Optimising Information Security Decision Making

Adam Beautement

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

April 2012

Abstract

The aim of the thesis is to investigate the relationship between human behaviour and effective security in order to develop tools and methods for supporting decision makers in the field of information security.

A review of the literature of information security, Human Computer Interaction (HCI), and the economics of security reveals that role of users in delivering effective security has largely been neglected. Security designers working without an understanding of the limitations of human cognition implement systems that, by their nature, offer perverse incentives to the user. The result is the adoption of insecure behaviour by the users in order to cope with the burdens placed upon them. Despite HCI identifying the need for increased usability in security, much of the research in the field of HCI Security (HCISec) still focuses on improving the usability of the interface to security systems, rather than the underlying system itself. In addition, while the impact of user non-compliance on the effectiveness of security has been demonstrated, most security design methods still rely on technical measures and controls to achieve their security aims.

In recent years the need to incorporate human factors into security decision making has been recognised but this process is not supported by appropriate tools or methodologies. The traditional CIA framework used to express security goals lacks the flexibility and granularity to support the analysis of the trade-offs that are taking place. The research gap is therefore not so much one of knowledge (for much of the required information does exist in the fields of security and HCI) but rather how to combine this knowledge to form an effective decision making framework.

This gap is addressed by combining the fields of security and HCI with economics in order to provide a utility-based approach that allows the effective balancing and management of human factors alongside more technical measures and controls. The need to consider human effort as a limited resource is shown by highlighting the negative consequences of neglecting this axis of resource measurement. This need is expressed through the Compliance Budget model which treats users as perceptive actors conducting a cost/benefit analysis when faced with compliance decisions. Through the use of the qualitative data analysis methodology *Grounded Theory*, a set of semi-structured interviews were analysed to provide the basis for this model. Passwords form a running example throughout the thesis. The need to provide decision makers with empirical data grounded in the real world is recognised and addressed through a combination of data gathering techniques. A laboratory study and a field trial were conducted to gather performance data with two password policies.

In order to make optimal use of this data, a unified approach to decision making is necessary. Alongside this, the usefulness of systems models as tools for simulation and analysis is recognised. An economically motivated framework is therefore presented that organises and expresses security goals with the methods required to fulfil them. The role of the user is fully represented in this framework which is structured in such a way as to allow a smooth transition from data gathering to systems modelling. This unified approach to optimising security decision making provides key insights into the requirements for making more effective real-world decisions in the field of information security and is a useful foundation for improving current practices in this area.

Acknowledgements

First, I would like to thank HP Labs and the ESPRC for providing the CASE studentship that funded this research. This studentship was attached the Trust Economics project that I was fortunate to be involved with throughout the duration of this project. Trust Economics provided an excellent environment in which to develop and explore the ideas presented in this thesis and I would like to express my gratitude to all its members. In particular I would like to thank David Pym for his support, insight and many excellent conversations. I would also like to thank HP Labs in general, and Martin Sadler and Simon Shiu, in particular, for offering me the opportunity to spend six months working at their facility as an intern. This was a very valuable experience which materially contributed to my research.

My supervisor Angela Sasse deserves special recognition for her unwavering support, belief and kindness over the last few years. It was Angela who encouraged me to start on this road and I will always be grateful that she did. Without her I would certainly never have reached this point and so to her go my most sincere thanks.

There have been numerous others that have offered their support and encouragement while I toiled at the workface and so I would also like to extend my thanks to Jude Beautement, Christine Wallace, Sacha Brostoff, Tom Wiggall, Giles Copp, Adam Shaw, Dougal Wallace and Andrew Shackleton. For a wide range of reasons this thesis would never have been completed without them.

Table of Contents

List of Figures

List of Tables

1. Introduction

1.1 The problem and thesis scope

The need to protect and secure valuable information has been understood for thousands of years. Julius Caesar is credited with the invention of the Caesar cipher in approximately 50 B.C., and over the intervening centuries, various methods have been used by countries and organisations to maintain their secrecy. The last century has seen many rapid advances in communications and computation, in particular the rise of the personal computer. Since its introduction into workplaces, companies have become increasingly reliant on computers for their day-to-day functions. Across large parts of the globe, millions of employees interact with a computer on a daily basis in order to communicate, share information, conduct business transactions of all kinds and to engage in social activities. This growth has afforded companies many additional opportunities to do business but new risks and threats, such as hacking and malicious software, have risen in parallel. These threats have driven the need for new methods of securing the information stored and transmitted in this digital space. The disciplines of information security and computer security have emerged as a result.

The aim of information security is to safeguard the confidentiality, integrity and availability of information seen as valuable to an organisation. In today's environment this practice is closely and inextricably linked to the use of IT systems. While there are academic disciplines that research the theoretical methods of information security, the prevalence of IT systems in the business environment also creates a requirement for the effective management of these systems. The relative short history of the information and computer security disciplines means that often these decision makers are working without firm scientific evidence or methods to support and enhance their understanding of the operational environment. In particular, there is a knowledge gap with respect to the behaviour and motivations of the users of these systems. New devices and software are released on to the market every year and the number of individuals using them continues to grow.

Nevertheless, many security decisions need to be made every day. Some are made by professionals, those responsible for shaping the security of entire companies. Many more are made by relatively untrained members of the public. These are most commonly employees in support of the security of the organisations they work for, or simply individuals, perhaps customers of the organisation, providing for their own personal security needs. The outcome of these decisions affects the security of us all but are they being made in an informed or even rational manner? Many security managers rely on their own experience and intuition when deciding which security option to take. In a world of reduced budgets and increased expectation this may no longer be enough. Without a systematic approach to security

decision-making managers may find it hard to justify their budget expenditure to a company financial officer. Worse, organisations or individuals without such expertise are left to make decisions based only on their own intuition and by following common 'best practice', or by following the advice of the security community and the vendors of security systems, who offer an ever-changing array of products. The inherent danger in this is that vendors are primarily motivated by product sales, rather than by delivering effective security, and so present themselves as experts. However they rarely have a solid empirical basis for describing the effectiveness of their products. Likewise, the lack of both consensus and a solid science of security in the research community can lead to contradictory or ineffective advice. Without a systematic approach to security the way is left open for regulatory pressures, pursuit of industry standards (which may be driven by vendors) and 'best practice' to govern security decision-making. This creates an environment that is far from optimally secure, but one that also contains opportunities for advancement and progress.

This is not to say that all or even most of current information security practice is in need of improvement. However, not all aspects of security are equally well developed. While the technical side of information security is relatively advanced, the so called 'human factor', and its impact on both security and productivity, is both less well understood and less studied. This is not because the importance of it has never been recognised. Auguste Kerckhoffs, one of the founding fathers of modern cryptography, wrote two journal articles in 1883 [45] that laid out six principles of secure communication. These principles are still referred to today. Of these six principles three refer to the usability of the system:

> *Principle 3: It must be easy to communicate and remember the keys without requiring written notes, it must also be easy to change or modify the keys with different participants;*
>
> *Principle 5: The system must be portable, and its use must not require more than one person;*
>
> *Principle 6: Finally, regarding the circumstances in which such system is applied; it must be easy to use and must neither require stress of mind nor the knowledge of a long series of rules.*

Over time, despite the application of his other principles, the necessity for usability in a secure system has fallen by the wayside. Over the past decade, however, the role of the human factor in security is starting, once again, to be appreciated. Information though is sparse on how to make best use of this understanding. Some practitioners have begun to apply economic thinking to technical security, to replace traditional military-influenced

thought processes such as the need for a centralised 'command and control' approach to security. Security researchers who understand and model security systems in this way have made significant progress [10, 13]. However, the work undertaken by such practitioners still abstracts away the role of the user in providing security. As a result, security policies are written which expect unrealistic feats of cognition and levels of effort from the people who are supposed to follow them, i.e., employees and consumers. Despite the advances in some of the specific usability issues in security, practitioners struggle to manage insecure behaviour and do not utilise knowledge of the human factor to improve security. This is the core problem area which this thesis will seek to address. The approach taken will be to utilise a cross-disciplinary method that combines economics, information security and HCI (Human/Computer Interaction). This train of thought leads to the following key research questions:

> **Question 1:** *What can we learn from economic theory that will move us toward solving the problems found in information security?*
>
> **Question 2:** *What promotes insecure behaviours among the users of IT systems and how can these behaviours be managed and minimised?*
>
> **Question 3:** *How can we maximise information security when the human factor is taken into account?*
>
> **Question 4:** *How can we include the human factor in a systematic approach to security decision making?*

The research presented in this thesis will primarily address a gap in the current state of the art – that being the provision of information on the full costs and benefits of security measures, in particular accounting for human factors. As such it may be useful for both security managers seeking to make more informed policy decisions, and for other academic researchers working on similar problems. It is intended to address the problems created by an insufficient understanding of the human factors associated with security implementation and management. While it may not be directly useful to many individual users, or those involved only in personal security, the results will yield an indirect benefit as security decision makers consider the wider impacts of their security policies. The purpose of seeking to combine HCI's understanding of user behaviour with computer security through an economic framework is to furnish information security decision makers, at a higher level, with a more viable model of human behaviour in the context of security. This will allow them to take the human factor into account when making security decisions. There will be a range of resultant benefits (such as increased usability, less wasted effort and greater

flexibility in the responses of the security system) passed down to the user indirectly as security-systems design begins to take in to account the associated user experience.

This research attempts to address the lack of metrics in security but the scale of the task exceeds that of a single PhD. The contribution does, however, provide a step towards this goal. The research presented here was conducted as part of the Trust Economics project [71, 36] and thus forms a part of a more comprehensive effort towards the goals of that project. The objectives of Trust Economics were:

1. *A study of the economics of information security policies, protocols, and investments. Our perspective is one of 'systems thinking' and, critically, our aim is to seek to integrate the following three perspectives:*
    a. *Modelling the behaviour of the users of systems, both internal and external (employees and customers), in the context of security policies and protocols;*
    b. *Mathematical modelling of systems, organizations, and networks, including the security policies and protocols which govern access;*
    c. *Economic modelling of the costs and value of security policies and protocols*
2. *Devise methodology that allows companies to make economically justified security decisions*
3. *Utilise Trust Economic modelling techniques within software tools. These tools can then be deployed for use within real-world environments.*

The results presented in this thesis should be taken in the context of the Trust Economics project and primarily address objectives 1 a), 1 c) and 2 of the project. Additionally, some of the key contributions were the result of collaborations with other members of the project, though I played a substantial and critical part in their development. The details of this research, and my specific contribution to it, are explained in Section 1.4. As this research was undertaken as part of a wider project it comprises several strands each of which offers insights in complementary but differing areas of the overall goals of the Trust Economics project.

This diversity has driven the need for this work to draw upon a variety of existing research fields including Information Security, Human Computer Interaction (HCI) and Economics. Each of these disciplines offers something to the research as a whole.

1.2 The research approach

The research strategy for this thesis is outlined in the diagram below. The nature of the problems being addressed necessitates a multi-disciplinary approach as no one field of expertise contains the information necessary for generating a solution. This thesis therefore follows three strands of research. Each of these differ in their approach to addressing the

same core problem - that of how to support information security decision making through understanding human factors - and focused on different but complementary aspects of the proposed solution.

The first strand of research investigated the causes behind user's unwillingness to comply with security policies. This involved analysing a series of interviews undertaken as part of the Trust Economics project, and extracting consistent themes and areas of cross over where participants made similar or related comments. From this analysis a theory, The Compliance Budget, emerged which drew upon the responses of the interview participants and combined them with key ideas from the field of economics. A second round of interviews, again conducted as part of the Trust Economics project, probed participants' attitudes toward security policy, allowing the theory to be tested with new data.

The second strand of research focused on gathering empirical data to support both the Compliance Budget, and the wider goal of the Trust Economics project to develop a more systematic approach to decision-making in security. The data gathering focused on passwords since these are a well understood security mechanism, and one that could be manipulated in experimental studies with relative ease. Additionally, the widespread and on-going use of passwords provided a real-world context for the experiments as well as a practical use for results found and recommendations made.

The password data were gathered using two primary systems. The first was an enhanced key logger that recorded keystroke timings within a controllable environment that allowed us to design experiments with some sophistication. Participants were required to repeatedly enter a variety of passwords under several different experimental conditions and their performance was recorded and analysed.

The second system was a server-based experiment management system that allowed participants to login and authenticate to the server when prompted by an email. Their authentication performance was again tracked. In contrast to the first lab-based experiment, the participants were required to authenticate with varying frequency over two two-week sessions. Facilitating remote logins also meant that the data collected were closer to the 'in the wild' performance of the participants; a point of key interest when seeking to make recommendations with regards to real-world policy.

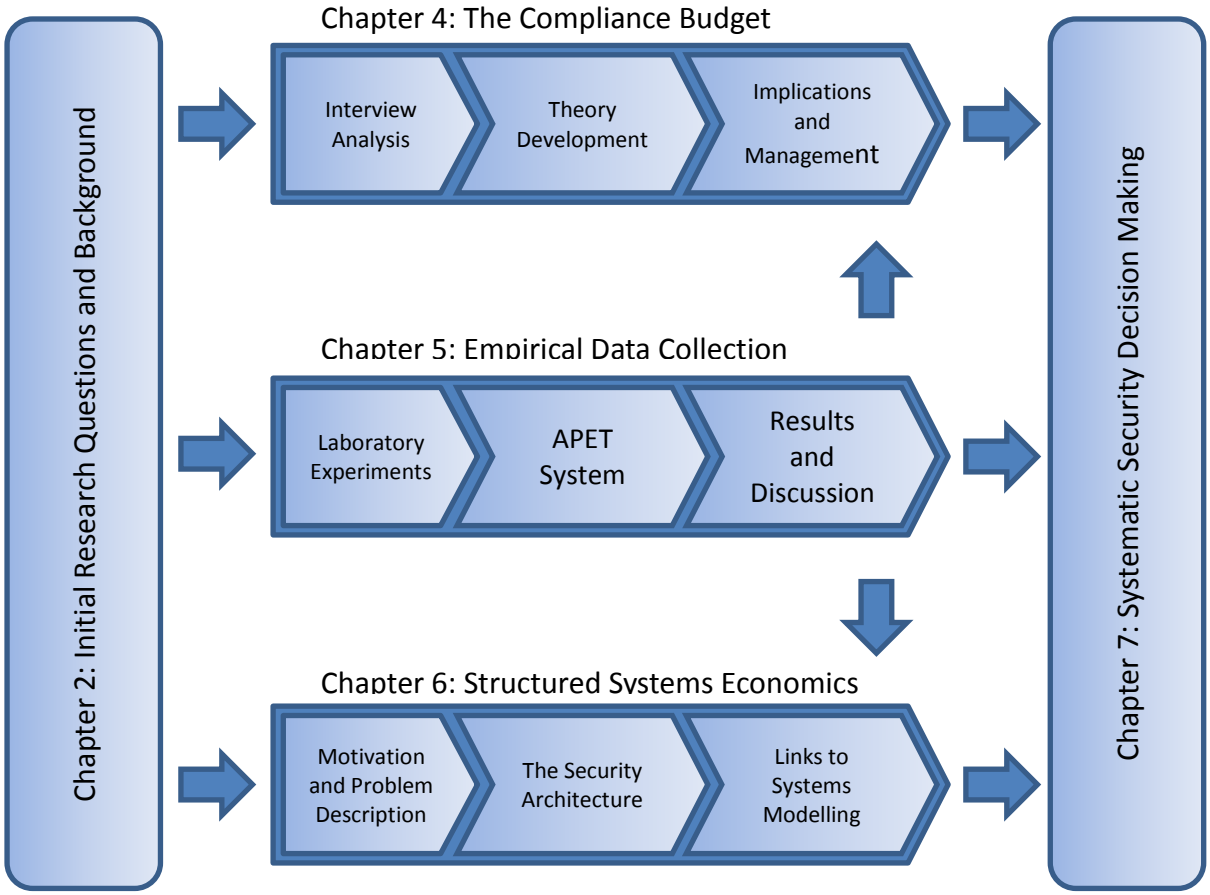

Figure 1: *Diagram outlining the primary research strands of this thesis*

Both the Compliance Budget and the security architecture and systems modelling strands depend on realistic empirical data to be of practical use to security decision makers. Therefore, a section on empirical data is included (see Section 5), and the results of that section feed back into the rest of the research.

The third research strand deals with the structuring of an organisations security goals and procedures. The aim here was twofold: first, to support the overarching goals of the Trust Economics project by providing knowledge to improve systematic security decision-making; second, to reorganise the information already available to organisations in a way that facilitates systems modelling. The approach here was to construct a logical multi-layer hierarchy describing the elements of a security system. The framework (see Section 6) was split into declarative and operational layers and populated with a variety of logical data constructs to meet the objectives of the research.

1.3 Thesis overview

The thesis is split into several chapters covering the three research strands outlined above. Chapter 2 reviews the key literature from the academic disciplines that provided the knowledge and methods on which this thesis builds: Computer Security and Security Objectives, Economics, Economics of Security and HCI. In each of the subsections, the relevant literature from one of these disciplines is covered with a critical commentary that illuminates its links to the thesis as well as to the other disciplines. The literature review takes a critical tone thus revealing the gaps in current knowledge that this thesis attempts to fill.

Chapter 3 reviews relevant methodologies for conducting empirical research on security performance and impact, and details those that are to be used to pursue the research aims. Both qualitative and quantitative methods are introduced including Grounded Theory, semi-structured interviews, observations and logging. The pros and cons of controlled and uncontrolled environments are explored and the role of modelling in decision making is considered.

Chapter 4 introduces the Compliance Budget. The supporting material for the Compliance Budget is laid out as background for the formation of the theory. The theory itself and its implications are laid out as well as some proposed applications and scenarios in which the Compliance Budget can be utilised. Additionally, the links between the Compliance Budget and subsequent publications drawing on its theme are highlighted.

Chapter 5 describes the empirical data collection process. The experimental designs for both sets of trials are covered here. In addition, the design and development of the proprietary software used in the experiment, although not a primary task for this thesis, is mentioned by way of background. Outlines of the results of the experiment are given here, as well as a discussion of how the results will be used in the wider context to which they relate. An outline of the broader range of experiments of which these form a part is also offered.

Chapter 6 covers a discussion on systems modelling and its use as a tool for systematic decision making. This section also brings to light the benefits of empirical data (of the sort gathered in Section 4) to systems modelling.

Chapter 7 details the development of the multi-layer hierarchy of roles. The motivation for its creation is discussed here and its function is related back to previous sections. The strong link between the hierarchy and the use of systems modelling is also made.

Chapter 8 draws together the findings of each of the chapters above to provide a set of conclusions and outcomes to the research.

1.4 A note on context and terminology

Considering all the different operational contexts of information security is beyond the scope of the thesis. In order to avoid confusion, and to allow the reader to focus on the concepts being discussed rather than trying to follow a changing narrative, a single context will be used throughout: that of a corporate business environment. Where it is useful to discuss other contexts it will be made clear in the text. The following terminology will have the following meanings when they are used in this thesis, unless explicitly stated otherwise:

> *Organisation:* used to describe a large (but not multinational) company operating in a corporate business environment, the purpose of which is to generate revenue for the company stakeholders. This company has no particular requirement for secrecy in its business operations, beyond that needed to remain competitive in the market place.
>
> *User:* refers to an employee of the organisation. This employee has no specific management or security responsibilities and can be considered to be a general office worker.
>
> *Primary Task:* refers to the business process of the organisation, and therefore the user. The primary task is therefore the productive process that generates revenue for the company.
>
> *Actor:* a specific piece of terminology used to denote the instantiation of a user (and other roles within the organisation) in the security hierarchy described in detail in Section 6.

1.5 Primary contributions

This thesis makes three main substantive contributions. For the parts where the research was carried out in collaboration with others, the author's specific contribution is identified. The primary contributions of this work are:

1. The development of a theory of user decision-making when faced with security compliance decisions. This theory, the Compliance Budget, describes the factors that a user is likely to take in to account when making a compliance decision and frames them in the form of an economic cost/benefit analysis. The notion of a *Compliance Threshold* is introduced, and the implications of breaching this threshold are discussed. This contributes to the fields of usability and security by drawing together research from both fields and expressing it a form that allows for an understanding of the user that goes beyond considering them as wilfully uncooperative or uneducated. By considering users as rational entities this theory prompts researchers to focus on the factors that contribute to insecure behaviour rather than ways to educate or train users to adopt better practices.

   This approach also offers security managers a better understanding the consequences of introducing security measures that burden users as part of their function. Techniques for the better management of user effort are discussed and advice suggested for the implementation of security policies that place more consideration on human factors.

2. A contribution to the body of empirical data required to produce metrics of user effort, which are required to support theories such as the Compliance Budget. While the wider aim of the data collection was also pursued by the Trust Economics project, the experiments detailed in this thesis are solely the work of the author. These experiments focus on user performance using passwords under various policy conditions and use frequencies. These data offers researchers and practitioners an example of how different password policies affect user effort. This serves to guide future research in this area; since users circumvent security policies and controls that require too much effort (as discussed in Section 2.2.1), we need metrics for the effort associated with security policies and mechanisms, not just for their security. The thesis makes an additional, methodological contribution by demonstrating methods of data collection that overcome some of the biases present in laboratory experiments, showing how researchers can effectively gather realistic data suitable for informing real-world decisions and strategies.

3. The development of a new framework for expressing security needs and requirements designed to allow a systematic approach to security analysis and the ordering of existing data in a form more readily usable by systems modellers. Researchers from the fields of systems modelling and information security, looking to collaborate, will benefit from this framework as it allows them to structure their information and requirements in a mutually supportive manner. The methods used in constructing this framework, while at the prototype stage, will be of use to security managers wishing to gain a more in-depth understanding of the relationships and dependencies inherent in security systems.

The research presented here has been published in three peer-reviewed conference papers, a conference workshop article and a technical report. These are listed below with details of the contributions I made to each.

1. **Modelling the Human and Technological Costs and Benefits of USB Memory Stick Security (WEIS 2008):** This paper was published as a collaborative work as part of the Trust Economics paper. My contribution was the analysis of a set of interview data. This analysis was the precursor to that used in the Compliance Budget (see below) and used the same data set. The results of my analysis were used in this paper to develop the scenario used in the systems models and also informed the behaviour of the users represented by the model. As such, this paper is an early example of combining empirical data with systems modelling. This approach is extended in Chapter 6 of this thesis.

2. **The Compliance Budget: Managing Security Behaviour in Organisations (NSPW 2008):** This work was published as a collaborative venture but was substantially my work. I conceived the original idea and introduced the notion of the decision as a cost/benefit analysis. I was also responsible for the primary analysis of the interview transcripts although I did not conduct the interviews. I identified and extracted relevant information from the interviews and inferred the behaviour of individuals (as represented graphically in the paper) from this information. I was also responsible for considering how an understanding of the Compliance Budget should affect the decision making of security managers. Some assistance was given in setting the Compliance Budget in the appropriate security context.

3. **Gathering Realistic Authentication Performance Data Through Field Trials (SOUPS 2010, Usable Security Experiment Reports Workshop):** This workshop paper outlines the APET system, a server-based tool for gathering performance data. While I did not write the APET tool myself I was principally involved in designing its requirements and acted as the supervisor of the project in which it was developed.

I also oversaw its implementation and use once the system went live and ran my own experiment on the system. As such, I contributed materially to its development and operation.

4. **Structured Systems Economics for Security Management (WEIS 2010):** This work was written during a six-month long industrial internship at the HP laboratories. I worked closely with a research partner throughout this time. I was responsible for the conceptual development of the framework and its components, the structuring of the framework and the inclusion of the OWL ontology language into the process. Following on from previous discussions regarding issues with the CIA objective set, I approached the problem methodically and set out to construct a conceptual framework that allowed expression of such security goals as a natural outcome of considering the requirements of the system. Assistance was received with formal logical expression of the concepts and the formalisation of the framework's relationship to systems modelling. The development and use of the running example used in the paper, that of baggage handling at an airport, was contributed to equally by myself and my research partner.

5. **Systematic Decision Making in Security Management Modelling Password Use and Support (HP Labs Technical Report 2011, now accepted as a paper at the QASA workshop, 2012):** This technical report was a second outcome from my internship at HP Laboratories. I took overall responsibility for collating the various sections of the report and writing the overall structure. In terms of the content I was responsible for designing the conceptual model and the scenario and structure of the systems model (though not its writing or execution). I also worked with David Pym to decide the content of the utility equations and with Simon Arnell to produce the results and analysis.

2. Background

Achieving information security is a large problem that crosses into many different areas of society. Most people regularly undertake security tasks, which do not necessarily involve advanced technology or take place in a westernised society. Security can be a tribal farmer ensuring his herd is protected from wild animals, a New York businessman ensuring that his wallet is tucked into a secure pocket and his car is locked, or a border guard maintaining the integrity of a national boundary. The protection of important assets is something that we are all concerned with. Information security deals with the security of data and, as more and more of the interactions of our society are handled digitally, many traditional security problems have become information security problems. For example, as banks move from branch banking to online banking the nature of their security requirements, and those of their customers, are changing. Now we have to be concerned with the risks of online banking and the security of our email accounts. Additionally, we are now burdened with a variety of new tasks such as remembering a wide variety of passwords. We must also learn to use an increasing number of technological devices and systems in contexts ranging from shopping for food to ordering a new passport.

The pervasiveness of technology means that information security now extends far beyond the remit of security and IT professionals. Gone are the days when only specialists were involved in delivering information security. Now, virtually everyone is a component in an information security system, whether they are an employee of organisations that provide security, or a consumer of the services offered. This means that individuals with a huge range of technical competencies are engaging with information security on a daily basis. In addition, systems are increasingly composed of actors with divergent interests. In an airport, for example, national and border security shares space with commercial flight operators, retail business in the airport terminals, and the security and passenger verification procedures of the airlines themselves. While all are concerned with the security of the aircraft using the terminal, the attention paid to profit and commercial opportunities varies widely between these stakeholders. Airlines know that passengers will choose other carriers if they feel the boarding process is too onerous or frustrating, yet they must follow certain procedures in order to meet regulatory standards. Even in this tightly-controlled security environment there is tension between business and security needs; in more relaxed corporate environments, with less obvious security requirements, this tension can be even higher. Where such divergent interests interact, incentives are becoming as important to dependability as technical design. Security failures are caused at least as often by bad incentives as by bad design [7]. Varian [73] observed that less money was being spent on anti-virus software than might be expected. At that time, viruses were enabling Denial of Service attacks rather than corrupting individual hard drives so consumers had little

incentive to spend $20 on anti-virus software to protect, for example, Microsoft's website. Furthermore, those individuals connecting insecure machines to the internet are not the ones bearing the full consequence of their actions.

Information security, and its core concepts of confidentiality, integrity and availability, also shares space with other highly relevant topics such as privacy and secrecy. While the lack of standardisation of terminology (discussed further in Section 2.1.2) makes defining such borders difficult, they clearly overlap in some way. Anderson [6] offers us this comparison:

- *Secrecy* is a technical term which refers to the effect of the mechanisms used to limit the number of principals who can access information, such as cryptography or computer access controls.
- *Confidentiality* (as a proxy for information security) involves an obligation to protect some other person's or organisation's secrets if you know them.
- *Privacy* is the ability and/or right to protect your personal information and extends to the ability and/or right to prevent invasions of your personal space. Privacy can extend to families but not to legal persons such as corporations.

Information security may be deployed in an attempt to ensure any one of the above such as a patient's right to privacy, the confidentiality of a bio-tech firm's intellectual property or the secrecy of a military force. To add further complications, we may also need to protect additional information associated with the data such as who spoke to whom or when a record was created or accessed. This is particularly true in the case of privacy where anonymity is required. Sweeney [70] found that 87% of the US population can be uniquely identified by gender, zip code and date of birth. This showed that simply removing the names from a set of records does not guarantee privacy. Thelma Arnold was identified from an 'anonymised' set of search queries collected by AOL by an analysis of her search history [53]. Information security then can be about protecting a wide variety of assets both physical and digital.

However, it would be a mistake to think of information security as a purely technical problem. With the rise of social engineering attacks, an adversary can gather valuable data simple by asking for it in the right way or under the right conditions. Phishing, the practice of falsifying websites in order to elicit secure information, has become one of the primary attack vectors of the last few years. This approach relies not on compromising servers through hacking or introducing malware into a target system but on simply asking a user for their information in a plausible way. As Mitnick [51] illustrates the likelihood of success can be increased by creating a sense of urgency or panic in the user and adopting a position of relative authority. No technical solution exists to defend against such an attack and we must look elsewhere for our security.

With all of these disparate goals and mechanisms, information security is no longer a problem of mathematics and technical solutions. Decision making in this complex environment creates a challenge for both individuals and organisations. To be able to conduct effective information security we have to be able to make the right decisions ourselves, and also understand the motivations behind the decision-making process of other stakeholders. To understand the role of perverse incentives, user vulnerabilities and other non-technical facets of security we must look to other fields. In recent years, economics, psychology and usability have all been profitably cross-linked with information security research. This thesis will draw upon all of these areas to tackle the information security problem it addresses. To do this it will be necessary to review these fields more closely, starting with the more traditional aspects of IT security.

2.1 CIA and security objectives

2.1.1 Overview

The primary objectives of information security are classically categorised as confidentiality, integrity and availability (hereafter referred to by the acronym 'CIA'). There are many varying definitions of CIA (investigated more thoroughly later in this section) although most of the variation is in the specifics. To begin our discussion, we take definitions of these words from *An Introduction to Computer Security: The NIST Handbook* [54]. These are:

> *Confidentiality:* A requirement that private or confidential information not be disclosed to unauthorized individuals.
>
> *Integrity:* NIST makes a distinction between data and system integrity:
>
> *Data Integrity:* a requirement that information and programs are changed only in a specified and authorized manner.
>
> *System Integrity:* a requirement that a system performs its intended function in an unimpaired manner, free from deliberate or inadvertent unauthorized manipulation of the system.
>
> *Availability:* A requirement intended to assure that systems work promptly and service is not denied to authorised users.

Most information security researchers believe that confidentiality and availability are linked through a relationship best categorised as a 'trade-off' (these are discussed in more detail in Section 2.3), in which as one property increases the other is necessarily decreased. Illustrating this through extremes; if we are concerned solely with confidentiality, then a valid strategy would be to bury our computers underground in a sealed concrete box with no network connections. The data contained on them, baring extreme circumstances, will remain confidential. However, it cannot be practically used either. In ensuring maximum confidentiality, we also reduce the availability of the data to zero. Conversely, making the data freely available and distributed through as many channels as possible guarantees we can makes best use of it wherever and whenever we please. Of course, so can anyone else and our information becomes public knowledge and manifestly not confidential. In between these two extremes exists a spectrum of states, perhaps favouring the one or the other, but still more balanced than our extremes. As we traverse this spectrum we will be trading-off confidentiality against availability at each stage with measures that improve one aspect cutting into the other; at least this is the current way of thinking.

However, in order to consider how confidentiality and availability interact, it is necessary to define what is meant by these terms. The literature contains many definitions associated with CIA and therein lies part of the problem. These are not only not standardised, but are expressed at a high level of abstraction, which does little to clearly delineate the boundaries between the conceptual spaces each component occupies. The effect of this is to leave gaps in the conceptual space in which CIA sits. Indeed, it is commonplace in the literature to find observations that CIA does not provide an adequate basis for practical, operational information security. The deficiencies in the framework have led many authors to suggest amendments, additions or adjustments to it, commonly adding such concepts as 'authentication', 'non-repudiation' and 'control'. One of the more developed examples of amendments to CIA is the 'Parkerian Hexad', created by Donn Parker, who in his account of security concepts [57] described the CIA framework as *"necessary but incomplete for modern use"*. To CIA he adds possession, authenticity and utility, defined by him as follows:

> *Possession* is the control over information, specifically preventing physical contact with data and preventing copying or unauthorized use of intellectual property.
> *Authenticity* is the correspondence to intended meaning e.g. avoiding fraud or nonsensical data entries.
> *Utility* is the usefulness of the information for specific purposes and aims to prevent against both conversion of the data to a less useful form and impenetrable coding of the information.

Parker loosely aligns possession with confidentiality, authenticity with integrity and utility with availability. With the addition of these new concepts, we can say that some improvements have been made to the CIA framework, but the Parkerian Hexad is by no means complete. Indeed, it introduces some new complications and errors. With the new concepts come more connections and relationships that must be examined. Within CIA there are three relationships, C-A, C-A and I-A. The Parkerian Hexad has fifteen, greatly increasing the complexity of our task. In addition, the Hexad brings in an example of a category error, something that plagues many attempts to improve on CIA. We will return to category errors shortly, for now let us focus on the issue with adding more concepts and relationships to CIA. Recall that our interest in it stems from wanting to understand the trade-offs between the various concepts.

While adding additional concepts to CIA can help us solve certain specific problems it does little to move us toward our real goal: a more systematic approach to understanding the relationships between security objectives. Indeed, by adding in more concepts, the situation becomes more, rather than less, confusing. To be able to understand trade-offs we need clear and concise definitions of the items we are examining. Even here we are thwarted as definitions of the core CIA concepts (let alone the wide range of possible additions) vary from author to author.

2.1.2 Definitions

Let us look at some additional definitions of CIA drawn from the literature.

Confidentiality:

- Protecting sensitive information from unauthorised disclosure or intelligible interception. [21].
- Is the concealment of information of resources.[14]
- Ensuring that information is accessible only to those authorised to have access. [40].
- The security goal that generates the requirement for protection from the intentional or accidental attempts to perform unauthorised data reads. Confidentiality covers data in storage, during processing, and in transit. [68].

Integrity:

- Safeguarding the accuracy and completeness of information and computer software [21].
- The security goal that generates the requirement for protection against either intentional or accidental attempts to violate data integrity (the property that data has

when it has not been altered in an unauthorised manner) or system integrity (the quality that a system has when it performs its intended function in an unimpaired manner, free from unauthorised manipulation [68].
- Refers to the trustworthiness of data or resources, and it is usually phrased in terms of preventing improper or unauthorized change.[14]

Availability:

- The degree to which a system, subsystem, or equipment is operable and in a committable state at the start of a mission, when the mission is called for at an unknown, *i.e.,* a random, time. Simply put, availability is the proportion of time a system is in a functioning condition [79].
- Refers to the ability to use the information or resource desired.[14]
- Ensuring that information and vital services are available to users when required. [21].
- The security goal that generates the requirement for protection against
    - Intentional or accidental attempts to (1) perform unauthorized deletion of data or (2) otherwise cause a denial of service or data
    - Unauthorized use of system resources [68].

There are differences here that substantially increase the levels of confusion surrounding these concepts. Let us turn to integrity, which NIST in their handbook say "*has been, and continues to be, the subject of much debate among computer security experts*". Bishop [14] states that integrity is concerned with the trustworthiness of the data, while Flechais [21] regards it as focusing on safeguarding accuracy and completeness. The differences in these definitions are sufficient to change the role of security in each case. Consider the following scenario: a fully authorised member of our organisation enters some information into the shared organisational network. No un-authorised users have gained access to the network and so we can trust that everything on there is as our colleagues meant it to be. However, due to some error the information isn't technically accurate. Bishop may consider that the integrity of the information has remained intact as no unwanted person has modified our data or inserted false data into the network; it remains trustworthy. Flechais, on the other hand, introduces accuracy into his definition of integrity thus making errors of this sort potentially the responsibility of the security department. The NIST definitions position CIA squarely as security goals and are highly data-focused. The NIST definition of availability also addresses unauthorised use of system resources. This is a different way of stating the objective. Rather than ensuring that resources are available, it aims to stop malicious users tying up resources. This definition does not, however, place problems such as insufficient provisioning under the aegis of the 'availability' objective.

Considering the question of definitions, and the meaning of CIA, is not just idle semantics. Kumar et al. [48] demonstrate, through rigorous modelling, the strategic need to consider confidentiality and availability differently. They point out that the countermeasures required to secure against availability losses are different to those needed to secure against confidentiality losses. Protection from availability losses can come in the form of firewalls and anti-virus software that guard against worms and malicious code, the agents of Denial of Service attacks. Cryptographic software and technologies such as virtual private networks are required to protect against the loss of confidential data. In addition, the costs associated with availability losses are not the same as those for confidentiality losses. When talking in general terms, the use of *confidentiality* and *availability* will suffice, but in order to develop a systematic account of a security system and its requirements such terms need to be rigorously defined, or an alternative approach found.

This variation in terms means that CIA as a framework is only suitable for discussing security in the most general terms. As the basis for a systematic approach to security with a view to making operational judgements it is entirely unsuitable. In attempting to move CIA toward a more functionally complete state many authors make category errors (see Section 2.1.3), rendering their additions more confusing than helpful. Clearly then, a new approach is needed to understanding trade-offs in security decision making.

2.1.3 Category Errors

A category error is a semantic or ontological error by which a thing is placed into a class in which is does not belong [80]. Bishop's definition of integrity illustrates this concept for us. From the excerpt given in Section 2.1.2 above the text goes on to state, *"Integrity includes data integrity (the content of the information) and origin integrity (the source of the data, often called* authentication*)."* Here we have a definition that includes both a property the data have (they are integral) and the means by which such a property may be ensured (some form of authentication). Now it may be the case that one of these is in the correct class but necessarily the other must not be as we cannot reasonably combine declarative and operational notions in the same definition. Declarative and operational must be treated differently and their combination here creates confusion when trying to understand what 'integrity' means a security objective.

Turning this same mode of thought toward Parker's Hexad, we find 'utility' included in among the security concepts. Utility is by definition an economic concept which, although it may be usefully applied in a security concept (and this thesis will argue for just that), does not belong in the same category as CIA.

2.1.4 Information security in practice

Despite these differing definitions, the primary focus of security professionals is still to protect assets valuable to the organisations for which they work. The decisions that are made and the policies that are put in place serve this purpose and, if considered in isolation, are often excellent technical solutions. Guided by the principals of CIA (in whatever particular form they adopt) technical solutions are deployed that offer functional security. The problem here is that the recommendations and advice offered by information security specialists too often *only* serves the purpose of security and neglects the context in which the measures will be implemented and the users that interact with the system.

Passwords are used through this thesis both as an example and a focus of the research. They also serve here to illustrate this point. Security practitioners and IT organisations encourage us to choose secure passwords and offer us advice to helps us achieve that aim. The following guidelines are taken from Microsoft's Safety and Security Centre website [50]:

- *"Whenever possible, use eight characters or more."*
- *"Don't use the same password for everything. Cybercriminals steal passwords on websites with very little security, and then they try to use that same password and user name in more secure environments, such as banking websites."*
- *"Change your passwords often. Set an automatic reminder for yourself to change your passwords on your email, banking, and credit card websites about every three months."*
- *"The greater the variety of characters in your password, the better. However, password hacking software automatically checks for common letter-to-symbol conversions, such as changing 'and' to '&' or 'to' to '2'."*
- *"Use the entire keyboard, not just the letters and characters you use or see most often."*

The suggestions, *"change passwords often"*, *"don't use the same password for everything"* and "*use eight characters or more*" are commonly repeated, in one form or another, by IT services that require passwords to be generated. However, the motivations behind such advice are often technical or based on out-dated concepts. The threats that a password conforming to the standards outlined above protects against, such as cracking and dictionary attacks, are not the predominant threats they once were and a lengthy and complicated password will not protect against a phishing attack.

The ability of information security professionals to protect against such threats is not in question. What is, however, is whether or not such measures are always the appropriate ones to take? In Section 2.2 I will introduce research from the field of HCI that conflicts with the above advice. This conflict between the advice offered by HCI and security, and the current lack of methods for resolving that conflict, lies at the core of the motivation and need for this research.

2.1.5 CIA and security objectives chapter summary

At the start of this research I began to examine the assumption that confidentiality and availability are connected by a trade-off relationship. The focus of my thought was on the following two questions:

1. Does increasing confidentiality *automatically and necessarily* reduce availability?
2. Are there are other similar relationships e.g. between confidentiality and integrity in the CIA framework?

Several difficulties immediately became apparent when approaching these questions. Initially the problem is with the definitions associated with CIA contained in the current literature. Additionally, as discussed above, CIA is an incomplete framework for expressing all possible security objectives. Let us consider a thought exercise to illustrate this point. Our organisation has suffered a data theft; a hacker has broken into our file system and extracted a copy of our most sensitive data. The hacker has not examined the data himself but is threatening to sell them to the highest bidder. In the aftermath, our security experts turn to the CIA framework to try to understand what has gone wrong. However, an interesting situation is revealed, reported as follows:

- Confidentiality has not been breached as the information remains secret; no-one outside of the organisation has read it.
- Integrity has not been breached as the information remains unchanged.
- Availability has not been breached as a copy of the information remains on our network and can be used and accessed at any time.

Clearly, a breach in security has occurred but the CIA framework is not sufficient to enable us to identify it. From this discussion it is possible to draw two main conclusions:

- CIA is an incomplete framework unsuitable for examining trade-offs in security decision making.
- Previous attempts to amend or add to the CIA framework have added rather than reduced confusion through category errors or inconsistencies in definitions.

These two conclusions each generated their own strand of research in this thesis. First, trade-offs were examined at a different level of abstraction. Rather than considering them at the level suggested by broad objectives, such as CIA, we consider the decision-making process of an individual user and the trade-offs they make when faced with security choices. This approach is covered in Sections 4 and 5 which provide the basis for a more bottom-up approach to tackling the question of trade-offs in security. In particular, I will return to the theme of the conflicting needs of security and HCI and approaches to a resolution.

Secondly, a hierarchical framework is created to express the components of a security system in an ontological fashion. This framework consciously separates the declarative and operational components of the system to avoid category errors; rather than attempting to trade-off between CIA directly, this allows trade-offs to be made between subcomponents in the system, the results of which reveal the system designers preferences within the CIA framework indirectly. This work is covered in Section 6.

2.2 HCI, productivity and security

2.2.1 Usability

Technological progress has long been considered one of the primary means of increasing productivity. The widespread adoption of the personal computer as a tool for business should therefore have led to spectacular advancements in individual effectiveness in the workplace. This, however, did not always turn out to be the case. The field of Human Computer Interaction (HCI) grew out of the need to understand why computers were not increasing worker efficiency in the service and manufacturing sectors by anything like the amount necessary to cover the cost of the adoption of the technology. Landauer [32] suggests that poor software or application design is at the root of the problem, because it is too often evaluated along technical lines (how fast or efficiently the application performs calculations for example), rather than to see how well it actually supports the individual's task and the business process. Several studies by Nielsen and Landauer in the mid 1990's showed that, even after just one cycle of empirical user-based evaluation and re-design, efficiency gains of 50% can be expected. The cost/benefit ratio of such an exercise can reach upward of 200:1 [34]. The problem then is not that it is impossible to design software that improves the efficiency of its users, but that such user-centric processes are not being utilised effectively in the software development process.

The requirement for systems to be usable was identified as early as 1983 by Norman [55], along with the need for systems to be designed with the user's mental models in mind.

Norman discusses the need for systems to be based on a coherent conceptual model and that this model is taught to the user. He gives three criteria that this model must have in order for the user to be able to effectively utilise the system:

- Learnability
- Functionality
- Usability

He goes on to say:

> *"What good is a conceptual model that is too difficult to learn? Or a model that has little functionality, failing to correspond to the system image or failing to predict or explain the important aspects of the target system? Or what of a conceptual model that cannot easily be used, given the properties of the human information processing structure with its limited short-term memory and limited ability to do computations?"*

Yet these failures are still present at the heart of many systems today, and HCI research has not always managed to address these concerns and improve system design and functionality as we shall see throughout this section.

The recognition of software usability as a problem has not been confined to the HCI community, or those producing commercial, business-orientated products. In 1999, Whitten and Tygar [78] evaluated the email encryption programme PGP 5.0 from a usability perspective. At the time PGP 5.0 was considered to have a relatively good user interface for a security application as it was the first version of PGP to ship with a graphical user interface, rather than the command line of previous versions. However, of the 12 participants in the study, only 4 managed to successfully send an encrypted email using PGP 5.0 in the 90 minutes allowed, and only 3 participants were able to decrypt an email in less than 10 minutes. Additionally, and perhaps more worrying from a security perspective, the majority of those who had failed to encrypt their email thought they had succeeded in the encryption task. The paper defines a system as usable for security if the people who use it:

- *Are reliably made aware of the security tasks they need to perform*
- *Are able to figure out how to successfully perform those tasks*
- *Don't make dangerous errors*
- *Are sufficiently comfortable with the interface to continue using it*

Following these guidelines the researchers concluded that "*PGP 5.0 is not usable enough to provide effective security for most computer users.*" A key problem identified in the study was that the visual metaphors used by the software in its interface did not correspond well to the mental metaphors the participants drew on. For example, even though signatures are created using keys in the protocol, the button to sign an email shows an image of a quill pen. This does relate to the act of writing a signature in general but is unrelated to the specific actions required in PGP – that of using a private key to generate a signature. This supports Norman's observations on usability.

While this study used a relatively small number of participants, which reduces the generalisability of the results (a common problem with lab-based usability studies), the experimental protocol did allow them a substantial amount of time in which to send the email, far in excess of what would be considered appropriate in the modern workplace. This adds weight to the conclusion that the system, as it was then, was unusable for practical purposes.

Sheng et al. returned to the topic of email encryption with a study in 2006 that evaluated PGP 9.0 [65]. Although only six participants took part – again too small to draw general conclusions from – it is worrying that none of them managed to successfully encrypt an email. The main problem was that *"no indication was given during message composition as to whether or not the outgoing date would be encrypted or signed".* No significant progress had been made in the usability design of PGP in the intervening seven years.

A lack of usability has implications beyond that of efficiency. In their 1999 paper *'Users are not the enemy'* Adams and Sasse [1] reported a study that looked into password use in two organisations. Rather than looking at a security mechanism in a lab context, they conducted interviews with more than 20 employees in each of the 2 organisations and then designed a questionnaire based on problems reported in these interviews. The questionnaire involved 139 participants, resulting in a data set of a more useful size. They set out to understand why employees of one of the organisations could not remember their passwords, a question prompted by the rising cost of provisioning help desks in that organisation. This makes the paper particularly relevant as examination of the security problem was pursued for economic reasons. The survey found that 50% of the respondents wrote their password down in one form or another. The main reason given for this was the requirement to use multiple passwords, combined with the fact that they were forced to change their password every month also being mentioned. In this case, users saw the password policy itself as the reason for their insecure behaviour. Recall from Section 2.1.4 that security practitioners explicitly advise that passwords not be shared between systems (creating the need for many passwords), and that they should be changed often. The security problem, that of employees

not being able to recall their passwords, was created by the practices and procedures suggested by the security community.

By shifting the responsibility for generating passwords from an automated service to the users, but not offering feedback on their choices, the security team allowed individuals to create their own rules for password creation that were often anything but secure. The cognitive demands exceeded users' capabilities; finding it impossible to comply, they created workarounds and security suffered as a result (this is discussed in more detail in Section 2.4.3). The security department noticed the user behaviour (workarounds), and formed the impression that users were 'inherently insecure', while users felt that the security mechanisms were a laborious and unnecessary overhead, and just designed to make their life difficult.

While this paper was the first to really gather concrete data on the real-world behaviour of users, it failed to communicate its findings to either the community or the users in a way that could affect change. Their task, in this sense, was more than challenging as the problem lies in the security domain and the analysis of the problem is to be found in the field of HCI. Adams and Sasse's work was a key step in bringing this tension to light, although, due to the lack of collaboration between the disciplines of HCI and security, little progress has since been made to resolve this inherent tension.

While the paper recognises that user behaviour is the product of a complex set of factors, and cannot be changed by simply threats of punishment, it did not to expose these factors. The problems resulting from security mechanisms being incompatible with work practices are properly expressed for the first time in this paper, but it did not provide solutions on which the security department or users could act. As such, its message went mostly unheard for over a decade.

Sasse and Weirich returned to this area in 2001 [77] and made progress in identifying the factors that influence user behaviour. Their study again involved gathering real-world data, this time through an interview process. With only 17 participants, the study was relatively small, but had a fairly heterogeneous set of participants: system administrators, research students and employees of a technology company. Through an exploratory interview process they found that most participants' internal mental view rendered them indifferent to attempts to persuade them to be more secure. Not only did the participants have a low fear of security threats, dismissing potential attackers as little more than technically smart kids rather than actual criminals, but participants also felt that they were not likely targets. This lack of fear makes it difficult to persuade them to see security as a priority. The study also discovered that security-conscious participants had a negative self-image, while there was a positive image of trust associated with people who shared passwords freely. These mental

constructs make enforcing security through fear-based strictures almost impossible and, while the more usable choice is to be insecure, this behaviour will continue to dominate. Not all the paper's recommendations are realistic and their effectiveness is not analysed here. However, the fact that his paper considers passwords (and by extension security mechanisms more generally) as part of a sociotechnical system, rather than just a technical one, is its main strength. It focuses on how to persuade users to be secure rather than force them through fear. This user-centric view supports the idea that security can be improved through working with the user rather than trying to control their behaviour.

Weirich's doctoral thesis [76] continues this line of research. He developed, through the analysis of interview data, a model of user choice in which users select their behaviours from a subset of all possible password-related behaviours allowed by the situation. This subset of available behaviours is defined by the actions the user is aware of and able to perform. As discussed above, many policies create demands that are outside the capabilities of human cognition, meaning that it may not be possible for a user to select a behaviour that is secure as there are nonesuch allowed by the situation. Weirich's model of user choice states that having made an estimate of the required security level of the system the user will choose 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.*

This model of choice is closely linked to the economic notions of utility and cost/benefit analysis that are explained further in Section 2.3 and form one of the themes of this research. Weirich's research did not explicitly link security, usability and economics, as this thesis does, but his insights provide a strong foundation for the need to recognise the relevance of user choice and the factors that contribute to it.

Even with the increased understanding of user behaviour and motivation offered by the research presented in this section, the problem of managing user behaviour remains an open one. Theresa Jones, a security manager at Dow Chemical is quoted by Johnson and Goetz [43] as saying, "*My biggest challenge is changing behaviour. If I could change the behaviour of our Dow workforce, then I think I've solved the problem".*

2.2.2 Problems with usability methodologies

The question remains why a more user-centric approach, as recommended by HCI researchers and practitioners, has not been applied to produce more usable security. Passwords, in particular, were identified as a problem area over a decade ago and researchers have developed a number of alternative, allegedly more usable, alternatives. However, but users still have to struggle with passwords today [22]. Part of the problem lies in the methods commonly used in usability studies, as can be illustrated through the history of Passfaces, a graphical authentication system.

Passfaces was studied by Valentine in 1998 [72] and the results seem to suggest it would be a viable replacement for passwords. Passfaces is a knowledge-based authentication system like passwords, but reduces the memory load by offering cued recall and recognition in the place of password's requirement for unaided recall. Users are show a series of four 3-x-3 grids containing headshots of people. To enrol users select one face from each grid and this sequence then allows them to log in. Valentine tested this system with 77 staff and students from Goldsmith College with participants being prompted to log in either every working day for two weeks, one week after enrolment or one month after enrolment. Users that logged in every day demonstrated 99.98% successful recall, and ever after a break of a month 84% of participants were able to successfully log in on. This is much higher than passwords which had only a 77% success rate after two weeks according to Zviran and Haga [82]. Valentine followed up this study in 1999 [73] when he contacted the participants after five months had elapsed and asked them to try again. An impressive 72% of those he contacted were able to log in after three attempts using Passfaces.

While this may seem to be an unqualified success Passfaces, was never widely adopted. In an attempt to find out why, Brostoff and Sasse conducted their own study into Passfaces in 2000 [15]. They point out that, in Valentine's Passfaces study, the system was used in isolation and participants were prompted to log in by the experimenter rather than being observed in the context of their normal activities. In addition, the password study they compared their results to was run with different participants over a different time-frame, which casts doubt on the viability of the comparison. Brostoff and Sasse conducted a within-subjects field trial over a period of three months with 34 student participants. The trial involved using both passwords and Passfaces to access a university system used for completing and submitting coursework assignments. In this way, the participants engaged with the task in its usual context (the first time this 'real-world' approach was taken in a usability study) and when they needed to use the system rather than when prompted to by an experimenter. The experiments tracked failed login attempts, and the reasons that logins failed. In keeping with Valentine's findings, students authenticating with Passfaces produced fewer errors. However, an analysis of use showed that students used the system

with Passfaces a third of the amount of the system with passwords. They also delayed an average of 4 days longer when handing in coursework. This is due to the fact that Passfaces log-in process took longer and felt more cumbersome. So, while Valentine's work showed that Passfaces were more memorable, it failed to demonstrate that they were more usable. This study also demonstrates the importance of the *primary task* to the user. Passfaces made executing the primary task – handing in coursework through the web portal – more onerous and so it was used far less.

The relationship between the primary task and security is expressed by Weirich and Sasse [77], who states that

> *"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 comers to reduce that extra load given a chance, unless they are motivated to make the effort to behave in a security-conscious fashion."*

Where this enabling task requires more effort than is necessary, the motivation to take short cuts will only be higher.

This conclusion is supported by Everitt et al. [19], whose 2009 study also moved away from the unrealistic assessment of single authentication credentials and examined the frequency of access to graphical passwords, and the interference resulting from utilising multiple graphical passwords. They recruited 110 participants who, as in Brostoff and Sasse, used Passfaces to control access to an online resource. Experimental conditions included using the same password throughout the study and the use of 4 different passwords through the five-week study period. The results of the study showed that participants using 4 different Passfaces passwords were 10 times more likely to completely fail to authenticate than those using a single Passface password once per week. In that most users make use of significantly more passwords than that each week, this study highlights a key problem with graphical passwords. This study is well designed, effectively choosing experimental conditions to explore interference, and with its large number of participants makes a strong case in support of Brostoff and Sasse.

These two studies demonstrate the need to conduct field trials which take place in the real task context to verify findings when research the usability of new systems. The group size of this study was too low to investigate some factors, such as whether requests for reminders differed between mechanisms, and the homogenous nature of the participants does not encourage generalisation of the result to more business-like environments.

2.2.3 Aligning security and productivity

Based on its older parent discipline, ergonomics, HCI has essentially two routes to improving performance on tasks undertaken within a system consisting of both technology and humans: fitting the task to the human, or fitting the human to the task. In the first approach, the tools (and in this context this often means software) used in the task are adjusted to more naturally fit with capabilities of the human operators, and the demands of the production processes they support. The alternative approach attempts to train or educate the user to effectively make use of the tools they are given. HCI considers that this second approach should be reserved for when it is not possible to design tools that fit human capabilities and processes (because the technology required does not exist, or is not affordable). A weakness of many usable security studies is that they look at the mechanisms in isolation without task context. Tests are also run under unrealistic conditions (such as enrolling and testing in same session), and do not consider the implications for security. These oversights create issues when aligning productivity with security, by creating pressure on the users of these systems.

A major weakness in current security research, as identified above, is that in most cases the role of the user is ignored, marginalised, or only ineffectually considered in simplistic terms. Security applications then would seem to be a natural focus for the HCI community, with its mandate for improving the usability of applications and the performance of human-technology systems. Many security implementations are presented to the user through software. However, the linking of HCI and security (HCIsec) has not been as productive as the above research might lead us to hope. This is because many HCISec researchers default to taking to the second approach and focus on educating or training users to be able to use security mechanisms more efficiently, or follow security advice more effectively. Despite studies such as *'Users are not the Enemy'* [1] the prevailing trend in HCI research (as indicated by the output of major conferences such as SOUPS and CHI) has been on fitting the user to the security process rather than designing security solutions that fit human capabilities and the needs of primary processes.

This trend exists in spite of previous research indicating that users are not fully to blame for their undesirable behaviours. Sasse et al. [61] recognised that user behaviour is a common cause of security failure, and that a key reason for this behaviour was security mechanisms placing demands on the user that exceed the natural limitations of human cognition. The paper highlights the fact the cognitive demands created by passwords and the policies proscribing how they should be used create impossible demands for human memory. Memory capacity is limited and decays over time, making items that are both regularly used and meaningful much easier to recall than those that are not. This makes frequency of use a highly relevant topic, and one that is again often ignored by usability studies.

Similar items also compete with each other, disrupting recall. It is interesting to compare these attributes with the common security advice for passwords which suggests that passwords should not be meaningful, should be not shared (requiring that many similar items by used), and should be changed often, thus cyclically removing the benefits of regular use. This results in a high level of recall failure, especially for systems that see only light use. Additionally, not aligning security with productivity – and instead leaving users to cope with conflicting security and production tasks – security managers are further compounding the problem. The paper argues that the solution to such issues will not come through blaming users, and that any attempts at education must take in to account the mind-set of the users targeted. They base their findings on a series of studies at British Telecom that included a questionnaire process in which 144 employees were asked to report the causes that led to password resets and their number of passwords at work, an analysis of six months of password reset logs, 17 in-depth interviews on password use on reasons for not following password policy, and an experiment involving a logging in to a web-based system. The study was the first of its kind to systematically study a population of real users. In line with the traditional HCI approach, the authors argue that educating users is not the way to deal with performance problems caused by demands that are incompatible with workings of human memory. A change in system design must be considered instead. This conclusion is reinforced by a discussion of the mismatch between password policies and the goals of users. Authentication is an enabling task that allows a more fundamental task to be completed, yet it requires a level of focus and accuracy more suitable for a primary task. This reinforces the view of security as being a burden and erodes security culture over time.

The paper also identifies areas of the user mind-set where user education could be effective - such as the belief that ordinary users are unlikely targets for attacks. The impact of socially accepted insecurity and informal work practises are also addressed. If education programmes are to take effect they first require a motivated workforce and should target such areas. While the recommendations this work makes are based on a solid real-world study they were not followed by the HCI community as a whole.

2007 saw the release of 'anti-phishing Phil', and online game that *"teaches users good habits to help them avoid phishing attacks"* [66]. While the aim is sound the underlying principle falls into the same trap as other approaches covered in section 2.3 – that being to consider user effort an unlimited resource. The attitude that motivates such research as the development of anti-phishing Phil is that users would be following security advice such as parsing URL's properly if only they were equipped with the correct knowledge. This goes hand in hand with the assumption that users have the time and effort to spare to both undertake the training and parse the URL's correctly. As we will see in the following sections this assumption is at odds with a more economically-based approach.

At the same time, other researchers were beginning to realise that educating users was not a viable solution. Working from an economic standpoint, Honeyman et al. [35] investigated the development of software improvements and fixes. They found that reliability and security errors were indistinguishable to a novice user. They also considered the cost of changing users from novices to experts and found it to be prohibitive stating that

> *"due to the high cost of information acquisition, users find it too costly to (i) determine which manufacturer's software component has caused a system failure and (ii) whether the system failure was caused by a reliability or security failure."*

Using an economic analysis designed to maximise profits they conclude that educating users is not a sound strategy from a business point of view.

2.2.4 A rational rejection of security

A new paradigm in security was thus starting to emerge. Researchers such as Microsoft's Cormac Herley picked up on the notion of rational users and in 2009 he published a paper [34] that posited that users rejecting security advice was not in fact a result of a lack of education or motivation but a predictable and rational response to the advice itself. Users are at many points presented with advice, guides and mandates as to how to protect themselves and their hardware. This paper examines the advice given about password creation, security certificates and phishing and in each case concludes that the money that would be spent (converted from time spent via a salary calculation) on security if the advice was followed accurately would be orders of magnitude greater than the expected losses that these measures seek to address. Thus, Herley argues, user tendencies to ignore, bypass or shortcut security procedures and advice is a reflection of misplaced and inappropriate security, rather than an indication that users are lazy and insufficiently educated.

Once we begin to examine security from this perspective it is not difficult to find instances where the systems in place have clearly been designed without regard to user behaviour, particularly in the field of passwords. Herley, in collaboration with Florencio, followed up his work with a study in 2010 of password policies of various online services [23]. He found that there was no correlation between the value of the assets being protected and the required strength of the password. There was also no correlation between the strength of the policy and the level of risk, or the number of reported security breaches. Instead services that were optional to the users, for example consumer services such as online retailers, had weaker policies than those for which users had no choice but to use, such as government services. They suggested the reason for this was that only those services with a guaranteed user base – ones where there was no alternative to the service – could afford to burden their

users with rigid policies. The vast majority of breaches reported in this study were a result of some other technical error. This demonstrates that in many cases user effort is being spent in a wasteful and unnecessary fashion.

2.2.5 Security and HCI today

So what has changed in the ten or more years since the HCI community has been pointing out issues with usability and recommending alternative solutions? Inglesant and Sasse presented in 2010 [39] a study which re-examined password practices in the workplace. 32 participants in a variety of job roles in two organisations kept password diaries for a week and were interviewed at the end to discuss the details. The results were all too familiar. The participants showed a willingness to be secure but the restrictions of the policies of the organisations they worked for, and the lack of alignment with their working practises led them to adopt insecure behaviours as coping mechanisms. Users felt burdened in particular by the need to constantly change their passwords forcing them to generate new passwords that were both strong and differed significantly from previous ones. In the organisation that required regular passwords changes over half the participants reported that they wrote the password down as an aid to memory – despite all the research to the contrary policies are still being set that place an excessive cognitive burden on the user.

It is clear that HCI cannot address the problem of usable security on its own. Undermined by misleading experimental results, and too focused on the problem of ease of use over functionality and applicability, HCI can offer only a partial solution to the problem. It also struggles to communicate to the security community the relevant information it does contain. Critically, the information produced by the HCI community does not in practice assist security managers to improve decision making by considering factors such productivity and usability. It has shown that significant problems exist, and explained some of the human factors associated with those problems, but has not offered a way of generating solutions that balance these issues with information security requirements. Too often, the choice appears to be usability OR security. The challenge then is to take the relevant knowledge from HCI and utilise it as part of a more unified approach that can deliver a productive form of information security that supports the business process without unnecessarily burdening the user. This statement of course implicitly acknowledges the need to *necessarily* burden the user, not remove all effort from them.

2.2.6 HCI, productivity and security chapter summary

The research gap that this thesis addresses is not simply a lack of knowledge. As this section shows there is a substantial body of literature that effectively examines the usability issues associated with security mechanisms by considering their use in context. Likewise, the security community has a strong understanding of the policies and technical mechanisms security managers can utilise to mitigate information security risks. As we have seen with passwords these two communities offer conflicting advice on how to approach information security. The central problem is then one of decision making. When is it optimal to pursue a technical solution? When should policies be relaxed in order to reduce the burden placed on the user? A new approach is needed to allow security managers to effectively make decisions by drawing on the existing knowledge in the fields of both security and HCI. Recent research has focused on combining economics with security in order to support decision making. In Sections 2.3 and 2.4 I will introduce the concepts of economics, and the economics of information security, and the many benefits such an approach can yield in terms of a systematic analysis of security policy and procedure.

2.3 Economics

A full review of economics is outside the scope of this thesis, however, as a discipline it contains several key concepts that are of use to us as security professionals. Section 2.1.1 introduced to us the notion of utility in the discussion on the Parkerian Hexad. A fuller discussion of the growing field of research relating security and economics will be given in section 2.3. This section will concern itself solely with giving an overview of the economic concepts used throughout this work.

Economics has two major branches [12]:

- *Macroeconomics* emphasises interactions in the economy as a whole. It deliberately simplifies the individual building blocks of the analysis in order to retain a manageable analysis of the complete interaction of the economy.
- *Microeconomics* offers a detail treatment of individual decisions about particular commodities.

Microeconomics also deals specifically with limited resource scenarios giving it a strong parallel with the environment in which security managers make their decisions. Thus it is the field of microeconomics that we will largely be focusing on. Within this field there are four areas of particular interest to us. These are trade-offs, utility, indifference curves and cost/benefit analysis. Each of these will be covered in its own section below.

2.3.1 Trade-offs

A trade-off occurs when increasing the amount of one good available decreases the amount of another. In economics trade-offs are more usually referred to as *opportunity costs*. An opportunity cost is the quantity of other goods that must be sacrificed to get another unit of the desired good [12]. While in a well-functioning economy the monetary cost and opportunity cost of a good will be closely linked this is not always the case [9]. The opportunity cost can be considered the true cost of a good as it looks past the simple monetary value of an item. Let us consider the cost of sending a student to University. Establishing the monetary cost would simply involve adding up all the bills and tuition fees they incurred over the years, but this does not tell the whole story. Calculating the opportunity cost would also involve estimating the money they would have earned if they had spent the time working instead and adding that to the total. Hence the relevant cost of any decision is the value of the next best alternative that the decision forces one to give up. It can then be said that a trade-off exists between the two alternatives as selecting one ensures that the other must be given up. If this relationship does not exist then the options should not impact each other's opportunity costs. Likewise it cannot be said that a trade-off exists between them.

In the field of information security such trade-offs exists in conceptual areas (as we have seen with confidentiality and availability) as well as in the decisions that have to be made at various stages of implementation. For example, (given a fixed budget) electing to employ more security guards in an airport will decrease the amount of automated equipment available to be used in the security process. Thus it can be said that a trade-off exists between manned and automated security stations. The trade-off between usability and security will also be explored in the course of this research.

2.3.2 Utility

Utility refers to the overall level of satisfaction an entity has as a result of consuming goods and services. It is assumed that a consumer will spend their resources in order to maximise their utility. While utility itself cannot be directly measured it is possible to construct a viable theory of consumer choice by framing the value of a purchase in terms of a second good. While it is not possibly to derive the exact utility a consumer gains from watching a film it is possible to ask them what they would give up in return for the movie ticket. This can either be done in purely monetary terms or with reference to other commodities. The totally utility of a good is the largest sum of money that the consumer will voluntarily give up in exchange for the good. However, we could also ask our consumer how many tickets to the theatre they would give up in order to obtain a cinema ticket. This then begins to give us an idea of their purchasing preferences. If they prefer plays to films they will be reluctant to

swap theatre tickets for the cinema ticket. However, if they are not a fan of plays but cannot wait to see the latest Darren Aranofsky film then they will happily exchange several theatre tickets for a ticket to the cinema.

As well as utility theory being based on the assumption that consumers will wish to maximise their utility it also includes the idea that the more of a good a consumer has the less utility is gained from the acquisition of further items of that good. The increase in utility the consumer gains by adding one more of a good to their purchases is called the *marginal utility* of that good.

When making trade-off decisions, a consumer will certainly consider the utility they will gain from each option. However, the changing value of the marginal utility of goods means that the rate at which options of goods are traded-off is not a constant value. To understand this we need to look at indifference curves.

2.3.3 Indifference Curves

Indifference curves are a way of representing consumer tastes. They show all the consumption options that yield a certain utility (see Figure 2 below). Thus at all points along the curve consumer utility is a constant. The shape of the curves are dictated by the nature of trade-offs and diminishing marginal returns. The lines must slope downward because otherwise the consumer would be able to select more of one or both goods and be better off. Additionally, the flattening of the curve as it moves to the right (and in fact the increased verticality of the line at the left) is a result of *marginal utility* and shows that additional items of good at high levels of that good are worth less.

Indifference curves are generated by asking a consumer to choose between various bundles of goods. Continuing with our example we might ask our consumer whether he would prefer to have two cinema tickets or five theatre tickets. All points where the consumer values the two options equally and is *indifferent* to the choice placed before them can be added to the indifference curve.

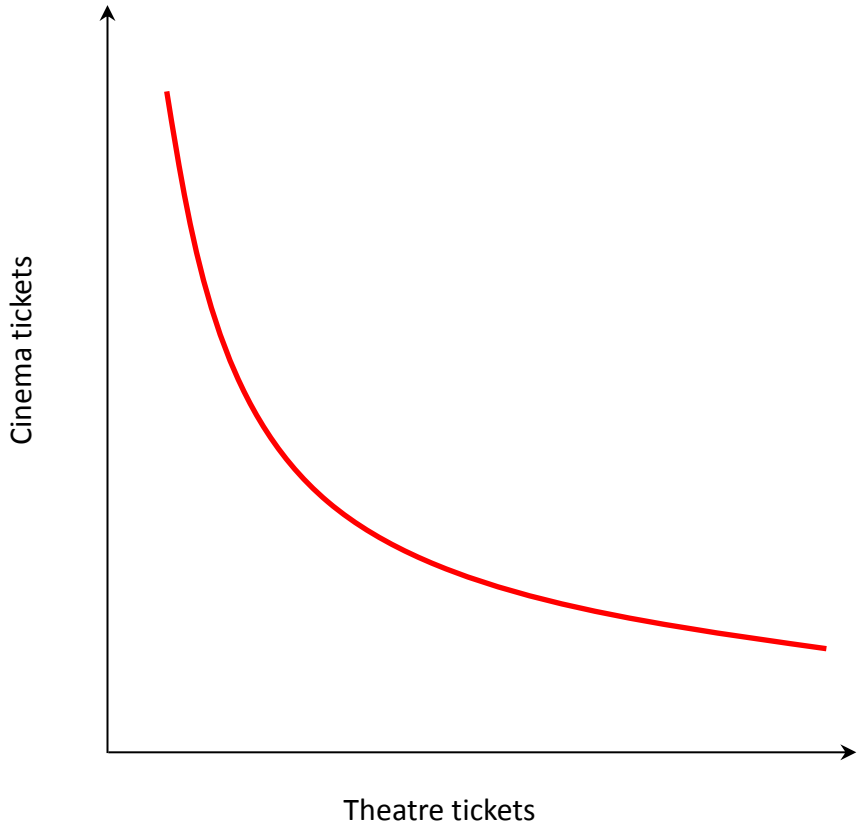

Figure 2: *Sample indifference curve showing consumer preference between theatre and cinema tickets*

Indifference curves reveal the rate at which consumers are prepared to trade-off between the goods in question. Transporting this into a security context allows us through a similar process to make explicit the security objectives of an organisation by examining the security measures they implement. The aim would be to expose the rate at which they would, for example, trade availability for confidentiality at various levels of each. However, to construct such a table in a security context we need to be able to talk accurately about what it is we are offering in these combinations. This requirement is addressed by the security hierarchy outlined in Chapter 6 which takes an ontological approach to developing security systems.

2.3.4 Cost/Benefit Analysis

A cost/benefit analysis (CBA) is simply a method of rationally making an economic decision. The process involves assessing the total costs and benefits inherent in one or more courses of action and deciding which to pursue. The more costs and benefits that are factored in to the decision the more accurate the final decision will be. For this reason when making a CBA appreciating the opportunity costs and trade-offs is highly relevant. The CBA process also highlights the importance of perspective when considering decision making as the decision

making entity can only factor in costs and benefits it is aware of. When making security decisions an organisation must therefore make the effort to be as fully informed as possible. Unfortunately, many organisations currently fail to account for the hidden costs associated with user behaviour when setting security policies, not realising that it detrimentally affects their performance. This research will reveal some of these costs and their potential impacts, thus increasing the effectiveness of security decision making.

2.4 Economics of security

2.4.1 Early beginnings

The economics of security is a recent but growing field of research. Pioneers in the field such as Gordon and Loeb [27], and Anderson [7] saw the potential of combining the two disciplines. Their primary reason was to bring the rational analysis of economics into the field of security to facilitate a financial analysis of security systems. No security manager has an unlimited budget (or often even a large one), so it is advantageous to be able to optimally deploy security measures within this budget. To do that an understanding of the costs and benefits involved is necessary. Additionally, the theories and principles of economics have been validated through many years of application and refinement. By drawing parallels between economics and security we can utilise proven theories with confidence when making security decisions.

The difficulty with this approach, however, is that economics deals exclusively with financial concepts; loans, interest rates and the acquisition and expenditure of capital. Critically these are measureable and easily quantified pieces of information. Security decisions, on the other hand, are not so clear cut, dealing as they do with intangible and difficult to quantify concepts. Information security seeks to protect assets whose material value can be very hard to calculate, such as the reputation of a company, or the potential value of a new piece of intellectual property. A financial trader who loses client data through a security breach will not be trusted, by both existing and potential clients, and will find it hard to engage new ones. Quantifying such a reputation loss is difficult at best. On top of this effective security means that breach events do *not* happen. Again how can we place a value on that which never happened? So with so little certainty available is not a fool's errand to attempt to apply economic theories to the field of security?

Despite these difficulties, in their book *Managing Cybersecurity Resources* [27] Gordon and Loeb argue that economic theories can and should be applied to security management. They discuss cost/benefit analysis and various financial decision making tools, such as Net Present Value (NPV) and Internal Rate of Return (IRR). The latter is of special interest

because the calculated result is compared against the organisations cost of capital to determine whether a purchase should be made. Cost of capital is the risk adjusted projected return on the capital investment. One of the major obstacles to applying economic theory to security is that it is difficult to calculate the benefits gained by introducing security measures. This is because the only way to gather data on the cost of the security breaches you hope to prevent is draw it from similar organisations that have had security failures. Many organisations do not report security failures, and those that do often downplay their severity. However, if we can factor some measure of risk into the decision-making process then we can make up for some of this data shortfall by using measures of risk which are more readily available. This cost of capital value can be scaled to take into account particularly high risk projects or where the cost savings are highly uncertain. Not only can we take steps to enable us to make at least good estimates of such quantities but we must also consider what our ultimate goal here is. We can avoid some of the uncertainty by tracking breach attempts and estimating the financial loss if they had occurred. Likewise we can review our business process and make an informed estimate of the likely loss of revenue through a reduction in our professional reputation. Most important though is to consider why we are interested in using economics in the first place.

2.4.2 A more systematic approach

In the past information security managers have often been recruited from the military or law enforcement. They relied heavily on their own judgement and experience to make security decisions and had less need than today to justify such decisions. However, things have changed in the last few years in two key ways:

1. The requirement to justify budget spending has increased in most organisations and intuition and personal judgement are rarely considered to be suitable evidence in a corporate context.
2. Not all organisations have access to a dedicated security manager, yet with the increase in cybercrime all organisations must deploy security measures in some form or another.

These two points highlight the need for a way of judging the effectiveness of security spending that does not solely rely on the expertise of an individual. Regurgitating 'best practise' without recourse to some system to measure the success of such practises is unlikely to provide many organisations with the security they need to safely operate within the modern threat environment. In short, a rational metric is needed to support systematic decision making.

With this in mind we can see that knowledge and tools that move security practice away from personal judgement and toward a systematic framework for decision making can be considered an improvement. Thus fully accurate financial figures are not necessary (although they remain desirable) and well informed estimates can still offer an improvement to the decision making process. Obviously with more accurate figures more effective decisions will be made, but less accurate figures do not undermine the usefulness of the methodology.

Some of the earlier works to analyse the link between security and economics did so in an attempt to find explanations for observed security behaviours that seemed to go against the perceived wisdom of the time. It is a common view that increasing the number of technical measures or the strictness of a policy results in better security. In 2001 Anderson [4] voiced a contrary view: security problems are at least as much the fault of perverse incentives and can often be explained clearly and convincingly through the language of microeconomics. For example, even though US banks spend less money on security to prevent fraud they spend it more effectively than European banks [3]. This is largely due a difference in liability. In the US banks are required to prove the customer wrong if they dispute a transaction. In Europe the burden of proof lies with the customer, an almost impossible task for them. As a result banks in Britain, Norway and the Netherlands became careless resulting in an epidemic of fraud [3]. The security behaviour of customers, attackers and the banks was influenced not by the technical systems in place but by the regulatory system surrounding them.

The negative impact of technical controls was also highlighted in Section 2.2.1. The requirement for a high number of passwords and frequent password changes at British Telecom resulted in the need for a costly IT helpdesk to cope with the volume of resets. Framing this problem in economic terms would have allowed them understand the costs of their policy and make a more effective set of decisions as a result.

In 2002 the Workshop on Economics of Information Security was held for the first time. Its call for papers stated that *"the confluence between information security and economics is of growing importance."* In that year Bruce Schneier [64] echoed Anderson and used economic thinking to explain why organisations and software vendors do not spend a lot of money on security, saying that *"the costs are significant: time, expense, reduced functionality, frustrated end users. The costs of ignoring security and getting hacked are small."* Anderson [5] added to his review of bank security indicating that despite issuing warnings against online use of credit cards in the 1990's they "*can see the transaction charges greatly exceed the fraud … and prefer to spend their energies quietly lobbying to have the risks of electronic fraud passed on to others*."

Gordon and Loeb also made the case for the need to take an economic approach to security in 2002 [25]. They point out that while there is a large body of work that focuses on the function and deployment of the technical defences utilised by the information security profession there existed little research to assist practitioners in deriving the optimum amount to invest in security. Without this information security is weakened, either through the deployment of insufficient security, or through too much, leading to overburdened users. The consequence of this latter situation has been discussed in Section 2.2.1.

These early insights showed that organisations practising security were not motivated by security for its own sake but rather by economic incentives to keep their costs down and their profits up. Where these incentives were not aligned with providing security (even when it would seem sensible to do so) the organisations were happy to minimise their security implementations.

2.4.3 Decision making and modelling

While this early research worked looked to economics to explain the behaviour of organisations practising information security later research began to look at ways of using economics techniques to improve security decision making. The motivation for this was a perceived lack of models and data to support security professionals. In 2002 Gordon and Loeb [26] investigated optimal levels of security investment. They tackled the question of whether or not an upper limit for security investments existed. In their paper they suggest that if such a level exists it would be approximately 36% of the expected loss. This value was based on a model of a firm contemplating the provision of additional security. The use of models to probe questions of security is an increasing trend in the field. Gordon and Loeb's findings were challenged in 2006 by Willemson [81] who demonstrated that it was possible to construct examples in which the required level of investment was as much as 50% of the expected loss. They also showed that by relaxing some of the original paper's requirements (such as making it possible to completely eliminate a threat through investment, something Gordon and Loeb did not allow) examples requiring investment of almost 100% of the expected loss were possible. The values manipulated in these models, such as investment levels, vulnerability and breaches, link financial input with a security outcome, bringing the two disciplines closer together.

Crucially, however, Willemson notes that *"There is a definite lack of general, reliable and rigorous models one could use in order to make such decisions."* So while the previous research has demonstrated the theoretical usefulness of such models they are not yet in a state where they can be used to support real-world decision making as "*in order to obtain any real results some concrete function must be plugged in.*" The production of data to allow the creation of such concrete functions is still a key task to tackle in this field.

Also in 2006, Cavusoglu et al. [17] looked at the economics of security patch management and came to a similar conclusion. Their paper showed that it is not always wise (deploying untested patches can open new vulnerabilities even as some are closed) or possible (150 vulnerabilities and fixes are released each week) to deploy every patch as it is released. Additionally, the downtime associated with patch deployment is expensive so a patching strategy is needed to optimise the process. The paper states that *"quantitative models are needed to help firms determine the optimal patch management strategy."* They go on to model the patching process for a single firm and piece of software. Damage is incurred through delayed patch release from the vendor or delayed patch deployment by the client. Patching has a fixed cost due to downtime. They then use game theory to determine the socially optimum (i.e. best for both the vendor and client) patching strategy. They conclude that this optimum strategy involves synchronising release and update cycles. The paper also discusses using cost sharing and liability to promote this synchronisation, concluding that vendors should share either the burden or the damage but not both in order to reach the social optimum. This work reinforces the earlier comments of Anderson and demonstrates that legislation not just technology can promote strong security. It also demonstrates the usefulness of well-grounded models for supporting security decision making.

Ioannidis et al. [41] also explore patch deployment through modelling, looking at the optimal patching strategies for two example types of organisation, which they characterise as *military* and *financial.* Interestingly they distinguish between these organisations by way of their confidentiality and availability preferences. This links back to the discussion presented in Section 2.1.2. They conclude that the optimal frequency for patch implementation is dependent on the attributes of the system that are preferred by the system's information security managers. Their results reflect the need to take in to account the security objectives of an organisation, and that optimal security decision making will therefore be situation dependent. They also state that *"Operationalising our approach will require more detailed models that more closely reflect the nature of specific organisations."* This highlights the research gap which the security hierarchy presented in Chapter 6 addresses, as well as the need to include the user in such models, without which the nature of an organisation cannot be closely reflected.

An increasing number of models employing economic approaches were by now being produced in the field, such as Hausken's work [30] on attacking and defending complex networks. Investment strategies are again considered here and utility measures are included as part of the evaluation metric. The Compliance Budget [10] is relevant to this theme and its contribution is discussed in full in Chapter 4.

When creating any model it is necessary to introduce abstractions and generalisations into the system to reasonably constrain the task. It is often the case in this field that the threat environment is abstracted away into a set of probability distributions controlling how often various attacks take place. However, if the attacker is modelled in more detail some interesting results can result. Huang at al. [37] did just that when considering security investment in the case of simultaneous attacks. The threat environment in this work is more developed than most and includes both distributed and directed attacks. They also consider investment in the context of both the vulnerability and value of the asset in question. They discover a minimum threshold of vulnerability below which they advise no investment at all. This is interesting because it is strongly counter-intuitive for the optimal security to be no security at all, yet – when considering the deployment of limited resources through an economic lens – Huang et al. prove that this is the case.

A second key point of interest in this paper is that they also consider how an organisation should shift their investment to defend against distributed or targeted attacks as the available budget changes. Specifically, they find that at low budget levels firms are better off protecting against either directed or distributed attacks but not both. Allocation shifts very fast at low budget levels too because a small portion of a low budget is ineffective at preventing attacks. This implies a threshold or 'critical mass of investment' below which it is not prudent to invest. The power of such models to locate and make explicit such thresholds is one of their strongest features. If as a decision maker we are aware of conditions under which investment is either worthless or particularly valuable we can maximise the gain from our resources and move towards a more efficient security solution.

More recent work by Baldwin et al. [24] has demonstrated the power of economically-framed systems models to change how security professionals make decisions. They combine a case study with their work on modelling to achieve these results. 12 security professionals were involved in the study. They were presented with a problem involving making improvements to the vulnerability and threat management for client infrastructure in a fictitious organisation. 4 decision options were presented and the participants were split in to 2 groups. Only one of these groups was guided to use the economic methodology, which was based on utility, and the systems model. This group produced justifications for their decisions that focused more on factors such as cost and productivity than the group that did not use the decision support tools. The study concludes that *"by externally representing alternative aspects of the problem and trade-offs between different factors our method can thus support decision making."* This paper represents a significant step forward in validating the use of economics of security and systems models as valuable methods for optimising security decision making.

2.4.4 Users: the missing link

So far, we have identified research that has demonstrated the use of economic modelling in security in supporting decision making. The requirement for solid data to ground the models has been noted, as well as an awareness of the fact that more detailed models yield superior predictions (albeit at the cost of increased complexity and difficulty). However, there is a key area that until recently was not well represented within the field of the economics of information security; that of user behaviour. While some papers do make reference to users it is most often in passing, and without the time and effort applied to either the attacker or the technological side of security. Kumar et al. [47] include users in their measures of cost, saying *"there are monetary costs of implementing both policy-based and product-based security countermeasures. In addition, firms need to consider the impact of such measures on end users."* While this demonstrates an awareness of the problem they go on to assume that *"a higher security level is achieved by stacking together security countermeasures."* This assumption does not take into account the effect on user productivity of having to cope with all these security systems and is based on the dangerous and misleading notion that user effort is inexhaustible – the issue of usability is not considered at all. Farahman et al. [20] also discuss the impacts of a computer breach stating that *"The dollar losses will depend on the type of computer affected and its relative losses. Security breaches pose various types of losses: (1) lost productivity (2) lost revenue, (3) clean-up costs and (4) financial performance impact, to name a few."* These loss categories were used by Herath and Herath [33], however, the loss of productivity in their model is still based solely on computer breaches. Lost productivity due to users being unable to fulfil their primary task is not considered, let alone the usability impact of task interruption. The significance of the primary task was discussed in Section 2.2.2 and is well documented in HCI literature. Its absence here highlights the need to integrate knowledge from these fields in order to make progress within the problem area. They do, however, attempt to cost more fully the types of losses due to breaches by considering the type of computer affected (web server, administrator's computer etc.). This is analogous to considering the actual role of the user involved in a security incident. This viewpoint is distinctly technical, however, and once more fails to fully consider the impact on the different users behind the machines. If a high-profile account is compromised and the information it contains deleted, or otherwise rendered inaccessible, not only is the data lost but so is the time of the user that produced it. As their time is likely to be valued at a higher rate than average (as reflected by their salary) this leads to an increased cost in terms of the time they will need to spend resolving the problem. Failing to consider the user leads to consistently underestimating the financial cost of breaches.

Johnson and Dynes [44] found *that "many of the large recent security breaches were not technical break-ins, but rather inadvertent disclosures."* The implication here is that users are not following security procedures properly, or that there aren't the right procedures in place.

Either way the security is compromised not due to malicious intent but due to absent minded or distracted users. In a previous paper Johnson noted that *"it is often not clear what security initiatives offer firms the greatest improvement"* [42]. This highlights both the lack of understanding available to security professionals on the role of the user in security and also the need for systematic metrics to evaluate the implementation choices available. Combining economics and security has proved effective in advancing our ability to evaluate the current state of affairs, but clearly there are still large gaps in the state of the art where decision making for the future is concerned. In particular the economics of security has failed to adequately include the human factor in security. This aspect will be covered further in Section 2.4.

This omission of the user is a trend that can be found in other works in the field, such as Chen at al. [18]. They support and extend the key contributions of the field with research that provides a system that links managerial decision making with technical implementations. The goal is to allow effective economic analysis through estimating security benefits. Their thinking supports the notion that increased systematic decision making support for security managers is necessary. They point out that without statistical data and tools security managers are left making decisions based on experience and judgement alone. However, they do leave the gap in that their system does not include a notion of user effort rendering their analysis incomplete and again highlighting the need for research in this area.

The application of economic concepts to information security has yielded many useful new results. In particular, this mode of analysis has revealed the perverse incentives responsible for many unwanted security behaviours. Additionally, the usefulness of models in understanding the possible outcomes of policy decisions has been demonstrated in many works. However, there are significant gaps that limit the usefulness of these results. These are:

- The lack of credible data on which to base the models being created.
- The abstraction, marginalisation or ignoring of the user when considering the costs and impacts of security

As the costs and benefits of different user behaviours have not been effectively considered in this field, the first point applies strongly to future attempts to model or quantify user data. Despite the strengths inherent in the field of economics of information security without models that include a notion of the user well-grounded in empirical data the full potential of the discipline will not be realised.

A similar line of reasoning has been developed by Pallas [56]. Like the research presented in this thesis his approach is based on the belief that a greater understanding of information security can be developed on the basis of established concepts from the field of economics. He identifies two strands which combine to create effective information security:

- Co-ordination – the identification of the optimal state of member behaviour.
- Motivation – the enforcement of the co-ordination outcome.

Where Pallas' contribution stands out is in his placing of the organisation *members* rather than technologies at the forefront of his model. It is important to note that the optimal state for information security is derived from the behaviour of these members, not from the technologies or systems implemented by the organisation. This approach ties in strongly with our findings that security and productivity (See section 2.2) are inextricably linked with user effort.

Pallas also discusses the '*meta-measures'* that influence human behaviour. He lists these measures as *architectural means, formal rule* and *informal rules.* He points out that architectural means, such as access control mechanisms, can only be circumvented by dedicated effort and their strictness leaves little room for trade-offs. By contrast formal and informal rules, being more flexible, are subject to trade-offs with other processes and can be ignored or bypassed. His key insight with respect to these rules is to consider the costs and motivations associated with them. Pallas informs us that while architectural means can be expensive to initially set up and have a moderate cost associated with them their rigidity means that they automatically enforce their rules and have no case-by-case cost. On the other hand formal and informal rules are enforced in a distributed fashion by the community the rules exist in and are dependent on information to be effective. That is to say if the chance of detection when breaking the rule is low, and/or the harshness of the sanction is not sufficiently high then the rules will be ineffective at enforcement. Given this it would seem that formal and informal rules are a poor choice for enforcement. Pallas explains that due to information asymmetries between the organisation and that arose with the advent of the personal computer (it became unrealistic for the centrally organised organisation to monitor the computing activities of all its personnel once such activities moved away from a mainframe) the use of such rules is actually an efficient delegation of the decision making process to those with more information about the activities – the users themselves. This work again highlights that the problem is one of decision making and ensuring that the correct information is available to those making decisions.

Economically speaking, architectural means also appear more expensive as they have high initial costs compared to the implementation of rules. This also contributes to the increasing use of rules for enforcement. However, as Pallas points out, the ultimate cost of enforcement

(due to needing to be enforced on a case-by-case basis and requiring larger investment to change the community culture if adjustments are required) is higher with formal and informal rules compared to architectural means. His discussion in this area also reinforces the need to fully cost the impact of decision making by considering costs that extend beyond the purely financial.

While Pallas does support the major themes of this research his is, by its own description, an *"abstract and theory-founded"* model. Linking such models to empirical data is a key contribution of this research.

2.5 Background: Chapter summary

The relationship between economics and security has allowed security professionals to gain many new insights into their discipline. It has offered a new perspective on evaluating and planning security. It has also revealed some unwanted results of security measures brought about by perverse incentives. However, the research focus of the economics of security is still largely on the costs and benefits of technical measures. Budgets are still largely considered in purely financial terms, and the role of user behaviour in security systems is largely seen as a hindrance to be overcome – preferably by ensuring compliance with existing policies. A recent trend in research that adds a more psychological approach to the mix of economics and security is yielding some interesting and valuable results. However, despite having identified that system design rarely takes the needs and capabilities of the user into sufficient account HCI and HCIsec is too focused on improving the user's interface with the system. While the underlying system remains unchanged this has the effect of trying to fit the user to the system. This is the picture revealed by a comprehensive review of the existing literature and is summarised in Figure 3 below.

This research contributes to, and extends, this recent trend of addressing system design directly. As stated above what is lacking in the state of the art is not necessarily the knowledge to do this. Each of the three fields covered above, information security, HCI and economics, contains valuable information for a security decision make, but in isolation they all offer an incomplete picture that struggles to achieve the goal of usable security. Information security experts know how to identify which assets the system should protect and the appropriate technical controls to do so. However, they lack knowledge of how these controls impact users and therefore how to ensure minimal disruption of the user's primary task. HCI researchers understand that the reasons for users' failures are not due to carelessness or idleness but because the demands of the systems are beyond human cognitive capabilities. However, they cannot link this to the systemic requirements and so are limited to making superficial changes i.e. through UI design. Economic theory and the

economics of security effectively models the incentives that drive decision making but does not know how to account for the role of the user.

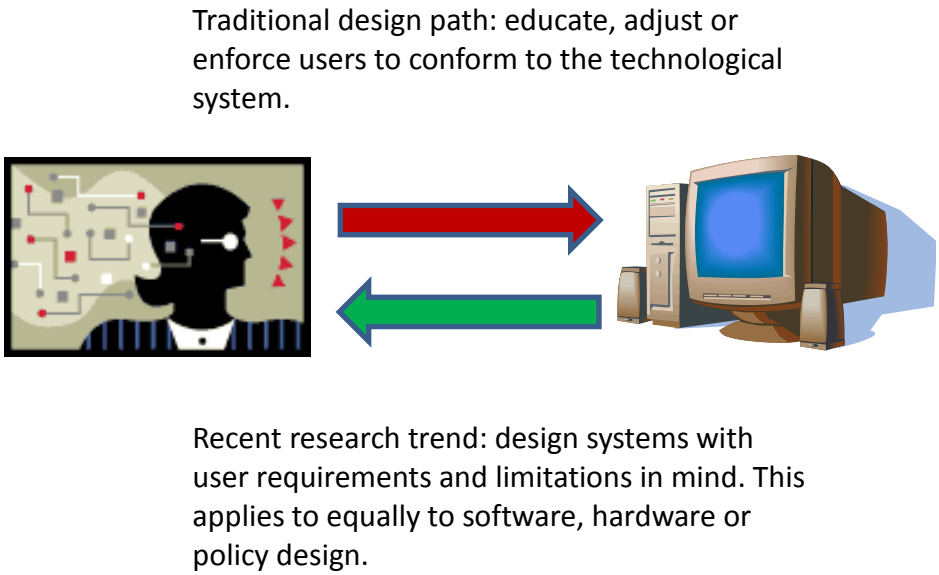

Figure 3: *Concept diagram for design priorities with respect to users and the technological systems they use*

The programme for this research is therefore to address the problem of how to effectively combine the expertise of these areas in order to develop a systematic approach that allows increasingly optimal decision making. The overall approach is to use economic analysis to understand the optimal trade-off between the needs of the user and the technical security requirements of the system and this concept is pursued in the following sections.

3. Methodology

The research disciplines of HCI, Security and Economics conduct experimental studies in different ways, using a variety of methodologies that support their aims. In research, data is collected to provide evidence to help decide research questions. As such, the research questions should determine both the type of data required, and the most suitable method for collecting them. Experimental design involves making choices between measurement and observation and the use of a controlled environment or gathering directly from the field. One must also consider the purpose of the data; if it is to be used to analyse functionality or usability, or support decision making through modelling and simulation. HCI evaluations, for example, can be *summative* or *formative.* In summative evaluations performance data is sufficient, as it will tell you which of two mechanisms is superior. A formative evaluation may seek to understand the reasons why a mechanism is not working, and how a solution may be generated. Here additional, explanatory, data would be required. These considerations will be discussed in this section. As this thesis is concerned with the collection of data on the usability, cost and organisational performance of security mechanisms, a review of those methodologies is necessary.

3.1 Laboratory experiments

Laboratory experiments form a significant fraction of the data gathering processes found in the HCI and HCIsec communities. Laboratory experiments allow the researcher to maintain control over all the variables present in the experiment. Experimental conditions can be enforced through manipulating experimental variables and monitoring control variables. These experiments commonly take place in a laboratory or similar environment, such as in [66], or are experimenter-led, as in Valentine's Passface study [72]. This means that the events that take place are scripted by the experimenter and, typically, participants are required to follow their instructions closely. The frequency and content of the experimental stimuli encountered by participants are determined by the needs of the experiment, rather than by the usual motivations of the users.

Practically speaking, controlled experiments require a significant commitment of time from both the participants and the experimenter, as both must be in the same location and actively involved in the task. The participant will often have to travel to reach the experimental location as well. These restrictions necessarily limit the experiment in a number of ways:

1. The number of participants is usually low, as each must be attended in person by the experimenter.

2. The duration of the experiment is limited – while participants may be willing to return to the experiment for a follow-up session it is unrealistic to expect this to happen over a period of weeks or months.
3. The recruitment process is usually local, resulting in a homogenous participant population.

These characteristics are demonstrated in the experiments reviewed in Section 2.2, in which there are usually fewer than 40 participants, and these are often drawn from the student population of the university at which the experiment is being run. This hampers the generalisability of the experimental results.

The advantage of controlled experiments is that they offer the experimenter the chance to orchestrate more closely the events of the experiment, and to interact in more detail with the participant. The data collected is typically quantitative and objective, for example whether or not a password was entered correctly, and can be analysed using statistical methods and tests. This makes controlled experiments a good choice when seeking to test performance on a task, where the factors that affect performance can be excluded or controlled. Often, however, the results of such experiments are generalised to predict real-world behaviours, and this can be a mistake. The small numbers of participants with heterogeneous backgrounds, and the artificial nature of the tasks (absence of a primary task, or any real security risk), mean that these results do not provide a suitable data set from which to generalise.

In this thesis, a laboratory experiment was chosen to examine performance with two different password policies. Although this work intends to support real-world decision making, and laboratory experiments can yield data that is disconnected from this operational context, such an experiment was necessary in this case. I was able to systematically alter the policy criteria and provide a stable environment, without which meaningful comparison would be difficult. The aim was to identify which elements of passwords construction caused users most effort. The experimental design and results for this study can be found in section 5.2.

3.2 Field Studies

Field studies allow the experimenter access to more naturalistic sources of data. However, there are many more variables outside of the experimenter's control potentially leading to a degree of noise in the data. Uncontrolled, in this sense, does not mean random. Unlike a laboratory, where outside influences are minimised, the experimenter accepts the presence of external factors in order to gather data in a more realistic context. In such studies the experimenter does not necessarily have to sit with each participant, making it easier to

involve a large number of participants. While experimenter intervention can take place, the events about which data is collected are not entirely orchestrated by the experimenter; their frequency, content and context are more informed by the usual activities of the participants. A software beta test is an example of a field trial in which the experimenter (in this case the designer of the software) can make modifications to the experimental environment (such as releasing a new build of the software) but the manner in which the testers interact with the software is left open. In some cases the participants in a field trial may not be aware the trial is taking place, such as when the provider of an online service like Google introduces new features and tracks user response. The collection of such data logs is usually done by a software or machine agent and can be reviewed at a later date. Information such as the number of password resets requested by the employees of an organisation can be accurately gathered through an automated logging process. Herley's study of web password habits [22] is based on the analysis of data gathered through automatic reporting, rather than through interaction with any specific users. Very large data sets can be gathered in this way from a potentially global population. Herley's large-scale study of web passwords collected data from 544960 clients via an optional component of the Windows Live toolbar. The automisation of data collection in this way extends the possible size of the data set far beyond anything that requires experimenter involvement. The downside of this type of trial is that explanatory data is absent, the *why* of user behaviour or intention is missing. Such explanatory data is critical to this research and so a trial of this sort would not be suitable.

An online survey is an example of field trial which does gather explanatory data. Such surveys need to be designed such that participants find it difficult to try to record the 'correct' answer, rather than responding truthfully. There is no way for the experimenter to guarantee the participants in such studies are responding accurately, or truthfully, but the potential rewards of gathering either much larger data sets, and/or more realistic data, can make such approaches worthwhile. Sasse and Inglesant used a 'password diary' method in their 2010 study on password habits [38]. These were filled in every day for a week by the participants, a level of engagement that would be virtually impossible in a laboratory engagement. This approach is also interesting because it collects data that is both quantitative and subjective, as only events that the participants recognise and remember to record will be included. Adams and Sasse [1] also used a survey methodology in their research and were able to recruit 139 participants, far more than would be feasible with a laboratory-based approach. A more recent, and as yet unpublished, survey I constructed for a recent project had over 1600 responses further illustrating the possible gains with this method.

Field studies are also more conducive to the running of experiments that involve repeated measurements, where data gathering is required over multiple days, or experimental sessions. This is rarely feasible in the case of laboratory experiments, in which the

participants must return to a specific location for each session. Field studies can run for days, or weeks, with far less intrusion into the daily routine of both participants and experimenters. In order to complement the laboratory experiment that I used to examine password performance, with respect to different policies, a server-based field study was used to test password performance over time and with different frequencies of use. Participants interacted with the server from their normal work context, making this data more directly grounded in the real world. I selected this method in order to develop a rounded picture of password performance that would be more useful in informing security decision making than a laboratory study conducted in isolation. The experimental design for this study can be found in Section 5.3.

While certain research questions can be answered solely with quantitative data, there is often relevant information that cannot be accurately measured. User perceptions and attitudes are a particularly germane example. In order to gather such information a qualitative approach is required. Although time-intensive, semi-structured interviews allow the experimenter to probe deeper into user motivations and thought processes. Subjective data such as this is highly necessary when examining decision making, as it is an individual's perception of their environment that they will rely on, rather than some more objective measure. Von Helverson et al. [31] reported that "*subjective but not objective effort has an effect on judgement formation"*. Thus subjective measures must be included, if an accurate model of user behaviour and decision making is to be developed. A semi-structured interview does not have a set question order but rather attempts to guide the participant into discussing the desired topics by using open ended questions. This method can be used as a primary source of data gathering, as in Sasse and Weirich, or as a follow-up to a more controlled study. Structured Interviews were used as the basis of significant portions of this work in order to answer research question 2, *"What promotes insecure behaviour among users"*. The motivations behind the adoption of insecure behaviour cannot be determined through remote observation or measurement. In order to answer this research question it was necessary to examine *why* users were following these behaviours, not just to observe that the behaviours were taking place and make assumptions as to the causes. As such, the structured interviews allowed me to gain access to the perspective of the users, a critical step in developing a unified view of the factors that a security decision maker must take in to account.

3.3 Data analysis and representation

It is not only important to collect the appropriate type of data to answer a research questions, but also to choose an appropriate method to analyse the data. Different methods are used to quantitative or qualitative, subjective or objective data.

Quantitative data is amenable to numerical analysis and so statistical methods are commonly used to analyse this form of data. Statistical tests can identify trends in the data, and the likely significance of these trends. Significance in this sense is a measure of how likely the observed event was to have occurred by chance as opposed to being the result of some causal effect. This value is particularly dependent on the sample size of the population.

The most common form of qualitative data is verbal data – transcripts from interviews or free-form answers collected in online surveys. The latter comes ready for analysis, whereas recorded interviews need to be transcribed first, which can be a time-consuming effort. Such data can be analysed in a quasi-numerical fashion, for instance by counting the number of participants mentioning a technology, problem or concept. However, to tease out deeper layers, such as the causes for reported behaviours, in-depth analysis methods are needed.

Grounded Theory [69] is a qualitative data analysis method widely used in social sciences, which allows identification of salient concepts and relationships between them. Over the past ten years, this method has been successfully applied to model user perceptions and attitudes in Human-Computer Interaction. Adams & Sasse [1] used this approach to identify factors that affect employees' perceptions of corporate security policies. Weirich & Sasse [77] used it to model employee decision-making around compliance with password security policies. In Grounded Theory, the data collected is marked with codes that indicate similar concepts. These can then be grouped into categories from which theories can be developed. In particular, Grounded Theory seeks to identify, through a process known as Axial Coding, why a particular event is occurring and how it is responded to. This is achieved by identifying various *conditions* that relate to these events. These can be:

1. Causal conditions: these directly influence the event.
2. Intervening conditions: which alter the impact of causal conditions. Unexpected events fall in to this category.
3. Contextual conditions: which are the products of causal events mitigated by intervening conditions.

Subsequently, the actions taken by individuals as a result of these conditions are identified, followed by the consequences of those actions. Grounded Theory was used in this thesis to analyses data gathered in interviews, and was instrumental in the formation of the Compliance Budget theory (see Chapter 4).

Various tools exist for representing category data, such as that generated by Grounded Theory analysis. Such representations are useful for sharing semantic data in a structured fashion through the creation of an ontology, which describes the concepts and relationships

relevant to a domain. The OWL 2 Web Ontology Language [52] is an open source ontology language that provides ontologies that include classes, properties, individuals and data values. Its open source nature makes it an available choice for use and the categories it contains serve equally well for representing Grounded Theory categories and the components of a systems model (see Section 3.5). These properties make it a suitable fit for use in this thesis. Additionally, there exists an open source editor, Protégé [67], which provides a usable interface for the construction of ontologies. These tools were used in section 6.4 to express the components of the security hierarchy developed in Section 6.

3.4 Modelling

The field of economics relies on the modelling of systems in order to determine the potential outcomes of a course of action. While no systems modelling has been undertaken as part of this thesis, neither the motivation for, nor the impact of, the research presented here can be understood fully without an understanding of systems modelling and its potential applications in the field of information security.

Systems modelling is a general term that is used across several disciplines (most commonly in the fields of business and information technology) to describe a process in which a system is represented, using some formal notation, (often mathematical but may be graphical or purely conceptual) in order to facilitate analysis or design.

Systems models are comprised of three principal components, *processes*, *resources* and *locations*. These can be defined as follows:

- *Processes* are the activities that occur in order to deliver the primary service of the system. In this context, the services delivered are security operations. Typically many such processes will execute, either simultaneously or sequentially, in order for the service to be delivered.
- *Resources* are the infrastructure of the system; the logical and/or physical entities that enable, and are manipulated by, processes in order to achieve their intended purpose.
- *Locations* are places in which events can take place within the model. Each resource must have a location and processes may require certain resources at certain locations to function. These locations tend to be highly distributed (logically and/or physically) and are linked by directed connections.

Each component of the system under analysis will be represented in the model as one of these types of conceptual objects. A fourth category, *Environment*, can be added to this.

Systems exist within external environments, from which events are incident upon the system's boundaries. Typically, the environment is insufficiently understood, and too complex, to be represented in the same explicit form as the system itself.

Of particular interest to this work is the use of such models as the basis for executable simulations. Various factors that affect the behaviour of the system are parameterised and, through adjusting these parameters, a range of behaviours can be observed through the simulation. This allows decision makers to predict the impact of decisions before they are made and thus select optimal strategies. Simulation modelling can be a powerful tool for reasoning about complex systems and exploring the decisions made within them. This has already been recognised within the information security community by work such as Beresnevichiene et al. [13]. This work recognises that "*the state of the art lacks a systematic methodology to support security investment decision making.*" The paper presents an approach that combines utility evaluation with mathematical systems modelling, a combination that has the potential to generate useful results with respect to managing trade-offs. Indeed, the study validated through a process of interviewing and collaborating with security experts, who perceived the process to be objective and rigorous. The model used in the case study represented the activities of IT staff responding to support-job requests. However, the behaviour of the IT staff was based on assumptions guided by the experience of the managers. This reflects the lack of available empirical knowledge regarding human behaviour available in the modelling community and therefore a research gap in the state of the art. The commercial possibilities have also been recognised by cutting edge technology companies such as HP, who now offer a Security Analytics service. This service uses model construction and exploration as a key component in supporting decision making. Supporting information security decision making through systematic methodologies is very much at the heart of the motivation behind this work and so it was natural to shape it with an eye on systems modelling. Additional work by Baldwin at al. [24] in this area added weight to these early findings and reported that, "*the system modelling tool was valued as a decision aid"* by the security professionals that participated in their study.

The involvement with systems models was present from the very beginning of this research, as well as the Trust Economics project. Understanding that a current weakness in the application of systems modelling was the lack of empirical data underpinning the models the Trust Economics project sought to address this. One of the results was a paper that reported on a model of the use of USB sticks, to which I contributed an analysis of interviews conducted within an information technology company and a financial services organisation. This same analysis went on to form the basis for the Compliance Budget. Understanding the need for better metrics in systems modelling was also a driver for the work presented in Section 6. The predictive power of such models is dependent on the quality of the data used to tune them. The principle of *garbage in, garbage out* is commonly

invoked in the fields of computer science, as well as in other fields such as engineering and information processing. It expresses the idea that bad inputs result in bad outputs, barring some specific design intervention [49]. *Problems of type* occur when the incorrect type of input is used. *Problems of quality* are the result of inputs of the correct type but which contain some defect. It is the second type of problem that can be avoided through improving the validity of the empirical data used in the systems model. As a consequence, decisions predicated on the results of the model can be made with more confidence.

The usefulness of such models is also linked to the level of abstraction from reality required to implement them. Certain assumptions must be made in the case of any systems model for it to be bounded and executed, within a reasonable time frame and budget. However, these assumptions must be chosen with care, likewise the transfer functions that govern the behaviour of the model.

Understanding the benefit of incorporating real-world data into a systems model does not always make that process straightforward. This problem becomes increasingly difficult when human factors are considered. Chapter 4 will demonstrate the need for a clear representation of the user in security planning and decision making, but this problem is typically abstracted away from in the more technical world of modelling. While this makes the creation of the model easier, it is not a viable long term solution if models are to remain a functional part of a new security paradigm, one that places an increased emphasis on the significance of the individual user. In order to address these issues I was motivated to consider new approaches to organising, and representing, information about security objectives and processes; approaches that would both facilitate systems modelling and incorporate a representation of the user alongside other, more technical, components of such systems. Chapter 6 details the exploration of this challenge.

3.5 A linked approach

There is a dichotomy at the heart of data collection that taps in to one of the key problems with progressing security decision-making. In order to be able to inform real-world decisions, data must be collected in conditions that are as close as possible to those found in the context in which the decisions are being made. However, in replicating this environment, or observing it directly, the experimenter runs the risk of gathering data in which any observed effects may be attributable to factors other than the ones of interest to the experimenter.

While HCI researchers favour controlled experiments, security practitioners prefer logging and economists rely on modelling and simulation, none of these methodologies on their own

offer the required tools to support efficient decision making. Just as theories and concepts from all three areas are needed, so it is with methods of data collection. Sasse and Brostoff demonstrated in 2000 [15] that field trials are necessary to understand the functionality and not just the usability of a security mechanism. The HCI literature review outlined the need for a user-centric conceptual model at the heart of security planning, and that requires qualitative data gathering methods such as interviews. Systems models are highly data dependent if they are to be usefully predictive and thus, where they seek to incorporate users, require collaboration with psychology and HCI.

In chapter 5this thesis lays out a range of approaches that attempt to overcome the inherent problems with both laboratory experiments and field studies, by combining the effective parts of both. Additionally, the use to which such data can be used to support decision making is discussed throughout. The focus is always on grounding any theory or model with a valid empirical data set. If a more systematic approach to security is be achieved, this combination of various data types, and their alignment with the appropriate conceptual models, is a mandatory prerequisite.

4. The Compliance Budget

4.1 Introduction and research background

The literature review, conducted in Section 2, concluded that information security mechanisms that negatively impact their users can create security weaknesses, or generate additional costs for the organisation they are trying to protect. While this problem has been recognised for over a decade there has been no tangible progress in reducing or minimising these negative impacts and moving toward a more usable state of security.

A key aim of this chapter is to propose a model that allows security managers to understand how the implementation of a security control leads to users adopting insecure behaviours. This will be done by demonstrating the link between the effort burden placed on the user by the security control and the user's decision making process regarding whether or not to comply with it. In this way the chapter will address research question 2, repeated here for convenience:

> *What promotes insecure behaviours among the users of IT systems and how can these behaviours be managed and minimised?*

Additionally, this chapter will shed light on the causes of insecure behaviour by showing that increasing the effort required by secondary tasks – such as security – beyond a user's tolerance or capacity for such tasks will drive them to adopt workarounds. These workarounds, which most often involve following an ad-hoc process decided upon by the user, are usually insecure. By highlighting the factors that go in to the decision to not comply and use a workaround the chapter will provide security managers with the required information to manage and influence such compliance decisions.

This chapter also contributes to answering research question 3:

> *How can we maximise information security when the human factor is taken into account?*

Without the co-operation of the users of a system the security of that system cannot be maximised. Therefore a discussion of how to manage compliance in order to maximise security will also be included. A utility-based economic approach is used throughout as a way of framing the problem in such a way that both business and security managers can access and understand the issue and its solutions.

In order to fulfil these aims it is necessary to understand the role of the user in information security. The practise of information security can be seen to be the practise of minimising security failures in the system being protected, but it must also serve the business needs of the organisation – if it costs more to implement and maintain than it saves in avoided breaches, then it is not of functional use. Where system failures are of a technical nature the tools exist to understand and remedy them with relative ease. However, it is widely acknowledged in security research and practise that many security incidents are caused by human rather than technical failures. In such situations the information security community is far less equipped to understand and solve the problem. Expertise does, however, exist in the field of HCI to understand the factors that underpin human error. These errors can have many causes but in the case of security the two most salient are difficulty and workload. Even technically-minded users with good IT skills can struggle to keep up with the additional complexity introduced into their working day by security mechanisms. Improving the UI and interface to such systems does not necessarily make them more usable, as was demonstrated with PGP 5.0 in Section 2.2.1. A more fundamental systemic change is required. While some failures can be attributed to security mechanisms whose execution places unrealistic cognitive demands on the user, this does not tell the whole story. Individual users have a choice about whether or not to comply with security policies and a choice to not comply may be indistinguishable from a failure event, thus weakening the system.

While the field of HCI has the knowledge to address this problem, it faces the difficulty of placing that knowledge within a framework that makes it relevant and available to business and security decision makers. Security decision makers have no systematic tools that allow them to incorporate information on user behaviour, limitations and requirements into their decision making process. This means there is a requirement for an approach that provides decision makers with the tools to balance the security and business processes with respect to the user with a view to optimising the overall utility of the organisation. Compliance is one area in which progress can be achieved. Non-compliance weakens a security system and implementing the monitoring and sanctions required to ensure compliance can be expensive for an organisation. However, not all instances of non-compliance are due to negligence. In some cases a non-compliant response is simply the logical choice due to the aforementioned cognitive demands of the system. It is necessary to understand the problem from a user perspective in order to eliminate such instances. This task is addressed in the remainder of this section.

When an organisation adds to or amends its security policy the opinion of the user population is rarely solicited, even though they have to deal with the workload and other implications of policy changes. Having identified from the comprehensive literature review in Chapter 2 that the role of the user in security planning is under-represented, the first stage

in understanding the impact of changes in security policy on user populations, and so on the effectiveness of an organisation at achieving its primary task, has to be gathering information on user opinions and attitudes. This is motivated by the understanding that effective security cannot be achieved without the consent and co-operation of the user. It has been pointed out that in order to achieve their security goals CISOs and other managers must consider that users both have a choice about whether or not to comply with security policy and that this choice is influenced by the individual's own perceptions and goals.

If the direction of security research is to be altered and a trend begun toward more user-centric design then it is first necessary to understand user's perception of and attitudes toward security. In particular an understanding must be gained of the decision making process of users when faced with security tasks. Implicit in this desire is the assumption that users are making decisions and are not blindly following policy either through coercion or a lack of awareness of any alternative.

As part of the Trust Economics project interviews were conducted with various employees of a major financial organisation and a technology research and development company. These interviews were designed to solicit user's views on how their working day would be affected by a change in security policy. During the interview they were specifically asked about mandatory encryption of USB sticks. When these interviews were analysed an interesting pattern emerged. Many of the participants stated that their reaction to the change in policy would not be to adjust their working practises to accommodate it, rather they would ignore it and/or seek a workaround. This indicated a more sophisticated decision making process was taking place. Considering research question 2, *'what promotes insecure behaviours among users?'* the Compliance Budget describes this process in a form that allows it to be manipulated so that compliance with policy can be increased. The result of this is a framework that supports security decision makers by providing them with guidelines to managing compliance through system design. The Compliance Budget seeks to express a new paradigm in understanding and managing user responses to security decision making. The core proposals of the Compliance Budget are:

1. That understanding the economics of compliance from an individual's point of view provides a better basis for influencing individual's security behaviour
2. That organisations should accept that the effort users are willing to spend on compliance with security policies is finite resource that needs to be carefully managed
3. That compliance decisions can be characterised as a rational cost/benefit analysis by the individual of the advantages and disadvantages of compliance

4.2 The empirical study

The empirical basis for this was model was a study run in two phases by the Trust Economics project to elicit factors that contribute to corporate and individual security cost and examine issues around understanding of and compliance with security policies. I was able to access this data through my association with Trust Economics, and chose to analyse it.

In the first phase 17 in-depth interviews were conducted with security staff, employees and managers in two major UK companies – a financial institution and the research lab of a technology company. The interviewees included the chief security officers in both companies, and 4 security researchers or operational security staff. The remaining 10 participants were staff working in the companies' main-line activities – financial services or technology research, and one of them was a lawyer. All participants had a university degree and at least two years' work experience.

The second phase of interviews was conducted on a much larger scale with over 140 participants. These were employees of a large multinational energy company who volunteered to take part in the study. They include members of a broad range of roles within the company from both the United Kingdom and the United States of America.

In both phases the interviews were semi-structured, exploring

- The tasks and responsibilities of interviewees,
- Their perception of the risks facing the company,
- Their attitudes to the company's security polices and security measures, and
- The perceived impact of security measures on individual's tasks and responsibilities, and company productivity.

While the interviews covered a range of security policies and measures, all interviewees were asked about their use of one specific bit of technology - USB sticks- and their awareness of security policies surrounding the use of these policies. My initial involvement with the phase 1 interview data was to support the work of Pym et al. during Trust Economics' modelling of USB use. The first pass of analysis focused on identifying and extracting responses dealing with USBs, specifically the costs and benefits to both the individual and the company associated with changing the security policy to mandate USB encryption. It was from the Grounded Theory analysis of this data that the pattern of responses leading to the Compliance Budget model emerged – participants showed a preference for workarounds rather than compliance in the face of the proposed policy

change. A second round of analysis was then conducted that focused on attitudes toward wider security issues, particularly with respect to compliance and decision making. The transcripts of the interviews were reviewed in their entirety, and analysed using techniques from Grounded Theory (see Section 3.4). The interviews using axial coding (the first stage of Grounded Theory) to produce an inventory of the individual employee's cost and benefit associated with security policies and mechanisms they came in contact with, and the cost and benefit for the organization. Data were coded by two researchers independently. An example of this process is given below. The following quote is taken from one of the interviews (A denotes the participant and B denotes the interviewer):

"A: But the biggest problem I have is probably on my windows laptop because it's Windows which is inherently a disaster waiting to happen. A lot of the like security stuff it's effectively damage limitation on Windows. The firewall and the other thing I really hate is Symantec anti-virus are both there for my own good you know.
B: Yes
A: However, I turn off the company firewall and I use the Microsoft one instead the one sitting in Microsoft XP because it permits me to open the network ports I need for my application while I'm outside the HP firewall (UI)
B: Yeah
The other problem is Symantec AV. When it's turned on our program takes an extra 20 minutes to build and test (UI) it takes 40 minutes. But with Symantec anti-virus turned on it takes 40 minutes. So that is basically tangible impact in my development time by running anti-virus stuff so right now I turn those things off when they annoy me."

Here a *category* of problem has been identified - a perceived negative impact on the user as a result of the company-installed virus checker and anti-virus software. The contextual condition is that these mechanisms limit or prevent the individual's ability to complete his work by blocking required ports or doubling the time taken to test a new build of their software. As a result, the individual takes the action of turning off these services and using only the firewall built in to Windows. The consequences of this action are that the individual is more productive but they (and by extension the company) are more at risk from the threats the full firewall and virus checker would mitigate against. Further conditions, actions and consequences identified in other interviews are added to this category until no new information emerges from coding.

4.3 Study results

The analysis of the interviews provided many examples of both compliance and non-compliance with security measures. When discussing their security behaviour individuals

frequently raised the issue of the costs and benefits of security policies. These costs and benefits are a complex web of perceived costs and benefits to the employee (both from a professional and individual point of view), and the perceived cost and benefit to the organisation. The individual employee's perception of cost and benefit is, however, largely determined by the impact of compliance on their tasks and responsibilities; the goal of their primary work task dominates their perspective whereas the security goal of the organization is a relatively small concern by comparison. A number of different scenarios were identified in which individual cost/benefit decisions were traded off against compliance with security policies. These are described in detail in the following section.

4.3.1 Example Cost/Benefit Scenarios

In this section five scenarios are presented, which illustrate individuals' cost/benefit perceptions of security measures. The quotes from the interviews illustrate the examples given.

4.3.1.1 Centrally scheduled maintenance tasks, such as weekly automated virus scan

Participants provided many examples of production tasks being disrupted by centrally scheduled maintenance tasks, such as virus checker updates, patches, and licence management software. The effect is either a complete interruption of the employee's production tasks, or their machines being slowed down to the point where they cannot work effectively:

P6: "when [the virus scanner] is turned on, our program takes an extra 20 minutes to build and test".

Most participants acknowledged that there was a purpose to those security measures. If they had the luxury of doing so, they would re-organise their work: *"go for a coffee"*, and resume when the maintenance task was completed. All participants felt it was justified to circumvent these security measures if they interfered with their ability to deliver their work on time, especially if they were up against a deadline. A minority, however, were more militant, feeling that no interruption of their work was justified. Each participant seemed to evince some awareness of the trade-off between personal and organizational good:

P3: "Anything that loses time is not good for the business."

Others clearly resented the loss of control over their own machines, and their ability to complete their work being at the mercy of *"whatever they think of next"*, especially if they could not perceive any benefit of those updates. In one of the organisations, several employees had permanently opted out of the updates:

P11: "next thing I know I'm installing some tiresome patch that I don't want to be installing. I used to circumvent that whenever I was in this building."

Employees did this either openly or covertly (some deliberately switched to an operating system or application software that was not centrally supported) - some with, some without, the knowledge of the security/IT managers. Most employees were aware that there must be organisational need for running these updates and that their *"opting out"* would undermine this:

P16: "different IT groups all want to scan these machines so I don't think they are too happy about that."

The prevailing attitude, however, was that - whatever the organisational need for these measures - it did not justify *"stopping people from getting their work done"*, and leaving several participants to wonder *"why they cannot run these things at night?"* This last example highlights a lack of awareness that there might also be conflicting goals at the organizational level, in this case the desire to reduce energy consumption versus effective security. This example also highlights the subjective nature of the decision making process, wherein the only costs and benefits taken into consideration are those directly relevant or visible to the decision maker.

4.3.1.2 Additional authentication

Several passwords need to be (re)entered, e.g. to connect to additional VPNs. At worst, failure to recall the correct password leads to failure to access (no availability), which in turn means inability to get work done.

P3: "The only [problem] I can think of is if either you lose a password or forget it."

Additional passwords increase cognitive load, and worry that e systems recall password accurately.

P10: "there are a lot of systems that need a lot of passwords and you know keeping a track of them is a bit of a pain."

Additional authentication hurdles cause delay in accessing systems and cause frustration at having to repeat a task.

4.3.1.3 Using encryption for data storage/transfer

The policy change hypothesised in the interview scenario was mandatory use of encryption for all data stored on USB sticks. For some participants, the risk to availability of data they needed was the key concern, but many participants balanced this against the confidentiality risk to their organisation. From the individual's perspective, the worst-case cost is permanent loss of data or lack of availability at a critical time for their tasks. The very possibility of losing access to data was an unacceptable cost to some participants:

P12: "I just feel a lack of control because most of the time the threat is unavailability created by the encryption system."

The fear of the risk to availability is deep, and several participants expressed a fundamental unwillingness to rely on a technology they did not understand.

P13: "I know very few people who run encrypted file systems on a laptop ... because they don't trust the file system. They want their data to be accessible."

In this context, data does not equal data – the fear is strongest for data that individuals have created or generated themselves; they felt strong ownership and resented this date being subject to a blanket company policy.

P10: "[USB encryption] would be irritating because much of the content is private."

Even if the data would not be permanently lost (e.g. because the file is still on the organisation's system) there a perceived risk to availability: not being able to access the data when it is needed. The result could be lost business opportunity, or embarrassment (looking incompetent):

P1: The only one I can think of would be the password on the device - I could forget it.
I: And what would be the result?
P1: I wouldn't be able to do my presentation. And that would be quite embarrassing.

The additional time to encrypt and decrypt data was a lesser perceived cost:

P10: "there is a cost to [USB encryption] in that it wastes my time."

These perceived costs made some participants feel entitled to break the policy:

P10: "if {USB encryption] was mandated I would probably work around it."

While others acknowledged the risk to the organisation, and the consequences that a breach of confidentiality would have:

P5: Well, for sure, a competitor could get hold of something and do something with it. The press could find out there was a breach and publicize that. And the client wouldn't be too happy. It's just bad at all levels if confidentiality were breached. [If] there was any perception of wrongdoing it taints the firm's brand and the value of that is huge."

Again we see strong evidence that it is the decision maker's knowledge that is relevant when considering how such policies will impact populations, not the knowledge of the security manager or his staff. While concerns that encryption software poses a threat to availability by somehow going wrong hold no water from a technical standpoint it does not mean that such concerns can be ignored. If this is a view that is held by the majority of the population being asked to use such software it will still influence their behaviour and compliance decisions.

4.3.1.4 Restrictive firewall

Firewall settings blocking access to a wide range of websites – from adult entertainment to social networking. In the worst case, participants reported that they could not get their work done, e.g. software developers could not access ports needed for running a virtual machine, and felt they had to break the policy:

P6: "the [company] firewall doesn't like those network ports being open. It lacks flexibility… By being too restrictive for my use they end up forcing people to cheat to disable things."

Other participants perceived the restrictions as more than an inconvenience – they felt it could damage business, and this, in turn would also affect their earnings:

P5: So how do you quantify that extra cost? … [it] can range from something as dramatic as business we won't win because we're not aware of information. And those fees [we miss out on] could be millions of dollars over one year or two. You know, if a company went public that we just had no idea existed, because we couldn't see the webpage.

Participants felt they were being forced into workarounds, so they work effectively.

P4: "I think people don't try to not follow the rules just for fun. When we do it's for a reason."

4.3.1.5 Over-zealous security classifications

Participants from one organisation stated that all data and documents tended to be blanket-marked as *confidential.* They reported that this led to unintentional breaches of the security policy, e.g. when sharing data with colleagues outside the organisation, even though the sharing was clearly needed to get their work done:

P10: "there is a kind of blanket imposition of inappropriate security restraints, sometimes on documents."

Not being able to share data was seen as resulting in lower productivity and freedom of thought/expression.

4.4 Perceived costs and benefits

As Weirich [76] points out, when an individual is presented with a security task, he has a choice of:

1. Complying and performing the required behaviour (or at least try to do so), or
2. Attempting to bypass the task.

The analysis of the interviews using Grounded Theory techniques revealed that users are aware of the costs to their time and productivity that result from security tasks. They also demonstrate an understanding of the need for security and in many cases state a desire to be secure where possible, in some cases reporting that they felt they were being forced into workarounds against their natural desire to be secure by the need to work effectively. This makes the above choice a non-trivial one. Users are not universally drawn to either reject or bypass security. This suggests that the decision – to comply, or not to comply - is the result of an internal cost/benefit analysis, in which the individual weighs up the advantages and disadvantages of complying or not complying; they do not regard compliance with security behaviour purely as a cost: the participants in our study value security, both for themselves and for the organisation they work for. Their disposition is often to comply with security policies - up to the point at which the perceived benefit of doing so is exceeded by the perceived costs associated with complying. (Or the point at which anticipated cost of complying – i.e. worrying about forgetting a password and the consequences of forgetting it – is perceived to be higher than the benefits of compliance.) Examples of costs perceived by individuals, from an HCI perspective:

1. *Increased physical load.* These are extra steps added to production tasks, increasing time and effort required to execute production tasks (such as a machine being slower because of security updates**)**, or new tasks added (such as additional authentication to enter building/log in to network, additional training to be completed for new security devices).

2. *Increased cognitive load* – additional information needs to be stored and recalled on demand (such as a password to encrypt/decrypt data). Research in psychology and human-computer interaction has established that human will try to avoid increased cognitive load even more than increased physical load.

In the interviews participants also discussed more intangible costs such as:

3. *Embarrassment* – security measures impacting on business practises may cause the individual embarrassment (such as not being able to open an encrypted file for presentation to a customer or audience) in addition to the inability to complete the task. Individuals also anticipate that the perceptions formed as a result of such an event could have a potential long-term negative impact on their career.

4. *Missed opportunities* – the increased time required to go through the steps of a security task, or the restrictions placed on data access by a security policy may mean a worker in a competitive environment misses an opportunity to effectively do business e.g. websites relevant to a deal blocked by firewall rules. As with embarrassment (3), individuals are concerned about the cost of long-term impact on their career.

These costs are supported by economic notions of utility. With an economic need to maximise profit and productivity, both for the individual and the organisation, security which threatens job security by preventing users from completing their production task will be seen to carry a high cost. The overall satisfaction of the individual is maximised by following a workaround and completing their primary production task with efficiency.

5. *The 'Hassle Factor'* –the perceived the cost associated with complying is often higher than the actual cost, because the knock-on effect on certain situational or contextual factors. For example, if an individual is under pressure to meet a deadline, waiting for a patch to install feels more onerous than if it happens over lunch or at a period of low activity. This cost is most apparent where the security task slows down or prevents the primary task

These costs also identified in HCI and economic literature. Previous studies (as shown in Section 2.2) indicated that users feel overloaded with the need to manage multiple passwords and are often already operating at their cognitive limits. This background context explains why users will seek to avoid additional cognitive load where possible. Framing this economically, whereby the individual seeks to maximise their personal utility, these costs clearly explain why non-compliance is in many cases the preferred course of action. Individuals that are directly motivated to be secure, i.e. those for whom security is their primary task, and the one from which they derive most utility, will prove to be compliant. However, such individuals are exceedingly rare and an individual's utility is far more commonly dependent on taking satisfaction in the successful completion of their primary task, a good salary and being able to provide for their family. In this context the behaviours reported during the interviews become a rational response. At the same time, compliance with required security behaviour can have perceived certain benefits:

1. *Avoiding the consequences of a security breach.* The most obvious benefit is avoiding the consequences of a security breach – such as losing data – which would have a negative impact on the individual's task. It is important to reinforce the point being made here; that users do value security, up to a point, and in most cases understand its benefits.

2. *Protection from sanctions* – there is no danger of being *"caught"* in breach of policy, and exposed to the sanctions stipulated by the organisation. Therefore the individual does not have to worry about being caught, or the sanctions.

However, the degree of worry depends on the individual's risk appetite and the likelihood of sanctions being applied. Adams and Sasse [1] and Weirich [76] point out that if policies are routinely breached, but stipulated sanctions are not applied, individuals quite reasonably do not expect any consequences from being caught in breach of policies. This highlights the need to have a policy that is both consistent and enforceable. Unrealistic policies that are then not enforced and left to slide are arguably less effective than a less restrictive policy that is realistic about what it can ask of its workforce and can enforce through sanctions and monitoring.

4.5 The Compliance Budget model

Security policies and mechanisms will not be effective without the co-operation of individual employees. The most obvious form of co-operation is compliance with security tasks. Most users will comply if this does not require any additional effort. When extra effort is required, individuals will weigh this extra effort against the perceived benefits for them,

in the context of their primary production tasks. It is this process of trade-off that is not well understood by either the HCI or Security communities. Security sees that there should be no need for a trade-off as security tasks are of prime importance. One security manager, when questioned about how users should cope with additional tasks, said *"they should come in to work 15 minutes earlier and leave 15 minutes later."* [59] Conversely, HCI sees the improvement of usability and the reduction of user effort as its primary objective. As the overall utility of an organisation or individual is dependent on a combination of these two factors a balanced view is needed. Thus, this trade-off is central to the decisions surrounding compliance and an accounting of the factors that influence it is necessary in order to support security decision making. This is the objective of the Compliance Budget model. If the user's goals are aligned with the organisation's security goals, there is no conflict, as the behaviour required of the individual translates into perceived gain for them as well as for the organisation. In rare cases the security task will be equivalent to the primary task, for example a member of airport staff manning the X-ray machine at a baggage checkpoint, and so no security-specific effort will be registered by the user. Thus, security policies are likely to be followed – at least by most individuals, most of the time.

Based on the results of Weirich [76] there will be some exceptions to this rule, because:

1. Some individuals have a high propensity for risk-taking
2. Some individuals may have other issues with the organisation, and contravene security policies as a way of expressing dissatisfaction or dissent.

Conversely, most individuals will not be inclined to choose the behaviour required by the organisation if there is a conflict between the security behaviour and their own goals. In this case, either some part of the joint set of goals will not be met, or the individual will need to expend effort *without gain* to help the organisation's security goals.

Compliance occurs if the individual chooses the behaviour required by the organisation, even though it makes it harder for them to realise their goals, or even prevents them from reaching them altogether. In the scenario mentioned above where the security task is in alignment with the primary task no true compliance decision has been made as there is only one reasonable course of action. Compliance can be seen as a kind of organisational altruism. From the individual's perspective, this is a situation of 'pain but no gain' (recall we are focusing on conflict situations where individual and organizational goals are unaligned). The 'pain tolerance' - the amount of extra effort an individual is prepared to make for no personal gain – is what is referred to here as the *Compliance Budget.* The limit of the Compliance Budget I have termed the *Compliance Threshold*; this being the point at which the individual no longer has the motivation to comply with official requirements. The closer an individual is to his Compliance Threshold the higher the cost to the organisation of

achieving compliance, as the perceived cost to the individual will also be higher. This is due to the 'hassle factor' – a task performed at the start of the day will seem less onerous than the same one executed after a series of security-related effort expenditures. In essence there is a cumulative cost associated with performing tasks in sequence; each new task has some residual level of effort from previous tasks added to it. Once the Compliance Threshold has been exceeded, there will be almost no way to achieve compliance except through heavy monitoring of individuals' behaviour and enforcement as the workforce will be continually looking to avoid security tasks to make their life easier, focusing exclusively on their primary production goals.

When an individual is faced with a compliance decision, the costs detailed in the results section will be weighed up by the individual (consciously or subconsciously) and measured against the benefits. As stated before the issue of compliance only comes into question when the individual is placed in a situation where there is a cost to him but no direct benefit. The decision to comply or not comply with a single task can be summarized with a brief formula. As discussed above the *hassle factor* is cumulative. This means that a steady erosion of the Compliance Budget will take place as tasks stack up, and the individual's tolerance for further tasks (or repetitions of the same task) will be reduced. This must be represented in the formula to take into account the cumulative effect of previous tasks already performed. We can express this decision as:

> Compliance if: (current task cost x *hassle factor*) + total cost of previous tasks < Compliance Threshold

In addition to the internal factors identified in section 4.2, there are external factors that impact both an individual's total Compliance Budget and the rate at which it is expended. External factors are created by the production task performed by the individual, and the organisational environment. It is factors in this second category that security managers can use to influence individuals perceptions of cost and benefits of compliance, and their security behaviour. The key external factors are:

1. *Design*
    a) The most direct way to influence cost-benefit perception is to reduce the actual mental and physical workload that individuals have to expend on compliance.

    b) Well-designed security seeks to minimize friction between security and business processes, and avoids putting individuals in situations where they have to choose between security goals and production goals. Improving the design of the security system will reduce the cost of each security action, meaning more tasks can be undertaken within the same Compliance Budget.

2. *Awareness, Training, Education*
    c) Effective training in using security measures can improve individual performance, which in turn reduces the cost associated with security measures. It can also build individuals' competence – which is a benefit to them, and their confidence – which can reduce the stress and anxiety associated with using security mechanisms.

    d) Raising awareness of the risks and vulnerabilities faced by the organization increases the perceived benefits of compliance.

3. *The culture of the organization*. The more security-minded an organization is, the less friction compliance causes. Weirich & Sasse [77] report that individuals' security behaviour is strongly influenced by behavioural norms – most individuals try to *"fit in"*, rather than seek conflicts with their colleagues. Building a positive and strong security culture reduces friction and perceived cost of compliance

4. *Monitoring*. The visibility of the organisation's compliance monitoring, and willingness to administer advertised sanctions, will determine how likely it is an individual thinks they will be detected and reprimanded if they do not comply with policy. This will in turn feed into their decision to comply or not.

5. *Sanctions*. Avoidance of sanction is a perceived benefit. Thus, for sanctions to be effective, they must be enforced, and seen to be enforced [77] used in response to a security failure will be factored in to the cost/benefit analysis of the individual. They influence the level of perceived benefit from the protection gained through compliance.

In general, each of the external factors affects either the total Compliance Budget, or the rate at which it is spent. This is an important distinction to make, as there is a finite improvement that can be made to the total Compliance Budget available through improving awareness, training and culture. Beyond this limit, further attempts will be counterproductive either in the business sense (time spent away from work) or in terms of the Compliance Budget (individual costs imposed through time taken up and attention demanded away from the individual's agenda). The implication here is that initial gains should be made through reducing the cost of each security task, and therefore slowing the rate of spending of the Compliance Budget. The most effective way to do this is to improve the design and integration of the security system. In order to do this effectively tools must be put in to place to allow security decision makers to systematically consider ways of making these

improvements. In particular data on how differing design and policy options affect user effort must be collected, and a representation of the system must be created that allows trade-offs between factors to be considered. These needs are addressed in Sections 5 and 6 of this thesis. The Compliance Budget therefore forms a key part of the motivation for the other sections of research.

The effect of the Compliance Budget and Compliance Threshold is to place a cap on the effectiveness of an organisation's security policy – the compliance limit. Once this threshold is reached, several outcomes can result, of varying significance. Firstly, adding in further security measures that require compliance from 'overspent' individuals are unlikely to improve security, unless they are heavily monitored and enforced. Secondly, in cases where adding a new measure breaches the Compliance Threshold, not only will the new measure have no effect, but existing security measures will be become significantly less effective as individuals lose patience even with security tasks they were previously willing to undertake.

4.6 Implications of the theory

The results of the empirical study allow four main conclusions to be drawn:

1. An individual's Compliance Budget sets a cap on the effecitveness of security practices they are involved in. This is because - once the threshold is crossed - the individual will chose work-arounds motivated by his or her own needs, rather than the more altruistic process of compliance.

2. The closer an individual comes to crossing the Compliance Threshold, the higher the organisational cost of achieving compliance will be. This cost may come in the form of an increased need for monitoring and sanctions, and/or a passive cost in the form of increased stress and reduced goodwill.

3. The organisaton can influence individuals' perception of where the threshold lies, and how fast that threshold is reached, through an understanding of the decision-making and feedback cycles associated with a compliance.

4. The organisation can employ economic reasoning to ensure the available budget is used to achieve compliance in areas where it matters most. A possible outline for these cycles is presented as part of this conclusion.

Figure 4 illustrates the effect of the Compliance Threshold on the effort required to ensure continuing compliance. The closer an individual is to his or her Compliance Threshold, in terms of effort already spent on secondary tasks, the higher the perceived cost of each task will be. This means that to ensure compliance, the organisation will have to expend more effort or resources on monitoring/sanctions etc. Once an individual is pushed passed his Compliance Threshold, it becomes very difficult to ensure compliance as the individual will be continually seeking ways to reduce the burden of compliance. The possible responses of the individual vary in severity and can be hard to predict, not least of all because the budget threshold is not a hard limit and reacts more elastically to pressure. For example, individuals may be willing to accept exceeding the Compliance Threshold for a short term if provided with sufficient motivation (increased remuneration, promise of future perks etc.).

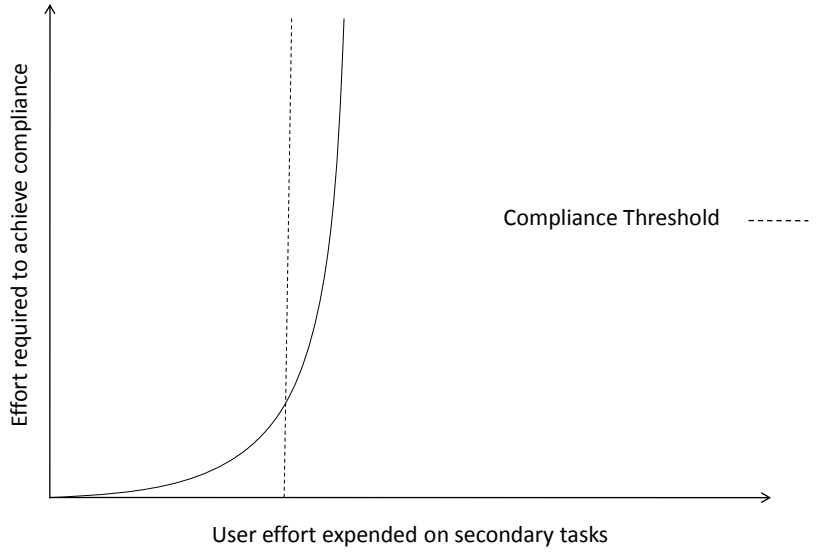

Figure 4: *The effect of the Compliance Threshold on organisational effort required to achieve compliance. Effort rises as we approach the threshold and rises dramatically once it is crossed*

If the budget is exceeded for longer they may continue to undertake their current compliance tasks, but begin to look for ways to undermine or get back at the organization (this may be termed 'spiteful compliance'). In extreme cases, the individual will effectively cease to comply whenever possible, and focus on his or her own goals entirely. His or her last resort is, of course, to simply leave the organization.

In Figure 5, it can be seen how the effectiveness of the security system changes with the costs placed on the individual. It is important to recognise that it is *perceived* cost that is important. The individual is making decisions from a personal perspective, so it is the impression of the costs placed upon him that is relevant. Adding new security measures will place some

burden on the individuals using or encountering them. This will increase security in the organisation if they are used effectively and in the intended manner. However, as the Compliance Threshold approaches, the rate of gain of effectiveness starts to slow as the individual perceives a higher cost to each new task. This implies a situation of diminishing returns when deploying security measures. Once we cross the threshold, the individual may choose to remove himself from the security process, at which point the effectiveness of the system drops dramatically. Now not only are new measures ineffective, but previously-performed tasks may be neglected. The level of effectiveness does not return to zero because in the course of normal business practice certain security measures will still be undertaken (such as using a password to log in to the network) even though minimal effort is being expended on security per se. Additionally, we are now operating in the grey area discussed above, where individual responses to exceeding the budget are hard to predict. The innate flexibility and mobility of the budget limit itself also contributes to the uncertainty.

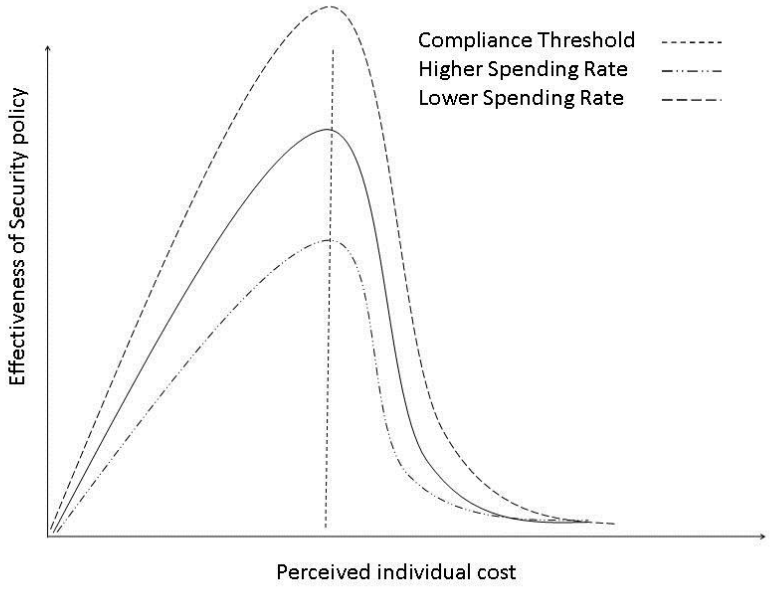

Figure 5: *How perceived individual costs relate to effectiveness of security. Alternative rates of compliance expenditure are also shown for comparison. Once the Compliance Threshold is crossed security effectiveness drops sharply as altruistic behaviour is increasingly exchanged for selfish options*

The initial gradient of the lines indicates how fast the security budget is being spent in each scenario. The more each task costs, the faster the budget is used up and the lower the gradient of the line. If each task costs relatively less we spend the budget slower and the gradient of the line increases. We can achieve better security behaviour, and a better return on security, for the same level of effort. By changing the rate of spending (for example by reducing the cost of each task through better system design), it is possible to alter the how

effective it is possible to make the security of the system before this breakdown takes place. This is illustrated by the 'lower spending rate' line in Figure 5. When compared to the 'normal spending rate' we can see that this line ascends more rapidly, reaching a higher level of effective security before reaching the budget threshold – we progress faster vertically relative to the horizontal consumption of compliance effort.

Spending the budget at a faster rate by implementing security tasks that conflict with the business process or are poorly integrated with the daily routines of the workforce lead to a lower maximum level of security (we are able to implement fewer security policies before the threshold is reached). This is shown by the 'faster spending rate' line. In this scenario security measures are more costly to the user and so at the point that the Compliance Budget threshold is reached, the effective security is below that of the other more efficient spending rates.

So how is the Compliance Budget set, and how can the rate of expenditure be influenced? In answering these questions we must be aware that these are not set values. They exist in a dynamic situation and are subject to constant change and revision. The results of previous compliance decisions, and the organisational environment, all contribute to the maintenance and expenditure rates of the Compliance Budget. This is illustrated in Figures 6 and 7.

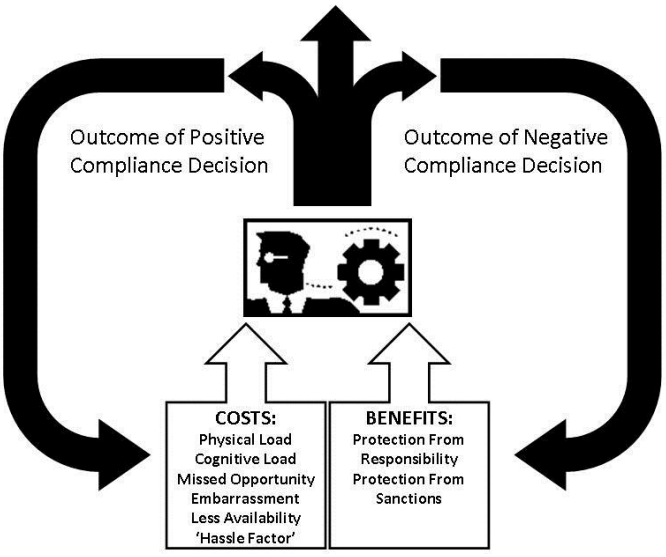

Figure 6: *Feedback cycle for the individual. Positive compliance decisions increase the weight given to individual costs. Negative compliance decisions increase the weight given to individual benefits*

The factors that may be considered in a single decision are shown in Figure 6. In figure 7 we can see how the organisational environment impacts on the cost/benefit analysis performed by the individual through those factors. Improving system design, and creating a positive

security culture, will simultaneously decrease the perceived cost of security tasks, and lower the rate of expenditure of the Compliance Budget. On the other hand, increasing sanctions and monitoring will give more weight to the benefits associated with compliance. Measures associated with either side of this cycle increase the likelihood of a positive compliance decision in the future.

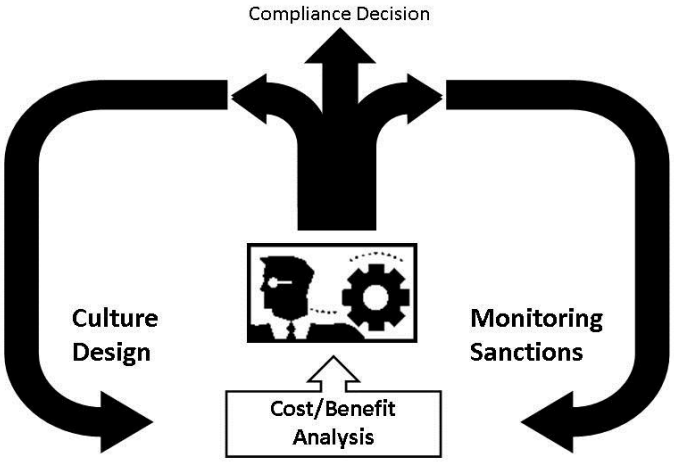

Figure 7: *How organisational factors can be used to increase the likelihood of compliance. Improving system design and creating a security culture decrease the perceived cost of compliance. Harsher sanctions and monitoring increase the perceived benefits of compliance*

4.7 Effective Compliance Budget management

Targeting research questions 3, *"How can we maximise security when human behaviour is taken into account?"* this section examines the implications of the Compliance Budget in the context of making security management decisions. In that security systems are, by and large, reliant on user effort and co-operation to drive them, the Compliance Budget places a cap on how much security can be supported by a user base. Spending user effort (and therefore compliance) in the most efficient way possible to maximise effective security before the Compliance Threshold is reached is critical. Some ways to approach this are discussed here.

The Compliance Budget is consumed when secondary tasks interrupt the user's primary task. While other secondary tasks can also be considered, the focus of this research is solely on security. In order to minimise the rate of spending of the Compliance Budget security managers should understand that for the vast majority of their users security is a secondary task. As such, any security task will be seen as an additional cost, and the degree to which

that task interrupts, delays or prevents the user's primary task determines how much perceived effort that task consumes.

Clearly not all such interruptions will be avoidable and the Compliance Budget, while finite, is not an empty resource. Effective management therefore involves selecting where and when these interruptions take place, rather than seeking to remove them entirely. Examples from both the literature and the empirical study outlined in Sections 4.2 and 4.3 illustrated that compliance issues most commonly arise when security mechanisms are perceived to prevent the user from being productive. Again the notion of individual perception is important. Understanding the core concepts of the Compliance Budget will allow managers to anticipate the existence of 'hotspots' of friction where the effort that following security policy requires is considered excessive by users. Data gathering techniques such as the semi-structured interviews used in the empirical study (see Section 4.3) that underlies the Compliance Budget can bring to light these hotspots. Once they are identified they can be managed. The appropriate management of these problem areas will depend on the nature of the hotspot in question. Let us consider some of the issues raised in Section 4.3 through the lens of the Compliance Budget.

Section 4.3.1.1 presented interview quotes relating to centrally scheduled maintenance tasks, such as automated virus scans and remotely delivered patches. The prevailing opinion in the interviews was that such processes detrimentally interrupt the production tasks of the organisation. The primary complaint was that the timing of the procedure was inconvenient, not that the procedure was unnecessary per se. The Compliance Budget tells us that perceived temporal load in this case is high – employees are frustrated by the *"tiresome"* nature of the task and are concerned that *"Anything that loses time is not good for the business."* Additionally, they are under the impression that multiple departments are running scans of some kind at different times, leading to inefficiencies. The goal of the security manager following Compliance Budget model is therefore to reduce the temporal burden of these managers by restructuring the way such maintenance tasks are scheduled. Removal of the tasks is not a requirement under the Compliance Budget.

Some of the employees mentioned that they would be happy to re-organise their work to accommodate the required tasks if there was more flexibility. A possible solution then would be to require the task to be initiated within a window of 4 hours but for the exact timing of the task to be at the user's discretion. If multiple maintenance tasks are required then bundling them together, where possible, would again minimise the amount of time a user loses control of their machine – another key factor in creating frustration. In this case the onus is on the organisation to effect change in order to remove the perception that such maintenance tasks are *"stopping people from getting their work done."* The Compliance Budget facilitates management decisions that achieve this aim by identifying the factors that drive

this perception and removing them. In particular in this case it has allowed us to identify concrete systemic factors that can be addressed structurally at the organisational level.

In contrast to the organisational-based solution that the Compliance Budget indicates would be optimal in the scenario above Section 4.3.1.3 presents a case in which a different management approach is required. The security control under discussion here is the use of encryption for data storage and transfer, specifically USB sticks. The feedback generated during the interviews indicates that employees feel that the perceived costs associated with encryption are related to a loss of availability with regards to their data. Comments such as *"I just feel a lack of control because most of the time the threat is unavailability created by the encryption system,"* indicate that a change in the system will not reduce the costs in the Compliance Budget model and therefore will not lead to a rebalanced cost/benefit calculation that favours compliance. From a technical standpoint there is no reason not to 'trust' encryption yet this distrust exists. As such, to manage this situation through the Compliance Budget rather than seeking to change the system the correct route would be to target the user's understanding of encryption. By targeting the distrust through education and training (one instance where this approach can be useful) the fear of availability loss can be removed and the cost/benefit equation altered in favour of compliance. Again the Compliance Budget allows us to understand the factors influencing the compliance decision and select a management strategy that works on those factors.

A more complex situation is presented in Section 4.3.1.4 in which the company firewall blocks ports necessary for software testing. Here the lack of flexibility in the configuration of the firewall left the employee feeling like they were justified in disabling the firewall because it so clearly conflicted with their primary task. In this case the manager could decide to remove the ability to disable the firewall. However, this will strongly negatively impact the primary task of that individual leading to a high cost to their compliance budget. It also removes control over the user's machine which as we saw above is a common cause of frustration. This has a knock-on effect on the user's Compliance Threshold due to this frustration. This can be likened to a lowering of their loyalty to the organisation as the user sees the organisation as a source of difficulty that hinders their work instead of facilitating it. Factors such as loyalty to the organisation and understanding of the threat environment affect the total Compliance Budget available. While enforcing the firewall rules may make sense from a purely security perspective the Compliance Budget model indicates that this will reduce the individual's Compliance Threshold as well as using up a section of the Compliance Budget.

Good management therefore involves ensuring that where possible secondary tasks conflict as little as possible with the primary task, and where this is not possible to ensure individuals have the highest available Compliance Threshold. The goal is not to weaken

security but to look for decision points where currently there is an unnecessary effort burden placed on the user and to mitigate that burden. The requirement to engage in trade-offs lies at the heart of this process of review. In the case of our firewall would unblocking the ports required by the software test open up vulnerabilities with a greater potential risk than the actions the user is taking to cope with them being closed? In that they are disabling the firewall altogether here the trade-off, seen through the lens of the Compliance Budget, seems clear, that they should re-enable the company firewall and the port should be opened. This view is simplistic and considers only one aspect but by iterating through this process a better solution can be reached. The mistake would be to try to force the user to conform to the existing policy. Not only does this contravene HCI principles but it will reduce their Compliance Threshold, negatively impacting future compliance decisions.

Failing to take the Compliance Budget in to account will introduce systemic vulnerabilities in an organisation's security. Workarounds developed by the user will be done so without expert security knowledge or a particular motivation to be secure. In most cases these actions will not be malicious but seen as a necessity by the individual. Users will seek to complete their primary task in some form. If they are pushed past their Compliance Threshold then the activities associated with their primary task will where possible take place outside the boundaries of the security system – an unwanted outcome produced by negative incentives. One of the primary objectives of the Compliance Budget is to demonstrate that such outcomes are not the product of poor user education or motivation but are the inevitable product of the security system itself. This understanding is critical if more optimal security decisions are to be made.

4.8 Operationalising the Compliance Budget

The discussion above shows how the thinking of a security manager in the case of centrally scheduled maintenance tasks, encrypted storage and firewall configuration should be influenced by the Compliance Budget. In each case a trade-off is considered and a decision made that seeks to find the optimal point along the spectrum of options. While the Compliance Budget in its current form can highlight the factors that influence a compliance decision, and give insight in to which management methods may be most appropriate what it cannot do is aid in the specifics of a decision. That is to say its contribution to security management is at a relatively high level. In order to 'operationalize' the Compliance Budget as a tool for supporting the real-world decision making of security managers an extra step is needed.

To optimise security within the Compliance Budget model managers should seek to spend user effort up to the Compliance Threshold as minimally and efficiently as possible. This should be done firstly by minimising friction with the primary process, typically by

choosing appropriate controls and strategies, and by implementing them in such a way as to avoid placing unnecessary burdens on the user (as discussed in Section 4.7). However, in order quantify the effort associated with each control, so that such decisions can be made, the Compliance Budget must be supported by empirical data. Specifically the burden placed on the user by various mechanisms must be examined empirically and with sufficient granularity that a range of policy options can be systematically compared. The volume of data required to fully characterise all the situations in which the Compliance Budget can be applied is vast; far beyond the scope of this thesis (although much of this burden can and should be shouldered by organisations wishing to pursues this management model). However, it is useful to examine the types of experiments that may be required and discuss how the results can be interpreted within the Compliance Budget framework. This objective will be addressed in Chapter 5.

It is also necessary to create a representation of the system in which the compliance decisions are taking place. The Compliance Budget offers security decision makers a model for user decision making that allows them to influence the outcome of compliance decisions. This is done through improving security design and managing the factors that impact on the user's perception of the costs and benefits associated with compliance. However, user behaviour is not the only component of a system. It is, as was demonstrated in the literature review (in particular see Section 2.4.4), a component that is both important and too often neglected, but it does not exist in isolation. Therefore, as well as a comprehensive data collection programme that addresses the need for metrics of user effort, a framework for representing the Compliance Budget model for human behaviour alongside technical considerations is also essential for systematic decision making. This framework should allow decision makers to consider the trade-offs between the business and security processes, and examine how policy changes will impact user behaviour. Chapter 6 presents an example of such a framework and thus represents a possible method for positioning the Compliance Budget alongside more traditional tools such as systems modelling.

5. Data Collection

In chapter 4 introduced the idea was introduced that user effort is a limited resource in the form of the Compliance Budget model, and that – to best serve the utility of his organisation – a security manager needs to manage this resource. This process involves making trade-offs at the policy level between various security mechanisms to ensure a good balance between the level of security mandated by the security policy of the organisation, and the burden that the policy places onto the organisation's employees. A policy that overburdens users will undermine security as they will adopt workarounds to cope with the pressure. It is necessary for security managers to consider how the requirements of the policy incentivise user behaviour and make the necessary trade-offs to ensure that users are not pushed into insecure behaviour, all the while meeting the security needs of the organisation. This is not an easy task however, and to understand and evaluate these trade-offs, security managers need access to more fully developed metrics. In particular the user effort associated with the various mechanisms they are considering is a requirement if a systematic decision making process is to put in place.

Such metrics must be built on valid empirical data – it can be described as a security version of the keystroke-level model in HCI [16], which comes with a catalogue of execution times for different interaction tasks, based on how practised a user is. Building such a database will require a huge data set, the collection of which is outside the scope of a single PhD thesis. Nevertheless it was useful to carry out data collection as part of this research in order to demonstrate the types of data that are relevant to problems addressed, and the associated methods of collection. The completion of this goal is the primary aim of this chapter. Data collection also serves the secondary purpose of contributing to wider data set necessary for supporting systematic decision making. Within the Trust Economics project, further data was collected on password habits and use [38, 39] and the data collected and presented here can be seen as a companion set to that.

While the conclusions drawn from such data are to some degree generalisable it is worth noting that the specific nature of the primary task and the context of use effect the impact that a particular policy or mechanism has on the user. Therefore larger organisations would gain more accurate measurements of impact by building their own specific data sets using the kinds of techniques discussed in this section.

These studies conducted to collect a set of example data focused on one security mechanism: passwords. Passwords are still in widely used, and there is no sign of this changing in the near future. This makes them a useful mechanism to target, because results will be relevant to security decision-makers in companies – a focus of this project. Additionally, passwords

have attributes, such as length and characters used to compose them, which can be directly controlled by the experimenter through the instructions issued to participants during the recruitment phase of the study. As discussed in Section 3.1 the ability to control variables is advantageous when wishing to assess task performance. Password entry is also an atomic security task that is generalisable, at least for non-disabled users on standard keyboards. This makes passwords a relatively simple focus of empirical experimentation.

In comparison to the other chapters of this thesis the work reported in this chapter focuses on a specific problem – the relationship between password length and complexity and the effort required to use that password. As such, a more focused research question is required:

> **Research Question 5:** *Does password length or password complexity have a greater effect on the user effort associated with remembering and entering a password?*

This question can be seen to underpin research questions 3 and 4, focusing on a specific aspect of those questions. In the same way this chapter underpins Chapter 4 and the Compliance Budget by investigating the impact of a single security policy decision. As was expressed in Section 4.8 in order to function as an operational tool to support decision making the Compliance Budget must be grounded with empirical data. The Compliance Budget states that user effort should be spent as efficiently as possible and here we can look at that approach in practise.

The security delivered by a password is linked to its length and complexity, against certain kinds of attacks (clearly a phishing attack, for example, is effective regardless of how long or complex the password is). Security policies have different length and complexity requirements. A policy that requires long but simple passwords may deliver equivalent security to a policy that mandates short but complex passwords. However, the two policies, given the nature of human memory, will place differing effort burdens on users. As was discussed in section 4.7 effective compliance management using the Compliance Budget model involves avoiding placing unnecessary burdens on the user. In order to know which policy should be selected it is necessary to determine whether password length or complexity burdens users more. The following sections will present methods of data collection, and an empirical data set, which will allow a security manager to effectively and systematically make that decision.

5.1 Data collection methods

An experimenter faces a set of decisions when considering the collection of empirical data. At one end of the spectrum, laboratory experiments offer precise control over the environment the experiment takes place in, and the stimuli and tasks presented to the

participants. This ensures that, as the conditions under which the data are collected are controlled, when a condition is varied by the experimenter and difference in the results will be a product of this variation. At the other end of his continuum, conducting experiments in real-world environments allows data to be gathered that most closely matches the day to day lived experience of the subjects. However, as not all the conditions of environment are under the control of the experimenter any differences observed in the results cannot always be assigned to the variations in the experimental conditions. Both approaches have drawbacks; laboratory experiments lack real-world validity and field studies will almost certainly involve dealing with confounding variables. Laboratory experiments also have a tendency to overestimate the performance of those that take part in them, because distractions available in the real world do not affect participants. For security experiments, there is also the question of experimental design: if the experiment only investigates participant performance with the security task, the data gathered might over-estimate performance because security is very rarely a primary task for those undertaking the activities it demands (as discussed in Sections 2.2 and 4.4). Data collected in real-world environments does not have this limitation – the participants will be involved in their normal routine, and the tasks studied should be afforded a similar level of priority to security tasks when they are routinely encountered. However, data collection under such conditions has many challenges: the experimenter has no control over environmental factors, and may not be able to capture all the factors and distractions that affect participants' performance.

With this in mind I set out to gather data in such a way as to minimise these limitations. In order to do this rather than focusing on a single avenue of data collection a mixed approach was taken. The purpose of these experiments was not to collect a single unified data set but to develop and refine techniques that could be used and extended as part of the Trust Economics project for wider data gathering efforts. The following sections describe the two experimental methods and the results gained.

5.2 Laboratory password entry experiments

The first set of empirical data gathered as part of this thesis was motivated by the need for supporting data on human effort to underpin the Compliance Budget. The Compliance Budget states that, faced with security tasks, users make an implicit cost/benefit decision. A key factor in this decision making is how close they are to their *Compliance Threshold*, the point at which their available effort for tasks outside their primary goals is exhausted. Pushing users over this threshold has strong negative consequences for the organisation because users who are over the threshold are less likely to comply, so it is in decision-maker's interest to avoid doing so. We need to be able to quantify the effort users have to

expend on security measures. This quantification will naturally be subjective; a measure of how the user views the task relative to other activities they may have to undertake and contextualised by the situation in which the task takes place. An objective measure of the effort involved is not useful in this context as objective information of this kind is not factored in to an individual user's decision-making process. Deriving a measure of absolute user effort in any context, let along information security, is a task sufficiently complex and challenging as to be rendered impractical. This is in a large part due to the lack of viable methods for directly measuring individual effort in any way other than purely physical terms (such as the work done in Joules). However, effort of the kind we are concerned with, although it relies on self-evaluation and is therefore no means objective, can be measured with sufficient accuracy through careful use of the appropriate techniques – such as carefully documented laboratory-based experiments. These measurements can be taken as a reasonably reliable basis for predictions of user performance and reactions. A relative notion of effort is an achievable goal and one that still facilitates the improvement of security decision making. For example, it allows the comparison of two different policies with a view to identifying which places the higher burden on the user.

The aim of the studies conducted was to obtain on how altering a password policy affects a user operating under that policy. The following stages would be needed to address this issue:

1. Identify variables that affect the effort required to use passwords
2. Determine from this list which variables can be controlled by policy
3. Select those variables that can be realistically and systematically manipulated for the study
4. Ascertain how user performance varies as the variables change

The outcomes of these stages are addressed in the sections below.

5.2.1 Experimental Design

5.2.1.1 Password-related effort

Passwords are in widespread and continual use in modern society and are found in both commercial and social environments. Anyone using a computer today is most likely to have to interact with multiple services that require a password to use with relative frequency. Florencio and Herley [22] report that users have an average of 25 password-controlled accounts, with password re-use being common. This implies that the number of passwords a user has increases the difficulty of use, leading to password re-use when multiple accounts are present. This then is one factor relating to password effort.

In addition the composition and structure of the password will affect its memorability, with more complex passwords being harder to recall. Thus to number of passwords in use we can add the following factors to the list of those affecting password-related effort:

- *Password length*
- *Complexity of composition*
- *Meaningfulness*

Complexity of composition here refers to the number of different character sets from which the password had to be drawn; the most complex passwords have to contain at least one from all sets (lower case letters, upper case letters, numbers and symbols). The vast majority of passwords are still entered via a keyboard, though with the rise of tablets and smartphones, the number of users who have to enter them via soft keyboards is increasing – and on soft keyboards, complex passwords require even more effort than on traditional ones. This highlights why different measurements are required for different combinations of users, tasks, systems and contexts – the use of a soft keyboard creates a different perceptual-motor task compared to standard keyboard entry, and this must be recognised experimentally.

Meaningfulness here refers to the degree to which the password is contains information that carries meaning for the user, and is thus easier to recall. Most users find it easier to recall passwords with meaningful content [61], though whether a password is meaningful is subjective, since some users have strategies for password generation that create meaningful passwords without using visible known words. For instance, a user may take a memorable quote or phrase, and take the first (or second, or last etc.) letter of each word as the password. For example, 'we who are about to die salute you' (the widely accepted, although historically questionable, phrase spoken by Roman gladiators) becomes 'wwaatdsy'. While this may appear the product of a nonsensical, or even random, process to an observer unfamiliar with the method of its construction, it retains a high level of meaning for its creator. Obviously using very well-known phrases for this method will yield weaker passwords as they can be derived in the same way by someone building up a dictionary of such constructions for use in an attack. A machine-generated password of randomly assigned characters would exist at one end of this spectrum while a straightforward dictionary word would mark an opposing point. Passphrases as described above would occupy some middle point.

Frequency of use also affects memorability [61, 76, 82]: a password that is used often will be encoded in long-term memory leading to automatic rather than conscious recall. As automatic recall is easier than conscious recall (with muscle memory also playing a part when entering passwords) a password that is used frequently will require less effort to use successfully than one that is used infrequently. This concept also links to how long a

password is in use: one that has been in use for some time and has been retrieved regularly will require less effort than one that has been in use for a short period of time. Thus, a new password will require more effort until it has become familiar enough to be accessed with automatic recall. From this I also hypothesise that changing a password more frequently will increase user effort as the password in use will spend less time in the automatic recall stage and more time in the newer learning stage.

The factors that contribute to password effort, as distilled from the literature [61, 76] are:

1. Number of passwords in use by the user
2. Password length
3. Complexity of construction
4. Randomness of Construction
5. Level of user input in generation
6. Frequency with which passwords are used
7. The duration for which passwords are kept in use

Of these, Factor 1 cannot be set as a matter of policy in absolute terms, although a security manager can control the number of passwords used within an organisation – for example by choosing between a single sign on approach or having one password for each service. As this study aims to provide data to support security decision-making, this factor is less relevant than the others, which can be set through a security policy. Additionally, this factor cannot be practically controlled for in a laboratory environment. While it would be possible to pre-test participants, ask them how many passwords they have, and from this filter them down based on the requirements of the study I considered this to be likely to overly limit the amount of participants available. As subject recruitment can be challenging (and potentially expensive if remuneration is offered for taking part in the pre-test) I did not consider this approach to be feasible for this study.

Factors 2 – 5 can all be controlled for in a laboratory-based study and a subset of these will form the basis of the variables used in the experimental conditions. Factor 6 will require a longer term study using a different methodology – we will return to this in section 4.2. Factor 7 could be controlled for in a very long term study but this is outside the scope and resources of this thesis and so will not be considered. The factors that will be used in the study are:

1. Password length
2. Complexity of construction

5.2.1.2 Experimental conditions

With a key goal of the study being to be able to compare the impact of password policies it was decided to ask participants to generate and register two different passwords with the experimenter, each one conforming to a different policy. They were also asked not to re-use one of their existing passwords. The policies were chosen in consultation with the industrial partners of the Trust Economics project who wished to compare two policies that were in use at the time with a view to ascertaining which was the more difficult for the workforce to use. The two policies were:

Policy 1:

Users must choose passwords that are at least six (6) characters in length and contain at least three (3) of the following categories:

- *Upper case letters (A, B, C, . Z)*
- *Lower case letters (a, b, c, . z)*
- *Numbers (0, 1, 2, . 9)*
- *Non-alphanumeric. The following non-alphanumeric characters are safe across all platforms:*
- *! (exclamation point), + (plus), ; (semi-colon), - (minus), % (per cent),*
- *: (colon), $ (dollar)*

Policy 2:

Password constructed MUST consist of the following characteristics:

1. *Minimum length of 8 characters*
2. *Chosen from any acceptable character sets available on the target system*
3. *Includes at least one alphabetic and one numeric character*

The wording of the policies was taken directly from the policies in use in the industrial partner organisations. Policy 1 requires a shorter but more complex password, whereas Policy 2 requires a longer but simpler selection. Performance and effort ratings would be gathered for both policies and compared.

5.2.1.3 Data collection tools

Two pieces of software were used during the study to assist in the gathering of performance and effort data. In order to gather accurate performance data I decided to use a piece of modified key logger software developed by Roy Maxion at Carnegie Mellon University. The Maxion Text Prompter (MTP) (as used in [46]) is a locally installed logging program that

stores precise keystroke information including accurate timings of when a key is activated and released. A series of prompts, the display of which is modified through a configuration file, are shown on screen and the user is required to enter the correct text string in order to progress. Using this software allowed me to assess performance in terms of the time taken to enter the passwords correctly and the number and type of errors made.

The MTP configuration files allowed close control over both what is displayed on the screen for the participants and the conditions under which the user can progress to the next screen. This can vary from a requirement to enter the prompted text exactly or to allow any text string as the correct answer. This allows MTP to be used to store subject data such as their name as well as for distraction tasks such as general knowledge questions where a correct answer is not a necessary part of the experiment. An example configuration file can be found in Appendix A.

Effort data was gathered through the NASA Task Load Index (TLX) [52]. This is a subjective workload assessment that derives an overall workload score based on a weighted average of ratings on six subscales: Mental Demands, Physical Demands, Temporal Demands, Own Performance, Effort and Frustration. The definitions given in Table 1 below were used during the experiment and were given to the subjects on a references sheet:

A software version of the TLX procedure was downloaded and locally installed for use during the study. The TLX scoring systems is based on two separate tasks, a *ratings tasks* and a *weights task.* This terminology is taken from the NASA TLX documentation and was used here in the same fashion, despite being arguably somewhat counterintuitive. This is because the descriptions of some of the factors do not seem to apply to the process of entering a password. The description for physical effort, *"How much physical activity was required (e.g., pushing, pulling, turning, controlling, activating, etc.)?"* is one example. The *weights task* is presented to the participants first: participants are asked to evaluate the contribution of each of the above factors (mental load etc.) to the workload of the task. They are presented with pairwise comparisons of the six scales, and asked to select the pair that contributes most to the workload of the task. This is repeated for each of the 15 possible comparisons of the scales. A score between 0 (not relevant) and 5 (most relevant) is then awarded to each factor representing a tally of how many times they were selected during the pairwise comparisons. The *ratings task* is conducted second: participants are asked to assign numerical ratings to each scale that reflect the magnitude of that factor in the task. Participants are presented with scale graduated into 20 equal intervals and marked 'high' and 'low' at the two end points. Responses are marked on the scales by the participants at the point they feel represents the magnitude of the task.

Title	Endpoints	Descriptions
MENTAL DEMAND	*Low/High*	How much mental and perceptual activity was required (e.g., thinking, deciding, calculating, remembering, looking, searching, etc.)? Was the task easy or demanding, simple or complex, exacting or forgiving?
PHYSICAL DEMAND	*Low/High*	How much physical activity was required (e.g., pushing, pulling, turning, controlling, activating, etc.)? Was the task easy or demanding, slow or brisk, slack or strenuous, restful or laborious?
TEMPORAL DEMAND	*Low/High*	How much time pressure did you feel due to the rate or pace at which the tasks or task elements occurred? Was the pace slow and leisurely or rapid and frantic?
EFFORT	*Low/High*	How hard did you have to work (mentally and physically) to accomplish your level of performance?
PERFORMANCE	*Good/Poor*	How successful do you think you were in accomplishing the goals of the task set by the experimenter (or yourself)? How satisfied were you with your performance in accomplishing these goals?
FRUSTRATION LEVEL	*Low/High*	How insecure, discouraged, irritated, stressed and annoyed versus secure, gratified, content, relaxed and complacent did you feel during the task?

Table 1: *Definitions of the subscales used in the NASA TLX Weights task*

The overall workload is computed by multiplying each rating by the weight given to that factor by the participant. The sum of the weighted ratings for each task is divided by the sum of the weights (15) to deliver a single final score for workload.

The TLX task was included as both a distraction task from the primary password entry tasks, and to measure how participants perceived the effort involved in using each password policy. The task was used without modification, as developed by the authors and presented in [28, 29] (With hindsight, this was a mistake, as will be discussed in Section 5.3).

5.2.1.4 Experimental procedure

In total, 21 participants successfully completed the experiment. Once they had registered two passwords that conformed to the policies in Section 4.1.1.2 participants were invited to attend a laboratory session in which they would use the registered passwords through interaction with the MTP software. They were asked to ensure they could recall the passwords without memory aids before they attended the session. This section of the study involved 7 phases. These were as follows:

1. Repeated entry of the policy 1 password: In order to familiarise the subjects with this password they were asked to enter the policy 1 password correctly 10 times. Incorrect entry attempts were logged by the MTP software but did not count toward the total. Once 10 successful entries had been achieved they subjects were moved on to the next phase.

2. Repeated entry of the policy 2 password: This phase functioned in the same way as phase 1 except with the second password. The order in which the subjects were presented with these tasks was varied systematically through the study so half the subjects were asked for password 1 then 2, and half for password 2 then 1.

3. General knowledge distraction task: Here subjects were asked to alternately enter the two passwords, with a general knowledge question presented after each password entry. An accurate password entry was required to move on to the question but it was not required to correctly answer the general knowledge question to proceed. The general knowledge questions were multiple choice and were taken from a Guardian newspaper article that gave the questions asked in a game of 'Who Wants to be a Millionaire'. There were 10 questions in total so each password was entered 5 times.

4. Policy 1 password entry followed by NASA TLX task: Subjects were asked to enter their policy 1 password then the MTP software was minimised and subjects were

shown the TLX software. They then entered their effort ratings for this password by answering comparative questions posed by the TLX system.

5. Policy 2 password entry followed by NASA TLX task: This phase followed the same procedure as Phase 4 except for the second password. Again the order in which participants were presented with these tasks was systematically varied between subjects.

6. Repeated entry of the policy 1 password: Mirroring phase 1 subjects were again asked to correctly enter their policy 1 password correctly 10 times. Whereas the phase 1 was seen as a familiarisation task this phase was intended to be a recall task that would more closely mirror the performance with a password that was already in regular use by the user.

7. Repeated entry of the policy 2 password. Identical to Phase 6 except with the second password. As before the order was systematically varied between subjects.

In all cases the MTP logged the timings and content of the keystrokes used by the participants. For each keystroke a 'make' event is recorded when the key is pressed and a 'break' event when it is released. Where a key is held down multiple sequential 'make' events are recorded.

5.2.2 Experimental results

5.2.2.1 Timings

The time taken to enter each password was determined by calculating the difference between the timestamp for the 'make' event for the first character and the 'break' event for the last character in the entered string. Incorrect password strings were included in the timing data. Due to the varying length of the passwords chosen by the participants, it would be meaningless to compare the total time to enter each password. Therefore the total entry time was divided by the number of characters to generate a per-character entry time. Although typing speeds varied greatly between participants I assumed that typing speed was constant within each participant's data and therefore typing speed would not affect the overall results. The following table is a summary of the timing data (all figures to 2 d.p.):

	Policy 1 character entry time (s)	Policy 2 character entry times (s)
Repetition Task 1	0.76	0.76
Distraction Task	0.75	0.75
Repetition Task 2	0.54	0.52

Table 2: *Per-character entry time data*

The data clearly shows that there is no appreciable difference in entry times, despite the differing requirements of the password policies. It is worth noting here that - although the policies had different composition criteria - many participants ended up selecting passwords according to a pre-existing generation scheme (following a favourite quote or film title for example), meaning that 25 out of the 42 enrolled passwords would have satisfied both policies, being both 8 characters and containing 3 different character types. This would also contribute to the similarities in timing for the two policies.

5.2.2.2 Errors and false attempt rates

When entering passwords there are two forms of mistake that are recorded by the MTP software:

> *Errors:* I define an error as an instance of the participant using the delete or backspace key in order to attempt to correct a mistake in the typing. The frequency of such events is determined by counting the 'make' events for the backspace and delete keys in the log files. Successful login attempts can therefore still contain errors as the backspace key is used to remove typographical mistakes before correct characters are entered.
>
> *False Attempt:* I define this as an instance of the participant pressing the enter key with an incorrect string in the field thereby attempting to submit a password and failing. False attempts can contain errors where the correction attempt was unsuccessful.

The mean frequency for such mistakes was calculated for each policy and task; the results are show in Figures 8 and 9 below:

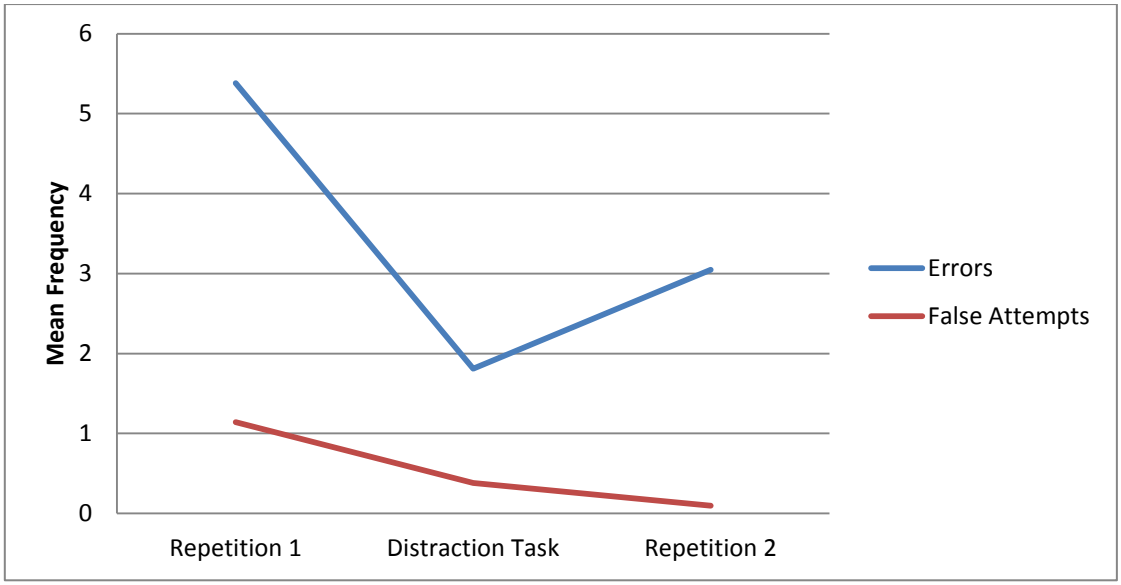

Figure 8: *Errors and False Attempts for Policy 1*

In both graphs, the patterns of errors are closely related. This implies that although errors are being recorded (i.e. participants are attempting to correcting typing mistakes) this process does not result in the elimination of incorrect strings being submitted. The exception to this is the repetition task for both Policy 1 and 2, both of which show relatively high error rates (and in the case of Policy 1 even an increase in errors compared to the distraction task) but almost 0 false attempts, implying that the correction process was successful in this task.

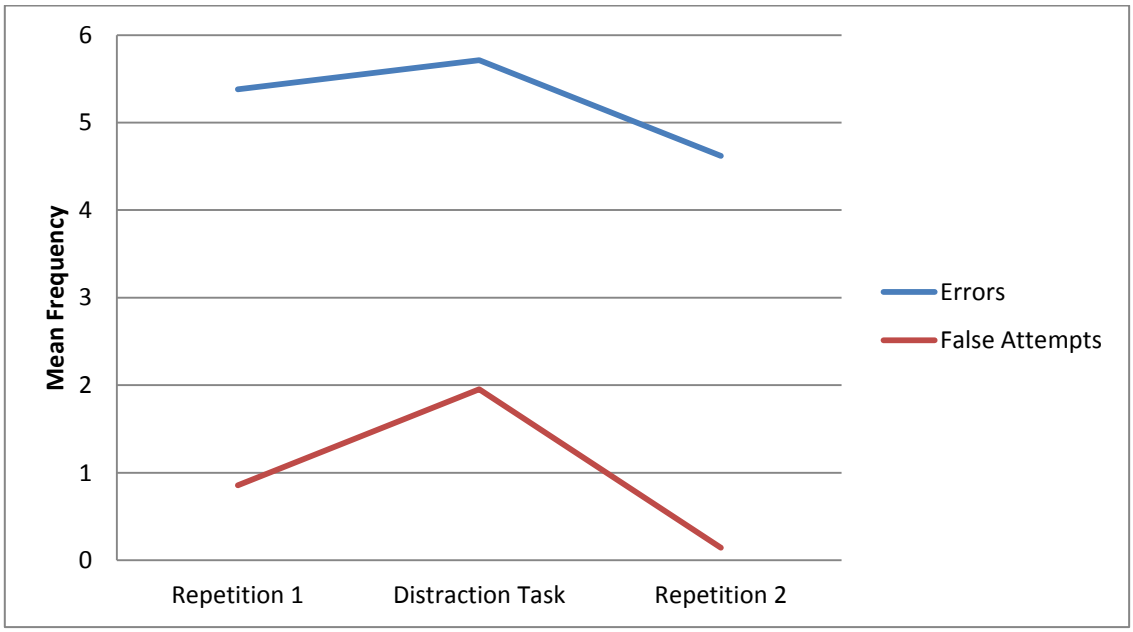

Figure 9: *Errors and False Attempts for Policy 2*

The difference in shape of these graphs indicates that while the per-character entry times for the policies were almost identical, the errors they generate are not. This shows that there is a difference in usability between the two policies. While initial error and false attempt rates are similar the effect of practise and distraction varied between the policies. By the end of the second repetition task both policies had achieved almost 0 false attempts this was at the cost of more errors for Policy 2. As this policy required longer (but potentially simpler) passwords this result implies that increasing the required length of the password creates more user errors than adding to the complexity requirement. This is especially noticeable with respect to the distraction task which for Policy 2 showed an increase in both error and false attempt rates after the first repetition task. Despite having a higher complexity requirement Policy 1 showed a drop in false attempt rate with practise, even during the distraction task. These results indicate that longer passwords receive less benefit from practise and are more vulnerable to errors caused by distraction.

However, as noted in Subsection 5.2.2.1 there was a tendency among participants to over-fulfil the policy requirements. This adds some confusion to the results and conclusions outlined above and deserves further attention. Due to this over-fulfilling the results do not indicate what the performance would be if participants followed the policy instructions to the letter. While it would have been possible to take greater control of the password enrolment process, for example by checking the enrolled passwords and rejecting those that over-fulfilled and asking the participant to generate a new password, I decided that it would be more interesting to take the opportunity to gather data on password creation. The results clearly illustrate that users do not follow password policies to the letter, but instead try to comply in a fashion that is workable to them.

Although we asked participants not to reuse passwords that were already in use there was nothing to stop them registering passwords that followed the same generation scheme that they already employed. Anecdotally, participants reported that the greatest difficulty in enrolling a password came from situations in which the policy they were trying to follow did not allow them to follow their usual generation scheme, some of which were quite creative. One participant registered the password 'TB's(x_x)Br', created from 'Tim Burton's The Corpse Bride'. Here the use of simple ascii art to represent the corpse creates a memorable, and strong password. While our policies allowed such creations the subject in question reported that not all policies do, and encountering one such always required extra effort.

So how does over-fulfilment affect our analysis? The two policies required either 6 or 8 characters and 2 or 3 different character sets to be used. In order to consider whether or not a password over-fulfils a policy it is necessary to characterise the passwords using the following criteria:

Short: The password is less than 8 characters in length
Long: The password is 8 or more characters in length
Simple: The password uses 2 different characters sets
Complex: The password uses 3 or more different character sets

From this list of criteria we can create three categories of passwords:

Short/Complex: A password that fulfils, but does not exceed the requirements for Policy 1. The lack of 8 characters makes it ineligible for Policy 2.
Long/Simple: A password that fulfils, but does not exceed the requirements for Policy 2. The lack of complexity makes it ineligible for Policy 1.
Long/Complex: A password that over-fulfils both Policy 1 and 2, satisfying both the complexity requirement of Policy 1 and the 8 character requirement of Policy 2.

A *Short/Simple* password would not satisfy the requirement for either policy and so has not been included in the list. Examples of each of these password categories were found in the data set. The table below shows the distribution of passwords across the two policies.

Policy/Password Category	Short/Complex	Long/Simple	Long/Complex
Policy 1	10	N/A	11
Policy 2	N/A	6	15

Table 3: *Distribution of password categories by policy*

As can be seen in Table 3, 26 out of the 42 passwords, or 62%, were *Long/Complex* and therefore over-fulfilled both policies. That the majority of participants decided to over-fulfil suggests that password policy may not be the driving force behind the participant's password creation process. This is supported anecdotally by conversations conducted after the experiment with the participants. Several participants reported that they had a scheme or method for creating new passwords (often based on famous people or meaningful phrases) and policies that prevented them using their personal methods were the ones they found hardest to use.

Policy 2 was more over-fulfilled, with 71% of passwords being *Long/Complex*, compared to 52% for Policy 1. This implies that participants were more reluctant to add length than they were to add complexity. We can also analyse the results graphically.

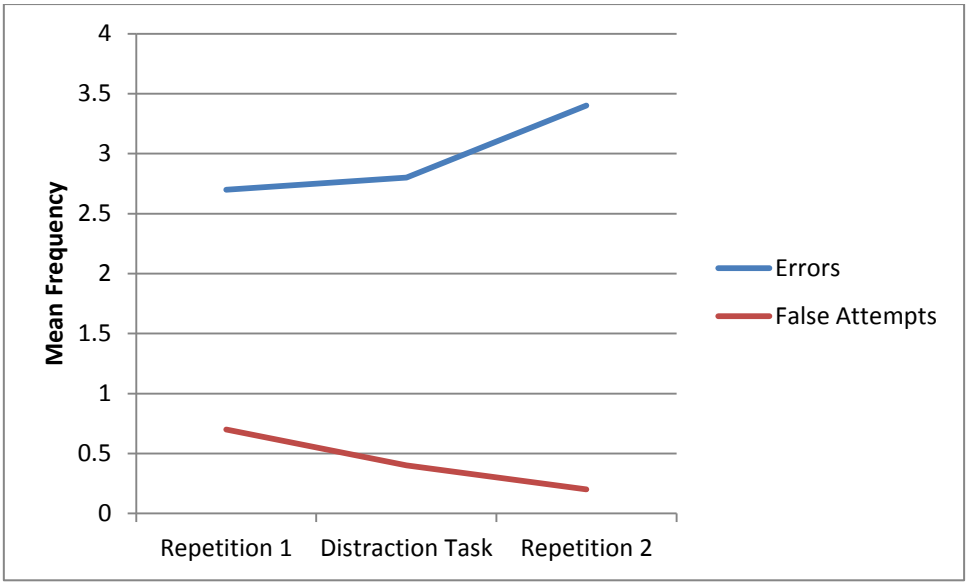

Figure 10: *Errors and False Attempts for Short/Complex passwords*

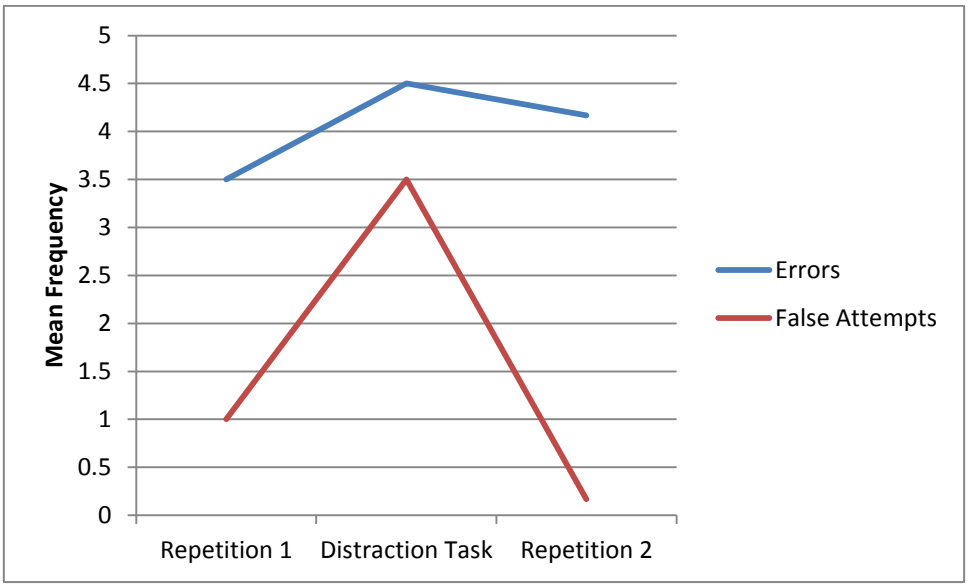

Figure 11: *Errors and False Attempts for Long/Simple passwords*

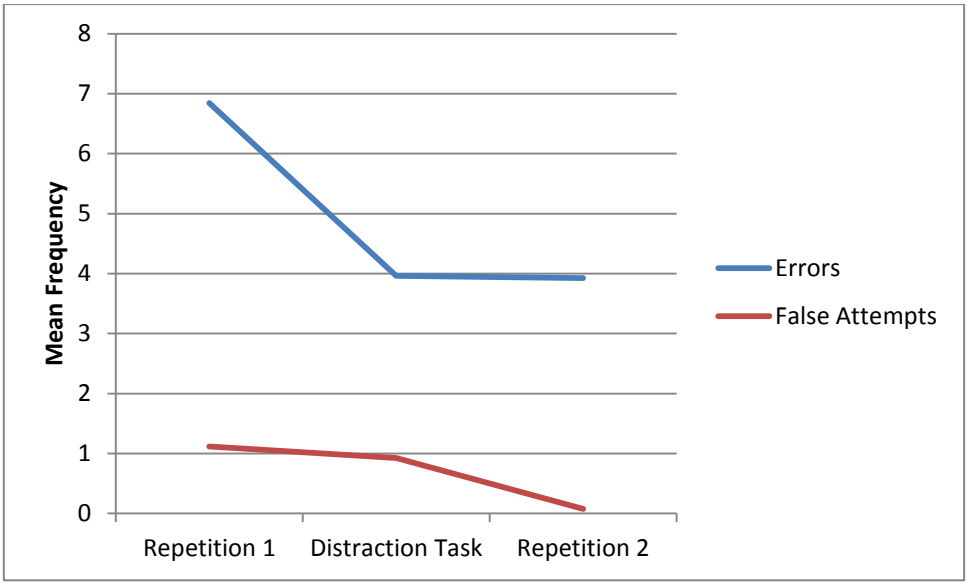

Figure 12: *Errors and False Attempts for Long/Complex passwords*

The relatively low number of data points for the *Short/Complex* and *Long/Simple* categories (10 and 6 respectively) gives a low confidence in the conclusions drawn from these graphs. In particular a single participant in the *Long/Simple* condition was responsible for 70% of the errors on the distraction task. It is worth noting, however, that all three password conditions ended up with similar frequencies of errors and false attempts by the end of Repetition task 2. In support of the previously drawn conclusions the *Short/Complex* condition was generating fewer errors, and comparable numbers of false attempts, than either of the longer conditions. This adds weight to the conclusion that adding length reduces performance more significantly than adding complexity. However, further experimentation and analysis with larger numbers of participants would be required to draw any firmer conclusion.

5.3 Server-based authentication experiments

To investigate Factor 6 – how password performance changes with frequency of use – a different approach was needed; one that did not require participants to visit a laboratory. In order to gather data relevant to actual password applications, a much longer time frame would be required than a laboratory session was capable of offering. To ask participants to return to the laboratory multiple times over a period of weeks would have been impractical in terms of both the logistics for both the experimenter and the participant and the budget required to remunerate participants on this basis. Additionally, I wanted to move away from the artificial nature of laboratory experiments. While such studies allow the experimenter to tightly control the environment in which the study takes place they do have drawbacks that mean that performance results gained under such conditions should be treated with caution, as discussed in Sections 3.1 and 3.2. Foremost among these is that the environmental

conditions in which the study takes place are markedly different from the norm, lacking the distractions that characterise a normal working environment, such as the demands of colleagues, email, telephones and the requirement to complete a primary production task.

To gain data that is both sufficiently objective and gathered under real-world conditions I conducted a semi-controlled field trial. In order to investigate how password performance changes with frequency of use, I needed record repeated password performance over time. I also wanted the participants in these experiments to be able to perform the tasks in their normal working context, and for performance data to be directly recorded, rather than subjectively through user diaries (See Section 3.2). The solution to these problems was developed as part of an MSc student project in which I was involved in an advisory role. The result of the project was a piece of software entitled the Authentication Performance Evaluation Tool (APET).

5.3.1 The APET software

Full credit for development of this tool goes to Alexandros Trepeklis, a former M.Sc. student at UCL), who worked under my supervision, and to a requirements specification and conceptual design drawn up by me.

APET is a web-based system that allows authentication experiments to be set up and managed remotely. Participants can enrol and take part in authentication trials over any Internet connection. The tool is split into two major components, the core system and the authentication plugins.

5.3.1.1 The Core

At the heart of APET is an experiment management system tied to a database of participants. A researcher can log into the management side of the software and set up an experiment. At this stage they can specify its duration, the authentication mechanism involved in the experiment (from the current set of plugins), the attributes they wish to log (from a selection specified by the authentication plugin, see below) and the participant groups they wish to be associated with the experiment. New participants can be added at this stage, by submitting their email addresses through a CSV file. Alternatively, the database of previous participants can be searched and filtered by a variety of criteria - such as age and gender - or by more experiment-specific criteria - such as which studies of authentication mechanisms they have previously participated in. Multiple groups can be assigned to each experiment. This allows different conditions to be applied to the same mechanism, without having to create large numbers of experiments.

Once an experiment is active, the researcher can email the participants of their experiments. Drop-down menus allow the experiment and participant group to be selected and an (editable) message is then sent to the specified group containing a link to the APET system. It is via this link that the participants themselves interact with APET. Using any web browser from any location, they can follow the link to a page that will ask them for their email address. Once they have confirmed their address (and thus that they are taking part in the correct experiment) they will be shown a screen containing whichever authentication mechanism is being tested. Their performance is then logged to a data file that the experimenter can download at the end of the experiment. The participant's interaction with APET is largely controlled by the authentication plugin.

5.3.1.2 Authentication Plugin

Each authentication mechanism used in APET needs to have its own plugin. This allows the system to be continually expanded as new technologies are developed without necessarily rewriting the core code. The plugin controls what the participants see when the follow the link emailed to them. For example, the password plugin can be configured to allow or disallow password resets, reminders or any other feature the experimenter wishes. Additionally, the plugin specifies which logging options the experimenter has. The logging options can be added or removed during the set up phase. The function of the authentication mechanism cannot be modified at this time and so to change the functionality of the authentication procedure a new plugin would need to be created. Typical logging options are such attributes as the number of attempts the participant took to successfully authenticate, whether they requested a reset or reminder, and the time and date of the authentication attempt. Any attribute can be logged assuming it has been coded into the plugin.

5.3.2 Experimental design

Using the APET software with a password plugin allowed me to design an experiment in which participants could enrol a password and login to the server when asked from any remote location when prompted by an email from the APET system. This removed the requirement for participants to attend laboratory sessions and returned the process to its more accurate status as a secondary task. Details of their login session were then recorded by APET removing the potential inaccuracies resulting from user reporting of data.

15 participants were recruited for the study. The group included both males and females in the age range 25-50, with a variety of occupations such as administrator, web developer and graphic designer. To demonstrate the fully remote nature of the experimental process two of the participants were located in America. I never met any of the participants face-to-face –

the entire process of recruitment and task completion was conducted via APET. This study was split into three phases – an initial enrolment task followed by two experimental sessions, the first of two weeks and the second of one. These sessions were separated by a six week gap. The participants were paid £20 for each successfully completed session.

In the enrolment phase the participants were sent a set of instructions (which can be found in full in Appendix B) and asked to generate a password that conformed to the following simple policy:

- Be at least 6 characters in length.
- Contain at least 3 of the following categories
    - Upper case letters (A, B, C, . Z)
    - Lower case letters (a, b, c, . z)
    - Numbers (0, 1, 2, . 9)
    - Non-alphanumeric characters (! + ; - % etc.)

This policy is identical to the first policy used in the laboratory experiment. As part of the enrolment phase participants also had to successfully complete the login task to ensure that they had correctly entered and remembered their password. This data was not used in the final analysis.

In the first phase of the experiment participants were send an email once a day every working day (Monday to Friday) asking them to login to APET. The daily timing of the email varied through the two weeks but always took place within normal working hours to ensure that the participants were engaged in their primary task. This was to avoid any sense of routine or expectation forming in the minds of the participants regarding they would next have to login. Email distribution was under the manual control of the experimenter (as opposed to being an automated part of the APET system) so the precise timing of the emails can be managed as needed.

During the second phase the participants were sent emails asking them to login three times a week (Monday, Wednesday and Friday). The purpose of this second phase was to examine the recall rate of a password that was previously well embedded but that hadn't been used for a relatively long period of time. Additionally, I wanted to see if the reduced frequency of use affected recall success rate.

In each case the number of attempts taken to successfully log in and the number of password resets requested were logged to form the data set for the experiment. For reasons of privacy and security the passwords used by the participants were not recorded.

5.3.3 Server study phase 1 results

One participant made no attempt to recall their password and simply reset their password on each of the 10 days of the experiment. This constituted a form of sanctioned non-compliance, in that participants were told in their instructions that *"If you cannot remember your password there is an option to reset it. Doing so is perfectly fine and should not be considered negative in anyway"*. The intent of this part of the instructions was that participants should not feel that being unable to recall was a failure. It was far more preferable for them to complete the task using a reset than to disengage from the task altogether. However, I did not anticipate that a participant would use the reset option instead of recall completely.

There was no sanction associated this form of non-compliance with the task. Participants were incentivised to log-in each day, not to successfully recall their password. The only other form of non-compliance available to participants was to not make an attempt to authenticate at all. However, participants were only remunerated if they participated in all the sessions, making this option unproductive. It was not possible for participants to enrol with a password that did not fulfil the policy correctly as the passwords were monitored by the experimenter before the sessions started.

The participant that only used resets created an interesting set of data for discussion. However I decided not to include them in the numerical analysis, as the data set is uniform – the participant contributed a single reset each day. Additionally, the experiment was designed to test user performance with passwords, including the recall phase, and so this participant does not provide relevant data for that aspect of the study. For these reasons I removed the participant from the data set for the following analysis. This participant again illustrates that users will interpret the rules in a way that works for them, and not necessarily as the policy author intended.

The remaining 16 participants all appeared to make genuine attempts to complete the task, although not all participants responded on all of the 10 days. This information is show in Table 4.

Days responded	Number of participants
10	11
9	2
8	1
7	1
6	1

Table 4: *Server study participant response rates*

The mean number of attempts used each day to successfully log in is shown in the Figure 13. Excluding the participant the used the reset function instead of recall only 1 password reset was requested during the 10 days of Phase 1 of this experiment. This reset was requested on the first day of the experiment after the participant had attempted to login 9 times. This is the cause of the higher average on Day 1. In all other cases once the password was successfully enrolled most participants successfully logged in on their first attempt. In 9 out of the 159 recorded login events the participant took 2 attempts to log in.

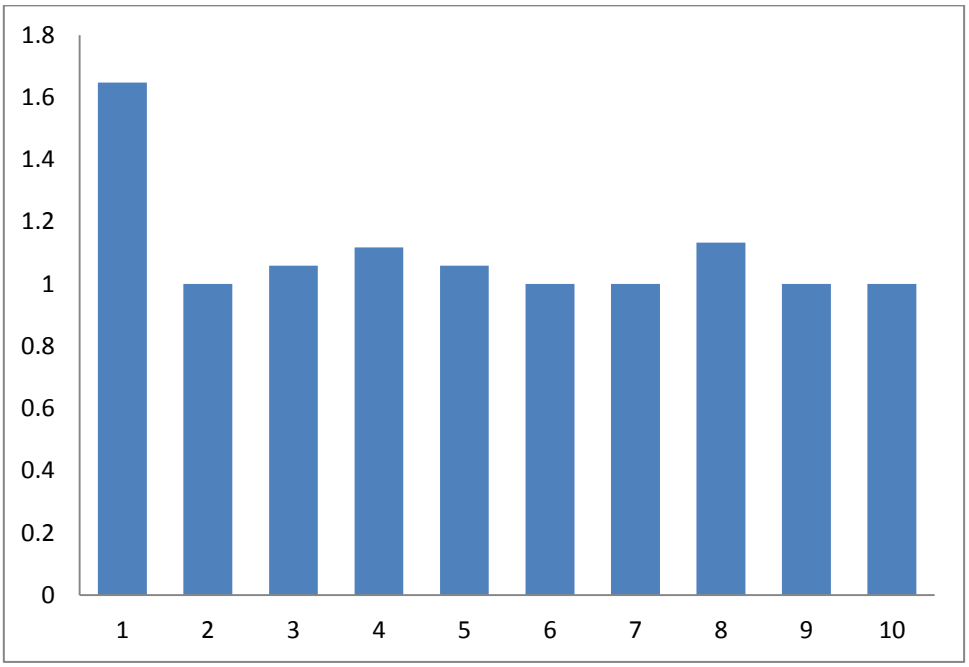

Figure 13: *Server study Phase 1 log in attempts*

5.3.4 Server Study Phase 2 results

In this phase again one participant (the same as in Phase 1) simply used the reset function rather than recall in order to complete the task. Their results are not included in the summary. In this phase all participants bar one responded on all three days. The chart below shows the mean number of attempts required to successfully log in on each day.

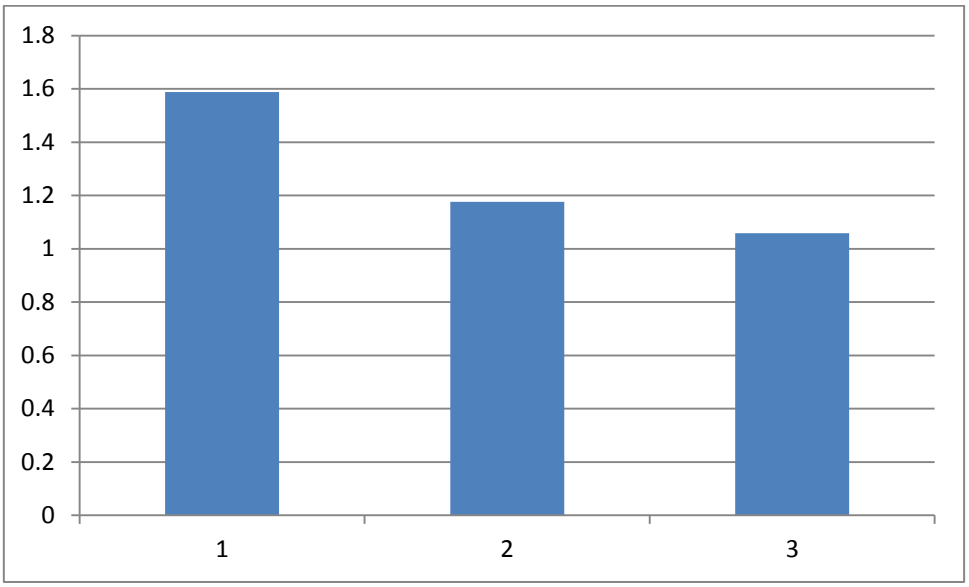

Figure 14: *Server study Phase 2 log in attempts*

Again the first day of use had a higher value for attempts required. In this phase, however, this was the result of more than one participant using several attempts to log in. One participant made 6 attempts then reset their password. Three other participants used 2, 3 and 4 attempts to login successfully without resetting.

The most notable difference between the phases was the use of resets to replace recall. Excluding the participant that used the reset exclusively through both phases there were 7 resets requested in Phase 2 spread across the three days, with a varying number of attempts made before each reset. The information relating to resets is summarised in Table 5.

It is interesting to note that relatively few login attempts are made before each reset and that there the number of attempts used before reset decreased as the experimental phase progressed.

Day	Attempts before reset

1	6
1	2
2	3
2	1
3	2
3	2
3	1

Table 5: Server study phase 2 password reset data.

5.4 Discussion

As stated previously in Section 3.6 the approach taken throughout this thesis is to combine different sources of information to form a coherent whole. As such, it is through considering the results of the laboratory and server studies together that a better understanding of how password composition affects user experience can be gained. The motivation behind collecting such data is to provide grounded empirical data for decision makers and modellers working in the field of information security.

The results of the password entry study showed that per-character typing speed was constant across different policies and password compositions. This means that password composition has no meaningful effect on entry time and that the length of the password will be the main contributing factor to the basic mechanical process of entry. However, while the per-character entry times showed that physical effort associated with passwords increases linearly with length, and be unaffected by composition, the usability of the password overall is affected by both frequency of use and composition. The password entry study also showed that longer passwords produce more typing errors. This results in an increased need for participants to correct their entry using the 'backspace' key in order to achieve the same level of successful login attempts. This increases the effort required by the participants.

In addition, 62% of participants in this study over-fulfilled the requirements of the policies, choosing to use passwords that were both 8 characters or longer and contained characters from 3 or more character sets. Working on the assumption that users will not deliberately choose a password that is difficult to use, this implies that user effort is linked to more than just the raw characteristics of password composition. In this case it was anecdotally reported by several participants that they used a personal algorithm for memorable password composition. Situations in which they were unable to use this method, for example where a policy places both a maximum and minimum length on passwords, were reported to be the most difficult to use. It would require a study focused on composition, rather than

performance use to explore this behaviour in more detail, but it is worth nothing that users do not necessarily minimally fulfil password policy requirements. This highlights that the situation is more complicated than simply saying short passwords are easier to use than long passwords.

During this study data was also gathered using the NASA TLX system. Unfortunately many of the participants felt that the definitions of the subscales did not really apply to the task in hand. This led to uncertainty when it came to assigning scores. Additionally, the participants were confused as to whether they were supposed to be rating the task with respect to the other password or against a more general idea of what a difficult task was in terms of the subscale in question. This made the TLX task a useful distraction tool as the participants were clearly not focused on password recall while they were engaged in it, but negates the usefulness of the data gathered for assessing task load. As such, it has not been included in this analysis.

The server study, which used Policy 1 and required only a password length of 6 characters, demonstrated that when used daily passwords are usually entered successfully on the first or second attempts. This agrees with the results of the laboratory study that showed that attempting to enter a false password happens with a frequency of between 0 and 1 per 10 successful entries. The exception to this was the performance of the Policy 2 passwords during the distraction task which peaked at 2 failure attempts per 6 successful entries. However, the server study also showed us that as frequency of use decreased, the use of resets increased, and that resets were requested after only 1 or 2 entry attempts. These results imply a pattern of user behaviour in which they will attempt to correct known errors in their password entry manually before submitting the password but that they will expend only a small amount of time and effort on repeated attempts if their first entry is not accurate. Instead of trying multiple times to successfully recall and enter their password they will simply reset it. The trigger point for password resets is therefore low in terms of user effort.

One participant used the reset function, rather than attempting to recall the password. They reset their password for each session and simply entered the new one. While this participant did not provide data on password performance they do highlight a relevant point, with respect to password use. Where systems allow a reset, the cost of a reset may be lower than recall. This is more likely to be the case with systems that are used infrequently as the recall effort is higher. If an organisation is experiencing high levels of password resets, when compared to recall and entry, this may be a reason. An alternative form of authentication may be preferable in such cases as password resets carry a risk of interception or abuse by malicious parties.

To put this discussion in the context of decision making, the information drawn from the experiments outlined in Chapter 5 can be used to inform password policy creation. It is important to know the context of use for the password, in particular how often it is likely to be used. For passwords that are used only a few times a week, or less, password policies that require 8 characters or more are likely to generate sufficient errors (and therefore user effort in terms of correction and failed entries) to push users towards resetting instead of recall, given their low tolerance for failure. Such a policy would therefore require more investment in IT support services to manage the increased volume of reset requests, or the utilisation of self-service resets which transfers the effort back to the users creating knock on effects elsewhere, as described by the Compliance Budget (Section 4.5). Resets also offer attackers increased opportunities for password capture and interception. In terms of increasing user effort, the results of the laboratory study imply that length has a greater impact on error rate than complexity and so where increased security is required, mandating more complex passwords will produce fewer errors than longer passwords. Passwords that are used daily, or more often, are less prone to this affect as the frequency of use compensates for the difficulty of successful entry and password resets are much less common.

This knowledge allows decision makers to decide if they wish to spend their resources on IT support infrastructure to cope with reset requests or on other forms of protection for assets that may be left vulnerable by shorter but more usable passwords. While this trade-off has been known to exist for many years data such as this allows the trade-off to be managed in a much more systematic fashion, with the likely impact (in terms of errors and resets) of changing password policy anticipated with greater success and fidelity. This type of information is critical if the understanding provided by the Compliance Budget (see Sections 4.7 and 4.8 for an outline of the advantages of security management through the Compliance Budget) is to be leveraged into a systematic decision making framework. The Compliance Budget tells us that improving security design to be more understanding of the user effort associated with security controls is one of the four primary methods for improving the likelihood of positive compliance decisions. In order to do this effectively the manner in which various controls impact the user must be studied and understood. Chapter 5 provides initial answers to the question of password composition and therefore functions as a sub-component of the Compliance Budget. Where the Compliance Budget outlines a model of user decision making with respect to compliance, Chapter 5 offers the information required to use that model to resolve a specific security management question.

The two studies offer an example of the data collection that needs to be carried out to inform decision making, and create a comprehensive understanding of this area. The two policies that were compared traded off increasing the length of the password by 2 characters with the requirement to have one less character type in the composition and the results and discussion are based on this single trade-off. In order to support a wider range of decision

making a series of studies that systematically varies password length and composition under laboratory conditions to assess error rates is needed. These studies would need to be then mirrored through live server studies in order understand how use patterns affect the results. As more data is added to the data set so to can a broader range of decision be supported. These studies offer an early contribution to this area and demonstrate both the usefulness of such results and the need for further work.

6. System design

The review of the existing literature (see Section 2.2 and 2.4.4) has revealed that it is a factor that is often neglected, or its costs not fully represented, in the field of information security. This thesis has suggested ways in which the human factor can be more fully represented, through the Compliance Budget model of user decision making, the details of which can be found in Section 4.5. The Compliance Budget provided the motivation for the empirical studies conducted and reported on in Chapter 5. The data gathered by these studies demonstrates types of data that must be collected in order for security managers to use the model as a tool for supporting decision making. However, the human factor is just one of many factors that security managers must consider. The purpose of the Compliance Budget is not to promote human factors as the single most important part of security, but rather to highlight their importance so they will be considered equally alongside technical mechanisms. It is also necessary to develop methods that allow this process of equal consideration to take place.

This section will address Research Question 3, which asks *"How can we include the human factor in a systematic approach to security decision making?"* Just as the Compliance Budget motivated the need for data gathering to support it, it provides the motivation to explore this question as well. It is necessary to recall at this stage that all the research in this thesis takes a utility-based approach to security decision making. This approach seeks to examine the trade-offs that exist behind all forms of decision making and provide security managers with the information required to make more systematic decisions. The Compliance Budget achieves this with respect to ensuring that user effort is spent as efficiently as possible on security tasks. However, this is only one facet of security decision making. Taking an approach to security management that focused primarily on human factors would produce results as far from optimal as those that marginalised it. What is required is a balanced approach where the human factor is included alongside other, more technical, factors to achieve the desired aims of the organisation as a whole. Without this process of inclusion the Compliance Budget cannot serve the goal of a systematic approach to decision making. There is a gap in the current state of the art, in that technical approaches do not incorporate an accurate understanding of human behaviour (where one is included at all), and HCI approaches focus too much on the usability of a system in isolation, ignoring the operational context in which it takes place. This neglects the primary purpose of security, that being to protect the primary operations of the organisation such that maximum productivity can be achieved.

The notion of the operational context is essential to a systematic approach to decision making. Compliance decisions of the sort addressed by the Compliance Budget do not take

place in isolation, nor are the issues raised during the interviews, and presented in Section 4.3.1, purely theoretical. More traditional methods of simulation and analysis, such as systems modelling, are more than capable of representing the technical mechanisms that exist in this context. In order to address Research Question 3 a framework is required that allows decision makers, and indeed systems modellers, to take the Compliance Budget model and include it in such tools. It is the combination of these approaches that offers the potential for creating a systematic approach to security decision making, which is the over-arching goal of this thesis. In pursuit of this goal I have developed a conceptual structure that allows for security objectives, human factors and technologies to be integrated in a framework. The key aim is to represent a security in its entirety so that the impacts security controls have on user behaviour, and vice versa, can be understood and explored. The intent is to formalise the understandings found in the Compliance Budget specifically, and in the utility-based approach to security in general, in a form that is accessible by security managers. If such decision makers are to make use of this information, and to take advantage of the insights in human factors that are now being produced, such a representation is essential.

In addition, this framework pays close attention to the requirements of system modelling. Having identified the practical uses of such tools (in Sections 2.4.3 and 3.4) the framework is structured in such a way as to facilitate an easy export of the data contained therein to a systems modelling environment. As such, conscious effort is made to represent the aspects of the security system in terms of processes, resources and locations. The whole structure sits within the same economic view of security that underpins and motivates this thesis in its entirety.

Where Chapter 5 provided a specific data set that allowed the Compliance Budget model to be used to solve a narrow question, the goal of this section is to provide a wider framework in which the Compliance Budget can be used to aid in understanding, and optimally managing, security at an organisational level. This sets the research presented here at a different level to the work in Sections 4 and 5, although still conceptually part of the same economically-motivated thinking, as show in Figure 15.

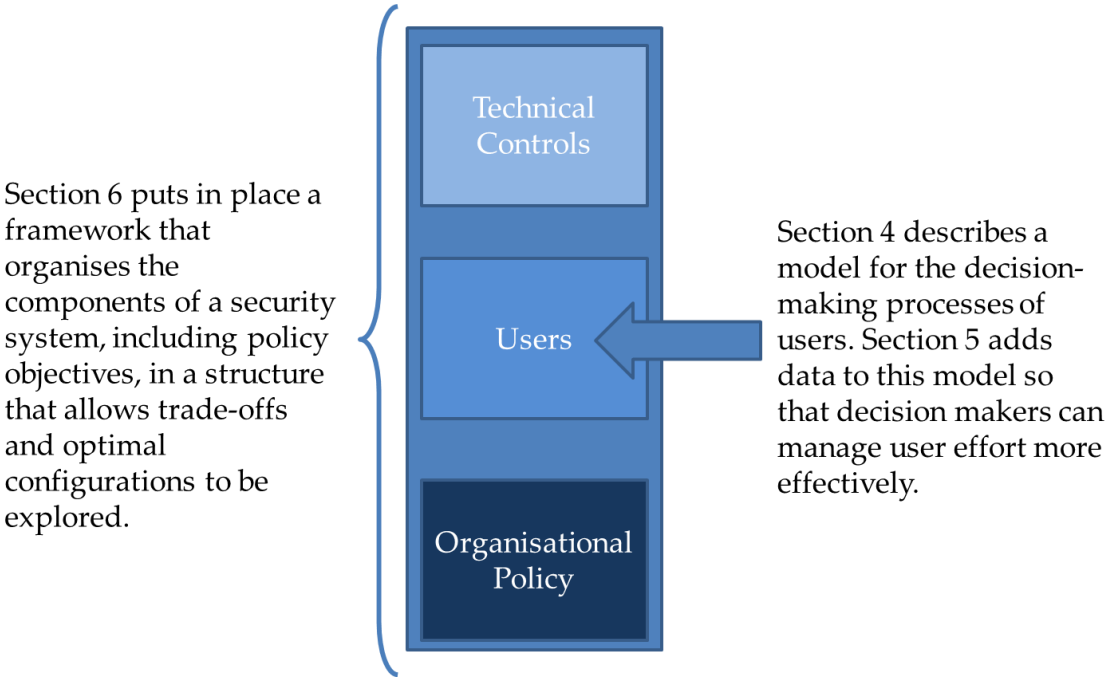

Figure 15: *Conceptual map of the research*

6.1 Background and motivation

As discussed in section 2.1 CIA, while providing a suitable framework for organising high level security objectives, lacks the range of concepts necessary for discussing a fuller implementation of information security. In particular, it does not lead to a systematic analysis of security requirements such as the level of confidentiality, availability or integrity required by a system. When taking an economic view of security we need to consider the trade-offs that are taking place, as security decisions are made by both organisations (as represented by their chief security managers) and the users of the system. These decisions can be at the management level, where policies are set, or at the individual level, where compliance decisions are made (as discussed throughout Section 4). In either case, the decision maker is trading-off between various options and projected outcomes and trying to maximise the organisation's utility, with respect to its business process. In the current state of the art the knowledge required to make such decisions is split between the HCI and security communities. Additionally, in the case of the individual user, non-security incentives, such as maximising productivity, the need to strike an appropriate work/life balance and the fundamental limitations of human cognition, drive behaviours that are insecure (although often without malice). A security decision maker attempting to systematically make optimal decisions is hampered by the lack of unification among these sources of information. To support decisions that move security in this direction an economic analysis of the situation, that accounts for all the costs and benefits associated with

a policy or mechanism, is required. The standard method of expressing security goals, the CIA framework, is unsuitable for this level of analysis due to its incompleteness and the high conceptual level of its definitions.

Furthermore, the analysis of the interviews (see Section 4.2) conducted as part of the Trust Economics project showing that the role an individual occupies within an organisation reflects their attitudes toward security, for example, their Compliance Threshold, with leaders and managers seeing more clearly the corporate need for security (although not always willing to follow it themselves). It is clear that any structure that seeks to represent security goals and objectives must be able to do so at all levels, and for all roles, within an organisation, while recognising that these goals and objectives will differ between the levels. CIA is too unwieldy a conceptual framework to provide the granularity need to work at these different levels. A new approach is needed if the human factors knowledge is to be integrated into a systematic methodology.

As was explained in Section 3.5, systems modelling can provide powerful tools that do allow such a systematic analysis. However, the strength of these models is dependent on the usefulness of the information used to tune them. Organisations already collect large quantities of data about their own activities that can be used for this purpose but it is not always organised or presented in such a way as to make this a natural process to follow. In addition, such models commonly abstract away from a notion of individual users and their behaviour. Any new approach suggested should also build on the principles of the economics of security discussed in Section 2.4. Thus an approach is needed which:

- Makes use of existing information but does not require organisations to gather large amounts of new data
- Organises this data in such a way as to allow a smooth transition into systems modelling
- Captures the tensions and trade-offs present in the system expresses them in terms of utility
- Allows a systematic analysis of the security system including the impact of human factors

To meet these requirements an ontological framework is proposed which integrates the security managers' preferences between investments to protect against confidentiality, integrity, and availability and the structure and dynamics of the system itself. The purpose of the ontological model is to represent security functionality in such a way that the model can be used as a tool to aid in both the design and subsequent audit of the security behaviour of an organization. It represents and expresses the security design and implementation choices of the organization in terms of a hierarchy of needs. It also aims to

reveal the tensions between the business and security processes, as well as the trade-offs that are taking place between the various security objectives of the organization.

This approach is not intended to function as a stand-alone technique, but as part of a security 'life cycle'; one link in a chain of technologies and processes that provide information security solutions. As methodologies such as HP's Security Analytics package become commercially available there is a need for techniques such as the one outlined here to be developed to facilitate the connectivity of such unified approaches to information security.

6.2 A running example

In that it is useful to have a more concrete example with which to illustrate abstract concepts such as these, one will be provided. While this thesis is largely concerned with information security, the principles outlined here are applicable to security more generally and so an example has been chosen that relates to both informational and physical aspects of security – the management of an airport's security process.

This is obviously a large and complicated problem, and I have chosen to focus on the set of processes involved with checking passengers and their luggage for their acceptability to fly. In order to be allowed to board the plane a passenger, with luggage, must navigate from the concourse of an airport's terminal building to a seat on an aircraft; that is, it is an access control process that is predicated on maintaining an integrity property of aircraft. This is achieved by maintaining that property for passengers and their baggage. This integrity property trades-off against other concerns, principally costs, incurred in providing security staff and equipment, and service availability. The passenger is subject to a range of security controls that are intended to ensure a certain integrity property—roughly, that certain dangerous or substances and objects are not present—of the aircraft. The manager of the security process has decided that, in order to access the aircraft, passengers must submit to the security process and must therefore sacrifice their confidentiality—the bags will be searched, there will be body searches. There are also non-trivial trade-offs between cost and availability, in that increasing investment in additional staff or scanning devices will reduce wait times at security checkpoints and improve the throughput of the system. The length of time a passenger must spend navigating the airport security procedures can be seen in this case as a measure of availability of access to the aircraft, with the extreme cases being unrestricted access to the plane on the one hand and security checks so lengthy and time consuming that passengers are unable to progress through them fast enough to make their scheduled departure time. These issues make airport security an interesting and relevant example that will be developed over the course of the following sections. In order to discuss

this, however, (and in keeping with the components required for systems modelling) a location map is needed. This is show below in Figure 16:

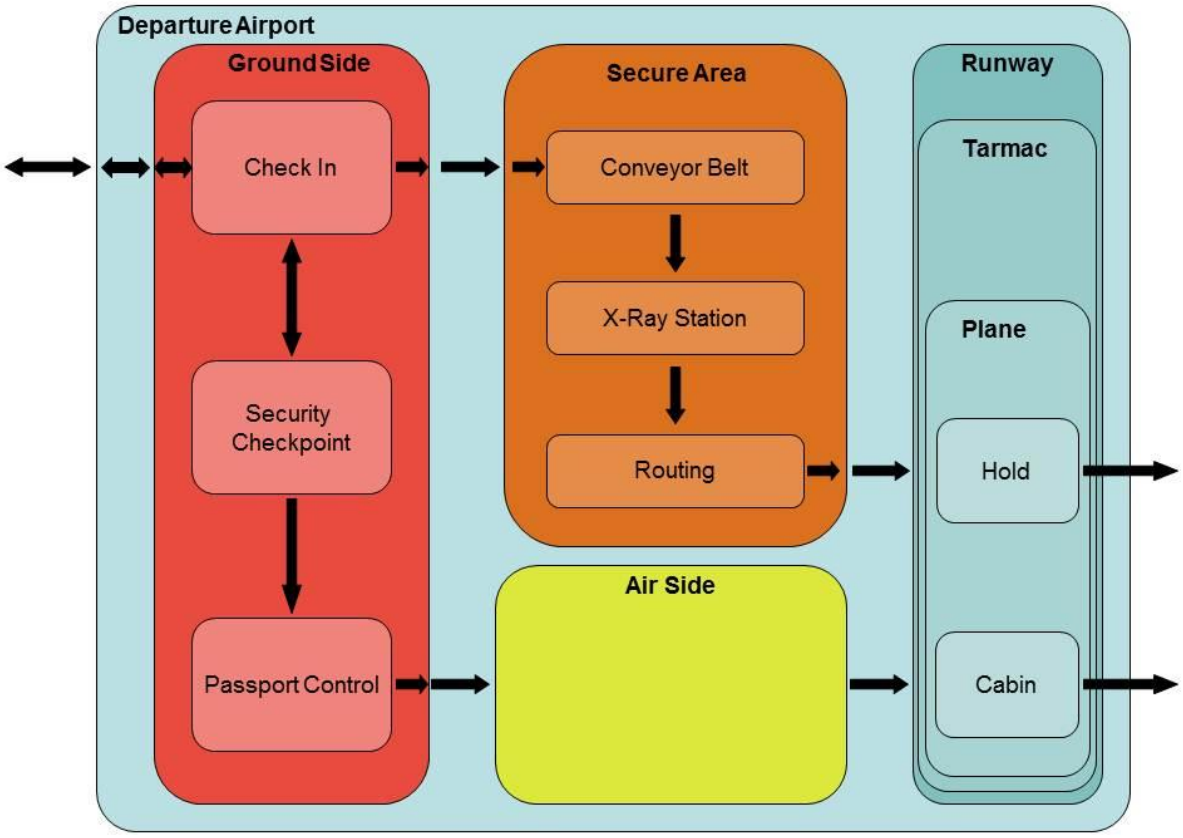

Figure 16: *Diagram showing the locations and associated connections and associated connections used in the Airport example*

6.3 The security architecture

This thesis is not the first work to identify gaps in the CIA framework and other authors have commented on its inadequacies. As discussed previously in Section 2.1 attempts have been made to suggest additions to CIA in order to rectify its perceived weaknesses. However, in many cases these additions suffer from category errors (see Section 2.1.3) and confuse operational and declarative concepts. The architecture of the approach suggested here is designed in such a way as to address this issue directly. In particular it is split into two sections, one containing the declarative and the other the operational components of the system. These are referred to as the *Framework Layer* and the *Instantiation Layer* respectively.

While the layers differ in terms of their internal components they share a commonality of structure across two key dimensions. 'Vertically' the framework is organised into a hierarchy of roles, with all relevant roles within the organisation being represented. Relevant in this case means roles that interact with the organisation's security system. This

serves two purposes. First, it reflects the findings from the Compliance Budget interviews that an individual's role within an organisation affects how they will interact with, and view, security systems. By creating a hierarchy of roles, within which security objectives and processes are represented, it becomes possible to differentiate the behaviour of the occupants of those rules with respect to the security mechanisms they encounter. More will be said on this topic in Section 6.3.2.2, which covers Actors.

Second, and as the term 'hierarchy' suggests, it allows an ordering to take place. The roles here are ordered by their ability to influence the security architecture of the system. In other words, they are classified by the toolbox that is available to them for modifying security objects (that characterize security tasks, defined in section 6.3.1). For example, the top level of the model might represent the strategic decision-makers of the organisation, such as an airport's security managers or their regulators, while the bottom level might represent an individual employee or user of the organisation, such as an airport's check-in staff or a passenger navigating airport security. The roles represent the possible positions individuals can adopt in the hierarchy. They do not represent any entity themselves. They are instead populated by Actors, which are another component in system and are described below. The roles represented in our airport example may include the airport security manager, security guards manning the various checkpoints, airline staff involved in check-in (where some security questions are asked) and the passengers seeking to board the plane. This ordering is necessary in order to logically construct the components of the security system. This will be explained in full in Sections 6.3.1 and 6.3.2.

This role structure is seen in each layer of the hierarchy but it is not the only area of commonality. Each hierarchy role or level is divided 'horizontally' into three sections representing the security dependencies, priorities and preferences of that level. The terms are defined as follows:

- *Dependencies* (strong requirement): Externally enforced requirements that Actors populating the role must meet (all of) in order to function within the model. Actors occupying this role have no choice in whether or not (and possibly even how) to meet these requirements regardless of how resource inefficient they are. Dependencies will often be informed by the environment within which the hierarchy exists (such as the Department for Transport regulations regarding random liquid testing);
- *Priorities* (weak requirement): Externally supplied tasks, as many as possible of which should be met by Actors in the associated role. Actors have some choice in which priorities to meet and how they are approached. In a limited resource environment, Actors can select the most resource efficient priorities and methods first. Priorities will often be informed by the role that the level represents (for

example while security guards have to search some items of hand luggage they will have some flexibility as to who and when is searched);

- *Preferences:* Actor-generated tasks that the Actor has decided are worth achieving from its own perspective. These can be generated by the Actor's inclusion in other hierarchies (passengers can for instance choose which security checkpoint they queue up for).

The purpose of this horizontal structure is to allow the Actors that populate the system to allocate resources according to the needs of their role. It also allows the system to be checked for functionality. If all the dependencies of the system are not being fulfilled then the system can be said to be failing – it will be in breach of the rules and regulations that apply to its area of operation. This will make the system illegal, and potentially dangerous. Where priorities are not being met the system can be said to be failing its primary process. The core objectives, or processes that would give the organisation a market edge, that the system designers would wish it to possess are not being achieved. Failure to meet preferences does not imply system failure in the same way. However, the Compliance Budget would indicate that a user base that is being forced to work in a way that does not meet its requirements may well be introducing costs elsewhere. A failure to meet preferences should therefore prompt a manager to assess the behaviour and compliance of the system's users more closely.

The overall structure of the hierarchy can be seen in Figure 17 below. This clearly shows the role-based nature of the hierarchy and the shared structure between the Framework and Instantiation layers. The nature and composition of these layers will now be discussed in more detail.

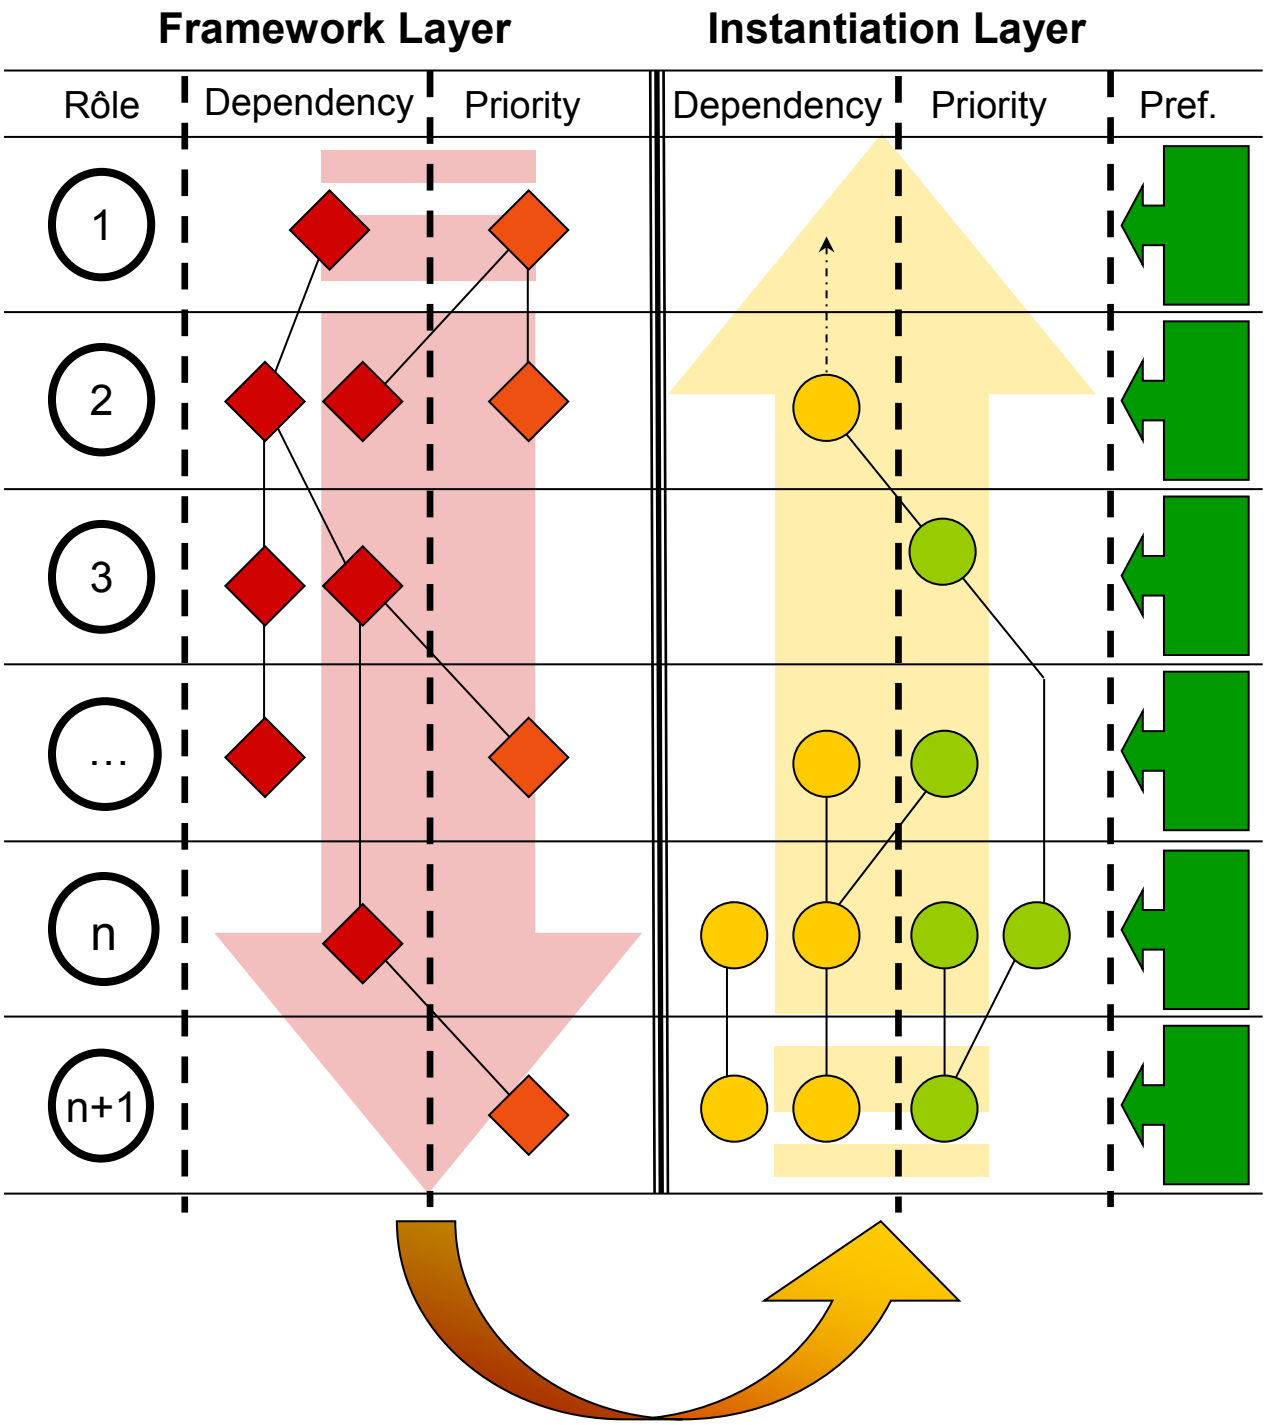

Figure 17: *Diagram showing the structure of the security hierarchy*

6.3.1 The Framework Layer

The Framework Layer represents the underlying structure of the system. It is static (in the sense that the processes and resources which would be represented in a systems model are not found here) and declarative but informs the construction of the more operational

Instantiation Layer. A completed Framework Layer consists of a hierarchy of roles with dependencies and priorities assigned to them. As preferences are derived from Actors (see Section 6.3.2.2) they do not appear in this layer, as Actors are found only in the Instantiation Layer. The dependencies and priorities assigned to the roles will each have at least one Security Object (SO) associated with them. SOs are a unique component of the Framework Layer and represent the security tasks which, if completed, will satisfy the dependencies and priorities with which they are associated.

For example, in the setting of the running example of airport security that that has been introduced, examples of SOs include the examining of checked luggage, the checking of hand luggage and passengers — to identify and so remove any prohibited contents — and the tracking of the relationship between passengers and checked luggage. These examples are developed below.

SOs are unconstrained with respect to their location within the hierarchy. SOs can only exist in one hierarchy and never populate multiple hierarchies. This is a key difference between Actors and SOs. One of the aims of this formulation is to improve the communication between different stakeholders and eliminate the duplication caused by the failure to understand the connectedness of security concepts between levels. In practice, that means a typical SO will exist at multiple levels and multiple sections (dependency, priority) in the Framework. It will commonly be the case that an SO created at a higher level will transition through and connect (or create) priorities and dependencies lower in the framework.

The Framework Layer is populated with dependencies, priorities, and security objects through an iterative process that requires design input from an expert source. The criteria under which security objects terminate, and by which the Framework Layer can be said to be complete, is the same for all frameworks created in this way. As indicated above, dependencies and priorities are externally generated. In practice, a hierarchy of roles will not encompass all possible contributors to the security system and will have been bounded at some sensible level. In the example, I have not represented any role higher than the airport security manager in the hierarchy. The creation and bounding of the hierarchy of roles is the first step in creating a Framework Layer. To populate a Framework Layer, it is necessary to determine the dependencies and/or priorities that the top role in the hierarchy will inherit from sources external to the hierarchy. At this stage, there exists a complete hierarchy of roles that is empty of dependencies and priorities except for the top layer. The next step is to assign security objects to these dependencies and priorities that will allow them to be fulfilled, as in Table 6.

Role	Dependencies	Security Objects

Airport Security Manager	Ensure no prohibited materials transit the airport	Scan checked luggage Scan hand luggage and passengers Track relationship between passengers and checked luggage

Table 6: *An example of security object construction*

At this point the iterative process begins. The construction of a framework always proceeds from top to bottom. At each iteration the following are checked:

- Are there any dependencies or priorities without an assigned SO?
- Are there any un-terminated SOs?

The construction of a SO terminates once it is possible to return a true or false value from its lowest point. If this is not possible, then the SO must be extended to the role below in the hierarchy, creating any necessary dependencies and priorities as it does so. The dependencies and priorities created will be informed by the role creating them; thus, as the SO descends through the hierarchy, it will become more detailed as the scope of the lowers is necessarily more limited. In the example above none of the security objects can return a value and thus need to be extended. Let us extend the 'scan hand luggage' SO. Table 7 shows the result of two iterations, one would generate the dependency from the SO above, the second would find a dependency without an assigned SO and create one.

Role	Dependencies	Security Objects
Airport Security Manager	Ensure no prohibited materials transit the airport	Scan checked luggage Scan hand luggage and passengers Track relationship between passengers and checked luggage
Airport Security Staff	Examine all passengers and luggage passing through security checkpoint	Identify contents of hand luggage and verify it is permitted

Table 7: *Output of iterated Security Object construction*

Note, for example, that the SO 'scan hand luggage and passengers' corresponds to a tree (red/orange in Figure 17) in the Framework Layer. At this point the SO can return a Boolean (true/false that the contents of the bag are permitted) and will terminate. The framework is not yet complete, however, as there are still un-terminated SOs at the manager layer. Iterating in this fashion would also close those at a suitable point. The final step in a SO is always a compliance step, which indicates that at this level and below, the roles in the hierarchy simply comply with the SO and are not involved in its execution. This would add the following line (see Table 7 below) to the framework:

Passenger	Comply with Security Object	

Table 8: '*Compliance step' line*

Once the framework is complete, under the criteria outlined above, the SOs will form a Boolean forest with the leaves connecting each dependency and priority in the framework. At this point the construction of the instantiation layer can begin.

The construction of the SOs can be likened to Schneier's attack trees [62]. Both SOss and attacks trees use a formal, methodical approach to describe a system. However, where attack trees describe, as the name suggests, the goals, and routes to that goal, based on their capabilities, of an attacker, SOs describe the security goals of an organisation. SOs are, in some sense, 'defense trees'. It may be possible, although not considered in detail in this thesis, to assess the security of a system my overlaying attack trees and the contents of the Framework and Instantiation Layers. Indeed if sufficient information about the threat environment was known then attack trees could be used to inform the construction of the SOs (and corresponding Security Components, for more on which see Section 6.3.2.1).

It is, however, worth noting the critical differences between attack trees and SOs. The Framework and Instantiation Layers separate the declarative and operational components of a system such that planning can take place more effectively. This separation also allows managers to ensure that the implementation they have chosen is actually fulfilling their objectives; it is providing security in the manner they require. Attackers are only required to plan a single successful attack. Defensively, security managers must provide protection against a range of attacks. In order for a system to be effective these objectives must work together in order to share advantages where possible, and certainly not contradict them. As such the hierarchy outlined here allows managers to view the whole system and ensure that the various security goals are being met simultaneously. Schneier's commentary on attack trees is however useful and many of the same lessons are relevant. The knowledge captured here is reusable (an important point when the scale of the task is considered) and both

require an iterative approach that becomes more effective as more refined information is added to the system.

6.3.2 The Instantiation Layer

Where the framework layer is static and declarative, the instantiation layer is dynamic and operational. Two new parts of the architecture are added during instantiation, Security Components (SC) and Actors. Actors will be discussed in more detail below; for now it is sufficient to know that they occupy roles and insert preferences into the hierarchy of roles at the Instantiation Layer.

6.3.2.1 Security Components

SCs combine together to form the operational counterparts of security objects. The instantiation layer is built in the same way as the framework layer. Again, an iterative process is adopted with certain termination criteria. The key difference here is that this layer is built bottom up. SCs lay out the processes and resources needed to perform the Boolean checks specified in corresponding SOs. SCs start at the final 'compliance' layer of the SO. Once the processes and resources required at this level are put in place a check is carried out to see if they are sufficient to complete the SO. If yes, then the SC terminates. If not then the procedure is repeated for the role above and additional processes and resources are added as needed. Again, this process repeats until all SCs are closed. At this point, the Instantiation layer is complete. A little more formally, corresponding to the slightly more formal view of SOs sketched above, I will describe how SCs are combined to instantiate SOs as follows:

- SCs are combined according to the and/or forest determined by the SO that they instantiate;
- Each SC implements a checking process that applies to Actors at the level below;
- SCs return Boolean values that instantiate internal nodes of the corresponding SO.

Working through our example again let us start at the passenger level and work upward until we have sufficient processes and resources in place to return a Boolean for the statement 'the passengers possessions and luggage are permitted'. The finished security component in this case would be as in Table 9.

Role	Dependencies	Security Objects
Airport Security Manager	Ensure no prohibited materials transit the airport	Provide resources (X-ray machine, metal detector, wands) Provide data on prohibited materials for X-ray comparison
Airport Security Staff	Examine all passengers and luggage passing through security checkpoint	Monitor X-ray machine and inspect results for prohibited items Hand-search suspect luggage Hand-scan suspicious passengers
Passenger	Comply with Security Object	Place luggage on scanner Walk through detector

Table 9: *A completed example of a Security Object*

In Figure 17, the SCs correspond to the green/yellow nodes in the Instantiation Layer
Note that whereas the SO terminated in a 'compliance' level the SC terminates at a 'provision' level when it reaches a role that can sufficiently provide the resources required to execute the SC without recourse to a higher role.

6.3.2.2 Actors

The final components of the Instantiation Layer (and the model) are Actors. Actors exist independently from any single security hierarchy. They represent entities that transition between hierarchies, this being the key difference between Actors and SOs. They can interact with any and all hierarchies present, simultaneously if necessary. Actors exist solely as a collection of tags corresponding to their attributes. Attributes such as an Actor's nationality, employer, age, and physical characteristics may all be included as an attribute tag. When an Actor interacts with a hierarchy and seeks to populate one of its roles, the hierarchy examines the Actors tags, some or all of which may be unreadable to the hierarchy. Our airport hierarchy, for example, will not be able to read tags relating to other hierarchies such as the security clearances assigned to the Actor by the organisation that employs it. Additionally, there may be tags assigned to the Actor that the actor does not know about. An example of such a tag would be one placed by a government agency to indicate the Actor

is on a terrorist watch list. The tags a hierarchy can see are stored as part of the role data of the hierarchy.

Having inspected the tags it can see, the hierarchy assigns the Actor to a role based on the information gathered. The Actor then inhabits that role for the duration of its lifecycle in that system. The Actor inherits the relevant dependencies and priorities of the role. The Actor brings its own preferences into the hierarchy. These preferences are used to generate *Preference Grids* which are covered in Section 6.3.2.3. Some Actors have no preferences. Such Actors represent items such as inanimate objects (e.g., a passenger's bags in the airport example) or data stores (e.g., a baggage handling label in the airport example) that can be passed between hierarchies (the departure and destination airports will have separate security systems and this will be considered separate hierarchies) but have no intentions of their own.

The Actor is now a full part of the hierarchy and can interact with its associated SCs, updating its dependencies, priorities, and preferences as necessary. The associated Preference Grids will determine how the Actor interacts with the processes and resources of the hierarchy. While the structure of each Actor in the hierarchy is the same it is the Preference Grids that determine the Actor's behaviour. In this way we can differentiate between an Actor representing a legitimate passenger, and one representing a terrorist. While they appear largely identical to the hierarchy they will have very difference Preference Grids, which will lead them to undertake very different behaviour as they move through the airport.

The status of the Actor may be changed by the hierarchy, and tags added or removed as a result. For example, a hierarchy designed to check the citizenship of an unknown Actor may return tags, which it did not previously possess, identifying it as a national, or an illegal alien. The introduction of Actors adds a third dimension to the hierarchy. Multiple Actors can exist in any one layer of the hierarchy (and, typically, the lower levels will have more Actors assigned to them). This also means SOs can span multiple Actors as well as multiple levels. To understand the governing mechanisms that inform the behaviour of Actors it is necessary to recall certain key points from the hierarchy conceptualisation.

First, that the lowest level in the hierarchy of roles is occupied by Zero Preference Actors (ZPAs). This level is essentially a (possibly extensive) list of the ZPAs allowed within, or understood by, the security hierarchy. In the case of the airport example this list would, for example, include

- Hand luggage

- Checked luggage,
- Passport,
- Ticket/booking documentation,
- Minor offensive items (such as liquids over the amount allowed in hand luggage),
- Major offensive items (such as weapons and narcotics).

Second, that the hierarchy of roles is ordered by the level of control each role has over the security system. This means that roles further up the hierarchy have authority over the lower levels.

Considering the question of how to model Actors, from a systems modelling perspective, the following is a useful guide:

- A model that is written from the point of view of a given Actor (e.g., passenger's view of his or her passage through the airport) would normally treat that Actor as a process.
- Other Actors (e.g., the airport security staff, as part of a model written from the point of view of a passenger) would normally be treated as resources — to be used by the subject, or processes called by the subject, in order to navigate the system.
-

6.3.2.3 Preference Grids

Within the hierarchy, then, Actors will often have cause to engage and interact with one another. Both to motivate and to mediate such interactions and determine their outcomes, it is necessary to represent within the system a notion of the intentions of the various Actors. This is achieved by using Preference Grids. As well as their tag clouds Actors have a set of Grids that relate to themselves and all Actors in the layers of the hierarchy below their own, including ZPAs. These Grids are tables containing a set of Preference Values representing the intentions of that Actor toward the other dynamic components of the security hierarchy. This allows Actors, for example passengers in the airport, to be constructed similarly regardless of their intentions. A family going on holiday and a terrorist will have the same structure as Actor. Their intentions, however, as represented by their preference grids, will be very different. Grids are location-dependent as the preferences of the Actors toward each other will vary as they move through the system. Referring to the airport location map in Figure 16, it is clear that a passenger will react very differently to a security guard attempting to search their bag at the check-in desk as opposed to the security checkpoint. This location map will be used as the basis for discussing the relationship between Actors and locations in the following section. Grids are, evidently, close related to utility, and I will return to this point in Section 6.3.2.4.

The Grids assign weightings to 'C, I, A', expressed as a value between -1 and 1. CIA is being used here as a conveniently accessible example. It is likely, and may be necessary, that in the specific context that other hierarchies are built in other classifications will be used, with finer granularity and more direct relevance to the context. The CIA dimension can be extended as necessary with other components, such as actions. Each Grid—for example, one expressing the intentions of a passenger toward hand luggage—is in fact a family of Grids representing the intentions of the passenger toward all instantiations of 'hand luggage' they have interacted with in the current state of the system. Where the Actor has an existing preference toward the Actor or ZPA in question these values will be used to populate the Grid. However, as not all interactions can be predicted in advance (and indeed to do so would bloat the information content of the Actor's tag cloud significantly) interactions can occur in which the Actor has no preference to draw upon. In such cases, a 'neutral' Grid will be generated from information contained in the Framework Layer. Neutral in this case does not mean containing values of 0 but containing values aligned with the goals of the security hierarchy rather than the individual Actor. In this sense, where the Actor has no preferences of its own, it inherits the preferences of the system.

Although it will be necessary to generate many of these Grids in a typical instantiation of the system, they are computationally light to handle and thus incur little overhead. These Grids, and the attitudes they represent, form the intention map for each actor. Do they want explosives to be in the passenger lounge or not? How concerned are they with the confidentiality of the contents of their luggage? As Actors can also have preferences regarding actors lower down the hierarchy and not just ZPAs interactions are allowed between groups of intentioned actors as well. For example, security guards may detain or search passengers according to their preferences. This allows the intentions of passengers, smugglers, terrorists and security personnel to be captured with equally fidelity using the same approach. The system can remain agnostic of the goals of each Actor and does not require special mechanisms to deal with 'enemy' Actors. This allows different types of actor to be created very easily by configuring their tag clouds to contain different preferences regarding the components of the system.

The inclusion of Preference Grids in this system as a method of representing the goal state and intentions of Actors allows user behaviour and choices to be modelled alongside other more technical considerations. This begins to offer up an answer to Research Question 3. How this works in practise will now be looked at in a little more detail.

To illustrate the flexibility of the Grid system let us pick up our running airport example and compare the Preference Grids of a passenger attempting to smuggle a disassembled ceramic firearm onto a plane and the airport staff present at the various locations the passenger passes through. Although a full system would require many more Actors and ZPAs to be

covered by the Grids for simplicity and space we will focus on the passenger and his hand luggage only. Additionally, we will assume that once the weapon is airside a security breach has occurred so the relevant locations for this example are the check-in desk, security checkpoint and passport control. It must also be noted that the values given here are a static snapshot of a dynamic system. As the Compliance Budget tells us, perceived costs are more important than some objective 'real' value. This means that the preferences expressed in the grids will vary, as the environment and status of the Actor changes.

At check-in, the passenger will not tolerate anyone examining his or her bag, and the check-in staff will have no interest in it at that stage, so their grids will be aligned. Moreover, the check-in staff, being concerned with the business processes associated with the airport, will actively want the passenger to be there. At this stage both parties are interested in the availability of the passenger and neither intend for the integrity or confidentiality of the bag to be breached. The Grids associated with check-in are shown below in Tables 10 and 11. The first column shows the factors to which the preference values are being assigned, in this case CIA. The column headings left to right show the members of the hierarchy of roles about which the Actor in question has preferences. ZPAs, being the lowest level of the hierarchy, will therefore be in the rightmost columns of a Grid. It is important to understand that the Preference Values contained in the Grid represent those of an Actor about the entities contained in the Grid column headings. Thus the first grid below shows us the preferences of the passenger in our example towards the CIA state of himself and his hand luggage while located at the check-in desk.

	Passenger	Hand Luggage
C	0.8	1
I	1	1
A	1	1

Table 10: *Passenger Grid at Check-in location*

	Staff	Passenger	Hand Luggage
C	1	0.8	1
I	1	1	1
A	1	1	1

Table 11: *Check-in Staff Grid at Check-in location*

The Grids for the passenger, x-ray operator, and bag searcher are as follows:

	Passenger	Hand Luggage
C	0.5	0.9
I	1	1
A	1	1

Table 12: *Passenger Grid at Security Checkpoint location*

	Staff	Passenger	Hand Luggage
C	1	0	0.6
I	1	0	0.6
A	1	-1	1

Table 13: *X-ray Operator Grid at Security Checkpoint location*

	Staff	Passenger	Hand Luggage
C	1	0	0.4
I	1	1	0
A	1	1	1

Table 14: *Bag Searcher Grid at Security Checkpoint location*

The lowered scoring for the passenger's preference toward his own confidentiality implies that he is willing to be searched during the process. His preferences for his bag are in alignment with the requirement of the x-ray operator, as discussed in the preceding paragraph, but conflicts with the bag searcher who wishes to reduce the confidentiality of his bag below the level he is comfortable with. The neutral weightings of the bag searcher and x-ray operator indicate that while they are not actively seeking to take items from the bag they have no interest in preserving its contents either. Their actions in this regard will be dictated by other factors, such as the legality of its contents. It is also interesting to note that the Grids allow us to express less directly relevant security concerns such as the x-ray

operator not wishing to have passengers in or around the machine as shown by the negative weighting for passenger availability.

The Grids can be used in this way to express a wide variety of behaviours simulating many types of Actor at the same level in the hierarchy. Intra-Grid interactions also reveal behavioural preferences. For example, if a passenger has higher integrity than availability at check-in it means he is unwilling to remove items from his bag (if it is overweight for example) even it if means they will not be allowed to check in. Preference Grids also have a key role early stage in systems models. The preferences expressed in the Grids are directly analogous to preferences for certain processes to be executed or resources allocated. By supplying these requests for resources and processes, Actors provide the dynamic force that drives the system model forward, and are usually as processes, though occasionally it is convenient to treat them as resources.

6.3.2.4 Utility revisited

The utility functions associated with economic approaches to information security, such as those described in Sections 2.4.3 can often be quite complex mathematical constructions, and it may not always, for a given system, be possible to formulate them. Indeed, it may not be necessary. Additionally, such mathematical expression is outside the scope of this thesis. A consideration of how such functions might be developed in this context can still be useful, even if they are represented in a highly simplified form. In the setting of the airport example, Preference Values can be seen as a crude representation of utility, encapsulating both weighting and deviation from target. For example, in the situations discussed above, we have:

$$U_{(C,I,A)} = P_C + P_I + P_A$$

Here, each of P_C, P_I , P_A is a Preference Value drawn from a Grid. Each of these values represents an expectation of an outcome. This can be compared to the actual outcome that occurs, and so a deviation from target is determined. Moreover, the size of the initial value can be interpreted as a crude measure of the weighting assigned by the Actor. A measure of an Actor's utility can be calculated using Preference Grids. By comparing the outcome of any action with the initial intention of the Actor, as expressed in its Grids, a value for how well the Actor's preferences have been met can be found. These values will be weighted by how strongly the Actor preferred the course of action in question.

These utility values will also serve as a measure of the friction between the primary process of the system (as represented by the intentions of the populating Actors) and the security

process. This illustrates one possible use for this approach; when used as a tool of analysis for an existing system this methodology serves as a warning system that is capable of identifying points of tension between various systems. Any actors with consistently low utility will likely be unable to complete goals they value within the current constraints of the system. Early identification of such individuals will allow the security manager to potentially address the incipient problem before it becomes endemic, or begins to undermine the security system, as users start to adopt insecure behaviours or other workarounds in an attempt to realise their goals.

For other types of Actors in the system, such as Smugglers and Terrorists, utility can still be calculated in the same way, except this time the resultants will measure how effective their organizations are at defeating airport security; their intentions (and therefore their measured utility) will be best served by circumventing the security measures with their contraband intact. This can be seen in our running example. Let us consider two scenarios, one in which the ceramic gun is successfully smuggled on to the plane and one in which he is caught at the security checkpoint. We can use the utility equation listed above to calculate the smuggler's utility over time as he moves through the airport.

First we need to construct outcomes for each location in each scenario. In our first scenario our Actor successfully passes through passport control with the gun undetected. The outcomes at each location are as follows:

> *Check-in* – Smuggler successfully checks in, answering the few questions the check-in staff present, and his hand luggage is not searched.
> *Security checkpoint* – The smuggler is not searched and his hand luggage is only x-rayed and the ceramic weapon is not discovered.
> *Passport control* – The smuggler and his luggage are subjected to only routine checks and make it through successfully.

At check in the Actor answers fewer questions than he is prepared to and so gains a positive deviation from his target confidentiality, otherwise he meets his targets for that area, giving us:

$$U(\text{Check-In, smuggler}) = 0.8(0.9 - 0.8) + 1(1 - 1)\ ) + 1(1 - 1)$$

$$= 0.08$$

U(Check-in, luggage) = 1(1 – 1) + 1(1 – 1)) + 1(1 – 1)

= 0

As the Actor's total utility is a combination of his preferences regarding both himself and his hand luggage we take the mean to determine this value. This is the simplest way of calculating utility, and in practise a more sophisticated weighted system will likely be prefereable.

U(Check-in, combined) = (0.08 + 0) / 2

= 0.04

At the security station again the smuggler was prepared to be personally searched but was not asked to so he gains a positive utility gain for exceeding his confidentiality target. His luggage also meets its targets in this area. So for the security checkpoint:

U(Checkpoint, smuggler) = 0.5(1 – 0.5) + 1(1 – 1) + 1(1 – 1)

= 0.25

U(Checkpoint, luggage) = 0.6(0.6 – 0.6) + 1(1 – 1)) + 1(1 – 1)

= 0

U(Checkpoint, combined) = (0.25 + 0) / 2

= 0.125

So we can see that the smuggler in this scenario maintains a positive utility throughout as he is searched less than he is prepared to be personally and the integrity and confidentiality of his bag remains intact. A scenario in which the ceramic weapon is detected would yield a negative utility as the smuggler's preference for the confidentiality of his bag would not have been met.

6.4 Relating the hierarchy to systems modelling and existing data

The relevance of systems modelling in the field of information security was discussed in Section 2.4.3, particularly its use in creating executable models of complex systems that can be used for simulation and predictive analysis. One of my aims when I formulated this approach was that it should not rely overly on the collection of new information but should make better use of existing knowledge. One of the ways in which this can be achieved is to place this knowledge in a form that makes it more readily available for use in the construction of the systems model. The structure and organisation of the hierarchical system, outlined throughout this chapter, has been created with this objective explicitly in mind. The components of the hierarchy have been intentionally designed to mirror those necessary for systems modelling.

At the highest level the hierarchy is split into two primary sections, the *Framework Layer* and the *Instantiation Layer*. This division places declarative concepts and components in the Framework Layer and operative concepts and components in the Instantiation Layer. As systems modelling aims to create an executable model this is immediately useful as the relevant parts of the system (from a modeller's perspective) will be found in the operative Instantiation Layer. This separation prefigures decisions that must be made when creating a systems model with regards to where to bound the model and where to abstract away detail. All of the components of the Instantiation Layer, being both functional and dynamic, should be found in a systems model in some form. The most common forms of components in the Framework Layer are Security Components and Actors. Returning to our encapsulation of the components of a systems model, that of *process, resource, location* and *environment*, let us consider how these relate to the components found in the Instantiation Layer. Examples will be drawn from our running example of Airport Security.

Locations can be said to exist outside of the Instantiation Layer as the locations relevant to the model are applicable to all areas of consideration. Locations influence both processes and resources, and indeed resources are directly tied to locations. As such, a map of the relevant locations and the connections between them (directed or otherwise) must be considered as one of the first stages of planning. The locations used in the Airport Security example are shown in Figure 16. Security Components make direct reference to these locations as part of their definitions; the specification of each Security Component contains a list of the locations it present or functional at.

Security Components define the activities that occur in order to deliver security operations, and the equipment and personnel required by these activities. As such, they represent both processes and resources. For example, in order to successfully scan hand luggage three security components are needed, one process and two resources, and one additional resource (in this case a zero preference actor). The first Security Component would be the scanning procedure itself, parameterised with relevant information such as tolerances and

lists of contraband. This would look for two resource Security Objects, a security staff member and an X-ray machine. Finally it would need the luggage itself. When these are all present the scanning process can execute. In our example security equipment such as boarding card scanners and metal detectors would be resources whereas the procedures put in place by the security policy (represented in the Framework Layer) would be modelled as processes.

Actors are also present in the Instantiation Layer. As mentioned above some of these will be zero preference actors, objects that transition between hierarchies such as luggage or boarding cards. The rest will be the human elements of the system, either users or managers of it. Like Security Components Actors can be characterised as either processes or resources, and this distinction will often be made by the perspective the model is built from. Typically the Actor (or group of Actors) from which the perspective of the model is take from will be treated as processes and the other members of the system will be resources.

Such conceptualisations allow the Instantiation Layer to function as a list of the Processes and Resources that a systems model needs to represent in order to represent the system effectively. In the creation of the Security Objects and Actors many decisions that prefigure those a modeller would need to make will have already been taken and so the Security Object specifications also function as a starting point for the design of the model.

This establishes how the structures of the Instantiation Layer both prefigure the decision making processes necessary for system modelling and also sets up convenient ways of passing data from one to the other but does not address how existing data can be easily imported into the Framework and Instantiation layer. The easy re-use of existing information is one of the key goals of this conceptualisation. Additionally, the system as outlined so far is likely to be to be relatively opaque to a pragmatic security manager. What is needed then is a template represented in an environment with which a security manager and his staff are likely to be familiar with.

The three key parts of this methodology are the Framework Layer, the Instantiation Layer, and the system model. Each informs the structure and population of the others. As described above structurally, the instantiation Layer is organized in such a fashion that it facilitates the exporting of data to the system model. As it is this layer that details the Security Components used to achieve the aims of the Security Objects, it is also the natural place in which to draw in existing system-specific knowledge. One possible approach would be to use an existing and widely used ontology language to create a template for the structure of the Instantiation Layer. This also as the advantage of standardizing this stage of the process, and provides a syntactic interface to security managers. I have chosen to use the Web Ontology Language (known informally as OWL) [75] here as an example, as it is open

source, and as such likely familiar to security managers or their staff. I will describe here how an OWL template, represented using the open source editor Protégé [67], can be used to describes a particular ontology for our Airport Security example. It is also just a small abstraction to see that this process illustrates a general form of such ontologies.

Such OWL templates retain the structure of the Framework and Instantiation Layers with classes for each role and category in the hierarchy. The notions of process, resource, location, and environment are also fully retained. By utilizing this structure, a fully populated OWL template of the Instantiation Layer functions as a map for the implementation stage of the system modelling process. Comparison between the system model and the OWL template will then act as a completeness/redundancy test.

Figure 18 illustrates several key areas of the OWL template. The upper left panel shows the structure of the OWL classes used to represent the hierarchy. This reflects the structure of the Instantiation Layer and holds all of the key concepts of the hierarchy. The lower left frame shows a sample of the relationships used to link entities in the ontology. These allow the creation of constructs within the ontology that are analogous to the Security Objects and Security Components of the Instantiation Layer. To the right of the figure are the members of the 'hierarchyCategory' class. These are an example of the individual entities that can be linked by relationships.

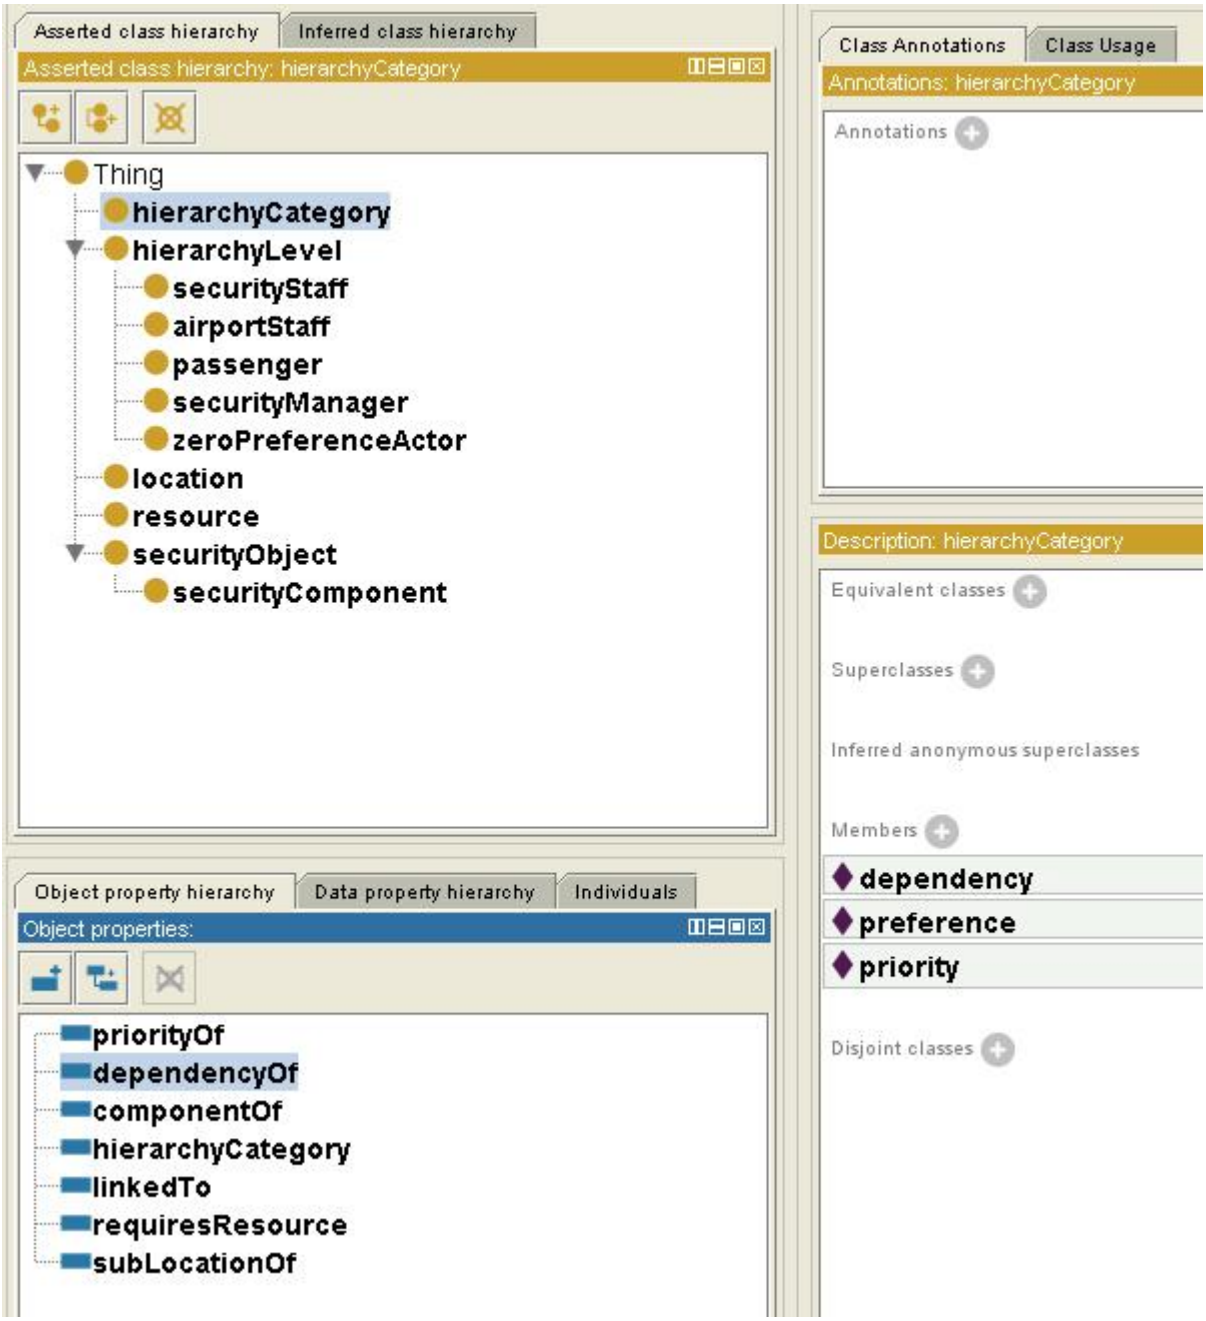

Figure 18: *Protégé screen shot showing the structure of the OWL classes used to represent the hierarchy*

Figure 19 illustrates in more detail one of the individual Security Components of the ontology. This component 'checkHandLuggage2DSC' is a member of the 'checkHandLuggage' Security Object, relates to Role 2 in the hierarchy (securityStaff), is a dependency, and operates at the security checkpoint location. This information is contained in the label of the Security Component and also in its associated properties. These can be seen in the right-hand panel of this figure and are made explicit through the use of the

previously defined relationships. Additional information that is relevant to the Security Component, such as its objectives, may be listed here.

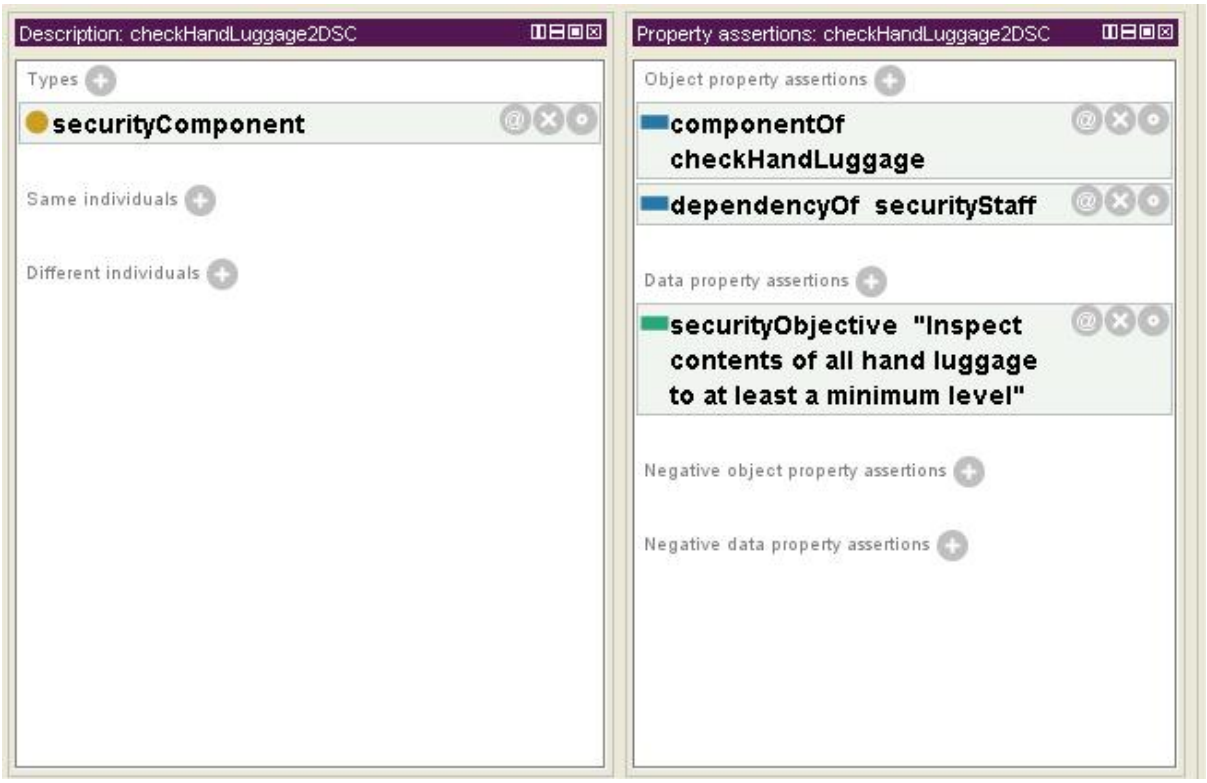

Figure 19: *Protégé screenshot showing the details of a security component*

The creation of a full OWL representation of the target security system (beyond the scope of this thesis) is a type of task with which security managers or their staff are likely to be familiar and so fulfils the needs of providing an environment into which existing data can imported. By structuring it in accordance with this template, the information becomes available for export to the system model.

This hierarchical approach to representing security objectives and information has therefore been designed from the ground up to support as smooth a flow of data from existing sources through to a systems model. This has been achieved by structuring its content in line with the requirements of this process.

7. Conclusions

Information security researchers and practitioners have traditionally looked to improve security by adding more technical controls - such as encryption and firewalls - and stricter policies – e.g. by requiring stronger passwords. This approach does not take into account the fact that security is a socio-technical system: it involves humans who have their own goals, and limitations. These users cannot overcome their inherent limitations, and will not comply with mechanisms that prevent them achieving their goals. While the discipline of HCI has produced substantive knowledge on limitations, and methods for goal-oriented design, this is currently not presented in a way is accessible to security decision-makers. Tools to support decision making – such as systems modelling and utility-based analysis – have been produced by researchers in economics of security, but the human element is currently not represented in these tools. Thus, there exists a gap in the state of the art: users are not included in security-related costing and the economic models to support security decision-making are not grounded in empirical data. The Compliance Budget, detailed in Section 4, makes the case for the need to include a user perspective in security decision making. Chapter 5 outlines empirical data collection methodologies capable of gathering data suitable for use in real-world models and Chapter 6 shows how the resulting data can be best structured to support a systematic decision-making framework. The work presented in these sections addresses the four research questions presented as the motivation for this thesis. These questions were:

> **Question 1:** *What can we learn from economic theory that will move us toward solving the problems found in information security?*
>
> **Question 2:** *What promotes insecure behaviours among the users of IT systems and how can these behaviours be managed and minimised?*
>
> **Question 3:** *How can we maximise information security when the human factor is taken into account?*
>
> **Question 4:** *How can we include the human factor in a systematic approach to security decision making?*

The research presented in this thesis combined knowledge from economics, HCI and security. This approach was taken to tackle the challenges presented by these questions. This has resulted in a linked series of solutions which, taken together, outline a structured approach to information security decision making. Each research question will now be looked at individually. An assessment will be made as to how well it was answered by this

research and conclusions drawn at the end of each section. A summarisation of the research including a discussion of open questions and projected future work will be given at the end of the section.

7.1 What can we learn from economic theory that will move us toward solving the problems found in information security?

The review of basic economic theory in Section 2.3 identified several useful tools and approaches that can be applied to the field of information security. Microeconomic concepts - such as preference and indifference, utility and cost/benefit analysis [9, 12] - provide useful frameworks for understanding the decision making processes in rational actors. These theories feed forward into the construction of both theoretical models such as the Compliance Budget and into the construction of systems models. In the case of the former they allow us to offer an explanation for the observed behaviour of individuals when presented with security decisions. Cost/benefit analysis tells us that a decision making entity will weigh up all factors relevant to the decision and execute the option that maximises their utility. This pushed me to consider as many relevant factors as possible *from the decision maker's perspective* when performing such an analysis. It proved necessary to draw non-security factors into the system to allow a proper costing under such conditions to be made. Such decisions should also not be considered in isolation. All micro-economic decisions take place inside a macro-economic framework which reflects the outcome of previous decisions. While an investigation and understanding of this wider context is far outside the scope of this work it is necessary to recognise that it exists and that security decisions are not made in isolation. In fact some of the costs and benefits factored in to the decision are the results of previous decisions (as represented either by internal structures or the environment in which the decision is made).

For example, the effort previously expended on security compliance by an individual is a factor when future compliance decisions are taken. Observed reasons for this include a cumulative frustration with additional security tasks when such tasks have previously been undertaken. One instance of this is automated virus scanning. In one organisations participants in the interviews analysed (see Section 4.3) reported that "*different IT groups all want to scan these machines*", leading to unhappiness in the workforce due to the repetition of the task. The fact that the repetition of the same task leads to increased frustration indicates that members of a workforce or employee population are aware of prior security tasks when considering the burden a current task places upon them.

Additionally, the timing of such tasks has an impact on their perceived cost and level of intrusion. Due to the additional load placed on computing resources by this process,

primary work tasks can be significantly slowed or disrupted – *"when [the virus scanner] is turned on, our program takes an extra 20 minutes to build and test"*. If the timing of this process is not under the control of the individual, then they cannot effectively organise their work to mitigate the loss of computing resources. If the scan then causes delays to the individual's primary task, as shown in the quote, the perceived cost of the task will increase. However, if they had the possibility of selecting the timing of the scan, they would re-organise their work and "*go for a coffee*", resuming their productive task once the maintenance task was completed. Obviously security decision making cannot be completely left to the users. In this example some users would simply postpone the task indefinitely. A compromise solution in which users have the flexibility, within a given time window, to schedule the task optimally with respect to their workflow would yield the maximum overall utility. Alternatively the system could automatically start the process once it detects the user is absent. Additionally, there may be other factors that the users are unaware of that will influence the organisations preferred patch timing as the next paragraph will show.

Macroeconomic concepts are not directly represented in the cost/benefit analysis undertaken by a user faced with a security compliance decision, but often influence the environment in which the decisions are made. One reason for virus scans and patching processes being run during working hours as opposed to overnight (as my initially seem more logical) maybe that the organisation is attempting to reduce energy expenditure. This indirectly creates a cost for the individual as the disruptive tasks are then present in their working day, but this decision to trade-off costs between the user and the organisation is not consciously made.

Economic theory therefore offers us boundaries and context alongside the tools it makes available to us. We can abstract away factors considered that can be considered macroeconomic and focus on the microeconomic inputs to the decision making process. Economics [9, 12] also informs us that the more accurately we represent the costs and benefits impacting a decision the more useful our result will be [13]. Factors such as timing of tasks (mentioned above) and productivity loss - which are not necessarily traditional security considerations but are of concern to the employees being asked to comply with security tasks - then become relevant when considering the effective deployment of security systems within an organisation. As discussed by Pallas [56] this approach treats security as a public good (in the economic sense) the production of which is centrally controlled and paid for by members of the organisation, irrespective of the cost of production for the individual being higher than the benefit gained. Pallas also points out that information security is delivered at least in part by formal or informal rules, the enforcement of which require the co-operation of the user population. This too introduces additional costs associated with the inefficient deployment of measures as a user 'paying' to support an inefficiently implemented centralised system of security is less likely to co-operate in the enforcement of

the formal and informal rules. This concept of user effort is central to the Compliance Budget (Section 4) and is discussed further below in Section 7.2.

7.1.1 Cost / Benefit Analysis

A cost/benefit analysis (CBA) is simply a method of rationally making an economic decision. The process involves assessing the total costs and benefits inherent in one or more courses of action and deciding which to pursue. The greater the number of relevant factors (those that impact the costs and benefits directly) that are included in to the decision, the more accurate the final decision will be. Framing decisions in this manner is a useful way of both understanding existing decision-making behaviour and predicting future outcomes [16]. Working from an assumption that decision making entities are, at least to some extent, rational, this approach provides a tool to identify the factors that have an impact on a decision. If when analysing a set of factors we predict a certain outcome, for example that users will comply with a security mechanism 80% of the time, but we observe only a 60% compliance rate in the real world, then we must conclude that, given our earlier assumption, the list of contributing factors is incomplete. We now know that we must re-analyse the situation to identify further relevant factors, or refine the weightings that are given to the existing list of factors.

User behaviour has often been described in information security as irrational but using techniques such as this highlighted that there may be other mechanisms at work, and encouraged me to look beyond the usual explanations for factors that were influencing behaviour. Additionally, once a decision making process is framed in this way, and grounded in a satisfactory set of inputs, it can be used to predict future actions. The impact of changing the elements affecting the decision is one such powerful use to which this tool can be put.

7.1.2 Utility

Utility is a key concept drawn from the field of economics [9]. Utility in economics is a measure of how satisfied an entity is with its purchasing decisions. An entity therefore seeks to maximise its utility through its choices. When considering this in the context of security two interesting concepts are revealed.

When considering a subjective concept such as utility the perspective we choose to take is of significance. I will discuss utility from two perspectives – that of an individual within the organisation and of the organisation itself.

7.1.2.1 Individual Utility

As revealed in the interviews that underpin the Compliance Budget (found in Section 4), an individual's utility, as considered by the individual themselves, is not focused on security (excepting the limited number of situations in which that individual's primary task is also security-related). While in an ideal world (from a security practitioner's point of view) individuals engaged in a cost/benefit analysis would be assigning said costs and benefits from the perspective of optimising the security outcome, the majority of the employees in an organisation will derive their utility more directly from their primary task, be this financial trading or technology research. When we think of an individual trying to maximising their utility we must do so in terms of this primary task not security.

Where security supports or does not directly conflict with this task, then an individual can increase or maintain their utility level while undertaking security tasks. However, where the security task comes into conflict with the production task, compliance with the security task can be said to *reduce* the individual's utility. It is in these instances that the difference in perspective between individuals and organisations can most clearly be seen. While in some sense both desire the organisation to be as successful as possible, the short-term goals of the individual may conflict with the operational and security policies of the organisation. When wishing to understand, predict or manipulate the behaviour of individuals in order to encourage a higher level of compliance with security policy it is necessary to work from their perspective. Utility gives us the tools to do this as the decisions made by the individual will be made in service of the goal of maximising utility not security. This is akin to an economic entity making purchasing decisions with its resources in order to gain the maximum satisfaction from the goods purchased. This process is one of trade-offs and balances where the expected outcomes of decisions are weighed against each other. When constructing a utility equation for an individual we may consider the following areas:

1. Level of productivity.
2. Level of security compliance.
3. Time spent on secondary tasks.

Taking the perspective of an individual user, their goal is to maximise their productivity while maintaining their level of compliance such that it is at the minimum level required to avoid sanctions (or termination). They also wish to minimise the time spent on secondary tasks. If a policy change requires more time to be spent on secondary tasks, then the individual has to either reduce their level of compliance or reduce their productivity to maintain their utility.

7.1.2.2 Organisational Utility

When considering an organisation's utility it is a flawed to think of 'security' as a separate concept. There are very few organisations for which security is their primary objective. Security then is one component of the utility equation but increasing security does not always increase utility. In that an organisation's security procedures are dependent on the co-operation of its members, and those members as we have seen above are factoring in more than simply security factors when making compliance decisions, we must also take in to account more than just a measure of 'how secure is the organisation' when determining its utility. Extending this thought further, we can imagine a restrictive security system that is sufficiently burdensome that it places heavy costs, both financial and temporal, on the organisation it supports. These costs are borne, at least in part, by the individuals that comprise the organisation's workforce. Such costs, especially time-related ones, negatively affect the productivity of those individuals, thus reducing the organisations profitability. Unless the protection granted by the security system (in terms of threats mitigated) increases faster than the impact of its deployment a point will be reached where the security system costs more to run than it saves through its mitigation. Diminishing returns and investment policies makes this point a virtual inevitability.

Here is where the perspective of organisational utility becomes interesting. If we take a purely security-centric view, then we are not terribly concerned when we reach this point – security does after all continue to rise, although at a slower rate with respect to resources invested. However, if we take an overall measure of organisational utility then this point is of critical concern. From that perspective our utility equation would report that our optimal strategy would be to *reduce* the impact of the security system, allowing the workforce to be more productive. This reduction could come through a decrease in the amount of security deployed, an increase in the efficiency of the existing systems or through some other approach.

This can be illustrated through the use of a simple example. Consider an organisation that wishes to strengthen its security policy. It decides to do this by increasing the minimum length of the password required to access its internal network. This password is used by all employees when signing on to a computer. While increasing the minimum password length will provide additional security with respect to some attacks it also has other impacts. In particular it increases the cognitive resources required to successfully remember and use the password. This increase means that individuals using the system will suffer failures in recall and entry more often. These failure events require a password reset to rectify. Resets are handled by the IT Helpdesk which must be provisioned appropriately to handle the expected volume of reset requests. In this small example we can see that changing the password policy in this way has potential impacts on three primary areas:

- **B:** The frequency of password-related breaches

- **P:** User productivity (a locked-out user cannot be productive).
- **R:** The resources required by the IT Helpdesk.

From this we can construct a simple organisational utility (U) equation:

$$U = P - B - R$$

Utility is maximised when productivity is high and breaches and resources spent are low. Returning to our proposed change in password policy, increasing the password length reduces breaches, a desired outcome. However, it also decreases productivity and increases the resources required by the helpdesk. Taking this viewpoint we can now say that the change should only go ahead if the cost of breaches is reduced by more than the losses incurred to productivity and helpdesk resources.

In a standard utility equation these variables are also affected by weighting functions that affect the relative significance of the factors in play. Our example equation should therefore be more correctly written as:

$$U_1 = w_1P - w_2B - w_3R$$

These weighting functions change the nature of the outcome and allow us to represent a range of organisational types. A very security-conscious organisation would have a strong weighting for w_2 so the negative impact of breaches on utility would be magnified. This would promote more resources to be spent on breach mitigation in order to increase utility. On organisation with a strong weighting for w_1 on the other hand would be very business focused and react badly to any loss of productivity.

Quantifying these changes is a vast challenge that is beyond the scope of this thesis but I have sought to illustrate here the benefits of taking an economic approach to security. Without it we would consider any gain to security to be worthwhile (within some reasonable budget constraints), failing to perceive the wider impacts and implications of our policy implementations.

7.1.3 Research question 1: Concluding remarks

The research question asked *'What can we learn from economic theory?'* The discipline of Economics of Information Security is a relatively new field that utilises well-established theories from economics to tackle security-related problems.

Economics provide security managers with established techniques for methodically and systematically considering all possible costs and benefits associated with a course of action. This can be particular effective when applied to human factors in security, an area previously not well understood by the security community.

Economics offers the security community tools for placing security problems in a rational framework that aids in optimising security decision making. By taking an interdisciplinary approach, it becomes possible to adopt a more scientific methodology to security, where before there has been an over reliance on the intuition and experience of high-level security managers and a following of established 'best practise'. How these economic tools can be incorporated into a security framework that includes human behaviour will be discussed in the following sections.

7.2 What promotes insecure behaviours among the users of IT systems and how can these behaviours be managed and minimised?

The security community's understanding of the role of the user in security has changed over time. From thinking of them as a hindrance to the effective deployment of security, attitudes have shifted to seeing them as an integral part of the security system – that effective security cannot be achieved without the co-operation of the population it seeks to secure. Schneier [63] said that *"security is only as strong as the weakest link in its chain."* Many security experts (both the research and security communities) have seized upon this quote and consider users to be that weakest link, regular citing Schneier to support their view. Instances of users adopting insecure behaviour – that is behaviour that does not support or undermines the security process – are well documented. From writing passwords on post-it notes to tailgating through security doors and barriers, it appears that users are often being deliberately insecure. Yet while some individuals do delight in in breaking rules, they are in a minority. Most people wish to feel safe and to conform to the normal rules of society, and most employees want their organisation to prosper, so they can remain gainfully employed. So why are such behaviours so common? Understanding why individuals are indulging in these negative actions is the first step in a chain that leads to what should be the prime goal of security practitioners - a co-operative workforce. By understanding the motivations behind non-compliant actions we can introduce changes into the system (which here refers to the total system including the internal states of the individuals) that counteract these motivations and lead to the reduction or elimination of these behaviours.

In the past such behaviours have been dismissed as the irrational actions of uneducated and unmotivated people. However, the research presented in this thesis was undertaken from the standpoint that, on the whole, individuals act largely rationally, albeit using incomplete

and subjective viewpoints on which to base their decisions. As such, I looked for underlying reasons for these behaviours. The data set that was used for this was a set of interviews undertaken at a technology research company and a financial institution. What became apparent from an analysis of these interviews is that (excepting the perverse few mention above):

> *Individuals adopt insecure behaviours not because they want to but because they feel they have to.*

This understanding then pushed me to look more closely at the factors that influenced these decisions. This core idea was then built into the theory of The Compliance Budget.

7.2.1 An economic re-framing

As we saw in Section 7.1, the discipline of economics offers us several valuable tools for framing models relating to behaviour and decision making such as utility and cost/benefit analysis. We can use these models to help answer the question of why users are adopting insecure behaviours. When asked during the interviews about complying with existing or new security policies participants mentioned several factors that would encourage them to comply or not to comply. The choice to not comply is a key form of insecure behaviour that will form the focus of this discussion. I decided to frame the decision to comply or not comply in terms of a cost/benefit analysis conducted internally by the individual in question. As discussed above in Section 7.1.2.1, the goal of such decision making is to maximise the individual's utility. As such, the factors that influence the decision and are weighed up in the cost/benefit analysis will be quantified subjectively from the individual's perspective.
However, while this framing is a useful way of examining the question it does not provide the answers we seek. Returning to our initial assumptions that users on the whole wish to be secure, and often have a vested interest in doing so, it seems more likely that they would decide to comply and that instances of non-compliant behaviour should be substantially rarer. The factors that contribute to such a decision were outlined in Figure 6.

A comparison of these factors would imply that while in some edge cases where security can be seen to be actively obstructing an important business transaction it would be rational to not comply with security in more minor or everyday situations the opposite would appear to be true. So far then this approach is not sufficient to explain the day-to-day issues of insecure behaviour such as writing passwords on post-it notes and tailgating through security barriers. There are two key factors missing that significantly influence decision-making processes. These will be examined in the following sections.

7.2.2 Primary task and context as factors

While security is regarded by many individuals to be important in an abstract sense, this does not necessarily translate into a motivation to behave securely in day-to-day compliance decisions. Where more direct motivators for behaviour exist, such as salary and task pressure from managers, security can and does take a back seat. The prime reason for this is that for the majority of individuals security is not the primary task for which they are financially (or otherwise) rewarded. When an individual seeks to maximise their personal utility they are not looking to make the organisation as secure as possible. Instead their utility is based on personal factors such as the successful execution of their primary task and earning money to provide for themselves and their family. While security may be a component of these it is not an explicit goal and this secondary status has a substantial impact on how it is perceived and acted upon as we shall now see.

7.2.2.1 Primary Task and Risk

Referring back to our economic framing of compliance decisions in Section 7.2.1, this substantially changes the weighting given to the factors involved. Those relating to the primary tasks and systems (the business task in our running example) will be considered to be more significant than those relating to secondary tasks, which for the majority of people include security. Figure 6 showed the various factors that an individual considers when faced with a compliance decision. These are summarised below, with an indication of whether they apply to the primary or secondary task, or are a general cost or benefit.

Costs:
- Physical load (general)
- Cognitive load (general)
- Missed opportunities (primary)
- Embarrassment (primary)
- Less availability (primary)
- Hassle factor (primary)

Benefits:
- Protection from responsibility (general)
- Protection from sanctions (general)
- Increased security (secondary)
- Protection from liability (secondary)

The majority of the costs associated with compliance are linked to the primary task – worries about missing opportunities or of having to take time out from being productive to complete the security task. Conversely, factors linked to the security process – such as being safer or reducing liability in the case of a breach – fall on the benefit side. This skews the cost/benefit analysis toward non-compliance for individuals for whom security is a secondary task

because the costs, being associated with the primary task, carry more weight than the benefits. Individuals pursuing goal-driven behaviour relegate secondary tasks to the status of distractions. This is compounded by the benefits of the primary task being concrete, expected and likely to happen in the short term. The benefits of the security task are diffuse and long term. Individuals of this kind will include most office workers, and other organisational and corporate employees.

Other factors also come in to play to push the imbalance further in this direction. A user's risk appetite will also affect their behaviour. The relevance of Adams' work on risk [2] became apparent as I considered the implications of the Compliance Budget. Adams expresses an individual's approach to risk through a model he calls the risk thermostat, shown in Figure 20. Like a mechanical thermostat this mental mechanism regulates behaviour by balancing a propensity to take risks with the perceived danger of those risks. Adams also tells us that there are three types of risks:

1. Directly perceptible risks e.g. crossing a busy road
2. Risks perceptible through science e.g. infectious diseases
3. Virtual risks (those risks that science has no firm consensus on) e.g. Global warming

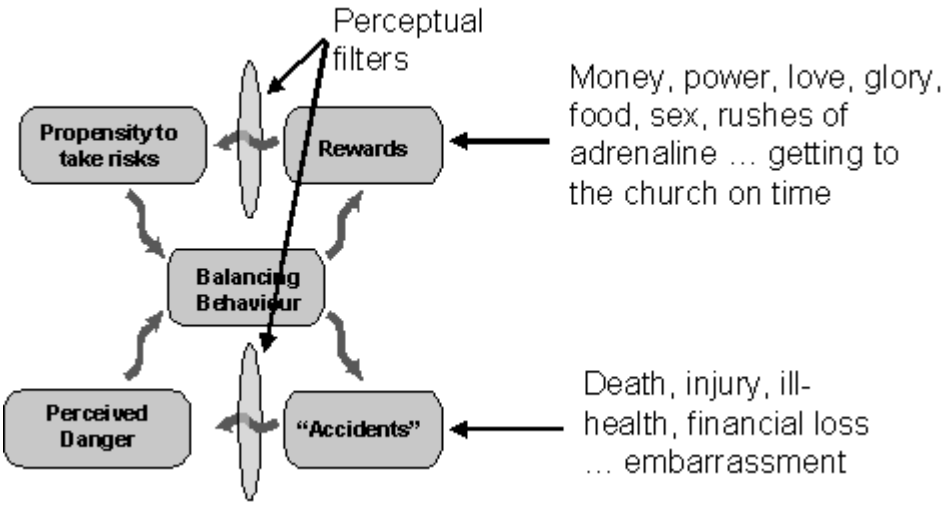

Figure 20: *Adam's risk thermostat*

The way we respond to risk varies depending on which category it is in. We manage directly perceptible risks intuitively and personally, overriding the judgement of experts in the risk with our own beliefs. However, for risks that are only perceptible through science we rely on institutional risk management. As security is not a directly perceptible risk (except for those that are unfortunate enough to suffer the consequences of a successful attack directly – again a minority) the employees of the organisation in our example will look to the

organisation itself to protect them from such risks. However, the weakness of this scientific risk management is that relies on an objective measurement of the safety and danger associated with the risk and a rational response to these measurements by those exposed to the risk. Security is notoriously hard to measure because 1) successful security results in events *not* happening and 2) much of the threat environment is opaque at worst, and visible but highly fluid at best to security professionals. This arguably pushes security risks closer to being a virtual risk which, without the guiding influence of scientific consensus, is then managed solely by an individual's beliefs about that risk. These factors make the effective transmission of security risks to individuals a very difficult task.

Returning to the risk thermostat we see that perceptual filters operate on rewards and accidents alike. In the case of compliance decisions the rewards are magnified (better productivity, less time spent away from primary task) and the perception of accidents, or breaches, is reduced as security risks are not directly visible to the individual. In contrast, the personal risks, missed opportunities, lack of productivity and other primary-task-related issues are directly perceptible by the individual, and so will be directly managed by their behaviour. This has the effect of increasing the individual's propensity to take security risks while decreasing the perceived danger of them thus leading to a set of balancing behaviours that condones risk taking with respect to security – the adoption of insecure behaviour.

So in our cost/benefit analysis we have a set of factors that are weighted toward making costs more prominent and undermining benefits. They do this:

- By their association with a task of secondary importance.
- By the application of perceptual filters that play down the dangers and make the benefits more attractive.

As these benefits are largely found in the 'costs' side of the analysis this exacerbates the bias already found there.

7.2.2.2 Contextual factors

So far I have demonstrated how individuals approach compliance decisions and how, using this approach, insecure behaviours can arise from a more rational and plausible chain of reasoning from the individual. The significance of this is worth restating – perceiving users to be insecure for better reasons than ignorance, indifference or laziness fundamentally changes how security managers need to address security design and incentivise compliance. There is more to add to the picture however. Such decisions are not made in isolation. With any conceptual model it is necessary to bound the inputs at some reasonable point. Bounding our analysis at this point in the conceptualisation of the decision-making process

would introduce significant errors into our understanding of the likely behavioural outcomes of any security policy. The reason for this is that a cumulative effect takes place as security procedures (or indeed any secondary task) follow each other. Time variations also take place and so the perceived burden of tasks varies based on when it is presented to the individual and which tasks have preceded it in the day.

This is understandable intuitively – the burden of a task presented to an employee during a moment of free time during the early part of the day will seem relatively light. However, if that same task is presented at the end of the day, when they are trying to finalise their primary workload so they can leave for home, the burden will feel much higher. Likewise, if they have already been presented with a series of tasks during the day and more keep arriving. This effect was reported in our interviews (in Section 4.2), in particular with respect to the timing of automated patching. Participants reported feeling very frustrated if patching took place when they logged in at the start of the day and were trying to begin their production tasks. They stated that if they could select the timing of the task manually, and set it running, for example, while they ate lunch, the frustration they felt with the patching process would be significantly reduced.

For some the logical next question would be '*Why should security managers care so much about how frustrating such tasks are when they are necessary?*' Indeed, one security manager with whom these concepts were discussed stated that if production tasks were delayed by security events then employees would *'come in to work 15 minutes earlier and leave 15 minutes later.'* [59] This type of attitude - that security is important and so can demand whatever amount of time and effort it consider necessary –is the start of the road to non-compliance. In the patching example above, it is true that - due to the automated nature of the process - participants had no choice but to let it run. However, there are situations where individuals do have a choice, to either comply or use a lower-effort but less secure workaround. The additional frustration caused by the poorly-timed patching process will make future compliance less likely to yield a positive outcome.

To take this in to account in our economic framing, we need to add in a 'hassle factor' to the costs of complying that accounts for the subjective burden perceived by individuals as a result of the timing of the task and the costs of other secondary tasks the individual has already performed. This 'hassle factor' adds further weight to the costs of compliance without adding any counterbalancing incentives to the benefits. With this our cost/benefit model is conceptually complete. While there may be specific factors unique to each decision or situation to consider all the core concepts required to answer our research question have been developed.

7.2.3 Research question 2: Concluding remarks

Our research question asked '*What promotes insecure behaviours among users?*' and this can now be answered using the model developed above based on the economic tools outlined in Section 7.1.

> *Insecure behaviours are adopted by individuals in order to ensure a positive outcome on their own cost/benefit analysis when faced with a compliance decision.*

The key insight here is that the decisions are made using a rational process but based on subjective, assumptions and perspectives. This addresses the criticism levelled at the rational actors on which most economic theories are based – that humans are not rational thinkers. What this research has attempted to show is that a version of rational processes such as cost/benefit analysis and utility calculations are taking place but they are often based on flawed information or from a limited and highly subjective perspective.

This insight relates directly to Kerckhoffs' principles [45] discussed in Section 1.1. Kerckhoffs understood that secure communications depended on the system being usable and three of his six principles reflect that. These principles are repeated below as a reminder:

> *Principle 3: It must be easy to communicate and remember the keys without requiring written notes, it must also be easy to change or modify the keys with different participants;*
>
> *Principle 5: The system must be portable, and its use must not require more than one person;*
>
> *Principle 6: Finally, regarding the circumstances in which such system is applied, it must be easy to use and must neither require stress of mind nor the knowledge of a long series of rules.*

Principles 3 and 6 reflect the need for a low cognitive load in the system. Principle 3 in fact hints at the result of an excessive cognitive load – that written notes will be required. This practise is often reported as a consequence of overly restrictive password policies [1, 38]. Principle 3 also expresses the need for responsiveness and agility within the system. Principle 5 relates to physical load and Principle 6 again expresses the requirement for an appropriate level of cognitive load. These costs are also reported as relevant today by participants in the Compliance Budget.

However, these principles have been neglected in security's recent history, likely due to the fact that the impacts of low usability on system security have never been well understood by the security community as a whole. This lack of understanding of the costs associated with low usability means that the incentives that are needed to drive the planning of usable security are absent. Work such as this thesis that explicitly shows the detrimental impacts of

neglecting usability on the security of a system provide a motivation for security planners and decision makers to return to Kerckhoffs' principles of usability and ensure they are followed.

7.3 How can we maximise information security when the human factor is taken into account?

The sections above have given us a useful set of tools and models with which to understand why individuals adopt insecure behaviours. The next question I wanted to address was how this information could be used to improve information security practice. Such models and theories are of no practical use if their insights do not offer up new approaches to security design. A strength of the above model is that it is based on a set of rational thought processes and has recognisable factors that feed in to the decision making process. This lead me to develop a more structured model that placed the decision making model in context and looked at how security managers could use the knowledge to maximise the effective security of their systems. This model was the Compliance Budget [11]. Rather than seeing it as an intractable or opaque problem the Compliance Budget treats compliance decision making as an understandable process based on rational principles as described above. Additionally, it recognises that the effort individuals are prepared to expend on compliance and overcoming the 'hassle factor' associated with it is finite. The implication of this is that user effort can and should be managed carefully as part of security planning and not squandered by inefficient implementations.

7.3.1 The Compliance Threshold

Understanding that individuals have a limited amount of effort to spend on secondary tasks is one of the key steps in understanding how to maximise security while taking human factors into account. The interview participants showed a range of reactions when asked how they would respond to the introduction of a new security measure, with some being considerably more impatient with security than others. From this it can be inferred that there is individual variation in where this threshold lies. Common factors became apparent that can and do influence this threshold. Loyalty to the organisation, technical ability, personal security culture and awareness of the threat environment all contributed to raising the threshold when high and reducing it when low. This gives security managers some approaches to maximising the Compliance Threshold of the individuals in their system.

Maximising security then means ensuring that where possible individuals are not being asked to work beyond their Compliance Threshold so as to minimise the pressure to seek workarounds to security tasks. Additionally, the organisation should seek to increase the

Compliance Thresholds of their employees where possible. Some thoughts on methods of achieving this will be discussed in a later section.

7.3.2 Rate of spending

It would be a mistake, however, to think that the onus is on the individual alone – that maximising security involves only pushing employees to work as close to their Compliance Threshold as possible with their Threshold set as high as possible. Individuals are pushed toward their Compliance Threshold by the effort required to complete secondary tasks. Maximising security therefore also involves ensuring that such tasks are implemented as efficiently as possible. Remember that user effort should be considered to be a resource that is consumed like any other. With this resource security managers are in essence purchasing security tasks. With tasks that push individuals over their Compliance Threshold being less likely to be completed effectively it is important to consider which tasks the users are asked to undertake.

In the case of financial resources, it is immediately obvious that spending money on over-priced goods is not an efficient way of conducting business. The analogy holds in this case, with inefficient tasks being analogous to the over-priced goods – the manager is paying more than he needs to in order to get the task done. In the same way that buying a series of over-priced goods is a poor use of financial resources asking users to complete overly-burdensome tasks is not a good use of their effort. There is a direct pressure on security managers looking to maximise security then to streamline tasks as much as possible with the primary business process of their organisation. Reducing the per-task effort means that for a given Compliance Threshold more tasks can be asked of the individual before that threshold is crossed. A second route to maximising security when taking into account human behaviour is therefore to reduce the rate at which user effort is spent.

7.3.3 The Operating Point

When seeking to maximise security through the Compliance Budget theory a somewhat counterintuitive result comes to light. The negative impact of adding in security measures once individuals are at or past their Compliance Threshold means that a point is reached where it is functionally counterproductive to do so. This creates a disconnect between the planned security, as set out in the policy, and the effective security of the system as it operates. Increasing required password strength is a good example to illustrate this point. For certain kinds of cracking attacks, increased password length and character variety add strength to the password, increasing security. A password such as ‘Ar5PwdIJC4tEx!’ provides good protection against such threats. However, such passwords are substantially harder to remember and input than ones created under a more lenient policy e.g. Jaunty89.

While in theory a policy that mandates the use of passwords such as 'Ar5PwdIJC4tEx!' provides better security the practise will rarely reflect that. Faced with the high cognitive burden of recall associated with such a password, the majority of users will use a memory aids in order to cope [1, 38, 39]. This memory aid, if the security manager is fortunate, will involve a discrete reminder system. However, passwords written on post-it notes attached to monitors, on whiteboards, kept in a desk drawer or under the keyboard are all often reported solutions. The reality, then, is not a strong password that prevents cracking attacks, but a systemic weakness where passwords can be freely harvested from the physical environment. The effective security of the system is lower than the planned security. This problem has been known to exist since at least 1999 (when Adams and Sasse published their work) yet little has been done to mitigate it. This may have been because the issue was seen to be intrinsic and unavoidable to the inclusion of users in a system – regardless of how the system was design these behaviours would still be seen. The Compliance Threshold offers a way forward; a method for managers to understand the impact of security processes and plan security in such a way as to find a more optimal balance between security and insecurity (or indeed any balance at all). By indicating that this limit exists, and that prior to this limit compliance is substantially more likely, the Compliance Budget makes it desirable to plan to operate at or under the Compliance Threshold rather than dismissing the problem us an unsolvable issue with users in general.

The negative outcomes described above are the result of pushing users of the system to operate at an unsustainable level of effort causing them to, quite rationally, adopt insecure behaviours in order to cope. The costs to the organisation go beyond just those inflicted on the individual users. When there are multiple security events in the system that create enough pressure to push individuals into using workarounds a culture of non-compliance can be created that will propagate over time. New employees joining the organisation will arrive and adopt the practises in common use that will include endemic use of security shortcuts and workarounds. This is analogous to Reason's argument for why repeated accidents occur in the workplace [60]. He states the existence of opportunities for accidents to occur is insufficient to explain my accidents recur. Reason explains that the driving force that accounts for this difference is the safety culture, and that this force is the result of the *"unsatisfactory resolution of the inevitable conflict that exists (at least in the short-term) between the goals of safety and production."* It is clear that 'safety' can be replaced with 'security' and the same sentiment be expressed. In this way an otherwise well planned security policy that doesn't take user effort into account can result in an insecure operational environment.

Following the inferences of the Compliance Budget would instead lead to a recommendation to identify the 'Operating Point'- this being the point at which users are spending their available effort on security but are not routinely being pushed over their Compliance Threshold into areas that encourage insecure behaviours. Counter-intuitively

this will at times involve *reducing* the demands of the security policy. Let us return to our example of the passwords outlined above. While allowing something as trivial as Jaunty89 is not advisable stepping back from a policy that mandates 'Ar5PwdIJC4tEx!' almost certainly is. Studies such as those conducted in Chapter 5 are key to approaching such decisions. Through analysing performance data with various password policies we can identify policies that decrease performance and increase the number of resets. It is these policies that will place the highest burden on users and thereby carry the highest risk of promoting insecure behaviour and workarounds.

By reducing the policy to a level that places a manageable cognitive load on the system's users the motivation to use insecure workarounds is substantially reduced, leading to a higher level of compliance and a more secure environment, with passwords not stored physically in the workspace. So although our policy on paper looks weaker our operating environment is actually strengthened.

7.3.4 Data-gathering support

Although this thesis has theoretically identified this Operating Point as the ideal point at which security policies should be balanced, the functional identification and implementation of these ideas will still require a great deal more work. Significant challenges are present in the understanding of how different mechanisms and policy variations burden users and how best to mitigate these burdens. This work does not claim to have the answers for such questions, simply to present a new way of looking at the problem of non-compliance and the issues of human factors in security.

However, one need that this research has started to address is the requirement for models such as the Compliance Budget grounded in relevant empirical data. Without such data, it is impossible to systematically fine-tune a security policy. Instead, any adjustments rely on the intuition and experience of security professionals alone. While such knowledge is not necessarily ineffective, it can be made more reliable by having valid data to work from. Such data is also necessary to unpin systems modelling work, which will be discussed in a later section.

Laboratory experiments can introduce biases into the conclusions drawn from their data simply because they do not take place in a realistic working environment. While in the laboratory the experiment becomes the participant's primary task. As has been discussed many times already in this thesis the nature of the task is critical to how it is weighed up in an individual's mind and thus the effort they devote to it. As security is almost always a secondary task laboratory experiments that focus on security technologies are rarely suitable for gathering the sort of data required to support theories such as the Compliance Budget.

The key advantage of laboratory studies, however, is that they take place in a controlled environment that can minimise the amount of external factors coming in to play, as discussed in Section 3.1. Gathering data from real environments comes with a host of problems including gaining access to such populations. In order to meet these challenges I developed a remote data gathering method utilising proprietary software developed at UCL. APET is a server-based experimentation system that allows participants to login over the internet from any location to complete a flexible range of tasks. This method offers a compromise solution to the problem of collecting realistic data. Each experiment is controlled by a software module within APET and so can be manipulated to create a systematic progression of experimental variables. This retains some of the pros of using laboratory-based experiment. However, as the participants login from their normal working environment to take part in the experiments the realism of the experiment being a secondary task is at least partly preserved and a more natural set of data is gathered. With experimental systems such as APET over time a database of information can be built up that will be sufficient to ground the Compliance Budget and related models in a valid empirical data set. This will allow them to become much more powerful tools for decision making support and prediction.

7.3.5 Research question 3: Concluding remarks

Our research question asked '*How can we maximise security when human behaviour is taken into account?*' by developing the economic framing of compliance decisions into a more comprehensive and contextualised model this question can be answered.

> *We can maximise security when human behaviour is taken in to account by minimising the friction between security and the primary business process and by managing user effort as if it were a finite resource that needed budgeting like any other***.**

Through seeking to extend the Compliance Thresholds of the individual users of the systems, and by slowing down the rate at which user effort is expended, security managers can aim to execute their policy at a functionally optimal Operating Point. At this point, the maximum number of security tasks are completed but individuals are not routinely pushed past their Compliance Threshold into a space where they are forced (according to their own internal cost/benefit analysis) into using insecure coping strategies to deal with the cognitive, physical and temporal burdens placed upon them by the security system.

7.4 How can we include the human factor in a systematic approach to security decision making?

In answering the previous research questions I have shown the importance of considering the human factor when implementing a security system. The requirement to correctly factor human behaviour into the decision-making process, and the recent advances in techniques and methods associated with this concept, can be seen as a new paradigm in security. However, this is not to say it deserves special attention above and beyond other existing factors. It is important to highlight it at this stage of security research because to date it has been relatively neglected due to the apparent intractability of the problem. This highlighting should not lead to a focus on human factors at the exclusion of all others. The final research question of this thesis asks *'How can we include the human factor in a systematic approach to security decision making?'* The focus on this section will be on looking at methods that allow human behaviour to be factored alongside technical considerations as an equivalently relevant part of a security system.

7.4.1 Systems Modelling

Systems modelling is a powerful technique for analysis and prediction that is becoming increasingly widely used in the field of Information Security. Services such as HP's Security Analytics seek to adopt a more scientific approach to Information Security and systems modelling forms a key part of their approach. In essence systems modelling creates an executable representation of a system that runs as a simulation of the real thing. The accuracy of a systems model is dependent on two things – how closely it represents reality and the quality of the data used to train and tune the model. With respect to the first point, a line has to be drawn somewhere. A model that fully represents all the variables present in the real system is indistinguishable from it. It is also unrealistic to consider that such a system may be built. Instead significant assumptions and abstractions are usually made when developing such a model in order to create the model, carefully chosen so as to preserve its usefulness. The data that is used to train the model should also be carefully gathered so as to be as realistic as possible. As the predictive and analytical power of the model is dependent on the accuracy of this data this is a crucial point to consider. It has been common in the past to abstract away human behaviour when constructing such models for three main reasons, two related to the points raised above. Firstly - and most fundamentally - until recently the importance of human factors in achieving effective security was not recognised by the research community. Secondly, the mechanisms of human behaviour have not been well understood by the security community and so representing it realistically within a model was not an option for security practitioners. Thirdly there has been a lack of available realistic data to allow any models that tried to incorporate a notion of human behaviour to be effectively tuned. As discussed above data from laboratory-based experiments contains biases that can make it unsuitable for real-world prediction.

As this thesis has demonstrated, not including a representation of human behaviour is not a viable decision if security operations are to be effectively maximised. This means that, if systems modelling is to be used as part of the toolset of security management, the issues above must be addressed. This thesis has taken the first steps to do just that. An understanding of the importance of human factors in security is one of the prime motivators for my research in this area. Additionally, through the Compliance Budget, a mechanism underpinning human behaviour with respect to at least one aspect of security has been identified and expressed as a theoretical model. This opens the door for such mechanisms to be represented in systems modelling approaches. Finally, the use of APET and other experimental techniques to gather more realistic real-world data contributes valuably to the pool of data that can be used to tune such models.

Yet, while the previous stages of research in some sense facilitate the use of human factors in systems modelling, they do not offer an integrated solution. I wanted to develop a conceptual structure that was capable of representing human factors alongside technical measures in a seamless fashion. The goal for this system was to allow a systematic approach to be taken to security decision-making so that security managers were less reliant on intuition and experience and could draw upon a more scientific methodology when approaching security policy and design.

7.4.2 Structured Systems Economics

The primary components of systems models are processes, resources, locations and environment. This organisational structure allows the model to be effectively and rigorously defined. By contrast the organisation of security objectives and mechanisms is often ad-hoc or follows standardised guidelines. Added to this there is a lack of rigour within the security community when it comes to expressing the requirements of a security system. This is best illustrated through an examination of what are seen to be the common goals of security – confidentiality, integrity and availability (CIA). It has been recognised by more than one security professional that while CIA represents a good basis for stating security needs and requirements there are gaps that it conceptually fails to address. In seeking to add to CIA most authors pay little heed to whether they are adding declarative or operational concepts; that is to say they do not differentiate between the desired properties of the system and the mechanisms used to create those properties. This sees 'authentication' listed alongside 'confidentiality' when the use of authentication is something that creates confidentiality.

I wanted to look at creating a systematic approach to security that addressed these issues as well as incorporating a notion of human factors. The result was Structured Systems Economics (SSE). This is a methodology for the organisation of security objectives and

information in such a way that the creation of a systems model from this data was not only facilitated but was the logical next step.

7.4.2.1 System Structure

SSE addresses the problems with CIA outlined above by intentionally separating the structure into two layers – the Framework Layer and the Operational Layer. The Framework Layer is declarative and expresses the desired security requirements of the system whereas the Operational Layer describes the processes and resources required to fulfil those requirements. The components of the Operational Layer are deliberately expressed as processes and resources so as to prefigure the creation of a systems model. In this way the description of the system also forms a blueprint for the necessary elements of the system model.

The population of the Operational Layer follows logically from the requirements expressed in the Framework Layer. Each requirement declared in the Framework Layer is addressed by one or more components in the Operational Layer. Components are added iteratively to the Operational Layer until all the requirements are met at which point population of the layer ceases.

This methodology ensures that the minimum of security is used to satisfy the system's requirements. However, SSE as described so far does not satisfy our research question in terms of incorporating human factors in a systematic fashion. The description of the system so far is a static one. It can be seen as the theoretical version of the system but it cannot tell us how the system will operate in a real-world situation. The solution was to incorporate a dynamic component in the form of Actors.

7.4.2.2 Actors

The structure of the SSE framework is separated horizontally as discussed above into declarative and operational layers. The vertical structure is defined by the roles available within the organisation the system represents, sorted hierarchically by their ability to influence the system. These roles are then populated by Actors, which can be seen simply as special processes that function along the lines of the Compliance Budget as described in Chapter 4. These Actors supply a dynamic function by attempting to complete the tasks associated with their roles in the organisation. As they interact with the security system following the economically-framed behavioural models described above they generate outcomes from which the overall utility of the system can be assessed. Again actors are listed as processes (although they may be seen as resources by Actors of a higher role) intentionally to frame them in a form usable by the systems model.

7.4.3 Decision support

The primary objective of this thesis in a wider sense and the creation of SSE more specifically is to develop a methodology for supporting security decision making through a systematic and scientific process. SSE is a key part of this as it structures existing information about the security system into a form that supports prediction and simulation. This allows security decision makers to explore the likely effect of policy changes before they are implemented. When used in conjunction with a utility-based framing of the system security can be optimised with respect to the organisations overall needs and not just improved with respect to itself alone. With a notion of human behaviour represented in the system by Actors SSE offers a more integrated approach to security management. Human factors are neither overly focused on nor abstracted away so as to become detrimentally absent. As part of a cycle of operations that includes effective data gathering and systems modelling SSE offers valuable decision making support.

The value of systems models in supporting decision making was demonstrated Baldwin et al. [24], working for HP Labs. Refining the models available by accurately representing human behaviour represents a valuable step forward in this area.

7.4.4 Research question 4: Concluding remarks

Research question 4 asked '*How can we include the human factor in a systematic approach to security decision making?'* Systematic approaches, such as those based on systems models [13] in general do not include a valid notion of the individual users of the system. Where such notions are included they tend to be simplistic, as in [10]. This makes such models unsuitable for supporting decision making as the responses of the user will not realistically match those found in a real operational environment. The answer to research question 4 can therefore be summarised as:

> *The human factor can be included in a systematic approach to security decision making by including a realistic notion of user behaviour in a framework that facilitates the use of predictive tools such a simulation through systems modelling.*

It is important to note that the representation of the user within the systems model must be sufficiently well developed and motivated as to generate realistic behaviours in response to the system's parameters. Framing the decision making processes of the user as a rational economic process (albeit one based on subjective factors) as is done in the Compliance Budget (Section 4.5) is one step forward along this route. Such models have to be supported by empirical data to allow the model to be calibrated correctly. The security architecture

outlined in Section 6.3 combines empirical data with an understanding of user decision making within a framework that also naturally expresses technical controls. The result is a unified and systematic approach to security decision making. An initial prototype of this method was explored during my internship at HP Labs [8]. This method used the results of the empirical study conducted in [39] to tune a systems model developed using the methods in [13, 41]. This model was able to balance trade-offs between breaches, productivity and help-desk costs within a fictional organisation in order to optimise the investments and security policy of that organisation.

7.5 Thesis summary and critical evaluation

In answering the key research questions posed at the start of this research this thesis has developed an integrated and systematic methodology for approaching security decision making that takes a full account of the influence human behaviour has on security. By framing the problem in economic terms, the methodology includes a fuller picture of the costs and benefits associated with security. This enables better decisions to be made. The optimisation of these decisions crucially is with respect to the overall utility of the organisation and so avoids the trap of implementing security for security's sake. Although there is substantial future work to be done in terms of populating the required data sets and understanding the specific issues associated with any one organisation the first steps toward a more effective system of security have been taken. The Compliance Budget and SSE combined with data gathered through the APET system form part of a security cycle that tackles some of the problems facing security managers today.

While this programme of research has yielded successes in terms of answering the research questions it is also important in any work of this kind to recognise limitations, not least of all so they can inspire and motivate future work. The collection and presentation of empirical data to support the Compliance Budget has taken a back seat in favour of developing other areas of the approach, and this is certainly one such weakness. Due to this lack the Compliance Budget remains at this stage a theoretical model only that is not ready to be used by security practitioners as a tool to support decision making.

Due to the necessity of developing a unified, multi-disciplinary approach to the problem much time has been spent on the presentation of the separate components of this approach. This has meant that the validation of the approach as a whole has remained outside of the scope of the work. This can certainly be raised as a criticism of the thesis and one that future work will address.

7.6 Directions for future work

The research presented in this thesis makes the case for changing the way we think about information security management, and this opens up new directions for information security research. In the Trust Economics project, research was carried out in close collaboration with industry partners, and the problems addressed in this thesis address problems observed in organisations today. Implementing and validating in an operational context an approach to decision making informed by the model of user decision making expressed in the Compliance Budget (Section 4.5), supported by grounded data gathering methods (Section 5), is a key future activity. Demonstrating that a utility-based approach to security, which includes costs associated with human factors, can help to identify changes and interventions that measurably improve security and productivity – and if not, what other factors have an influence - would further the development of theories of information security management that are grounded in the real world, rather being disconnected from it. The development of such theories still requires further building blocks.

The Compliance Budget describes compliance decision-making in terms of a cost/benefit analysis. Throughout this thesis, the importance of fully costing the impacts of a security policy or mechanisms has been stated. To move from being a theoretical model to a practical tool that can be used by decision makers, the Compliance Budget framework needs to be populated by a more complete range of factors than those shown in Figure 6. 'Cognitive load' can be broken down into factors such as recall, password generation, and interference (the effect of using multiple instances of the same information type – such as having to manage 25 or more passwords). As these factors are *subjective* this list can only be populated through discourse with the relevant users, although some likely candidates can be suggested, such as the impact of task interruption and the pressure of the business process. Additionally, the weighting and importance of, and the relationship between these factors, needs to be researched in more detail. As mentioned above pressure from the business process will affect the cost associated with missed opportunities and any delay or blocking of the primary task. The dependency of such costs on the business environment is one such relationship that future work may focus on. One approach would be to elicit a deliberately wide range of factors such as those mentioned above, but extended to include contextual factors like the type of organisation and personal factors based on an understanding of individual users' subjective experience, through consultation with expert security researchers. Due to the high cost (in terms of both time and the resources needed for incentivisation) of approaching users directly this approach would likely be the most suitable. It leverages the existing knowledge of the research community to form a starting point for what will, due to the necessity of including feedback from multiple stakeholder perspectives, be an iterative process. This list could then be presented to users through a weighting and comparison task in order to rank the factors. Links to industry and public sector organisations are necessary in order to provide a working organisational environment

so that the results are grounded in an appropriate context. A year-long project involving government and industry partners that pursues this research direction is currently underway. This project aims to combine data gathered in the commercial organisations linked to the project with expertise from the security community to develop guidelines and recommendations for the more effective management of compliance. One of the goals of this project is to modify the NASA TLX system (described in Section 5.1.1.3) in to a Security TLX by replacing the factors used in the NASA TLX with those that are more relevant to security. As described in Section 5.3 the use of NASA TLX in a security context can lead to confusion among participants and this motivates the need to re-design it for future use. The factors used in the re-design will be determined by the research approach outlined above. This tool can then be used to measure effort associated with various tasks, allowing us to better understand the links between user effort and compliance.

Models such as the Compliance Budget (found in Section 4.5), and the systems models suggested as an executable outcome of constructing the systems structure presented in Section 6.3, need valid empirical data to produce reliable predictions, as expressed by the principle of 'garbage in, garbage out' [49], discussed in Section 3.4. A key direction for future work is the continued development of methodologies for reliable and efficient data gathering to support security decision-making that fully accounts for human factors and the formulation of these models in particular. A cross-disciplinary approach that combines expertise in instrument design (e.g. from HCI) with security expertise is most likely to produce the required tools. This research direction should pursue two objectives in parallel:

1. the collection of an empirical data set grounded in an operational context, and
2. the development of a set of methodological instructions that can be exported from this research context to other researchers or interested parties to allow them to follow the same process without intervention from myself (or associated colleagues).

The continuing goal of this research is to offer support to security decision makers. The usefulness of such an empirical data set has been outlined already. However, not all organisations will have the resources or expertise to manager their own data collection. The data gathering and analysis approaches contained in this thesis require them to be tailored to each organisational context in which they are implemented. For example, the SCs and SOs (see Sections 6.3.1 and 6.3.2.1) used to populate the security architecture discussed in Section 6.3 must contain specific information about the system they are instantiating. The scale of this task, and thus the practicality of this method, has not been tested. The management effort required to map a large organisation and each of the SCs and SOs it would contain would necessarily be huge. The biggest barrier to this operation would be the time investment required to successfully achieve this goal. At minimum this would require a period of 6 months, and realistically could take several years. Obviously this would be

affected by the size of the organisation and the number of staff assigned to the project. It is not unknown for organisations to invest several years in rolling out new system architecture and I envisage that the mapping of the organisation into SCs and SOs would take place as part of a project of such a scale. The managers of an organisation wishing to pursue this approach would have to be prepared to commit to a long-term project that changes the way they approach security. It is not something that can be achieved in the short term. Until refinements to the method are implemented this will likely make it impractical in its current form for the majority of organisations.

That said, the approach that is being suggested here represents something of a paradigm shift in how organisations approach security decisions making. While the investment will be large the benefits have the potential to exceed that. The challenge for researchers is to demonstrate the gains that can be achieved through smaller scale interventions first. It is not necessary for the principals to be applied organisation-wide in the initial case. The recent work of my research group has focused on exploring how to identify the compliance issues with specific mechanisms within organisations, such as access control, and how to optimally select interventions based on an analysis of the user population.

If a meaningful change is to be effected within the security community then it is not enough for such tools to only be available to a single researcher (or research group). The approaches contained within this thesis need to be incorporated in to tools that can be used by those without specific human factors expertise, or developed into training and services that impart the relevant expertise in a concise fashion. Security practitioners would be the focus of such information dissemination, although researchers could also use the knowledge to direct future work or their own.

A necessary first step in this direction would be to build up a database of instances of non-compliance within an organisation, and gather data on the behaviour and attitudes of users that lead to these non-compliance events. One possible approach is to use semi-structured interviews to identify common areas of friction between the business and security processes, as well as the behaviours and strategies used by the users to cope with them. These situations can then be developed into a scenario-based survey using the psychometric design principles of *integrity testing* [24]. These principles allow tests to be created that make it difficult for participants to 'game the system' – that is try to record the answer they think the experimenter wants to hear, rather than what they honestly believe. This is particularly relevant when questioning participants about security as the fear of the consequences of admitting to insecure behaviour may drive them to falsehoods. Participants will be asked to rank the severity of the non-compliant behaviours presented (that do not include a 'correct answer as the presence of such would undoubtedly bias the results) as well as the likelihood that they themselves would use those behaviours.

I recently conducted a study of this kind with an industry partner, as a post-doctoral researcher. The survey was formed of 10 scenarios that presented situations in which the business process and security policy came into conflict. 5 of the scenarios assessed behaviour and 5 looked at attitudes. For each behaviour scenario we created 4 possible non-compliant actions that would allow the actors in the survey to complete their business task at the expense of security. The scenario topics themselves were generated from two sources. The first was a set of 127 interviews conducted during an early study with the organisation's employees in both the UK and the US. These interviews explored attitudes towards security and knowledge of the organisation's security policy. Common issues were extracted and expanded into approximately half the scenarios. The others were developed through consultation with the organisation's security personnel and focused on particular areas of concern the organisation had. The results of the survey allows us to identify 'hotspots' where options that were rated as severe were still commonly selected, indicating that non-compliance was motivated not by a lack of knowledge but by other factors, such a lack of agility in the security process. Additionally, we were able to detect a cultural difference between the US and UK divisions of the organisation. Again the purpose was to provide security managers with the tools to identify key areas of insecurity and the information needed to implement an improved solution that takes in to account the human factor. In this case the survey results directly informed the content and targeting of an awareness and training campaign that is soon to be launched within the organisation.

Any future work should retain the principle that a unified approach is needed to address information security problems. Framing such problems in the economic notion of utility serves the purpose of creating a unified environment in which the various strands of research can be pulled together. For example, Pallas [56] uses principles from the field of *New Institutional Economics*, such as transaction costs and the cost of information asymmetries, to frame his insights. An additional area of interest would therefore be to provide a utility-based system that would allow security managers to explore the potential impacts of their decisions through modelling and simulation. Parkin et al. [58] provides a prototype of how such a tool might be implemented in practise. The creation of a similar tool where in the executable portion is based on the structured system presented in Chapter 6 remains an outstanding goal. The effectiveness of such approaches in changing the way security professionals think about and justify security decisions is being explored by work such as Baldwin et al. [24]. This study extends the results found in [13] and discussed in Section 3.5. They found that methods based on security economics changed the decision making process of security professionals, indicating the potential of future work in this area.

References

1. A. Adams, M.A. Sasse. Users are not the enemy: Why users compromise security mechanisms and how to take remedial measures. *Communications of the ACM, 42 (12), pp. 40-46* December 1999.
2. J. Adams, 1995. *Risk.* Routledge, 1995.
3. R. Anderson, 1994. Why cryptosystems fail. In: *Communications of the ACM* Vol. 37 No. 11 pp. 32-40, 1994.
4. R. Anderson, 2001. Why information security is hard – an economic perspective. In: *17th annual proceedings of the Computer Security Applications Conference*, 2001.
5. R. Anderson, 2002. Unsettling parallels between security and the environment. In: *Proceedings of the second Workshop on the Economics of Information Security.* Berkley, USA 2002.
6. R. Anderson, 2008. *Security Engineering: A guide to building dependable distributed systems.* Wiley Publishing Inc. ISBN 978-0-470-06852-6.
7. R. Anderson, T. Moore. 2009. Information Security: Where computer science, economics and psychology meet. In: *Philosophical transactions of the Royal Society*, Vol. 367, No. 1898, July 2009.
8. S. Arnell, A. Beautement, P. Inglesant, B. Monahan, D. Pym and A. Sasse, 2011. Systematic decision making in security management modelling password use and support. *HP Laboratories Technical report No. HPL-2011036,* 2011.
9. W. J. Baumol, A. S. Blinder, 2009. *Economics; Principles & Policy.* South-Western Cengage Learning. ISBN 978-032458206.
10. A. Beautement, R. Coles, J. Griffin, C. Ioannidis, B. Monahan, D. Pym, A. Sasse, and M. Wonham, 2008. Modelling the Human and Technological Costs and Benefits of USB Memory Stick Security. In: M. Eric Johnson, editor, *Managing Information Risk and the Economics of Security*, pages 141–163. Springer, 2008.
11. A. Beautement, M. A. Sasse and M. Wonham, 2008. The Compliance Budget: Managing security behaviour in organisations. In: *Proceedings of the 2008 Workshop on New Security Paradigms*. ACM, Lake Tahoe, California, US, 2008.
12. D. Begg, 2008. *Economics.* McGraw-Hill Higher Education. ISBN 978-0077117870.
13. Y. Beresnevichiene, D. Pym, and S. Shiu, 2010. Decision support for systems security investment. In: *Network Operations and Management Symposium Workshops (NOMS Wksps), 2010 IEEE/IFIP, pp. 118 -125,* 2010.
14. M. Bishop, 2006. *Computer Security: Art and Science.* Addison-Wesley ISBN 0-201-4409-7.
15. S. Brostoff, M. A. Sasse. Are Passfaces™ more usable than passwords? A field trial investigation. In S. McDonald, Y. Waern & G. Cockton [Eds.]: People and Computers

XIV - Usability or Else! *Proceedings of HCI 2000 (September 5th - 8th, Sunderland, UK), pp. 405-424.* Springer, 2000.

16. S.K. Card, T. P. Thomas and A. Newall, 1983. *The Psychology of Human-Computer Interaction*. London: Lawrence Erlbaum Associates, 1983 ISBN 0-89859-243-7
17. H. Cavusoglu, H. Cavusoglu, J. Zhang. Economics of Security Patch Management. In: *The Fifth Annual Workshop on Economics and Information Security.* Cambridge, UK 2006.
18. Y. Chen, B. Boehm, L. Sheppard, 2007. Measuring security investment benefit for off the shelf software systems – A stakeholder value driven approach. In: *The Sixth Annual Workshop on Economics and Information Security.* Carnegie-Mellon, USA 2007.
19. K. Everitt, T. Bragin, J. Fogarty and T. Kohno, 2009. A comprehensive study of frequency, interference, and training of multiple graphical passwords. *In ACM Conference on Human Factors in Computing Systems (CHI),* April 2009.
20. F. Farahmand, S. Navathe, G. Sharp, P. Enslow, 2004. Evaluating Damages Caused by Information Systems Security Incidents. In: *The Economics of Information Securit*y, J. Camp and R. Lewis (eds.), Kluwer, 2004, pp. 85-94.
21. I. Flechais, 2005. *Designing secure and usable systems*. Ph.D. University College London.
22. D. Florencio and C. Herley, 2007. A large scale study of web password habits. *In proceedings of the 16th international conference on the World Wide Web,* 2007.
23. D. Florencio and C. Herley, 2010. Where do security policies come from? *The sixth Symposium On Usable Privacy and Security,* 2010
24. A. Furnham and J. Taylor, 2011. *Bad apples: Identify, prevent and manage negative behaviour at work.* Palgrave Macmillan, March 2011. ISBN-10: 0230584748
25. L.A. Gordon and M.P. Loeb, 2002. The Economics of Information Security Investment. In: *ACM Transactions on Information and Systems Security, 5(4):438–457,* 2002.
26. L. A. Gordon, M. P. Loeb, 2004. The economics of information security investment. In: *ACM Transactions on Information and System Security*, 5:438{457, November 2002. Reprinted in *Economics of Information Security*, 2004, Springer, Camp and Lewis, eds.
27. L. A. Gordon, M. P. Loeb, 2006. *Managing Cybersecurity resources: A cost-benefit analysis.* Mcgraw-Hill. ISBN 978-0071452854.
28. S. G. Hart, 2006. NASA-Task Load Index (NASA-TLX); 20 Years Later. *Proceedings of the Human Factors and Ergonomics Society 50th Annual Meeting, 904-908. Santa Monica: HFES,* 2006.
29. S. G. Hart and L. E. Staveland, 1998. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. *In P. A. Hancock and N. Meshkati (Eds.) Human Mental Workload. Amsterdam: North Holland Press,* 1998.
30. K. Hausken. Strategic Defence and Attack of Complex Networks. In *The Sixth Annual Workshop on Economics and Information Security.* Carnegie-Mellon, USA 2007.

31. B. von Helversen, G. H. E. Gendolla, P. Winkielman, R. E. Schmidt. Exploring the hardship of ease : subjective and objective effort in the ease-of-processing paradigm. *Motivation and emotion, Vol. 32, H. 1. S. 1-10,* 2008.
32. M. G. Helander, T. K. Landauer and P. V. Prabhu. (eds.), 1997. *Handbook of human-computer interaction*. Elsevier, 1997.
33. H. S. B. Herath, T. C. Herath. Cyber-Insurance: Copula pricing framework and implications for risk management. In *The Sixth Annual Workshop on Economics and Information Security,* 2007.
34. C. Herley, 2009. So long, and no thanks for the externalities: the rational rejection of security advice by users. In: *Proceedings of the 2009 New Security Paradigms Workshop,* 2009.
35. P. Honeyman, G. A. Schwartz and A van Assche, 2007. Interdependence of reliability and security. In: *Proceedings of the Sixth Annual Workshop on Economics and Information Security.* Carnegie-Mellon, USA 2007.
36. HP Labs, 2001. *Exploring Information Stewardship with the Cloud Ecosystem Model.* [online] Available at <http://www.hpl.hp.com/news/2011/oct-dec/Final_Report_collated.pdf> [Accessed 31st August 2012].
37. C. D. Huang, Q. Hu, R. S. Behara, 2006. Economics of information security investment in the case of simultaneous attacks. . In: *The Fifth Annual Workshop on Economics and Information Security.* Cambridge, UK 2006.
38. P. Inglesant, M. A. Sasse, 2010. Studying Password Use in the Wild: Practical Problems and Possible Solutions. *Symposium On Usable Privacy and Security (SOUPS) 2010, July 14-16, 2010, Redmond, WA, USA,* 2010.
39. P. Ingelsant, M. A. Sasse, 2010. The true cost of unusable password policies: password use in the wild. *In Proceedings of the 28th international conference on Human factors in computing systems. Atlanta, Georgia, USA: ACM, pp. 383-392*, 2010.
40. International Standards Office, 2005. ISO/IEC 27002 Risk Analysis.
41. C. Ioannidis, D. Pym, and J. Williams, 2009. Information security trade-offs and optimal patching policies. *Manuscript*, 2009.
42. M. E. Johnson, E. Goetz, 2006. Embedding information security risk management into the extended enterprise. *IEEE Security and Privacy*. June, 2006.
43. M. E. Johnson and E. Goetz, 2007. Embedding Information Security into the Organisation. *IEEE Security & Privacy May/June 2007 pp. 16 – 24,* 2007.
44. M. E. Johnson, S. Dynes, 2007. Inadvertent disclosure – Information leaks in the extended enterprise. In: *The Sixth Annual Workshop on Economics and Information Security.* Carnegie-Mellon, USA 2007.
45. A. Kerckhoffs, 1883. La cryptographie militaire. *Journal des Sciences Militaires, pp. 5-38*, 1883.
46. K. S. Killourhy and R. A. Maxion, 2009. Comparing Anomaly-Detection Algorithms for Keystroke Dynamics. In: *International Conference on Dependable Systems & Networks*

(DSN-09), pp. 125-134, Estoril, Lisbon, Portugal, 29 June to 02 July 2009. IEEE Computer Society Press, Los Alamitos, California, 2009.

47. V. Kumar, R. Telang, T. Mukhopadhyah, 2006. Enterprise information security: Who should manage it and how? In: *The Fifth Annual Workshop on Economics and Information Security.* Cambridge, UK 2006.
48. V. Kumar, R. Telang, and T. Mukhopahhyay, 2007. Optimally securing interconnected information systems and assets. In: *Proceedings of the Sixth Workshop on the Economics of Information Security,* Carnegie Mellon, USA 2007.
49. W. Lidwell, K. Holden and J. Butler, 2010. Universal principles of design: 125 ways to enhance usability, influence perception, increase appeal, make better design decisions and teach through design. *Rockport Publishers,* 1 Jan 2010. ISBN-10: 1592535879
50. Microsoft, 2012. *Create strong passwords.* [online] Available at <http://www.microsoft.com/en-gb/security/online-privacy/passwords-create.aspx> [Accessed 31st August 2012].
51. K. Mitnick, 2002. *The Art of Deception: Controlling the Human Element of Security.* Wiley Publishing Inc. ISBN 0-471-23712-4.
52. NASA, 2012. *NASA TLX: Task Load Index.* [online] Available at <http://humansystems.arc.nasa.gov/groups/TLX/> [Accessed 31st August 2012].
53. New York Times, 2006. *A Face is Exposed for AOL Searcher No. 4417749.* [online] Available at <http://www.nytimes.com/2006/08/09/technology/09aol.html> [Accessed 31st August 2012].
54. NIST, 1995. *An Introduction to Computer Security: The NIST Handbook*. [pdf] USA: NIST. Available at <http://csrc.nist.gov/publications/nistpubs/800-12/handbook.pdf>
55. D. A. Norman, 1983. Some Observations on Mental Models. In Gentner, D.A. & Stevents, A. A. [Eds.] *Mental Models.* Hillsdale, NJ: Erlbaum, 1983.
56. F. Pallas, 2009. *Information security inside organizations*. Ph.D. Berlin University. [online] Available at <http://opus.kobv.de/tuberlin/volltexte/2009/2320/pdf/pallas_frank.pdf> [Accessed 31st August 2012]
57. D. Parker, 1992. *Fighting Computer Crime: A New Framework for Protecting Information.* Wiley Publishing Inc. ISBN 978-0471163787.
58. S. Parkin, A. van Moorsel, P. Inglesant and M. A. Sasse, 2010. A stealth approach to usable security: Helping IT security managers to identify workable security solutions. *Proceedings of the 2010 Workshop on New Security Paradigms. ACM Press, New York, US, pp. 33-49*, 2010.
59. Personal conversation with Robert Coles
60. J. Reason, 1998. Achieving a safe culture: Theory and practice. *Work and stress, vol 12 (3), pp. 293 – 306*, 1998.

61. M. A. Sasse, S.Brostoff, and D. Weirich, 2001. Transforming the "weakest link": A human-computer interaction approach to usable and effective security. *BT Technology Journal, Vol 19 (3) pp. 122-131*, July 2001.
62. B. Schneier, 1999. *Attack Trees.* [online] Available at <http://www.schneier.com/paper-attacktrees-ddj-ft.html> [Accessed 31st August 2012]
63. B. Schneier, 2000. Secrets and Lies: Digital Security in a Networked World, *John Wiley & Sons*, 2000. ISBN 0-471-25311-1
64. B. Schneier, 2002. Computer Security: It's the Economics, stupid. In: *Proceedings of the First Workshop on the Economics of Information Security.* Berkley, USA 2002.
65. S. Sheng, L. Broderick, C. A. Koranda and J. J. Hyland, 2006. Why Johnny still can't encrypt: Evaluating the usability of email encryption software. *The second Symposium On Usable Privacy and Security,* 2006.
66. S. Sheng, B. Magnien, P. Kumaraguru, A. Acquisti, L. F. Cranor, J. Hong and E. Nunge, 2007. Anti-phishing Phil: The design and evaluation of a game that teaches people not to fall for phish. In: *Proceedings of the third Symposium On Usable Privacy and Security,* 2007.
67. Stanford, 2012. *Protégé.* [online] Available at <http://protege.stanford.edu/> [Accessed 31st August 2012].
68. G. Stoneburner, A. Goguen, A. Feringa. Risk management guide for information technology systems: Recommendations of the National Institute of Standards and Technology. *NIST Special Publication 800-30* 2002. UK: ISO.
69. A. Strauss and J. Corbin, 1998. Basics of qualitative research: Techniques and procedures for developing grounded theory. *Sage Publications inc. 2nd Edition,* 1998.
70. L. Sweeney, 2000. *Uniqueness of Simple Demographics in the U.S. Population, LIDAPWP4*. Pittsburgh, PA Carnegie Mellon University, Laboratory for International Data Privacy.
71. Trust Economics, 2008. *Trust Economics.* [online] Available at <https://wiki.cs.ncl.ac.uk/trusteconomics/TrustEconomics> [Accessed 31st August 2012].
72. T. Valentine. An evaluation of the Passfaces™ personal authentication system *(Technical Report) London: Goldsmiths College University of London,* 1998.
73. T. Valentine. Memory for Passfaces™ after a long delay. *(Technical Report) London: Goldsmiths College University of London*, 1999.
74. H. Varian, 2000. Managing online security risks. *The New York Times*, 1 June 2000.
75. W3C, 2012. *OWL 2 Web Ontology Language Document Overview.* [online] Available at <http://www.w3.org/TR/owl2-overview/> [Accessed 31st August 2012].
76. D. Weirich, 2005. *Persuasive password security*. Ph.D. University College London, 2005.

77. D. Weirich, M. A. Sasse. Pretty Good Persuasion: A first step towards effective password security for the Real World. *Proceedings of the New Security Paradigms Workshop 2001 (Sept. 10-13, Cloudcroft, NM), pp. 137-143.* ACM Press, 2001.
78. A. Whitten and J. D. Tyger, 1999. Why Johnny can't encrypt: A usability evaluation of PGP 5.0. In: *Proceedings of the 8th conference on USENIX Security Symposium, Volume 8*, 1999.
79. Wikipedia, 2012. *Availability.* [online] Available at <http://en.wikipedia.org/wiki/Availability> [Accessed 31st August 2012].
80. Wikipedia, 2012. *Category mistake.* [online] Available at http://en.wikipedia.org/wiki/Category_mistake> [Accessed 31st August 2012].
81. J. Willemson, 2006. On the Gordon & Loeb model for information security investment. In: *The Fifth Annual Workshop on Economics and Information Security.* Cambridge, UK 2006.
82. M. Zviran and W. J. Haga, 1993. A comparison of password techniques for multilevel authentication mechanisms. *The Computer Journal 36 (3), pp. 227-237*, 1993.

Background Reading

1. R. Anderson and T. Moore. The economics of information security: A survey and open questions. *Science, vol. 314, pp. 610-613*, Oct. 2006.
2. S. Chiasson, A. Forget, R. Biddle, P. C. Van Oorschot. Influencing Users Towards Better Passwords: Persuasive Cued Click-Points. *Proceedings of HCI*, 2008.
3. M. Collinson, B. Monahan, and D. Pym. A logical and computational theory of located resource. *Journal of Logic and Computation, 19(6):1207–1244*, 2009.
4. M. Collinson, B. Monahan, and D. Pym. Semantics for structured systems modelling and simulation. *In Proc. Simutools 2010. ICST: ACM Digital Library and EU Digital Library*, 2010. ISBN: 78-963-9799-87-5.
5. R. De Neufville and A.R. Odoni. Airport systems: planning, design, and management. *McGraw-Hill Professional*, 2003.
6. Peter Neumann. Computer-Related Risks. *Addison-Wesley*, 1995.
7. S. Parkin, A. van Moorsel, P. Inglesant, M.A. Sasse. A Stealth Approach to Usable Security: Helping IT Security Managers to Identify Workable Security Solutions. *NSPW 2010: Proceedings of the 2010 Workshop on New Security Paradigms. ACM Press, New York, US, pp. 33-49*, 2010
8. Schneier, Bruce. Beyond Fear: Thinking Sensibly about Security in an Uncertain World, *Copernicus Books*, 2003. ISBN 0-387-02620-7
9. K. P. Yee. User Interaction Design for Secure Systems. *In L. Faith Cranor & S. Garfinkel [Eds.]: Security and Usability: Designing secure systems that people can use 2005, pp 13-30.* O'Reilly Books, 2005.

10. M. E. Zurko, R. T. Simon. User-Centred Security. *New Security Paradigms Workshop,* 1997.
11. M. Collinson, B. Monahan, D. Pym. Located Demos2k - Towards a Tool for Modelling Processes and Distributed Resources. *HP Labs Technical Report No. HPL-2008-76,* 2008.

Appendix A: Sample configuration file for the Maxion Text Propmpter

```
################################################################################
#
#  UCL1 - May 09
#
#  (MTP Configuration File)
#
################################################################################

################################################################################
# Configuration
################################################################################

# This UCL1 config file for MetriTextPrompter demonstrates
# how you can set up a simple password experiment in MTP.

# The name of the experiment
[Set ExptName = UCL1_Experiment]
[Set DataDirectory = Data]

# Prompter window properties
[Set AutoMaximize = True]
#[Set windowwidth = 700]
#[Set windowheight = 700]

# Prompt and Box Font properties
# "Prompt" refers to the label where MTP displays prompts to the participant.
# "Box" refers to the box supplied by MTP for the participant to type in.
# You can set the font face, size, bold/not bold, and font color for both
# the Prompt and the Box.
# The font is any valid font name on your Windows machine.
# The size is in points.
# The color is in hexadecimal BGR format (see above for details).

[Set PromptFont = Arial]
[Set PromptFontSize = 15]
[Set PromptFontBold = true]
[Set PromptFontColor = &h000000]

[Set BoxFont = Arial]
[Set BoxFontSize = 15]
[Set BoxFontBold = true]
[Set BoxFontColor = &H000000]

# Each line that is not a command will be displayed to the participant,
# and a single-line textbox is provided type into.

################################################################################
# Start experiment
################################################################################

## Begin by asking the subject to type their first and last name (to
## confirm that there is no mixup between subjects.)
```

[Set ErrorCheck = None]
[CheckNumChars 3]
[Set CheckNumCharsErrorMsg = Please type your full first and last name.]

Please type your first and last name.

Next, tell the subject what they have to do.

[Message]
TASK 1 _
[br]_
[br]_
Please type in the policy 1 (minimum of 6 characters) password you chose. _
[br]_
If you try to enter an incorrect password you _
[br]_
will have to type the text from the beginning. _
[br]_
[br]_
To complete this task you will need to enter the password correctly 10 times. _
[br]_
[br]_
Please press the enter key to start typing.

[Set ErrorCheck = WholeString]

[Set HidePromptBoxText = True]

[Set PromptFontColor = &hE8E8E8]

G08ildA
G08ildA
G08ildA
G08ildA
G08ildA
G08ildA
G08ildA
G08ildA
G08ildA
G08ildA

[Set PromptFontColor = &h000000]

[Message]

TASK 2 _
[br]_
[br]_
Please type in the policy 2 (minimum 8 characters) password you chose. _
[br]_
If you try to enter an incorrect password you _
[br]_
will have to type the text from the beginning. _
[br]_

[br]_
To complete this task you will need to enter the password correctly 10 times. _
[br]_
[br]_
Please press the enter key to start typing.

[Set PromptFontColor = &hE8E8E8]

T1h9e4o6xarh
T1h9e4o6xarh
T1h9e4o6xarh
T1h9e4o6xarh
T1h9e4o6xarh
T1h9e4o6xarh
T1h9e4o6xarh
T1h9e4o6xarh
T1h9e4o6xarh
T1h9e4o6xarh

[Set PromptFontColor = &h000000]

[Message]
TASK 3 _
[br]_
[br]_
You will now be asked to alternately enter one _
[br]_
of the passwords you have just typed and answer _
[br]_
a series of general knowledge questions.

[message]
Now type your policy 1 password. _
[br]_
[br]_
Please press the enter key to start typing.

[Set ErrorCheck = WholeString]

[Set PromptFontColor = &hE8E8E8]

G08ildA

[set ErrorCheck = None]

[Set HidePromptBoxText = False]

[Set PromptFontColor = &h000000]

Complete this phrase. As sick as a.. _
[br]_
[br]_
A. Penguin _
[br]_
B. Parrot _

[br]_
C. Puffin _
[br]_
D. Partridge

[message]
Now type your policy 2 password. _
[br]_
[br]_
Please press the enter key to start typing.

[Set ErrorCheck = WholeString]

[Set HidePromptBoxText = True]

[Set PromptFontColor = &hE8E8E8]

T1h9e4o6xarh

[set ErrorCheck = None]

[Set HidePromptBoxText = False]

[Set PromptFontColor = &h000000]

Which legal document states a person's wishes _
[br]_
regarding the disposal of their property after death? _
[br]_
[br]_
A. Will _
[br]_
B. Shall _
[br]_
C. Would _
[br]_
D. Should

[message]
Now type your policy 1 password. _
[br]_
[br]_
Please press the enter key to start typing.

[Set ErrorCheck = WholeString]

[Set HidePromptBoxText = True]

[Set PromptFontColor = &hE8E8E8]

G08ildA

[set ErrorCheck = None]

[Set HidePromptBoxText = False]

[Set PromptFontColor = &h000000]

Complete the title of the James Bond film The Man With The Golden.. _
[br]_
[br]_
A. Gun _
[br]_
B. Tooth _
[br]_
C. Delicious _
[br]_
D. Eagle

[message]
Now type your policy 2 password. _
[br]_
[br]_
Please press the enter key to start typing.

[Set ErrorCheck = WholeString]

[Set HidePromptBoxText = True]

[Set PromptFontColor = &hE8E8E8]

T1h9e4o6xarh

[set ErrorCheck = None]

[Set HidePromptBoxText = False]

[Set PromptFontColor = &h000000]

In which sport do two teams pull at the opposite ends of a rope? _
[br]_
[br]_
A. Ice hockey _
[br]_
B. Basketball _
[br]_
C. Tug of War _
[br]_
D. Polo

[message]
Now type your policy 1 password. _
[br]_
[br]_
Please press the enter key to start typing.

[Set ErrorCheck = WholeString]

[Set HidePromptBoxText = True]

[Set PromptFontColor = &hE8E8E8]

G08ildA

[set ErrorCheck = None]

[Set HidePromptBoxText = False]

[Set PromptFontColor = &h000000]

Where would a cowboy wear his chaps? _
[br]_
[br]_
A. On his hands _
[br]_
B. On his arms _
[br]_
C. On his legs _
[br]_
D. On his head

[message]
Now type your policy 2 password. _
[br]_
[br]_
Please press the enter key to start typing.

[Set ErrorCheck = WholeString]

[Set HidePromptBoxText = True]

[Set PromptFontColor = &hE8E8E8]

T1h9e4o6xarh

[set ErrorCheck = None]

[Set HidePromptBoxText = False]

[Set PromptFontColor = &h000000]

Sherpas and Gurkhas are native to which country? _
[br]_
[br]_
A. Ecuador _
[br]_
B. Morocco _
[br]_
C. Nepal _
[br]_
D. Russia

[message]

Now type your policy 1 password. _
[br]_
[br]_
Please press the enter key to start typing.

[Set ErrorCheck = WholeString]

[Set HidePromptBoxText = True]

[Set PromptFontColor = &hE8E8E8]

G08ildA

[set ErrorCheck = None]

[Set HidePromptBoxText = False]

[Set PromptFontColor = &h000000]

Whose autobiography has the title 'A Long Walk To Freedom'? _
[br]_
[br]_
A. Ranulph Fiennes _
[br]_
B. Nelson Mandela _
[br]_
C. Mikhail Gorbachev _
[br]_
D. Mother Teresa

[message]
Now type your policy 2 password. _
[br]_
[br]_
Please press the enter key to start typing.

[Set ErrorCheck = WholeString]

[Set HidePromptBoxText = True]

[Set PromptFontColor = &hE8E8E8]

T1h9e4o6xarh

[set ErrorCheck = None]

[Set HidePromptBoxText = False]

[Set PromptFontColor = &h000000]

Complete this stage instruction in Shakespeare's The Winter's Tale: _
[br]_
"Exit, pursued by a .." _
[br]_
[br]_

A. Dog _
[br]_
B. Tiger _
[br]_
C. Bear _
[br]_
D. Clown

[message]
Now type your policy 1 password. _
[br]_
[br]_
Please press the enter key to start typing.

[Set ErrorCheck = WholeString]

[Set HidePromptBoxText = True]

[Set PromptFontColor = &hE8E8E8]

G08ildA

[set ErrorCheck = None]

[Set HidePromptBoxText = False]

[Set PromptFontColor = &h000000]

The young of which creature is known as a squab? _
[br]_
[br]_
A. Pigeon _
[br]_
B. Salmon _
[br]_
C. Octopus _
[br]_
D. Horse

[message]
Now type your policy 2 password. _
[br]_
[br]_
Please press the enter key to start typing.

[Set ErrorCheck = WholeString]

[Set HidePromptBoxText = True]

[Set PromptFontColor = &hE8E8E8]

T1h9e4o6xarh

[set ErrorCheck = None]

[Set HidePromptBoxText = False]

[Set PromptFontColor = &h000000]

Which of these zodiac signs is not represented by _
[br]_
an animal that grows horns? _
[br]_
[br]_
A. Taurus _
[br]_
B. Capricorn _
[br]_
C. Ares _
[br]_
D. Aquarius

[message]
TASK 4 _
[br]_
[br]_
You will now be asked to complete two tasks, a weights task and a ratings task. _
[br]_
Please ask the experimenter to start these for you. _
[br]_
When both tasks are complete please click the button below.

[Message]
TASK 5 _
[br]_
[br]_
Please type in the policy 1 (minimum of 6 characters) password you chose. _
[br]_
If you try to enter an incorrect password you _
[br]_
will have to type the text from the beginning. _
[br]_
[br]_
To complete this task you will need to enter the password correctly 10 times. _
[br]_
[br]_
Please press the enter key to start typing.

[Set ErrorCheck = WholeString]

[Set HidePromptBoxText = True]

[Set PromptFontColor = &hE8E8E8]

G08ildA
G08ildA
G08ildA
G08ildA
G08ildA

G08ildA
G08ildA
G08ildA
G08ildA
G08ildA

[Set ErrorCheck = None]

[Set PromptFontColor = &h000000]

[Message]

TASK 6 _
[br]_
[br]_
Please type in the policy 2 (minimum 8 characters) password you chose. _
[br]_
If you try to enter an incorrect password you _
[br]_
will have to type the text from the beginning. _
[br]_
[br]_
To complete this task you will need to enter the password correctly 10 times. _
[br]_
[br]_
Please press the enter key to start typing.

[Set ErrorCheck = WholeString]

[Set PromptFontColor = &hE8E8E8]

T1h9e4o6xarh
T1h9e4o6xarh
T1h9e4o6xarh
T1h9e4o6xarh
T1h9e4o6xarh
T1h9e4o6xarh
T1h9e4o6xarh
T1h9e4o6xarh
T1h9e4o6xarh
T1h9e4o6xarh

[Set ErrorCheck = None]

[Set PromptFontColor = &h000000]

[message]
The experiment is now over. _
[br]_
[br]_
Thank you for taking the time to participate.

Appendix B: Instructions for Server Study participants

Thank you for agreeing to take part in the HCS password server study. Please read the following instructions carefully. You will shortly receive an email from 'APET Project; called 'APET Project Password Study'. The experiment server is called 'APET Project' so all communication during the study will come from this email. **It is very important that your email is not pushing these messages into your junk folder.** Please check carefully that this isn't happening and adjust your filters appropriately if it is.

The first email will contain a link to the server itself, as well as some relevant information. The first page you will see after following the link will ask you to enter your email address. Then you will be asked to register a username and password. Your password should conform to the following rules:

- Be at least 6 characters in length.
- Contain at least 3 of the following categories
    - Upper case letters (A, B, C, . Z)
    - Lower case letters (a, b, c, . z)
    - Numbers (0, 1, 2, . 9)
    - Non-alphanumeric characters (! + ; - % etc.)

All passwords should be created for this study. Please do not re-use any passwords existing passwords you may already have. This is both to allow us to collect more accurate data and also for your security. Once you have registered you simple need to wait for further emails from 'Apet Project'. Each one will contain a link to a login screen. Simply login using the details you registered. If you cannot remember your password there is an option to reset it. Doing so is perfectly fine and should not be considered negative in anyway.

Some final points:

- Please check your email regularly (at least once a day) between now and December the 18th
- Please DO NOT write down your password and login details or use any other memory aids. If you can't remember it just use the reset function.
- You will be paid £20 at the end of the session. This will be in the form of an Amazon voucher which will be delivered by email.

If you have any questions please feel free to ask.