Aiding Information Security Decisions with Human Factors Using Quantitative and Qualitative Techniques.

James Turland

School of Computing Science

Newcastle University

This dissertation is submitted for the degree of

Doctor of Philosophy

June 2016

I dedicate this thesis to my friends and family for their love and support throughout my time at university. I would also like to make a special mention with respect to my partner Emma, whose emergency meals, cups of tea and late night proof reading has made this possible!

Declaration

I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. This dissertation is my own work and contains nothing which is the outcome of work done in collaboration with others, except as specified in the text and Acknowledgements.

James Turland
June 2016

Acknowledgements

First and foremost I would like to offer my thanks and sincerest gratitude to my supervisor Professor Aad van Moorsel for this opportunity, as well as his guidance and tutelage throughout. It has been a pleasure working with him and I have many fond memories of our travels and meals out over the years!

I would also like to thank the various colleagues I have had the pleasure of working with over the past 6 years spent at Newcastle University. In particular, it has been a great experience working in the RISCS project where I have gained valuable research skills and become friends with some great minds.

Finally, I wish to thank all the staff from the Computing Science Department that have in some way been a part of my academic life. I would particularly like to mention the work of the administrative staff and their role in looking after me through my years of service!

Abstract

The Information Security Decision Making Process is comprised of an extremely complex and dynamic set of sub-tasks, sub-goals and inter-disciplinary practices. In order to be effective and appropriate, this process must balance both the requirements of the stakeholder as well as the users within the system. Without careful consideration of users' behaviours and preferences, interventions are often seen as obstacles towards productivity and subsequently circumvented or simply not adhered to. The approach detailed herein requires an intimate knowledge of both Information Security and Human Behaviour.

An effective security policy must adequately protect a given set of assets (human and non-human) or systems as well as preserve maximal productivity. Companies rely on their Intellectual Property Rights which are often stored in a digital format. This presents a plethora of issues regarding security, access management and locality (whether on or off the premises). Furthermore, there is the added complexity of employees and how they operate within this environment (a subset of compliance, competence and policy). With the continued increase in consumerisation, more specifically the rise of Bring Your Own Device, there is a significant threat towards data security that persists outside of the typical working environment. This trend enables employees to access and transfer corporate assets remotely but in doing so creates a conflict over identity, ownership and data management. The governance of these activities creates an extremely complex problem space which requires the need to balance these requirements relying on an accurate assessment of risk, identification of security vulnerabilities and knowledge pertaining to the behaviour of employees.

The risks to company assets can be estimated by the analysis of the following issues:

- **Threats to your assets**. These are unwanted events that could cause the deliberate or accidental loss, damage or misuse of the assets.
- **Vulnerabilities**. How susceptible your assets are to attack.
- **Impact**. The magnitude of the potential loss or the seriousness of the event.

The ability to quantify and accurately represent these variables is critical in developing, implementing and supporting a successful security policy. A methodological based approach

is an effective way to design, simulate and perform impact analysis of a potential new policy adoption. Being able to identify both the cyber and human threats is itself a difficult task furthered by the necessity to accurately populate their data sets.

This thesis documents a methodology towards aiding the policy decision making process. In doing so, we introduce several experimental design methods aimed at understanding user behaviours with respect to 'nudging' in the field of Information Security. Specifically, we contribute the following:

- **Survey & Literature Review:** We conduct a multi-disciplinary review of the relevant literature pertaining to the study of Human Behaviours within Information Security. We examine the 'state-of-the-art' practices with respect to both academic research and industry practices.

- **Propose a Methodology Based on Empirical Data:** We design and investigate the effectiveness of an iterative methodology based on empirical data as opposed to simulated, from two separate investigations (the IRIDIUM Study and our pilot study). Our methodology employs user feedback, CISO interactions and the formulation of bespoke experiments that aim to identify and improve our understanding of specific behaviours highlighted within our pilot studies.

- **Behavioural Interventions - Nudge:** Our research contributes to the field of Human Computer Interaction and Behaviour Psychology by investigating the first application of nudging [228] within an Information Security setting.

- **Bespoke Experimental Design:** We design and investigate three separate user behaviours identified from our CISO interaction and pilot study in an effort to improve our understanding of the problem space. Specifically, we examine the role of nudging and behavioural interventions with respect to Wi-Fi selection, error-reporting, and cookie acceptance. Additionally, the work conducted with respect to these studies required the development of tools that are adaptable to other scenarios, we feel that these are a valid contribution to the literature.

- **Modelling:** We extend the BPMN 2.0 formalism through the introduction of new nodes which aid the expressiveness of the modelling notation. Our contributions aim to improve the precision in which a CISO can model the current environment and design new policy using empirical data obtained within the field regarding user behaviours and practices. The application presented here is somewhat abstract and this is featured in our limitations. We feel, however, that this is a valid approach to formulating a solution within the Information Security domain.

The dissertation is structured as follows. Chapter 1 provides an abstract overview of the problem space and highlights our aims, objectives and publications. Chapter 2 details an in-depth literature review of the cross-disciplinary problem space. This involves both the analysis of industry standards, practices and reports as well as a summary of academic literature pertaining to theoretical frameworks and simulations for discussion. Chapter 3 introduces our problem space and documents the rationale for designing our methodology. Each successive chapter (4, 5, & 6) documents a separate investigative strategy for populating specific data sets with respect to the behaviours and practices highlighted from our pilot study and CISO interaction. This provides the rationale behind each approach as well as a documented implementation and evaluation of our experimental design with reference to publications in the field. Chapter 7 documents our modelling strategy and highlights the extensions we propose to the BPMN 2.0 formalism. Chapter 8 concludes our work with reference to our contributions, limitations and the direction of future study.

Table of contents

List of figures

List of tables

Chapter 1

Introduction

Information Security Policies (ISPs) are designed and enforced to protect Intellectual Property Rights (IPR), assets, and employees. Failure to protect these often has profound implications, resulting in financial and reputation loss. The need for a robust Decision Making Process to formulate such policies is of great importance and must be appropriate to the environment in which it is intended for. Lobban [98] states "about 80 per cent of known attacks would be defeated by embedding basic information security practices for your people, processes and technology", a figure that supports the notion that humans are the biggest threat to security. We must therefore ensure that user practices are understood, supported with appropriate tools, educational material and technologies as well as having effective remediation processes in place.

Attack vectors and the complexity of the attacks are growing exponentially, "what was considered a sophisticated cyber attack only a year ago might now be incorporated into a downloadable and easy to deploy internet application, requiring little or no expertise to use" [98]. With a continually evolving threat environment and a multitude of vulnerabilities and access points it is imperative to devise a strategy to manage cyber risks.

To do this we require a robust understanding of our environment and the critical factors that govern the behaviour within it. Without this intimate knowledge it is impossible to formulate an appropriate strategy for defence that fully reflects the needs of the user and the organisation.

Figure 1.1 presents a high level overview of a typical working environment in the modern age. The abstraction aims to highlight the numerous working environments along with the data that traverses these environments. The figure also depicts the wealth of attack vectors that are exploitable by malicious outsiders and insiders.

The birth of consumerisation and Bring Your Own Device (BYOD) has enabled users to break the spatial divide and be able to operate remotely (access corporate data off-site).

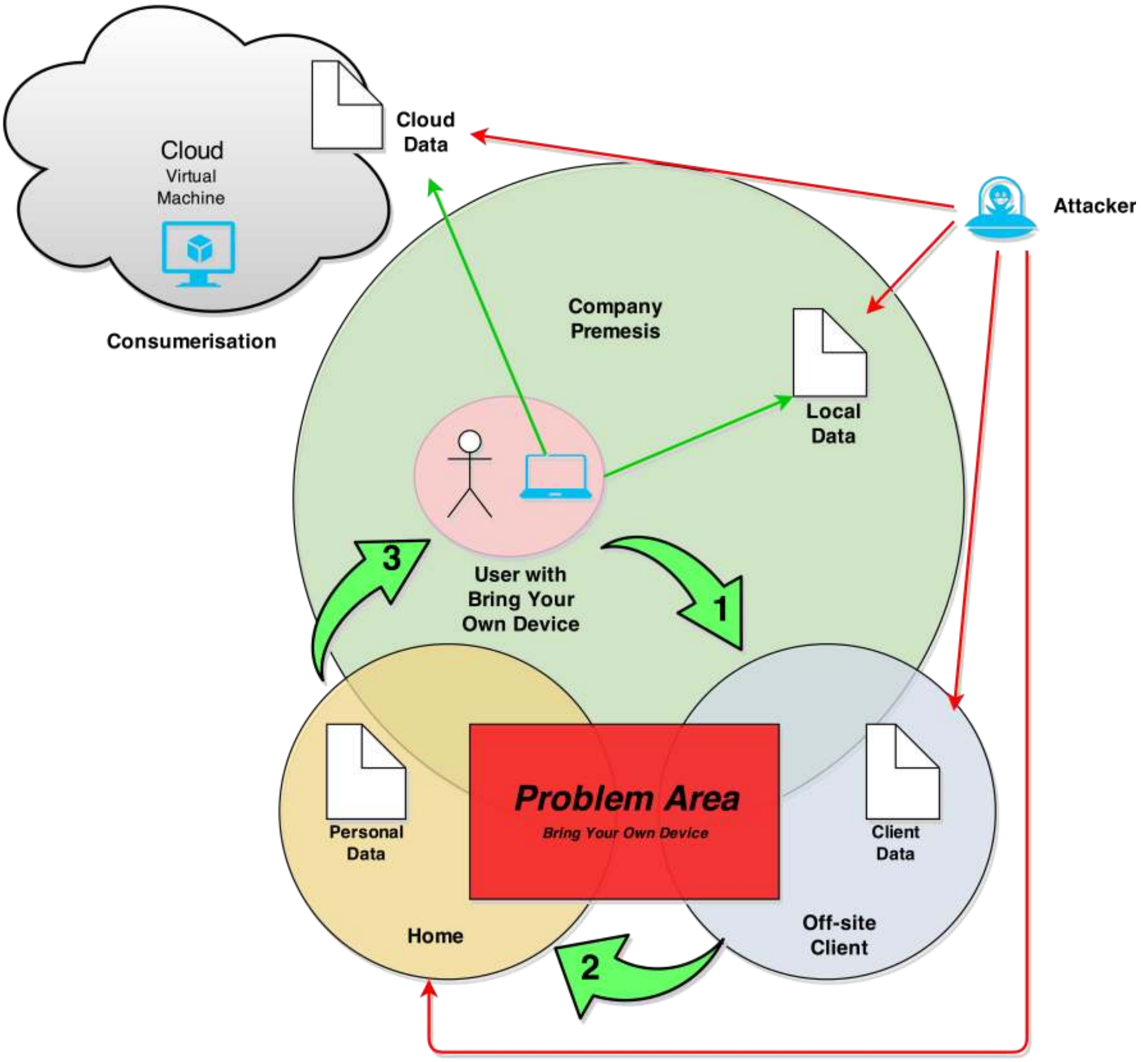

Fig. 1.1 Data movement within the modern working environment

This is both beneficial and problematic. BYOD can be seen to boost the productivity of users as well as the profitability of the company as the potential client base is expanded and the number of effective working hours are increased. Users benefit from being able to work during transit as well as being able to collaborate off-site (internationally as well). There are, however, numerous issues regarding end-point security, access management and the vulnerabilities regarding the access mechanism.

Let us consider a typical working day involving off-site activities (see figure 1.1). At arrow position 1 the user is located on the company premises and is equipped with their own laptop and smartphone that is allowed to be used in the corporate (allowed to access company resources locally and cloud-based) and non-corporate environment. A typical day may involve the employee working remotely or travelling to a client to conduct off-site activities, requiring remote access to corporate data and cloud-based materials. Any work is conducted on the user's device and uploaded to central/cloud-based servers whenever possible (perhaps passively via a VPN or file-sharing platform). The user finishes their working day and returns home for the evening. Their device may then be used in their home environment and not necessarily restricted to a single user.

The above scenario boasts a wealth of critical security issues that require an effective security policy to manage. The issues are not only from the computer security discipline, but also require an understanding of Risk Analysis and Human Behaviour. In this example, it is effective to envisage the vulnerabilities from an attacker's perspective allowing one to identify weaknesses within the system.

In an abstract view, without focusing on a specific method of attack or exploit, it is clear to see where the vulnerabilities reside. With the user's device being transient and user managed, the very notion of a changing threat environment creates a difficult problem space to assess and quantify from both the user's and security officer's perspective. Threats are numerous and mutate to match the vulnerabilities in each given environment, but may also lay dormant on the user's device until they enter their specific targeted environment. This means that an attack may be on-going and persist between locations (a vulnerability may be exploited in one location to impact a flaw in another). The fact that the device is user managed and operating outside of a controlled environment inevitably increases the likelihood of a successful attack.

Changing threat environments and individually managed devices require effective security policies. Before consumerisation and BYOD, a purely technical 'blanket' solution could effectively have solved many of these issues. A strict information security policy of no data transfers (it is not uncommon for media drives to be banned within organisations) as well as stringent governance of physical security and staff practices would mitigate many risks, as demonstrated in [98]. In this scenario, however, the policy must not only govern the local premises, but also cater for movement between the home/public and other corporate environments.

This dissertation addresses these highlighted issues of data security in a mobile working environment and presents methods and tested solutions towards the better understanding of the threat environment, understanding the users who operate within it, and the ability to identify and populate data sets. This work aims to empower the decision maker by aiding in the development of appropriate policies that reflect the needs of the user and the organisation.

1.1 Example Applications

This section details the example scenarios that are investigated within the dissertation. These scenarios are referred to with respect to the methodologies and solutions provided.

Designing New Policy - Chapter 3:

Universities present a difficult security environment to manage. Creativity, novel solutions and continued collaboration play a key role in many disciplines. Maintaining this freedom whilst understanding the needs of different user groups provides a ideal example for detailing how to develop an effective policy decision making process.

Data Transfer and Network Selection - Chapter 4:

User X belongs to company Y where BYOD is allowed. User X is working in a public, off-site location and must submit urgent documents to fellow colleagues. There are no company approved networks that can be utilised and the company does not have a Virtual Private Network infrastructure. User X must decide which wireless network to connect to in order to send the necessary corporate documents.

Understanding Security Behaviours - Chapters 5 & 6:

Users are required to complete a specific task. During this task users are interrupted and must make security decisions based on their knowledge and the task at hand.

Extending the BPMN 2.0 Formalism - Chapter 7:

Thrombolytic stroke victims require urgent medical attention in order to minimize further medical complications. A hospital administrator requires a decision making support tool to optimize patient throughput with a finite budget and time frame.

1.2 Aims and Objectives

1.2.1 Aim:

- ‘To design, develop and validate with scientific rigour, a robust methodology for aiding the chief information security officer in the policy decision making process’.

1.2.2 Objectives:

- To understand the current ('state of the art') practices through a review of literatures and best practices.
- To accurately interpret and document the issues and concerns from the CISO and narrow the problem space into a more focused investigative avenue.
- To identify and formulate solutions to specific problems with respect to CISO and user requirements.
- To develop tool support and experimentation to aid in the assessment of these problems.
- To evaluate the effectiveness of such tools and experimentation.
- To reflect on the overall impact of the approach, address its shortcomings and outline the direction of possible future work.

1.3 Thesis Contributions

A cross-disciplinary background and literature review of the factors influencing effective policy design

The literature review provides an extensive overview of the core concepts and methodologies from multiple disciplines. The review focuses primarily on the individual aspects of the environment which must be understood in order to develop an effective policy. Understanding the threats within these environments and detailing user behaviours is a critical component of this.

The literature review includes an in-depth survey of academic and industry led research on user behaviour, user management and information security, as well a formal documentation of many practices within the workplace. Psychology literature provides a necessary understanding of how users behave under certain conditions and which stimuli are most influential with respect to interventions. We include this so that a comprehensive impact analysis can be conducted.

The benefit of a cross-disciplinary approach allows for a more detailed analysis and understanding of the problem space. Typically, poor security decisions lack the understanding of how users behave, comply with, and are impacted by interventions. As the discipline matures, users requirements and behaviours are increasingly becoming a core component of the policy decision making process.

A methodology for investigating stakeholder and user requirements

The proposed framework encapsulates the problem space from the perspective of the stakeholder, and critically, the user. This approach differs from traditional work in the field of Information Security which often focuses explicitly on technological solutions. This is often to the detriment of the overall validity or applicability of the proposed solution. User requirements and behaviours are important and are often ignored through forced, blanket security that specifically dictates allowed user actions through policy. This often has profound impacts on user compliance and productivity and ultimately impacts negatively on the organisation itself.

Our explicit implementation is based on empirical observations from two separate studies (the IRIDIUM study, pilot study), benefiting from genuine interaction and feedback from users with respect to practical tasks within the environment. This provides numerous benefits over simulated data sets that are often based on assumptions, as we are able to articulate solutions that specifically match our problem domain.

The benefits of including user requirements enables the CISO to develop a more effective, holistic set of policies that aim to support 'good' user behaviours whilst also maintaining data security. This promotes compliance and productivity by working with and not against users.

Experimental tool design and implementation

Through experimentation in chapters 4, 5, 6, and 7 we develop and evaluate tools designed specifically to investigate the highlighted user behaviour. The tools we build are multi-use by design, allowing simple modification for use in subsequent studies. Specifically, we develop an Android Application for Wi-Fi selection (chapter 4), two web and database frameworks (chapters 5, 6), and an extension of the BPMN 2.0 modelling formalism (chapter 7).

An iterative methodology for policy design

This methodology facilitates the creation and adoption of policy through an iterative design process. This differs from previous work which typically aims to improve the decision making process through a single investigation, problem or observation. Critically, this methodology understands that policy design is an iterative process and one that requires repeat investigation to remain relevant and meet the changing requirements of users.

The value of the proposed methodology stems from the iterative process which assumes that subsequent investigation of the problem space will ultimately yield a more intimate understanding of the environment variables. This allows for solutions to be augmented based on newly discovered variables or behaviours that arise from repeated investigation. Subsequent solutions will be more effective as they are built from a more accurately quantified variable set.

Modelling-based approach towards policy testing

This modelling process improves the decision making process by enabling policies to be simulated on test populations before implementation. These populations can be either simulated or based on empirical data from our pilot studies that will more accurately reflect the specific nature of the target environment.

This methodology aims to specifically address the problem of policy evaluation after adoption. Typically, policies are adopted and then empirically evaluated. The problem with this method is that it relies on the first implementation of the policy being correct. By a process of iterative investigation into the environment, and subsequent modelling and simulation before implementation, the design process is improved as we enable a more appropriate policy to be designed. The resulting policy implementation process is therefore optimised as the policy is more likely to reflect the needs of users and the company.

1.4 Publication History

This dissertation includes work that has previously been, or is in the process of, peer review. These publications are referenced to at relevant points throughout and form the scientific basis of several chapters.

- James Turland, Lynne Coventry, Deborah Jeske, Pam Briggs, Aad van Moorsel. ‘Nudging towards security: Developing an application for wireless network selection for Android phones’. British HCI 2015.

- Daniel Nesbitt, James Turland. ‘BPMNdm – Extending the BPMN Formalism to aid the Decision Making Process’. IEEE Services 2016. Emerging Technologies track on Formal Methods in Services and Cloud Computing.

- Iryna Yevseyeva, James Turland, Charles Morisset, Lynne Coventry, Thomas Groß, Christopher Laing , Aad van Moorsel. ‘Addressing consumerisation of IT risks with nudging’. CENTERIS 2014.

- Iryna Yevseyeva, Charles Morisset, James Turland, Lynne Coventry, Thomas Groß, Christopher Laing , Aad van Moorsel. ‘Consumerisation of IT: Mitigating risky user actions and improving productivity with nudging'. CENTERIS 2014 - Conference on ENTERprise Information Systems / ProjMAN 2014 - International Conference on Project MANagement / HCIST 2014 - International Conference on Health and Social Care Information Systems and Technologies.

- Deborah Jeske, Lynne Coventry, Pam Briggs, James Turland. 'Less redundancy and increased relevance: Encouraging technical and security error reporting'. International Journal of Human-Computer Studies. (In preparation).

- Debora Jekse, James Turland, Pam Briggs, Lynne Coventry. 'Personality and Framing Factors in Privacy Decision-Making: A study on cookie acceptance'. ACM Transactions on Computer-Human Interactions (March 2015). (Submitted)

Acknowledgements & References

- Lynne Coventry, Pam Briggs, Debora Jeske, and Aad van Moorsel. 'SCENE: A structured Means for Creating and Evaluation Behavioural Nudges in a Cyber Security Environment'. Design, User Experience, and Usability. Theories, Methods, and Tools for Designing the User Experience. Springer Publishing International. 2014.

  This paper is included for its relevance to our approach with respect to user behaviours and intervention methods. Specifically, the cyclical nature of improved, repeat investigation. The paper is also referenced throughout the thesis.

- Robert Cain, Aad van Moorsel. 'Optimization of data collection strategies for model-based evaluation and decision making'. DSN 2012.

  This paper is included for contributions made to the 'Trust Economics' project in conjunction with The University of Illinois. The rationale for the study and methodology contained compliment our approach with respect to finite budget allocation and the necessity to obtain accurate variables with respect to our environment, a phenomena all security policies are confined to.

- Deborah Jeske, Lynne Coventry, Pam Briggs. 'Decision Justifications for Wireless Network Selection'. STAST 2014.

  This paper is included for contributions with respect to chapter 4 and the examination of how users justify their decisions in this context.

- Deborah Jeske, Lynne Coventry, Pam Briggs. 'Exploring the relationship between impulsivity and security'. Under review with Computers and Security. 2014

1.5 Collaboration

Throughout this thesis, reference to the above papers and their respective authors is made.

Chapter 2

Background and Literature Review

This section details a background and literature review within the subsequent fields of enquiry. Specifically, it focuses on the multi-disciplinary nature of aiding the Decision Making Process. This requires in-depth study of the many intricacies that formulate the problem space.

More formally, this section identifies the relationships between the different study areas and highlights the inter-connected nature of the multiple disciplines. This highlights the complexity and the difficulties faced when designing policy by exploring the plethora of variables that constitute to various phenomena. Once defined, these phenomena are then examined and further understood through an analysis of the relevant literature pertaining to the understanding and summation of the problem space. The principle disciplines examined fall under Computing Science, Human Behaviour and Economics.

This approach differs from traditional studies that focus on a single discipline. The inherent value of this approach is observed due to the more accurate representation of the problem space from the inclusion of multiple disciplines. This provides a theoretical and practical base for investigation that considers the technological, social and economic aspects and ultimately aids the CISO in the Policy Decision Making Process.

The structure of this chapter is as follows. Section 2.1 provides an overview of Information Security. Section 2.2 investigates the role of Social Phenomena and Human Behaviours. Section 2.3 investigates the rise in the Consumerisation and the BYOD trend. Section 2.4 discusses Trust Economics. Section 2.5 identifies the role of Trust within Information Security. Section 2.6 details Uncertainty and Probability Theory. Section 2.7 investigates the role of modelling and its applicability.

2.1 Information Security

Information Security is required to protect an organisation from unauthorised access of its assets. The importance of such a practice is highlighted by "A National Cyber Security Association [168] survey of small business in the US, conducted in 2012, suggested a cyber security disconnect where 47% of companies believed a data breach would have no impact on their business, yet 87% did not have a formal written Internet security policy and 69% did not even have an informal one" 4][144]. More recently in 2016, the same National Cyber Security Association published that these figures were still alarmingly high with 77% not having a formal ISP [215] indicating that this is still a serious issue. This realisation sparks the necessity for a thorough, methodological approach to Information Security and the Policy Decision Making process.

Information Security is the practice and theory of defending data or information systems from unauthorised or unintended access with the goal of preventing destruction, disruption and tampering [234]. Conceptually, the practice can be split into 6 considerations as denoted by Parker's Hexad of Information Security [66]:

- **Confidentiality:** "the assurance that information is not disclosed to individuals or systems that are not authorised to receive it" [234].

- **Possession or Control:** denotes the loss of control of a possession or asset but does not include a breach of confidentiality.

- **Integrity:** "the assurance that information cannot be modified by those who are not authorised to modify it, or that any such modifications will not pass undetected" [234].

- **Authenticity:** the strength of the claim of origin, identity and authorship.

- **Availability:** "the assurance that information is available when it is needed, and that mishap or malice cannot affect the ability of the systems to provide information when requested" [234].

- **Utility:** the usefulness of the data or asset in question.

These measures are necessary to protect assets. Asset management is typically divided into an organisation's "tangible and intangible assets" [107] defined as an item that "has physical form ... does not have physical form" [107] respectively. Intangible items are diverse and can range from "intellectual property to goodwill" [107]. Tangible items are typically easier to track in smaller organisations but are much more difficult in larger organisations.

Asset management requires a financial attribution (asset valuation). To calculate, one can use:

$$\frac{Cost - Salvage\ Value}{Useful\ Life} = Yearly\ Depreciation \tag{2.1}$$

A temporal aspect is required within the equation as business capital expenditure is approved based on an assessment of return on investment. The longer a product is viable, the greater the return on investment, and the less likely it is that further expense is required.

Intangible items are much more difficult to value often owing to the variables one associates with value. As such, one must ask "what would you pay for the asset if you did not already own it?" [107], "what revenue will this asset bring to the organisation in the future?" [107]. With this, it is important to understand the following four methodologies. These methodologies "have become the most important, whether for transaction, tax, or litigation purposes, or whether in an ongoing concern valuation or liquidation" [18]:

- Cost approach

  Historical cost basis

  Replacement or reproduction cost

- Market approach

- Income approach

  Future income stream

  Duration of the income stream

  Risk associated with the generation of the income stream

- Relief-from-royalty approach [18]

With this knowledge it is possible to utilise the "single loss expectancy (SLE)" formula "measuring the specific impact, monetary or otherwise, of a single event" [18]. Understanding this likelihood forms part of the rationale behind the policy development process and ensures that protective technologies and intervention strategies are appropriate to the value of the asset at risk.

$$Asset\ Value \times Exposure\ Factor = SLE \tag{2.2}$$

The exposure factor (EF) is a combination of the threat and vulnerabilities within the environment. It is important to stress that this an estimate however, as "it cannot include everything possible because we do not know all of the possible exposures" [18]. The EF is "the percentage of loss a realised threat event would have on a specific asset, that is, the consequence" [18]. The EF can be a large number owing to the nature of the event and the method of storage. For example, "a major event such as a fire or a small number like the loss of a hard drive" [18]. This can also be expressed as a percentage if necessary if for instance "a virus brought down your Web farm, this may cause a 75% loss in the Web farm's functionality" [18].

The exposure can be furthered by adopting an annualised rate of occurrence (ARO) which details the frequency that the exposure is expected to occur. The "ARO is not a definite number and can be subjective" [18]. It is based on empirical data and utilises an assessment of the "organisation's metrics on hardware, software and past threats" [18]. With this, we can formulate that if a = average attempt of unauthorised access and b = number of employees:

$$a \times b = ARO \tag{2.3}$$

Being able to quantify and identify risk vectors it is necessary to define security systems designed to combat risk and alleviate unwanted exposure. ISO 27001 "specifies the requirements for establishing, implementing, maintaining and continually improving an information security management system within the context of the organization" [123]. This standard specifies the methods in which security systems can be designed and implemented to address one or more of the following controls:

- Physical controls - walls, locked doors, guards. [234]
- Procedural controls - managerial oversight, staff training, defined emergency response processes. [234]

  Social behaviour

- Regulatory controls - legislation, policy, rules of conduct. [234]
- Technical controls - cryptographic software, authentication and authorization systems, secure protocols. [234]

With any intervention method it is essential to conduct a survey on the control mechanisms to ensure validity, assess impact and apportion budget. These measures should reflect the current threat environment. For example, if "we have an SLE of £58,000; but if we are spending £100,000 a year to protect it, we are spending more than we need and new controls should be selected" [107].

With the above defined, it is now possible to quantify the return on security investment (ROSI). The ROSI is calculated by:

$$Annual\ loss\ expectancy\ (ALE)\ -\ Current\ cost\ of\ control\ (CCC)\ =\ ROSI \tag{2.4}$$

[107]

Equation 2.4 aids in budget calculation. The following table details the entire process:

Table 2.1 ROSI for Propriety Confidential Data [107]

Steps			**Formula**
Asset identification and valuation	**Asset**: Proprietary confidential data	Valuation: £5,000,000	
Threat and vulnerability exposure factor (EF)	**Threat**: Disclosure of data	**EF**: 90%	
Determine the single loss expectancy (SLE)	£5,000,000 × .90 =	**SLE**: £4,500,000	Asset Value × Exposure Factor = SLE
Annualised rate of occurrence (ARO)	Based on observed data, the probability is 1 in 20 years	**ARO** = 0.05	
Compute the annual loss expectancy (ALE)	£4,500,000 × .05 =	**ALE**= £225,000	Single Loss Expectancy (SLE) × Annual Rate of Occurrence (ARO) = ALE
Survey controls	Current controls are costing £95,000	**ROSI** = £130,000	
Calculate ROSI	£225,000 - £95,000		Annual loss expectancy (ALE) - Current Cost of Control (CCC) = ROSI

Use of the ROSI 2.4 is not without drawbacks and is contested in its' validity owing to the assumptions one must make in assessing the given variables. One argument "is that valuing the ROSI lacks precision and is based on approximations" [107]. A solution to this problem, however, is through continued data collection of the given variables; "as more data is collected ... the picture will become clearer, much like insurance actuarial tables can predict the probabilities of certain events" [107]. Another criticism one must consider is that "the ROSI is immutable; but if it is made a part of the annual review process, this should not be the case" [178].

2.1.1 Policy Decision Making

Having defined the above formulae regarding asset valuation and ARO, it is possible to begin developing policies to combat and mitigate risk. In addition, having determined the financial cost of intervention systems utilising the ROSI 2.1 we must focus on the users within our environment. Policy making is intrinsically a simple process (locate and protect assets), however, designing and implementing a policy that is sustainable and appropriate for user's needs and behaviours is highly complex. Often there is a need to balance productivity, availability and security to promote user compliance. This notion is cemented in [198] which states that "literature agrees that the major threat to IS security is constituted by careless employees who do not comply with organisations' IS security policies and procedures". It is unfair, however, to blame the user entirely as the notion of compliance contains a plethora of variables. It is likely that the policy itself and the method in which it is enforced and disseminated amongst staff (i.e. staff training) is to blame as it does not accurately reflect the working environment of the users and thus constitutes a more systemic failure. Providing an holistic policy that compliments user's actions in a secure manner is critical and supported by "the fact that if information security is not addressed in a holistic and comprehensive way, taking all its dimensions into account, real risks exist preventing a really secure environment" [244].

2.2 User Behaviour and Cyber Security

The importance of the user in cyber security cannot be disregarded. In a survey conducted in 2013, "93% of large organisations reported having a security breach in the previous year, and 87% of small businesses" [144]. Of these, "36%" [144] of the worst breaches were attributed to "inadvertant human error" [61] resulting in the loss of confidential information. Furthermore, "87% did not have a formal written Internet security policy ... 69% did not even have an informal one" and "18% said they would not even know if their computer network was compromised" [168]. This naturally presents a significant risk area where users are in essence acting in their own free will without controls, management or monitoring. Further problems arise with the continued adoption of BYOD which further blurs the line between device ownership and the policies that govern its use, covered explicitly in section 2.5.

Research highlights that "employees seldom comply with the guidance outlined by their information security policies" [214] resulting in "billions of dollars lost annually" [46]. The result of such findings has issued users the moniker of "the weakest link in information security" [163], [246]. For this reason, it is essential that we understand how to work with and understand user behaviour to create meaningful, effective security that users trust and are willing to comply with, thus transforming users from "the biggest information security vulnerability to the first line of ISP compliance defence" [170].

Organisations have typically relied on technology-based solutions [260] in an attempt to improve information security. This approach, however, is seldom sufficient to eliminate risk [53]. Reports [11], [219], provide empirical and anecdotal evidence that the number of information security incidents is increasing "even as organizations invest more in technology-based solutions" [44]. This requires a shift in the approach towards information security; a new paradigm that requires "organizations to invest in both technical and socio-organizational resources" [44].

2.2.1 Human Behaviour

Human behaviour in IS has typically been studied by two streams of behavioural research:

- **threat-coping perspective**. This approach aims to depict how humans assess and cope with threats. Drawing on protection and coping theory, the suggested models explain individual user's responses to perceived threats [94], [131], [151], [256]. We use coping theory to explore an underlying relationship between employee stress caused by burdensome, complex, and ambiguous information security requirements (termed "security-related stress" or SRS) and deliberate ISP violations.
- **policy-compliance perspective**. This approach draws on theories that explain why humans do or do not comply with organisations' IS security policies [94], [44], [167], [179], [206]. It aims to examine how employees' intention to comply with policy is driven by cost–benefit assessments, personal norms and organizational context factors. Studies indicate that employees' compliance intention is the result of competing influences of perceived benefits, formal sanctions, and security risks.

Both of these approaches are built upon studies that define IS security policies as sets of rules "describing acceptable and unacceptable technology usage" [94]. These rules must be kept in mind by employees while they perform their day-to-day work tasks with technology. By understanding how threats are assessed along with how users comply with mitigation strategies we benefit from being able to design better ISPs.

Decision Making

The method in which humans formulate and act out a decision has been theorised for decades. If we can begin to understand how we think and subsequently act, we can begin to predict and in certain cases apply interventions that influence the process and subsequent decision. This is especially important in IS and compliance as we can modify user behaviour towards a more secure practice (see section 2.3.2).

Kahneman [135] acknowledges and summarises the many works conducted in decision making and proposes a "Two Systems" approach, a phrase first coined in [213].

- **System 1** operates automatically and quickly, with little or no effort and no sense of voluntary control.

- **System 2** allocates attention to the effortful mental activities that demand it, including complex computations. The operations of System 2 are often associated with the subjective experience of agency, choice, and concentration.

Typical examples of each system are as follows:

- **System 1**:

  Detect that one object is more distant than another.

  Orient to the source of a sudden sound.

  Understand simple sentences.

- **System 2**:

  Monitor the appropriateness of your behaviour in a social setting.

  Check the validity of a complex logical argument.

  Fill out a tax form.

The distinction between the two systems is important and enables a deeper understanding towards how we think. System 1 describes how we "effortlessly originate impressions and feelings that are the main sources of the explicit beliefs and deliberate choices of system 2" [135]. Principally, "the automatic operations of system 1 generate surprisingly complex patterns of ideas, but only the slower system 2 can construct thoughts in an orderly series of steps". System 2 therefore requires us to pay greater attention to the given task and requires a higher cognitive load.

As detailed in section 2.2.2, user compliance is paramount to any successful ISP. Understanding System 1 and System 2 furthers our knowledge through the understanding of "the busy and depleted system 2" [135]. Within this theory, it is determined that "both self-control and cognitive effort are forms of mental work". This is important when designing ISPs and implementing DLP strategies as we know that requiring users to perform additional tasks increases their cognitive load (which has a finite capacity [32]) and increases the likelihood of non-compliance. This phenomenon is studied in depth within the field of Psychology and is known as ego depletion.

Ego Depletion

Fischer [92], investigated "how the availability of self-control resources affects risk-taking inclinations and behaviours" and concludes that "risk-taking is increased when individuals find themselves in a state of reduced cognitive self-control resources (ego-depletion)".

The foundation of this paradigm is self-regulation. To "self-regulate is to exert control over the prepotent (i.e., automatic) psychological and behavioural responses" [92]. In doing so, an individual is continually consuming a limited resource, terminating in a state known as "ego depletion". In this state, "users show reduced performance at other tasks that draw on the same self-regulatory resource of self-regulation" [242]. Most importantly, this resource appears to be "global in the sense that this negative effect holds in domains as diverse as basic level intellectual performance and reduced social self-presentation abilities" [92], [242].

There is a positive correlation between an ego-depleted state and increased risk taking, whether through the necessity to complete the task more quickly (minimize suffering), or through the lack of fully understanding the task and risks associated. Risk taking "refers to one's purposive participation in some form of behaviour that involved potential negative consequences or losses (social, monetary, interpersonal) as well as perceived positive consequences or gains [36]. [92] demonstrates through 4 separate studies that sensation seeking, "which is one of the primary and most prominent determinants of risk-taking behaviour" [156], reduces self-regulation resources resulting in increased risk-taking.

Ego depletion can be seen as a highly relevant problem for organisations from an IS perspective. Reducing a users self-regulatory ability through the introduction of inadequate security measures will increase the likelihood that users will violate policy or behave in a non-desired fashion. The ability to combat such a scenario is critical and has been discussed in several studies.

The methods with which to remedy ego depletion have been studied extensively [230]. Throughout 4 separate studies it has been proven that "inducing positive emotion can counteract the effect of ego depletion" [230]. More specifically, the studies involved users mixing tasks by creating an "initial state of self-regulation" on one task, and then asking users to perform another different task. The self regulation was measured between these two tasks and it was noted that "for some participants, positive mood was induced in between the two self-regulation tasks. The positive mood resulted in improvement in self-regulation in all four studies, as compared to participants who performed the same self-regulation tasks but did not have positive mood induced" [230].

Stress and Security Related Stress

Stress is a complex concept that has typically been defined and operationalised in terms of stimulating conditions that produce stress reactions [71], [187], [188]. In IS, the literature "provides the technostress creators construct, which delineates five stress-creating aspects of organizational IT usage: overload, invasion, complexity, insecurity, and uncertainty" [71], [189], [222].

Security related stress (SRS) is a common cause of IS violations [71]. Within psychology there are numerous literatures that detail empirical studies that have shown that negative work "stressors predict a variety of undesirable employee behaviours" [99], [195]. Specifically for IS literature, "research indicates that employee stress-related to the use of information technology influences a number of IT and non-IT-related cognitions and behaviours" [189], [222].

SRS is conceptualised in terms of "overload, uncertainty and complexity dimensions" [71] and is similar in principle (with respect to a finite capacity) to ego depletion.

- SRS Overload: is a term related to scenarios where "requirements increase workload for employees, and as a result, create added time pressure for them to complete job duties" [71]. A common example includes users requiring administrative access on their computers to complete a task, a process that often requires additional paperwork and valuable time to complete.
- SRS Complexity: defines situations where security requirements are viewed as complex and thereby force employees to expend additional time and effort in learning and understanding procedures. For example, where "security policies involve multiple contingencies or contain technical jargon, employees will have to devote greater time and effort toward understanding the appropriate policy and deciding how to act" [71].
- SRS Uncertainty: refers to an organisation's continual update of job-related security requirements (whether internally driven or as a result of government or industry regulation. ISO 27001 [123]).

Coping Theory

Coping theory [188] provides a usable framework for understanding how employees respond to SRS. Coping theory traditionally describes cognitive and behavioural processes to "manage psychological stress, of which SRS can be considered an example" [71]. Importantly, this theory states that individuals go through two interrelated forms of appraisal, "primary and secondary" [71] when determining whether a particular situation is stressful. Primary appraisal refers to the evaluation of a situation and whether or not it is benign or stressful (stressful situations are defined as harmful, threatening, or challenging [188]). Secondary appraisal refers to the individual's control of the stressful situation. It is important to note, however, that these two appraisals often "operate in unison" [188] supporting the belief that SRS is an outcome of this combined process.

Defining the primary and secondary appraisal approach enables the understanding of coping efforts that aim to alleviate stress. The definition of coping has undergone numerous iterations (often specific to the scenario at hand) but the most common distinction is between "problem-focused and emotion-focused coping" [71].

- Problem-focused coping: defines the direct effort made to manage or alter the stressful situation. In a working environment this could constitute as evading obstacles that hinder productivity (possibly increasing the likelihood of non-compliance) or engaging in knowledge building exercises that enable the user to complete a task more easily.

- Emotion-focused coping: refers to the manner in which an individual feels towards this stressful situation. Specifically, it refers to the way in which the user employs "cognitive processes (e.g., reappraisals, distorting reality) directed at reducing emotional stress" [71]. The likelihood of this coping method is increased when there is little that can be done to address the problem directly (via problem-focused coping). The opposite is true in high controllability situations where problem-focused is more a more probable application.

The role of coping theory in IT is witnessed through various processes. Technological change [71] requires user adaptation, a concept similar to coping (see above). Much of the literature within IT [31], [151] identifies several coping strategies including "mental relaxation techniques, modifying work tasks, and reinventing and adapting the technology" [71]. Importantly, user adaptation research indicates that "when the expected consequences of an IT event are appraised as a threat of personal or professional relevance, and users feel that they have limited control over the situation, their adaptation efforts will be mainly emotion-focused" [31], [151]. This is of great significance for policy design as it reinforces the belief that security is an obstacle in the users' mind and therefore a mitigation strategy is required to either educate the users (to combat problem-focused coping) or to implement security measures in a non-intrusive, undetected fashion.

Moral Disengagement Theory

MDT enhances our understanding of SRS. Though they share many parallels related to stress literature, MDT provides a thorough and detailed explanation of the cognitive disengagement process within a theoretical framework [71]. Specifically, MDT can be seen as a method in which to extend the understanding of emotion-focused coping.

Studies [99], [195], have shown that negative stressors predict undesirable employee behaviours such as "counterproductivity and deviance". MDT is grounded in social cognitive theory [27] and provides "eight interrelated cognitive mechanisms, conceptualized as three broad categories (reconstruing the conduct, obscuring or distorting sequences, devaluing the target)" [71] that allows users to "disengage the internal self-sanctions that govern their behaviour" [28].

Within the field of social cognitive theory, Bandura [28] discusses emotion-focused coping in terms of moral disengagement mechanisms that "cognitively restructure the meaning of stressful situations" [71]. Importantly, this provides evidence for a "theoretical linkage between the emotion-focused aspect of coping theory and MDT" [71]. However, as stated in [71], the question remains as to how moral disengagement can serve as emotion-focused coping in an organisational context. Several independent research studies have shown that negative stressors (including technology characteristics and aspects of the IS environment) "produce strain on the employees" and "foster negative emotions and affect" [24], [148], [183]. Moral disengagement can be enacted to address this strain and negative emotion in an attempt to "restore emotional stability and reduce the tensions emanating from stressful work conditions" [71]. In this frame, one can conceptualise that moral disengagement may serve as an "instrumental coping function that mitigates the negative effects of workplace stress on subsequent strain" [71]. In such instances, moral disengagement may also be motivated by "a desire to cope with uncontrollable stressors in the work environment (e.g., SRS) such that MDT's cognitive rationalizations allow employees to assert and regain a degree of psychological control" [70].

Table 2.2 defines the cognitive mechanisms and self-sanctions that can be deactivated along with their IS consequences. It provides a useful reference for identifying and adapting policies towards user behaviours.

Table 2.2 Moral Disengagement Mechanisms [71]

Category	Mechanism	General Description	IS Security Policy Context
Reconstruing the conduct	Moral justification	Reconstructing harmful conduct as personally and socially acceptable by portraying it as serving worthy or moral purposes; that is, service of a greater good.	Employees may justify an ISP violation in the name of getting the job done more efficiently or meeting a particular deadline, whether it is for personal accomplishment or because they feel they are doing a service to the organization.
	Euphemistic labelling	Relabeling harmful conduct through sanitized or convoluted language or concepts to make it sound benign. For example, terrorists label themselves “freedom fighters” and in the business context laying people off is referred to as “downsizing."	Employees may euphemistically label certain ISP violations as “no big deal,” not such a bad thing, or an inevitable reality in the workplace.
	Palliative comparison	Considering harmful acts as acceptable by contrasting them with more reprehensible behaviors.	Employees may justify a seemingly innocuous ISP violation such as password sharing or failing to logoff a workstation by comparing it to a more severe policy violation such as stealing company information.
Obscuring or distorting consequences	Displacement of responsibility	Viewing harmful acts as stemming from the social pressures or dictates of authority rather than being one’s own responsibility.	Employees may deny responsibility for an ISP violation due to perceived work overload or a lack of alternative methods for getting the job done (both of which are the fault of management).
	Diffusion of responsibility	Diffusing responsibility across a collective (i.e., division of labor) rather than holding oneself personally accountable for harmful conduct.	Employees may perceive management, the IT department, or other employees as more responsible for IS security than themselves. An employee may also perceive that other employees are violating policy, which limits his/her overall responsibility for security.
	Distortion of consequences	Cognitive efforts to ignore, minimize, or distort the harmful consequences of one’s actions.	Employees may distort the consequences of ISP violations by deeming them as not hurting the organization, at least not directly. This is plausible given that the negative effects of many ISP violations are often not directly seen or experienced by employees.
Devaluing the target	Dehumanization	Divesting the target(s) or victim(s) of harmful conduct of human qualities.	Harm resulting from an ISP violation primarily affects the organization (and not humans), and so violations may occur if employees view the company as bureaucratic, lacking emotions, or not being people oriented.
	Attribution of blame	Ascribing harmful conduct to compelling circumstances outside of one’s control, such as the environment or surroundings, rather than a personal decision.	Employees may attribute ISP violations to the strictness or unreasonable nature of policies.

With the culmination of the above theories (SRS, MDT, Coping Theory), D'Arcy [71] proposes the following model of the influence of SRS on employees' deliberate ISP violations. It is important to note that the model depicts the ISP violation intention rather than actual behaviour, a stance that is consistent with several security compliance research studies [70], [206] and is "driven by the difficulties in obtaining actual policy violation instances" [71]; a problem often associated with the reluctance of organisations to share their IS policies for fear of future attack (exploitation) or culpability (financial loss).

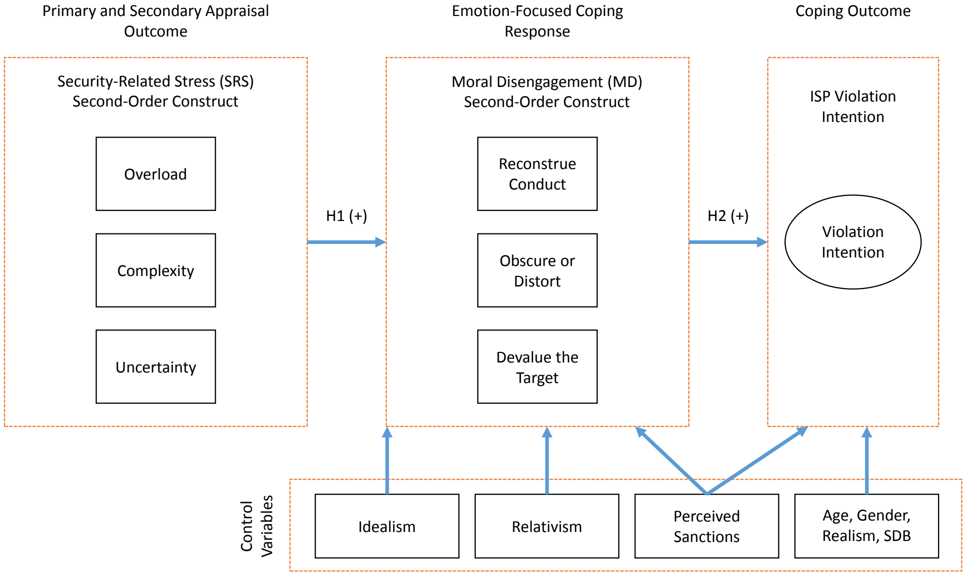

Fig. 2.1 Influence of **SRS** on employees deliberate **ISP** violations

Acknowledgement of this model is important with respect to identifying and rationalising potential violations and forms an intrinsic component of our study design. Knowledge of how user's formulate their non-compliance behaviours enhances our problem identification process. This in-turn enables the targeting of these behaviours along with bespoke experimental design aimed at mitigating the likelihood of occurrence, such as nudging.

2.2.2 Compliance

There are a multitude of threats (viruses, malware, corporate espionage, etc) "to the confidentiality, integrity and availability of organizational information and information systems" [22]. Whilst there are a plethora of security mechanisms to attempt to control and mitigate the risks posed by these vulnerabilities, "it is often incumbent upon users to utilize the technologies for them to be effective" [22]. In order to understand user compliance, it is necessary to understand hows users behave with respect to information security procedures. This requires an understanding of the "theory of planned behaviours" (TPB) [22] which itself is an extension of the theory of reasoned action (TRA) [14]. The TPB dictates that human behavioural intention to perform an action is driven by perceived behavioural control, attitude towards the behaviour and subjective norms [12], [13]. [13] states that Subjective norms are beliefs about the normative expectations of other people that result in perceived social pressure. Furthermore, an employee's attitude towards a behaviour is determined by their belief that performing (or not performing) the behaviour will lead to certain consequences [44]. Lastly, perceived behavioural control refers to a person's perception "about the presence of factors that may facilitate or impede the performance of a behaviour" [13], [147].

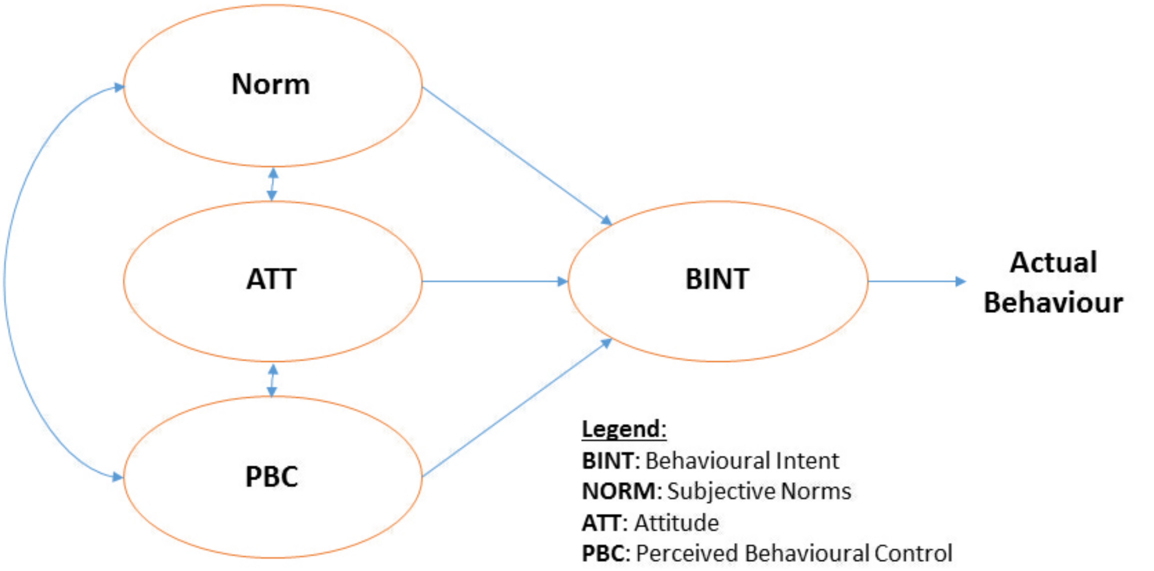

Fig. 2.2 Theory of planned behaviours model

The use of 2.2 is documented throughout many core literatures in the field of ISP compliance [44], [170], [225], [70] in areas such as insider security contravention [255] and computer abuse [147]. [22], however, argues that without a "common theoretical framework, and use of consistently operationalized constructs, comparison of past studies is hindered and focus for future research is obscured". To remedy this, [22] produced the following ISP behavioural compliance framework Fig: 2.3 based on "empirical studies that were published in high-quality, peer-reviewed journals".

The critical contributions of this composite ISP behavioural compliance framework can be seen when compared to the previous version in 2.2. [22] discusses the advantages, stating that the "value of the composite framework for behavioral compliance with information security policies is twofold. First, the framework synthesizes the results of related recent

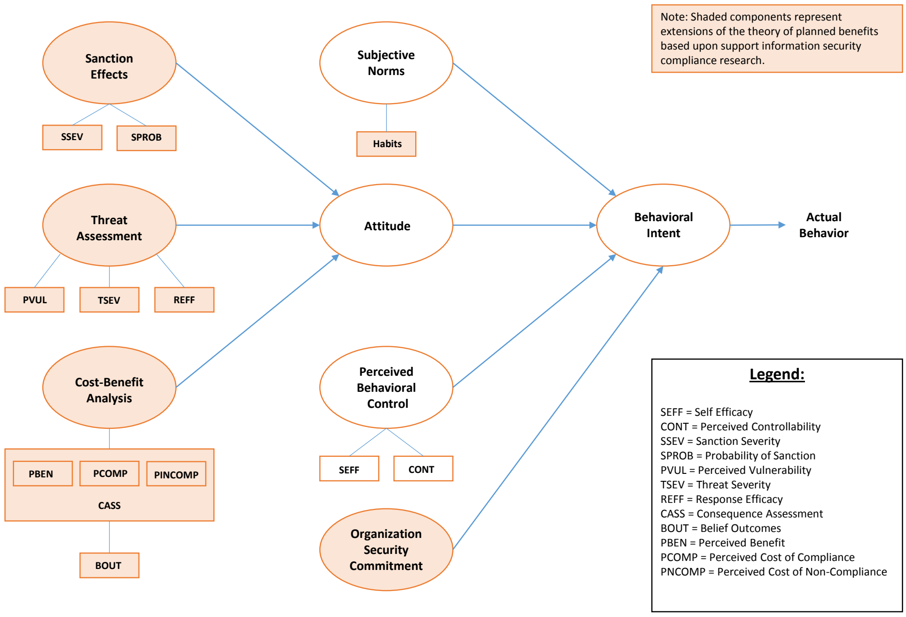
Sanction Effects
SSEV
SPROB
Subjective Norms
Habits
Note: Shaded components represent extensions of the theory of planned benefits based upon support information security compliance research.
Threat Assessment
PVUL
TSEV
REFF
Attitude
Behavioral Intent
Actual Behavior
Cost-Benefit Analysis
PBEN
PCOMP
PINCOMP
CASS
BOUT
Perceived Behavioral Control
SEFF
CONT
Organization Security Commitment
Legend:
SEFF = Self Efficacy
CONT = Perceived Controllability
SSEV = Sanction Severity
SPROB = Probability of Sanction
PVUL = Perceived Vulnerability
TSEV = Threat Severity
REFF = Response Efficacy
CASS = Consequence Assessment
BOUT = Belief Outcomes
PBEN = Perceived Benefit
PCOMP = Perceived Cost of Compliance
PNCOMP = Perceived Cost of Non-Compliance

Fig. 2.3 Composite ISP Behavioural Compliance Framework

research to produce a more complete, yet still parsimonious, model based largely upon the theory of planned behaviours". By adding a "core theoretical extension of organizational security commitment" a more complete model of the problem space is constructed allowing practitioners to "better focus their security education, training efforts and technology to maximize ISP compliance".

The use of figures 2.2, 2.3, provide a complex, detailed model of user behaviours (specific to compliance) which allow for bespoke policy implementations. This benefits the policy decision making process but it does not ensure that such policies are implemented or enforced correctly (itself a product of the environment in which it is deployed), thus compliance cannot be guaranteed. Simply understanding the user's compliance preferences is insufficient, it is necessary to adopt and fit technological solutions that are effective .

The complexity of the issue is further demonstrated when perceived beneficial technological implementations are in fact further hindrances towards security. This is the result of user intervention as "problems of appropriate response to cyber incidents are exacerbated when security technology is perceived as an obstacle to the user" [199] as they are often overwhelmed by difficulties in security implementation, or may mistrust, misinterpret the security" [199]. Whitten & Tygar [253] examined these difficulties from a Human-Computer Interaction (HCI) perspective and determined that "many human failures are caused by security mechanisms that are too difficult for a non expert to use. Even users with good technical skills, such as system administrators and software developers, often struggle to keep up with the increased complexity and workload created by security mechanisms". This suggests a disconnect between users and security system design requiring a fundamental reassessment of usable security. HCIsec is a relatively recent venture into understanding the importance of usable security [134] with the goal of providing "security tools that the intended users can operate correctly and complete a security task" [32].

To "achieve effective security from an organisational point of view, security designers and managers need to consider that:

1. Individual users have a choice on whether to comply with security policies, and
2. This choice is influenced by the individual's own goals, perceptions and attitudes, and norms which govern the individual's behaviour." [9] [248] [249].

The importance of user choice is something which cannot be ignored. Even if strict policies and practices are adopted, a user may simply decide to circumvent them or not fully comply. Managing employees' security behaviour is still a major challenge:

> "My biggest challenge is changing behavior. If I could change the behavior of our Dow workforce, then I think I've solved the problem." [130].

It is clear that changing behaviour is important, but "needs to be accompanied by changes in the security tools and models used" [32]. The best method of adoption, however, is somewhat convoluted and there are contrasting views on how to approach it. Whilst current attempts focus on "placing more responsibility on line managers, in some cases even imposing financial penalties on them if an employee they are responsible for causes a security breach" [130] there is also significant research that shows that "negative reinforcement - which should include financial sanctions for security transgressions – used in isolation, are as ineffective in changing security behaviour as they have been in changing behaviour most other areas of life" [249].

Bélanger et al [33] examined "resistance behaviour" in mandatory password change and discovered that "even when passwords were changed as required, the changes were intentionally delayed and the request perceived as being an unnecessary interruption" [199]. In fact, "people are concious that a password breach can have severe consequences, but it does not affect their attitude toward the security policy implementation" [33]. Furthermore, "the more technical competence respondents have, the less they favour the policy enhancement. ... In a voluntary implementation, that competence may be a vector of pride and accomplishment. In a mandatory context, the individual may feel her competence challenged, triggering a negative attitude toward the process" [33]. This finding is critical and further demonstrates the importance of the user in the field of Information Security. If we cannot find a secure, holistic approach that does not inconvenience the user the likelihood that any security measures will be adhered to are very low.

Cost Benefit Analysis in Compliance

The perceived cost and benefit from a user's perspective towards performing a task plays a critical role in compliance. Numerous studies [9], [248], [249], support the notion that to "achieve effective security from an organisational point of view security designers and managers need to consider that:

1. Individual users have a choice on whether to comply with security policies, and
2. This choice is influenced by the individual's own goals, perceptions and attitudes, and norms which govern the individual's behaviour." [32].

This belief that "individuals and organisations place different values on the cost and benefits of behaviours associated with security policies" [32] is supported empirically but requires a finer granularity with respect to the user's motives. It is stated that "[the users] choice is not an entirely selfish one - individuals consider the cost and benefits to both the organisation and themselves - but:

1. The perception of cost and benefit is centred on the individual employee's immediate work context,
2. There is a limit to the amount of effort individuals are prepared to expend on compliance unless there is a perceived benefit to balance it,
3. Cost-benefit imbalances accumulate until the compliance limit is reached" [32].

Understanding this user requirement and knowing that users act in this fashion greatly aids towards ISP creation and management. Simply designing and implementing a system does not necessarily mean that users will comply, regardless of how secure that system is. This reinforces the importance of the user in security 2.2. Thus, the 'compliance budget' is seen to addresses the necessity to understand the user's needs and behaviours and accounts for their behaviours by tying them to their perceived benefits. Understanding the 'compliance budget' provides "2 key benefits:

1. Organisations can focus their effort available on key security tasks to maximize their return on investment and avoid wasteful expenditure on less critical measures.

2. It is possible to determine the cost of achieving an employees' compliance with a security measure, and to include this cost in models of cost and benefit of security measures." [32].

This section 2.2.2 highlights some critical findings but also introduces difficulties towards how best to solve them. Users work and behave independently and of their own free will presenting a difficult problem to manage. Users effectively behave according to their own perceptions of the cost and benefits of a given task. This is problematic as it suggests a disconnect between the security policy (what the organisation feels it must protect) and the user (what they feel is important to adhere to). This requires education of the users to effectively understand the purpose of the ISP in place, and readdress their assumption of perceived cost and benefits.

2.2.3 Educating the User

Traditional literature has often fallen into two distinct categories; "strict technology-based controls of computer-based human behaviour" [199] along with "comprehensive education and training of system developers and users" [199]. Through subsequent research, neither has been found to have been particularly beneficial as extremes rely too much on depth rather than breadth of the problem space. Ofsted [235], conducted an investigation throughout 35 schools that supports this claim. They chose institutions "where the provision for e-safety was outstanding, the schools had managed rather than locked down systems. In the best practice seen, pupils were helped, from a very early age, to assess the risk of accessing sites and therefore gradually to acquire skills which would help them adopt safe practices even when they were not supervised" [235]. To summarise, "the most successful security behaviours were exhibited in schools where students were taught appropriate behaviours and then trusted to behave responsibly" [199]. The belief behind this notion is to combat the undermining of user's abilities (as noted in [33]) and instead cement trust and belief in user's actions with the knowledge that they have been given the necessary tools and received sufficient information to appropriately decide their actions and then to be trusted with their cyber security decisions. Through the continued implementation and support of such an approach, it is possible for it to become culture or common practice.

These findings are complimented in [20] where Julie Peeler (Director of the (ISC)^2) states that "organisations need to think beyond IT when planning IT security awareness training, and tackle it from the bottom up, as well as the top down". This involves identifying the "movers and shakers" within the organisation to better understand the formal and informal communication channels. Doing so will "help to identify the best way of putting security messages across and who needs to be involved and onside to ensure the programme is effective".

Peeler discusses the top down approach stating that it is essential to include company executives in any new security practice as "their actions, intentional or unintentional, can have a greater impact and be more difficult to uncover" [20]. She states that you must begin with the Chief Financial Officer (CFO) "who will know the potential cost of a breach", as well as the compliance team, "who will know the potential liabilities" (both critical in the calculation of the ROSI 2.1). In order to boost awareness, "security professionals need to stretch their leadership skills across the organisation and form partnerships at all levels". Most critically, "anyone putting together an IT security awareness programme should use as many of the ways people learn as possible and plan to reinforce the messages continually to ensure IT security becomes part of the way the organisation operates".

2.3 Nudge

A nudge is defined as:

> "any aspect of the choice architecture that alters people's behavior in a predictable way without forbidding any options or significantly changing their economic incentives. To count as a mere nudge, the intervention must be easy and cheap to avoid. Nudges are not mandates. Putting fruit at eye level counts as a nudge. Banning junk food does not". [228].

2.3.1 Choice Architecture and the Choice Architect

A choice architecture refers to the choices we have available within a given environment, at a given time with the necessity to make a choice. It encompasses all eventualities at a given time whether it is choosing what to eat at a canteen or which clothes to purchase in a shop. By its' very nature, a choice architecture is therefore highly random.

A choice architect is someone who modifies this environment and presents choices in a multitude of possible ways. Essentially, a choice architect has the "responsiblity for organizing the context in which people make decisions" [229]. Continuing the previous canteen analogy, consider the following:

Figure 2.4 highlights the plethora of choices available and the power that the choice architect holds. Each choice is different and fits to serve a different purpose. For example;

Synopsis: *The director of food services for a large city school system runs a series of experiments that manipulate the way in which the food is displayed in cafeterias. Not surprisingly, she finds that what the children eat depends on such things as the order of the items. Foods displayed at the beginning or end of the line are more likely to be eaten than items in the middle, and foods at eye level are more likely to be consumed than those in less salient locations. The question is, What use should the director make of this new found knowledge?* [229]

1. Arrange the food to make the students best off, all things considered.
2. Choose the food order at random.
3. Try to arrange the food to get the kids to pick the same foods they would choose on their own.
4. Maximize the sales of the items from the suppliers that are willing to offer the largest bribes.
5. Maximize profits, period.

Fig. 2.4 Choice Architecture - A Worked Example [229]

- **Option 1:** appears to be the most appealing and logical at face value. One might argue, however, that it seems somewhat paternalistic and undermines ones' sense of choice or preference.

- **Option 2:** makes sense with respect to overall fairness. It aims to favour no one but perhaps is somewhat chaotic for staff and customers alike (it make no sense to have different parts of a salad scattered along the length of the canteen).

- **Option 3:** is perhaps the most honourable attempt although the merits of such an approach are limited in their feasibility. For instance, perhaps younger students are unable to make sensible choices as to their meal as they are not old enough and thus ill-informed to make appropriate choices. Another problem arises with trying to please too many people that will all have different preferences. Finally, what does it mean to try to determine what the students would choose? Furthermore, no matter which you choose it will be impossible, with physical objects, for there to not be some kind of organisation or arrangement.

- **Option 4:** may attract the attention of a corrupt cafeteria manager that is operating for their own personal gain at the expense of others.

- **Option 5:** seems logical from a financial perspective where budgets and targets are required.

It is therefore extremely difficult to balance these seemingly insignificant choices without having far larger implications (in this case significantly modifying people's behaviours). As Thaler suggests [229] it is better to assume "everything matters" when altering the user environment.

2.3.2 Nudging

Nudging can be applied by manipulating the choice architecture. To be successful, nudges should:

- Nudges maintain freedom of choice: in this sense, a nudge can be seen as a "soft form of paternalism" [218] with the goal to steer an individual without limiting the overall choice. A GPS routing system is a good example.
- Transparency and effectiveness: nudges should be open rather than covert. This is increasingly important in governmental nudges (automatic pension enrolment for example). Disclosure of information can increase performance, combat inefficiency and help to reduce corruption.
- Evidence and testing: the most effective nudges tend to be based on empirical evidence and draw heavily on behavioural psychology [218]. Some nudges seem promising in the abstract but tend to fail in practice. Empirical, randomised controlled trails provide a rigid testing platform for designing, testing and implementing successful nudges.

2.3.3 Defaults and Least Resistance

"Treading the path of least resistance" refers to people's natural desire to adopt the easiest method of task completion (a process often involving some attribution of "laziness, fear and distraction") [229]. It implies that if given a task (where a default option is available - an option that is automatically selected if the user does not respond), we can expect the majority of a population to end up with this choice irrespective of whether or not it is best for them. The impact of such a default is magnified in its' effectiveness if there is some form of explicit or implicit suggestion that it reflects normative behaviour or the recommended choice.

A default is therefore an extremely powerful component of the choice architecture owing to its unavoidable, ubiquitous nature. In presenting a choice, there must always be a rule that governs the system's behaviour if the user does not choose an option. The importance of defaults is further stressed through the common belief that "if I do nothing, nothing changes" [229]. This is highly problematic and completely open to manipulation from a corrupt choice architect.

In Computing, the role of defaults can have profound impacts on IS. During software installation, many install wizards often contain additional pieces of software that are defaultly selected for installation (that are typically not from the same vendor) in order for them to profit from additional software sales or advertisement revenue. This software is from an unknown source which may or may not be trustworthy. Once installed, the software may then execute whichever commands it has been programmed to do so.

Defaults are also controversial in nature. Most common to cause complaints are those defaults that relate to financial transactions or cause some form of loss to the user, recurring subscriptions for example. Other examples include the necessity to manually opt-out of default charges (environmental CO_2 off-set, or default travel insurance).

To those who oppose the idea of defaults, Thaler [229] suggests the idea of "required choice, or mandated choice", a concept that is clearly illustrated through the example of organ donation [129], a decision that often sparks considerable debate especially given the pretence that many people are willing to accept an organ but far fewer are registered donors. This has similarities with Information Security decisions where users often view their personal sensitive data with higher regard than corporate data, and thus are more likely to take proactive measures to protect it. Within Thaler's study, it is noted that some countries have an opt-out policy with respect to organ donation, a term coined "presumed consent". This approach inevitably increases the number of people who implicitly agree to organ donation. The initiative has had strong negative feedback, however, as people feel that the government has no right to presume anything with respect to their organs and body after death. To remedy this, a mandated choice could be adopted. Thaler [229] discusses such an approach with respect to driving licence renewal in the state of Illinois, where upon renewal drivers are asked whether or not they wish to become organ donors. Such a policy has had a considerable effect on the number of donors resulting in a rise to 60% adoption from the national average of 38%.

2.3.4 Feedback

An essential component of any effective choice architecture is the ability to provide feedback to the user [229]. Examples include digital cameras that allow the user to instantly review a photograph they have taken (allowing the user to identify any errors) rather than having to send the film to be developed and finding out that the photograph is imperfect when it is too late.

Other useful feedback mechanisms can be visualised in the home energy market with the recent addition of smart meters. These smart meters aim to give a visual representation (through the use of colour coding) as to the amount of energy you are currently using in your home. similar visual stimuli can be observed in the automotive industry with the inclusion of miles per gallon displays which aim to push drivers into adopting a more economical, environmentally friendly driving style. (Defaults will be further discussed in 2.3.4).

MINDSPACE

MINDSPACE is a framework for understanding "nine effects on behavior operating largely on the “automatic” system: **m**essenger, **i**ncentives, **n**orms, **d**efaults, **s**alience, **p**riming, **a**ffect, **c**ommitment, and **e**go" [84], [83], [82]. Dolan et al [84] use the following diagram to highlight the links between the respective nodes.

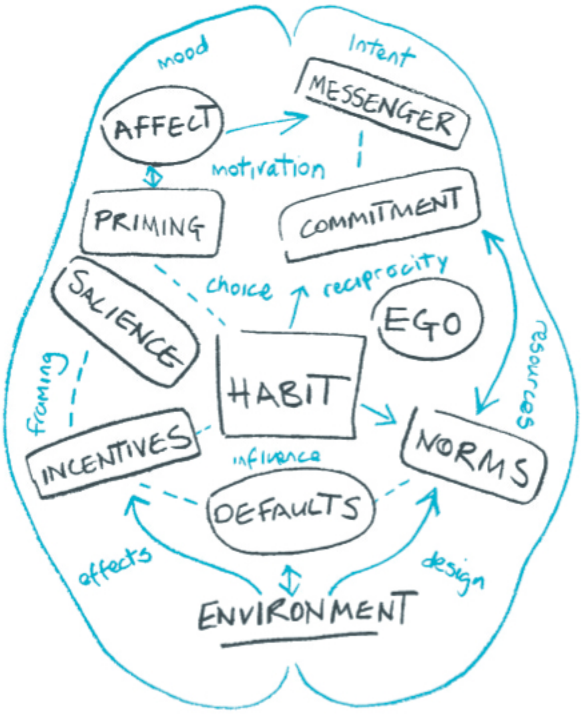

Fig. 2.5 Mindspace Diagram [84]

MINDSPACE is important in understanding how nudges are effective as it describes specific behaviours and how they can be modified. It is a tool for changing behaviour. To effectively change behaviour one must combine MINDSPACE with the 6 E’s of learning [84]. This process consists of understanding how users react to MINDSPACE and delivering our nudges in methods that conform to the 6 E's, both of which have been extensively documented within literature. Delivery must be holistic to working practices in order to improve the effectiveness of our nudges.

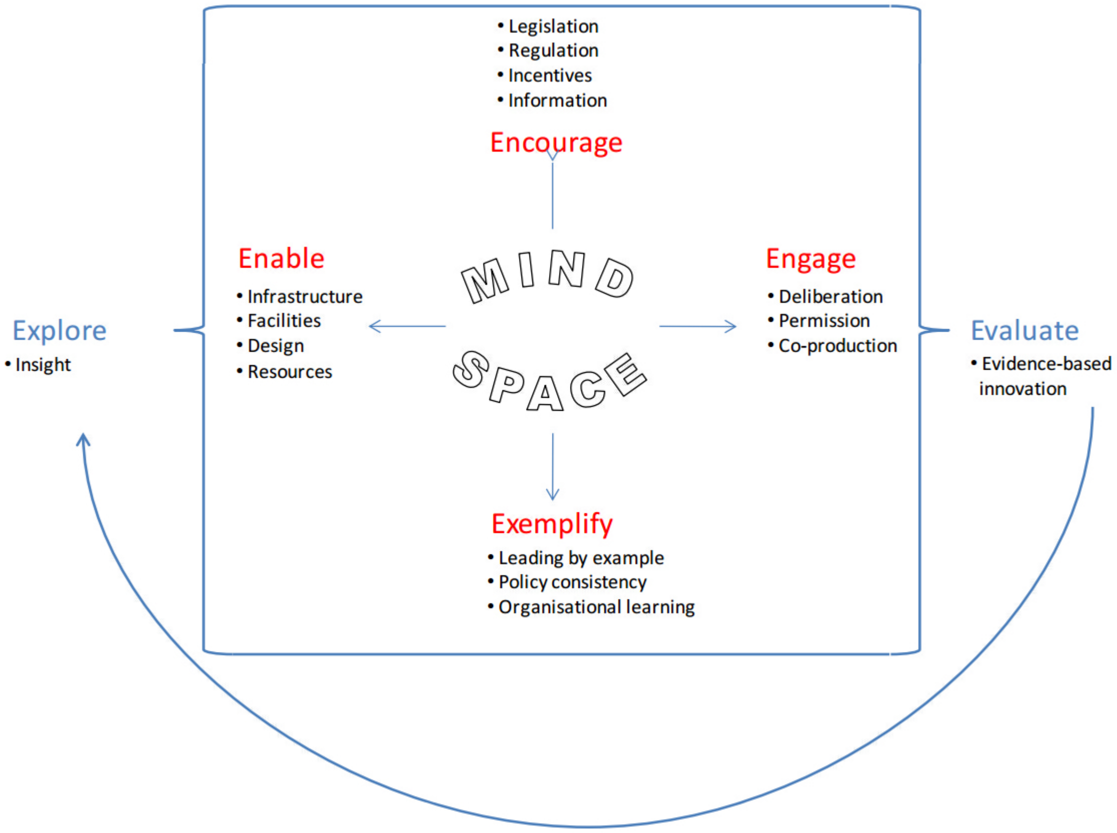

Fig. 2.6 Mindspace Diagram with 6 E's [84]

Defining the specific elements of behavioural change and how they may be altered is a powerful tool as it allows focused attempts to modify specific behaviours. Such a process enables the creation of highly effective nudges that can have a significant impact on the way in which people behave in a given environment.

Nudges are designed to exploit a given behaviour. Understanding the intricacies of MINDSPACE allows for the specific targeting and deployment of such nudges. Sunstein [218] provides a detailed overview of the ten most important factors with respect to nudges:

1. Defaults: see 2.3.3.

2. Simplification: complexity is a serious issue in both rich and poor nations. Programmes should be easily navigable and simplification in all forms should be a high priority. Overly complex procedures deters participants from important processes.

3. Use of social norms: social norms aim to highlight normative behaviour, essentially they reinforce the user that their behaviour is in line with most other users within the system. The use of social norms can be effective in reducing crime and promoting a healthier lifestyle (smoking and alcohol campaigns are highly prevalent).

4. Increases in ease and convenience: people often tend to choose the easiest option and thus a good strategy to adopt is "make it easy". If the desire is to promote a certain behaviour then it is a good idea to remove barriers that prevent adoption. Resistance to change is often a result of perceived difficulty or complexity rather than disagreement or scepticism.

5. Disclosure: openness of data repositories promotes trust through partnerships (opengovernmentpartnership.org for example). Simplicity is of the utmost importance in any strategy. Sunstein [218] states that "disclosure can operate as a check on private or public inattention, negligence, incompetence, wrongdoing, and corruption".

6. Warnings, graphic or otherwise: If a risk is serious, the best nudge might be to issue a public warning. "Large fonts, bold letters, and bright colors can be effective in triggering people's attention" [218]. The central belief to this theory is that "attention is a scarce resource" and warnings are attentive to that fact. Warnings aim to counteract the natural human tendency to towards unrealistic optimism and "simultaneously increase the likelihood that people will pay attention to the long-term". Warnings can be discounted by users in the belief that it will not happen to them, but can be combated by informing users (via descriptions) as to the relevant strategies to mitigate risk (if I do *x* I can combat *y*).

7. Pre-commitment strategies: denote a person's actions with respect to their predetermined goals. In essence, people with goals (i.e., quitting smoking) are more likely to behave in a manner that will help them to accomplish this goal rather than actions that will prolong its fulfilment. Such behaviour typically reduces procrastination.

8. Reminders: Reminders serve to combat non-fulfilment of a time dependant task. Whether it is through reluctance, procrastination, competing obligations or any other reason, reminders can have a profound impact on completion. Timing is of great importance with reminders as there needs to be sufficient time to complete the task at hand.

9. Eliciting implementation intentions: "people are more likely to engage in activity if someone elicits their implementation intentions" [218]. As a result, a simple question,

"do you plan to vaccinate your child?", can have significant impact on the user's choice and likelihood to act.

10. Informing people of the nature and consequences of their own past choices: Public and private organisations often have detailed information on their customer's past financial transactions (bank statements, energy bills), often having a more detailed account than the customers themselves. Being aware of one's past choices has a profound impact on future choices and can be a useful tool for behavioural change.

2.3.5 Ethics of Nudging

The ethics of nudging have been debated within academic literature for several years now. As such, it is an important and highly popular issue as users are increasingly aware of its adoption (many people are aware of Thaler's International Best-Seller 'Nudge' [228]). For this reason, it is important to understand why we should be concerned with respect to its implementation, and how these ethical concerns manifest themselves. As many of the studies included in this thesis investigate such interventions, it is important to discuss where ethical concerns may present themselves.

The ethics of nudging is a complex issue that stems from the belief that no matter how subtlety affected, a user's choice is being influenced and thus the decision is not entirely of their own free will. Ethics are closely related to the core elements of nudges and thus require analysis in a similar light.

Further complexity pertains to the application of the nudge and the context in which it is adopted. For example, one may argue that in a physical security context, a shove is indeed the correct manner in which to influence choice as the alternative may be detrimental to the user's safety or well-being. Within the context of IS we can envisage the process of securing sensitive data where the consequences of not securing it may outweigh the SRS implications related to shoving. There are several instances where data cannot be accessible to non-permitted users.

The discussion follows the 9 components of the MINDSPACE framework 2.3.4.

- **Ethics of Incentives:** Careful consideration is required before adopting incentive based behaviour change. It is necessary to determine:
  - the amount of incentive offered [38]. If the amount is set too high it may be considered coercive and present the user no realistic alternative (a "shove" rather than a nudge).

  - whether the incentive will disadvantage the people most in need [202]. Schmidt [202] likens this to the current highly incentivised models adopted in Germany where the "participation rates among people in the top socio-economic quintile are close to double the rates among those in the poorest quintile".
  - whether the incentive will result in the group that fails to meet the criteria for receipt being treated unfairly [202]. Perhaps it is better to have different completion requirements for different groups of people.
  - whether the incentive will harm the patient-physician relation [202]. For healthcare it is important that medical staff are not visualised as police causing patients to withhold important medical information. Volpp et al [243], found that enrolling patients in a lottery after successfully detailing the specifics of their medical history resulted in a drop from 22% to 2.3% in cases where the wrong medication had been prescribed.
  - and whether the incentive is fairly directed [38]. In the above example, the physician clearly plays a key role in determining the correct course of medication, but so does the patient.

- **Ethics of Defaults:** Defaults are a powerful tool that can have a significant impact on user behaviour. It is important to consider the following:

  - it must be easy to opt out [38]. This is necessary to preserve choice and requires that users are aware of the existence of the default, knowledge regarding how to opt-out and can opt-out without significant burden.
  - the harms and benefits of the default nudge [38]. It is important that we default people towards the most beneficial decision for them. Whilst seemingly obvious, Hanssesns [109] describes the process of opt-out HIV testing where it is beneficial physically to determine whether or not you have HIV, but the psychological harm can be significant.

- **Salience and Effect**: The method with which choices are presented has significant ethical effects that trigger certain methods of thinking and system 1 responses.

  - Does the use of salience affect an individual's autonomy [38]?
  - If so, is this ethical (manipulation can be ethically justifiable in some scenarios [105])?
  - Is the data presented accurate and true or is it exaggerated and misrepresented [38]?

- **Norms and Messenger:** The people that we listen to and gather advice from are not always the best role models.
  - we often nudge people towards bad decisions due to our "herd mentality" [38]. Obesity in the United States can be seen as a clear example of this process. Many celebrities and television personalities that are held with high esteem in the public eye often exhibit traits that are not in an individual's best interest to emulate.
  - there is a temptation to falsify accounts of users' actions [38]. For instance, "we could tell them that five out of seven people get screened [cancer] - but this would be a lie".
- **Priming:** Subconscious priming is conducted with good intentions but requires ethical considerations.
  - As such, the decision should be based on empirical evidence [38]. For instance, exercising promotes a healthy lifestyle and is compatible with the values of the individual.
  - Is it still easy for the individual to make their own decision rather than the decision they are being primed towards?
  - Is the priming ethically justifiable? In many cases, a user is unaware of the priming they are subject to. This manipulation must be used to promote "benefits that outweigh risks" [38].
- **Commitments and Ego:** The belief is that people wish to be consistent with public "promises and commitments and act in ways that make them feel better about themselves in order to nudge them toward healthier behaviours" [38]. The main ethical issues are:
  - whether ego is used for "good ends and good reasons" and whether this approach is preferable to reasoned argument. Spellecy [211], dictates that commitments based on intentions are "reason-centered commitments" and thus deserve more weight than desires.
  - whether bypassing this reason is done for "good ends" (e.g., not selfish ones) and for "good reasons" (people are harming themselves) [38].

The ethics of nudging are therefore a highly complex set of individual concerns which relate to specific components of the nudge. A detailed summary follows:

Table 2.3 Summary of Recommendation [38]

Nudge mechanism	**Ethically relevant considerations**
Incentives	•The amount and kind of incentives used. •Whether the incentive plan will disadvantage those most in need or result in the group that fails to meet criteria for receipt being treated unfairly (e.g., cost-shifting to those who fail, leaving those who fail by the wayside). •Whether the incentive plan will harm the patient–physician relationship (e.g., through actual or perceived monitoring). •Whether the incentive is fairly directed (e.g., at patients as opposed to or in addition to their physicians if the patients themselves are the ones who improved their health).
Defaults	•Whether people are aware of the existence of the default and whether it is fairly easy for people to opt out. •Whether the expected benefits of the default outweigh any anticipated harms, where harm is construed not just physically but also psychologically, socially, and financially. •Whether there are injustices or harms brought about to vulnerable or marginalized populations by the default (e.g., presumed consent for organ donation exploits the homeless who do not have easy opportunities to opt out/dissent) and whether attempts have been made to mitigate those effects.
Salience and affect	•Whether what is being represented saliently is true and accurate, as opposed to exaggerated or misrepresented. •Whether the use of salience and affect techniques will be perceived negatively by those it is directed toward. •Whether bypassing people's capacity for reason is done for good ends (e.g., not selfish ones) and for good reasons (e.g., people are harming themselves). •Whether there is a justification for using salience and affect instead of rational argument.
Norms and messenger	•Whether the information about what "most people are doing" is true and accurate. •Whether the use of comparisons and norms will do more good than harm in light of the fact that "what most people do" is often unwise. •Whether the power differentials between messenger and recipient have been considered.
Subconscious priming	•Whether it is fairly easy for people to go in a direction other than the one in which they are primed. •Whether subconscious priming is done for good and evidence-based ends. •Whether there is a justification for using subconscious priming instead of rational argument.
Commitments and ego preservation	•Whether ego is used for good ends and good reasons and whether there is a justification for using ego instead of rational argument. •Whether the person is making a commitment to self-destructive ends. •Whether the commitment is to long-term preferences or fleeting ones.

2.4 Human Computer Interaction in Security

Human Computer Interaction (HCI) governs the methods in which we interact with technology. The discipline has undergone significant transformation as our dependency increases and technology advances, generating new problems and the necessity for new technological solutions.

Carroll [50] defines HCI as "one of the first examples of cognitive engineering", a term coined within cognitive science that presents "people, concepts, skills, and a vision for addressing such needs through an ambitious synthesis of science and engineering". The inter-dependencies and multi-disciplinary nature of the subject can be visualised below (see 2.7):

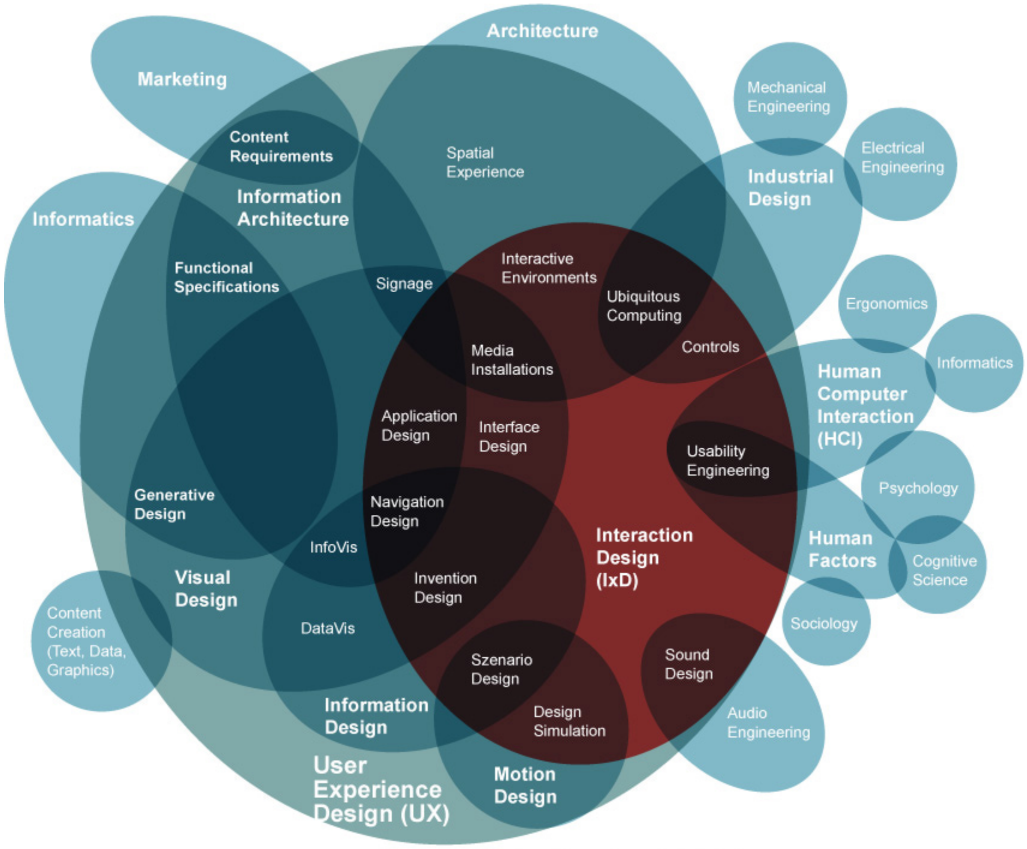

Fig. 2.7 The variety of disciplinary knowledge and skills involved in contemporary design of human-computer interactions [50].

The "original and abiding technical focus of HCI was and is the concept of usability" [50]. Usability accounts for:

- well being

- collective efficacy
- aesthetic tension
- enhanced creativity
- support for human development

A more dynamic view of usability is one of "a programmatic objective that should and will continue to develop as our ability to reach further toward it improves" [50]. HCI has traditionally been associated with the ’desktop’ environment but in recent years has expanded to encapsulate many new areas of technology (smart devices, information systems) which are typically more mobile.

To design effective HCI solutions, a cyclical design process is adopted known as the "task-artefact cycle" [51], 2.8.

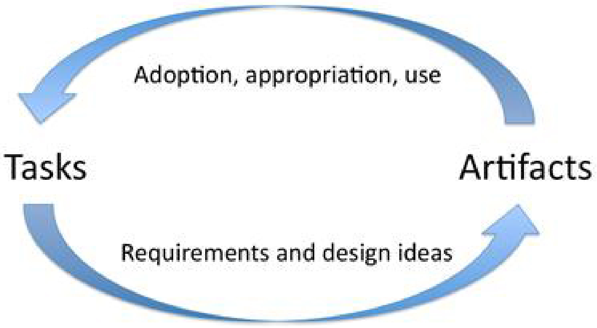

Fig. 2.8 Task Artefact Model [51]

The Task-Artefact Cycle relates to the "co-evolution of the activities people engage in and experience, and the artefacts - such as interactive tools and environments - that mediate those activities" [50]. HCI is therefore a process of critically evaluating the interactive technologies people use and their user experience. It is also about understanding how those interactions evolve with the adoption of new technologies and also how users’ knowledge, expectations, skills and visions expand. It is this assessment of such variables that drives forward the realisation of new devices and systems.

Carroll [50] states that the "dialectic of theory and application has continued in HCI" (approximating perhaps a dozen "currents of theory") that can be grouped into the following three eras:

- theories that view HCI as information processing,

- theories that view interaction as the initiative of agents pursuing projects,
- theories that view interactions socially and materially embedded in rich contexts. [50]

The change in paradigms throughout these era is represented in 2.9:

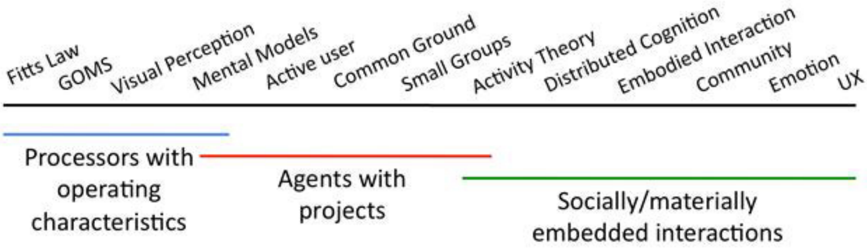

Fig. 2.9 A Series of Theoretical Paradigms Addressing the Expanding Research Ambitions of HCI [50].

To design effective **HCI** one must follow a strict set of criteria (successful HCI design will require an iterative approach often requiring repeated steps). It is essential to:

- Establish the requirements [196]. This requires careful consideration as to how the users interact with the product and what it is they are trying to achieve.
- Determine the alternatives [196]. This denotes predicting what alternatives are present and why they may be suitable or indeed preferable to the user. Assimilation of features may be necessary.
- Prototype [196]. Any implementation requires a testing phase. This allows for a visualisation of the solution and user testing.
- Evaluate [196]. Analysing the prototype and user feedback is critical to assessing whether the user requirements have been met. If not, this requires the repetition of previous steps.

2.4.1 HCI and Behaviour Change

The role of HCI is important when one considers the effectiveness of a given program or technology device in influencing behavioural change. Numerous studies have focused on health behaviour change encouraging physical activity [[17], [57], [58]], healthy diet [[75], [97]], glycemic control in diabetes [[157], [209]] and self-regulation of emotions [165].

Based on extensive empirical research, Prochaska et al [185] conclude that "for behaviour change to truly stick, a person has to maintain the target behaviour for several years". The problem associated with change is evident when one considers relapses or setbacks. For example, behaviours are often entrenched within an individual's routine requiring significant effort to alter. Cycling to work may seem a healthy alternative but may require the person to arrange child care, shower on arrival, find somewhere to secure their bike and accommodate bad weather [142]. Neglecting these prerequisites often has a negative impact on the user (changing routine can often affect social and workplace relations) causing resistance and resulting in failure with respect to behaviour change. There is also a budget (similar to ego depletion), where a user has a limited capacity for change and after a point, a user will typically revert back to previous behaviours. This often occurs during unexpected disruptions; catching a cold may dissuade a user from cycling to work.

It is therefore a complex process to ascertain the impact that technology plays on behaviour change. To be able to rule out "renewed commitment, social pressures, the effect of participating in a study" one must adopt a control group of "hundreds or even thousands of people and a matching control group" [142].

Efficacy Trials

Randomized control trials (RCT), a technique traditionally adopted during medical drug trials, are increasingly being utilized within the assessment of technological interventions. van der Berg et al [237] noted 10 cases where "internet-based interventions for promoting physical activity" had used RCTs exclusively as a method of trial. This process also featured heavily in testing mobile phone applications.

It is clear that RCTs play a vital role in the testing and evaluation of behaviour change systems but it is important to state their limitations. RCTs are often not feasible for smaller groups with early prototypes. Large trials require significant financial and time contributions which may not be viable at an early stage. RCTs often "reveal little about why the technology under evaluation is or is not effective" [142].

Hurling et al [119] recently adopted such a trial strategy to evaluate the effectiveness of a physical activity application that used accelerometers to assess physical movements. The data was subsequently viewed on mobile and web-based systems. The approach itself involved numerous behaviour change strategies: "self-monitoring, identification of barriers to change, planning, problem-solving, public commitment, and customized feedback". This approach was cross-analysed with a control group that used the accelerometers but did not receive feedback (N=77).

The results showed that there was no significant difference between groups "in overall physical activity" but "indicated that the intervention group increased their amount of leisure time activity more than the control group" [142],[119]. Whilst these results seem compelling, Klasnja [142] aruges that there are 4 main issues that "limit the usefulness [of the study] and suggest why RCTs should not be seen as the only valid model for evaluating ... especially in the early stages of research".

1. Sample Size: Although Hurling et al's [119] study was large by HCI standards, it was very small in relation to RCTs. For this reason it is impossible to rule out any outside changes that may have had an effect on the group. In addition, the results were deemed strongly suggestive and not the conclusive, a factor which can only be combated by a much larger RCT, something that typically involves hundreds if not thousands of participants.

2. Multiple Behaviour Strategies: As the study involved numerous behaviour changes, it is unclear to which had the most effect or indeed which changes affected which behaviour. To fully assess this, it would require an even larger RCT than the one previously mentioned for sample size.

3. Qualitative Data: The study lacked qualitative data that is needed to make "a thorough analysis of how participants perceived the system". Hurling [119] concluded that intervention was effective but did not know how or why. From a HCI perspective, this is highly significant as we cannot design a better system without knowing which parts of the system worked and which did not [142].

4. Time and Cost: Due to their size and cost, "efficacy trials typically evaluate complex systems that combine many intervention strategies to maximize effectiveness" [142]. Therefore, Hurling et al's [119] study of N = 77 over a 3 month period are short compared to typical RCT studies. Klasnja [142] concludes that "the resources and effort required to run true efficacy trials make evaluations of innovative technologies that embody early-stage, high-risk ideas simply infeasible".

These limitations are important to acknowledge with respect to our bespoke studies that typically include similar small-scale samples. It is important that the design of our studies are sensitive to such limitations, specifically in relation to the effects of testing multiple behaviours simultaneously. We do, however, combat the lack of qualitative data that is highlighted as we specifically assess the impact of our interventions upon the user.

Klasnja [142] offers novel solutions to these problems. Based on over 40 years of research it has been shown that (with respect to self-monitoring) "simply keeping track of a behaviour changes the frequency of that behaviour in a desired direction" [143], [169]. In light of this, a self-monitoring intervention should therefore evaluate:

- The rates of the target behaviour increase from their baseline levels prior to the interventions [142].

- Whether after the intervention is stopped, the rates of behaviour begin to go down again [142].

It is therefore necessary to tailor evaluations to specific intervention strategies. This enables HCI researchers "to show that their systems are doing what they are supposed to be doing, without requiring a full-blown demonstration of behaviour change" [142]. Such a strategy is beneficial as it allows for direct comparisons of the same intervention strategies across different implementations. This leads to being able to determine how "the design of a technology for behaviour change affects the technology's use by its target audience in situ" [142].

2.5 Consumerisation & BYOD

The inclusion of Bring Your Own Device (BYOD) in the workplace is no longer a new phenomenon, it is an everyday occurrence. The consumerisation of the workplace is partly due to employees expecting "to be able to use all the innovative new devices and tools at their disposal, both to do their jobs and to maintain their always-connected lifestyles while being able to work whenever and wherever they need to" [37].

The trend in BYOD is controversial as it introduces both a number of benefits and risks to the individual and organisation. In summary:

- Benefits: increased employee efficiency, improvement of employee satisfaction, and eventually lower IT costs [37].

- Risks: multiplicity of devices, the right to use those devices freely, greater access to the company network can degrade security, reduce productivity, increase support costs, and expose companies to compliance and reputational risks [37].

To integrate users' devices within the workplace there are generally two methods. Firstly, they can "bring in" employees under the "corporate umbrella" allowing them to use corporate devices as if they were their own. They can install custom software, use the device for social and out of work activities and loosen restrictions on web access. Alternatively, they can "reach out" allowing employees to use their own devices to complete work, perhaps with the aid of virtual clients [37]. Fig 2.11 identifies the full choice matrix.

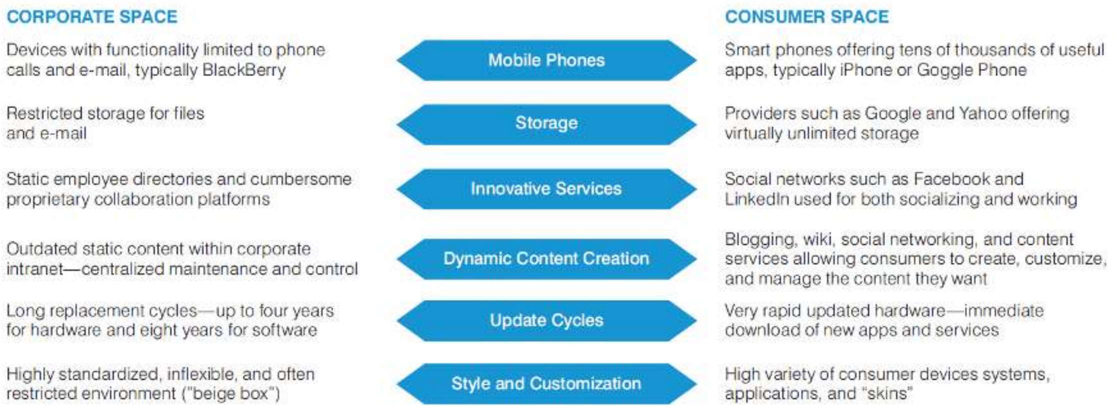

CORPORATE SPACE		CONSUMER SPACE
Devices with functionality limited to phone calls and e-mail, typically BlackBerry	Mobile Phones	Smart phones offering tens of thousands of useful apps, typically iPhone or Goggle Phone
Restricted storage for files and e-mail	Storage	Providers such as Google and Yahoo offering virtually unlimited storage
Static employee directories and cumbersome proprietary collaboration platforms	Innovative Services	Social networks such as Facebook and LinkedIn used for both socializing and working
Outdated static content within corporate intranet—centralized maintenance and control	Dynamic Content Creation	Blogging, wiki, social networking, and content services allowing consumers to create, customize, and manage the content they want
Long replacement cycles—up to four years for hardware and eight years for software	Update Cycles	Very rapid updated hardware—immediate download of new apps and services
Highly standardized, inflexible, and often restricted environment ("beige box")	Style and Customization	High variety of consumer devices systems, applications, and "skins"

Fig. 2.10 Consumerisation: Corporate vs Consumer [37]

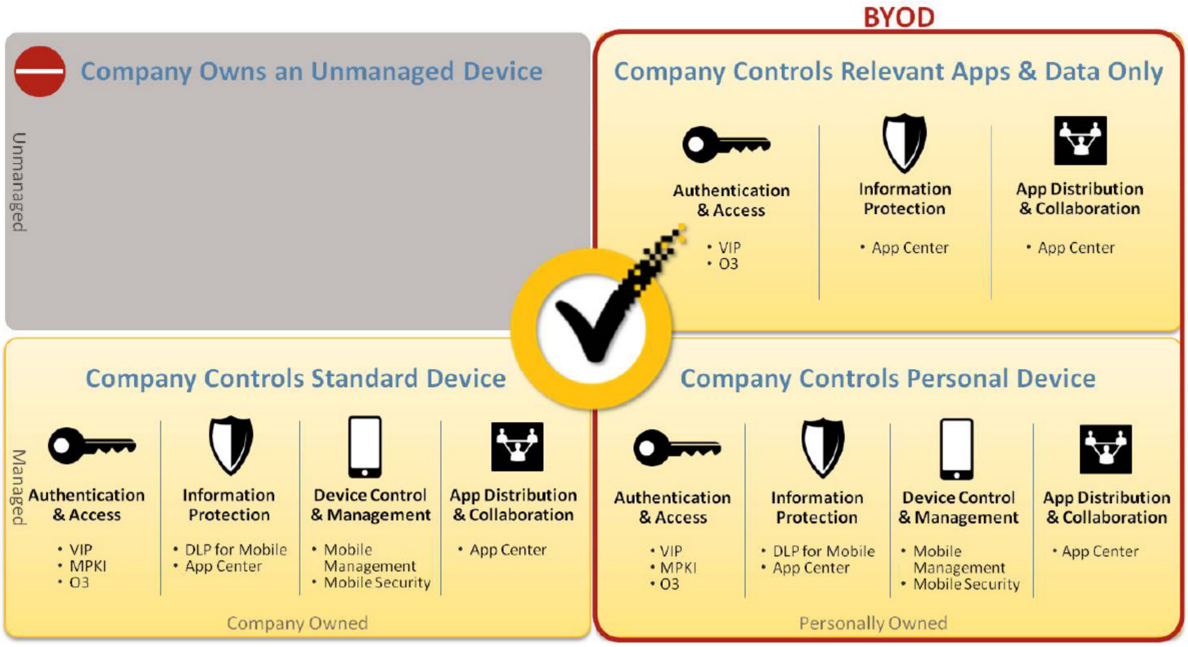

Fig. 2.11 Bring Your Own Device Matrix [221]

The scale of the problem is highlighted in a survey conducted by Harris Interactive and ESET (Nov 2012) [48] where it was discovered that "more than 80% of employed adults use some kind of personally owned electronic device for work-related functions". Further investigation also shows that:

- 47% of employees use personal desktop computers to access or store company information
- 41% of employees use personal laptops to do this

- 25% use smartphones
- and 10% use tablets

Morrow [166] argues that as "half of these devices" are not protected by even the most basic security measures, organisations may begin to feel that the "security challenges associated with BYOD far outweigh the benefits".

One of the most difficult challenges that organisations face is the knowledge that corporate data is being accessed and shared with devices which are not within the control of the IT department. This has profound impacts on the likelihood of "data leakage, data theft, and regulatory compliance" [166]. Morrow [166] further highlights the issues of BYOD through analysing the semantics of the acronym itself. He states that the **D** (device) includes more than just smartphones. It also includes:

- Employees logging into web applications such as Outlook Web Access and SharePoint.
- SaaS applications such as CRM systems.
- Healthcare billing applications hosted in the cloud services.

and concludes that "laptops, smartphones and tablets that connect to corporate networks significantly increase threats to sensitive data". Organisations must be aware and should be concerned about the security risks these endpoint devices pose. The ability to steal confidential information and share via cloud based storage solutions (Dropbox, YouSendIt, E-mail) is significant and companies must do more to "control the data after it's delivered to the device in order to prevent accidental or intentional loss by careless of malicious end users" [166].

Loss of devices is also of critical importance in the realm of BYOD. Ernst & Young [259] define loss as one their "five basic concerns" along with physical access, the role of ownership, always on with increased data access, and lack of awareness. They speculate that approximately 22% of the total number of mobile devices produced will be lost or stolen during their lifetime and over 50% will never be recovered. Although the motives are often opportunistic, and focused on the value of the hardware itself, "a growing amount of lost or stolen phones have their content accessed by someone other than their owners" [259]. This highlights the importance of "basic security features such as password protection, encryption and robust procedures to wipe the device once lost".

In light of the above 'five basic concerns', Ernst & Young [259] propose the following with respect to securing employees' devices:

- Evaluate device usage scenarios and investigate leading practices to mitigate each risk scenario.
- Invest in a mobile device management (MDM) solution to enforce policies and monitor usage and access.
- Enforce industry standard security policies as a minimum: whole-device encryption, PIN code, failed login attempt actions, remotely wiping, etc.
- Set a security baseline: certify hardware/operating systems for enterprise using this baseline.
- Differentiate trusted and untrusted devise access: layer infrastructure accordingly.
- Introduce more stringent authentication and access controls for critical business apps.
- Add mobile device risk to the organization's awareness programme.

2.5.1 Mobile Devices

It is estimated that there is now in excess of 7.4 billion mobile devices [122] and 3.7 billion "unique mobile subscribers" [122] with an annual growth rate of 6.1% and 5.35% respectively. A study conducted by Infonetics [121] found that almost all enterprises they surveyed had instances of malicious apps being downloaded onto a devices. Furthermore, 64% reported that users' devices "containing sensitive or proprietary data had been lost or stolen" but there were very few enterprises who had security measures in place to protect those devices [166], [121].

Symantec [220] claims that during 2011, within the Android environment, "more than half of all Android threats currently collect device data or track users' activities" and almost a quarter of mobile threats are designed to send information such as personal data (see Fig 2.12). A common example of such a practice is the 'free-to-play' game marketplace (these apps are consistently within the top 10 most downloaded apps [103]) where users often ignore permissions or do not fully understand the implications of the software. When using the device in a corporate environment this may lead to inadvertent leaks of confidential company information simply by saving email attachments or other media to the local storage [166].

In a similar fashion, the rise in Quick Response (QR) codes has enabled the rapid spread of malware. These codes are simple to use (take a photograph or actively scan using your camera enabled device) and provide a hyperlink to website. The problems occur with respect

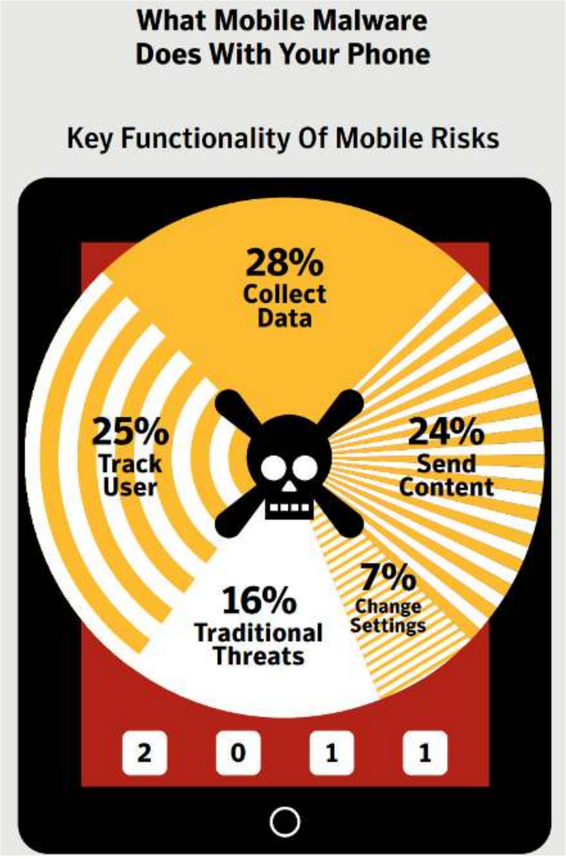

Fig. 2.12 The Purpose of Malware [220]

to the obscurity of the URL and the pending web page or download. It is difficult if not impossible to fully verify the address as they are often adopt URL shortening, in an attempt to mask their malicious nature, and often open by default on the mobile operating system. This may result in malicious software being downloaded under the guise of an official product. The user must then act upon the download or web page for the attack to be complete or indeed it is possible for actions to be automatically taken if the user has decided to automatically ignore prompts. The device is then open to data leakage and malicious activity [220], [221].

2.5.2 Device Vulnerability Management

Tenable [226] reported in 2012 that device vulnerability management was a top concern for security professionals. Their study surveyed attendees at the RSA Conference 2012 and discovered that nearly 70% of people believed that mobile device vulnerability management was ’very important’ when compared to other security avenues. Furthermore, almost all participants believed that mobile devices posed a significant threat to their businesses security, yet 68% said they currently have "no way of identifying known mobile device vulnerabilities that could be affecting their network and 67% said they either have no controls in place for mobile device usage on their network, or their employees simply ignore existing mobile device usage policies" [166], [226].

Although the majority of corporate data we access, and material we consume via the internet is encrypted and tunnelled via Secure Sockets Layer (SSL) encryption and Transport Layer Security (TLS), the end point device used to access this content poses a significant security threat if unsecured. In BYOD this end point is an unknown (with respect to non-managed user devices). IT professionals are unaware as to whether the device has software defences (virus/malware scanner) and even whether this is up to date. Furthermore, a 2012 study conducted jointly by Skype, Norton and Tom Tom [208] found that 40% of the respondents admitted they don't upgrade software when they should leaving them open to many cyber attacks and malware.

Whilst virus/malware scanners are somewhat effective tools for mitigating malicious software acquisition, they do not prohibit man-in-the-middle attacks (or more appropriately malware specific to mobiles, man-in-the-mobile), or man-in-the-browser attacks [166]. This observation is important for numerous reasons:

- Browser-based sharing is a popular and highly problematic activity with respect to BYOD. (See Fig 2.13).
- Such software does not provide security against such an attack.
- IT professionals cannot control which browser the user chooses to use (each have separate vulnerabilities), or indeed which plug-ins and security patches they have.
- It is impossible to know whether or not the device is currently infected.
- IT staff cannot access the devices' cache, password storage, web history etc. Simply copying and pasting may cause sensitive data to be compromised by malware scraping cache files.
- Browser-based file-sharing is set to become more popular and require ever closer management (see Fig. 2.14).

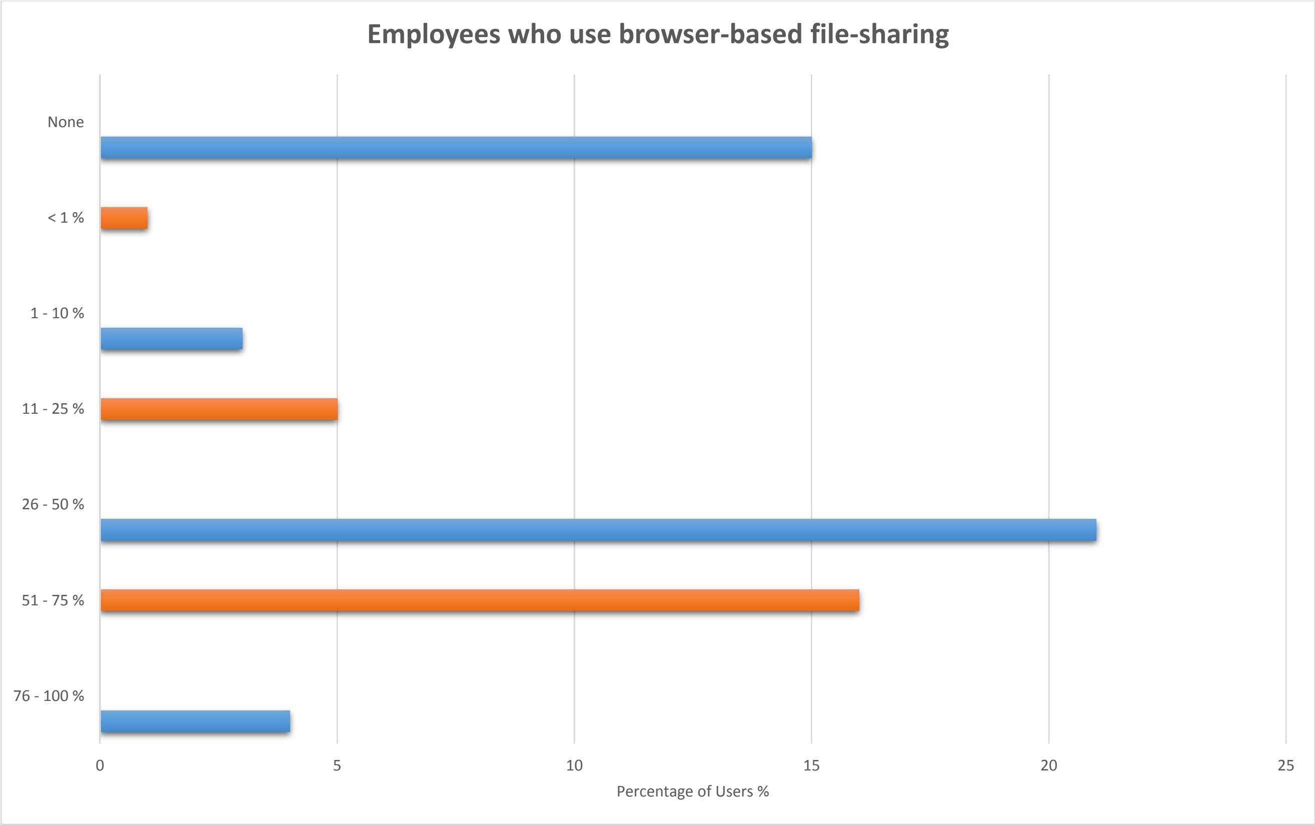

Fig. 2.13 The percentage of users who use browser-based file-sharing [184].

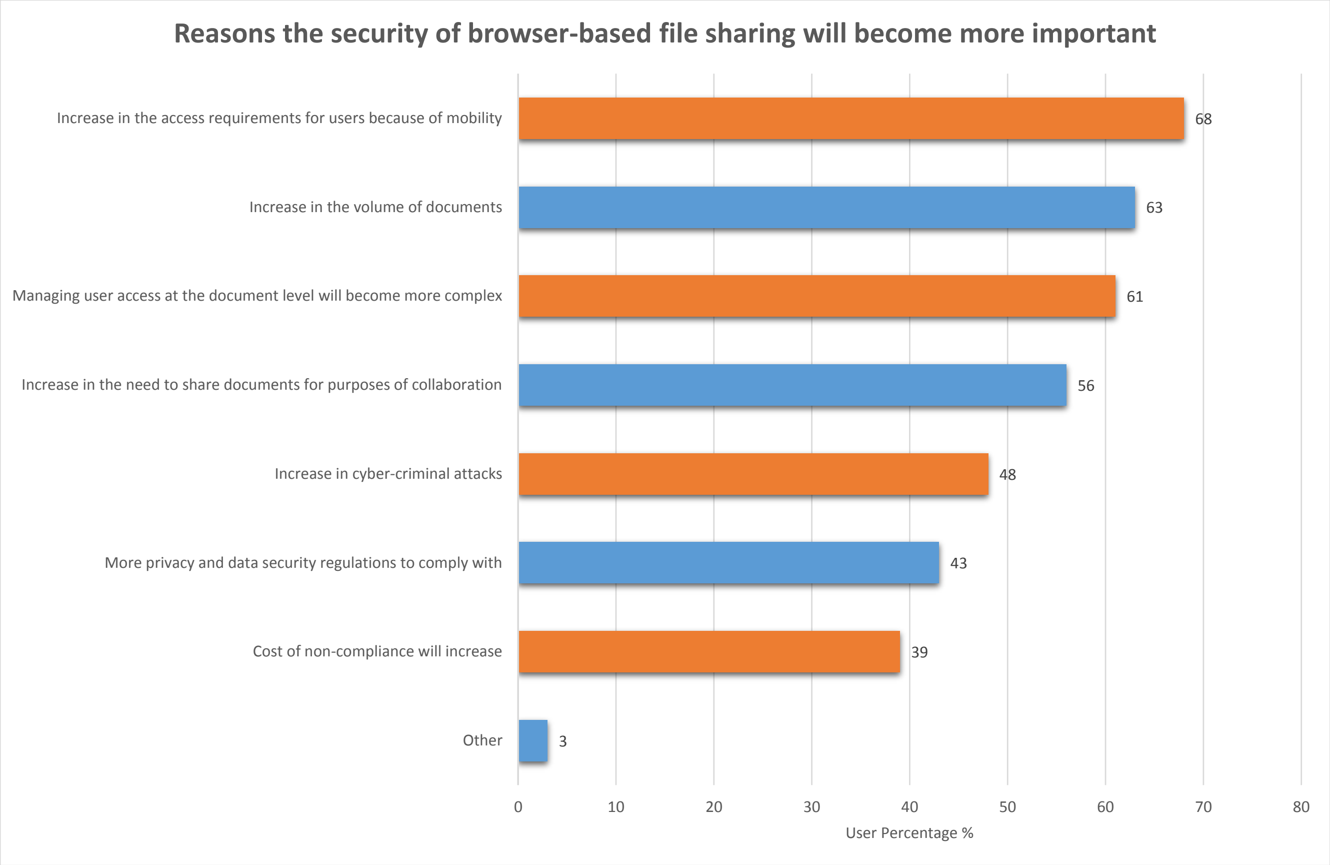

Fig. 2.14 Reasons the security of browser-based file sharing will become more important [184].

From the previous discussion and Figs 2.13, 2.14, it is clear to see why BYOD introduces many significant risks with respect to information security management. It also shows how this trend is set to continue (2.14) and how at present there are few security strategies in place to mitigate the risks. Furthermore, many IT professionals are either unaware, or simply do not know how to calculate this risk, let alone manage it. Morrow [166] argues that there is no single solution to securing your network from the vulnerabilities and risks that BYOD introduces. Instead, he proposes that "to counter these sophisticated threats, organisations should employ a layered security strategy that provides necessary access to corporate information while minimising risk and maintaining compliance". To go further, one must place importance on more than just authorised and unauthorised access and must ensure that content delivery is secured from transport to delivery and subsequent end-point access. To do so requires the dismissal of archaic network visualisations, a BYOD device is a part of your network that needs to be protected. Employing strategies such as "compartmentalising access to sensitive data", employing better "auditing logs" and log analysis, and deploying strategies that are actively engineered to address BYOD are all required to reduce data loss.

Education can also be seen as a critical factor in BYOD security (see section 3.4.2). Enabling users to be able to distinguish between the use and appropriateness of such devices in the workplace will help to reduce accidental, careless leaks. Ensuring that employees are familiar with security policy and procedure will also be of benefit [166].

2.6 Privacy

The role of privacy is important to our investigations as it contributes towards a user's decision-making process. Specifically with respect to data, privacy is often a key concern for users. In the digital age, inter-connectivity and the BYOD trend presents many challenges in areas such as location tracking, search history, and cookie usage. As such, it is necessary to understand exactly how privacy is defined, how it can be related to specific problem domains, and ultimately how it impacts the decision-making process.

Privacy is a "state in which one is not observed or disturbed by other people" [79]. Information privacy refers to "the user's ability to control when, how, and to what extent information about themselves will be collected, used, and shared with others" [162].

Figures [2.15, 2.16] [162], denote which information users share, and to whom they share it with respectively. The likelihood of information sharing and the willingness to share with a given party reduces radially meaning that users are more likely to disclose sensitive personal information with somebody they know and trust rather than a large organisation with whom they do not.

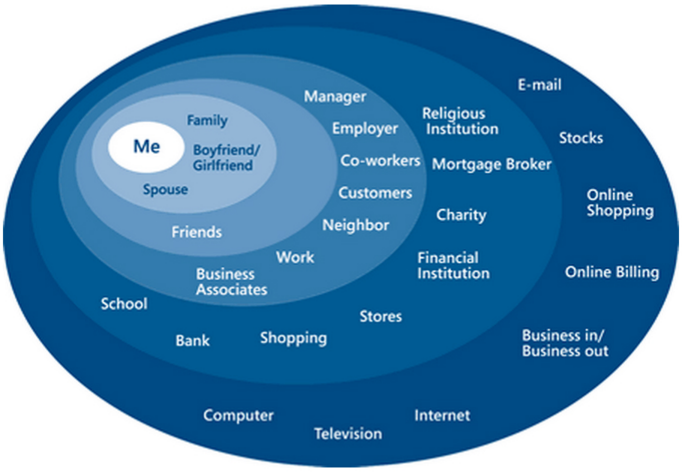

Fig. 2.15 Who users share their information with [162]

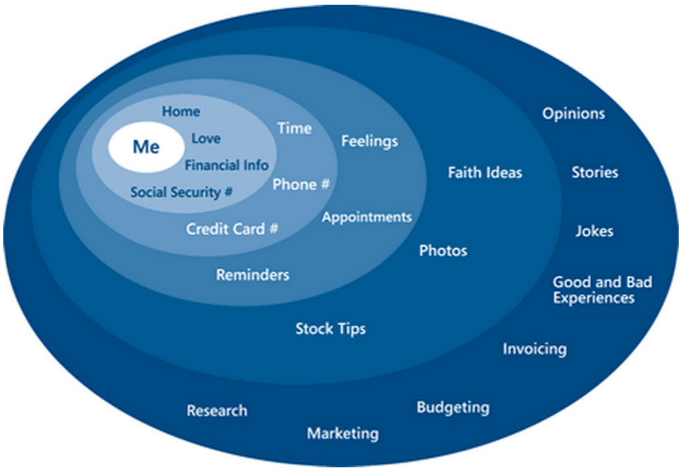

Fig. 2.16 The information users share [162]

Privacy was far simpler to preserve before the digital technology age. Records were physically based, and observations were physical and impossible to record in real-time. Since then, the birth of modern technology, specifically computer-based technologies, has allowed for autonomous data collection and mass surveillance. No longer are records solely physical, or observations simply literal recollections; there is now a digital footprint. The impact of this footprint is significant and the rights to control and the methods in which it is controlled and secured is a topic that is highly controversial and not fully understood. In essence, the problem has arisen far quicker than a remedy can be doctored.

Acquisti [2], highlights the complexities we face with modern day privacy and states that "several technological approaches have been proposed to solve the problem of personal privacy" and that "in almost any conceivable scenario, when making purchases, ... the identity of the individual can be disassociated from the rest of the information revealed during the transaction". Furthermore, he states that "comapnies based on such technologies (preserving anonymity) ... have struggled to balance the differing needs of the various parties in the privacy equation ... failing to gain to widespread adoption". This is important as it implies that identities are requested outside of necessity, eluding to the presence of external influences.

This idea is supported through the knowledge that whilst "privacy and security of personal information remain a concern for many, the economic incentives have not generated widespread adoption, and government intervention has increased the responsibilities for companies to collect personal information, without determining their liabilities for misuses of those data" [2]. For this reason, it is common to view privacy from an economics perspective.

One of the major difficulties underlying preserving privacy lies with the very ambiguity of the phrase itself and hence, "protecting privacy is a vague concept" [2]. This is exemplified through the notion that not only do different parties "have opposite interests and views about the amount of information to disclose during a certain transaction" but that the individual may also "face trade-offs between [their] need to reveal and [their] need to conceal different types of personal information" [2].

2.6.1 On-line and Off-line Identities

People often have separate on-line and off-line identities that are utilised to differentiate which information they share and how they share it (ultimately affecting the manner in which they can be identified). From an economic perspective:

- **On-line:** typically includes information such as purchase history, browsing behaviour, **IP** address, and cookies. Economic models would differentiate users by adopting types that aim to categorise users based on their preferences and behaviours. For example, "when I log into Amazon.com with a Hotmail.com email address ... I am revealing my on-line identity" [2].

- **Off-line:** differs in that it actually defines the identity of an individual. Identifiers include personal attributes such as credit card numbers, National Insurance Numbers. When completing the same transaction on "Amazon.com with my personal credit card, I am revealing my off-line identity" [2].

The problem with maintaining this difference in identities is exacerbated by the dependency on several "legacy" processes and existing infrastructure that is still in operation [2].

Re-identification often requires simple cross-analysis between records. For instance, when purchasing from Amazon cookie data may be cross-referenced with credit card information sharing online behaviour and spending habits with other third parties.

It is important to state, however, that not all outcomes are negative. Several studies have identified numerous scenarios where it is beneficial to have your identity known:

- "allowing firms to use cookies make society better off, because the buyer can benefit from the seller knowing him better and thereby providing him target services" [6].
- "sharing information between sellers reduces the distortions associated to asymmetric information between buyer and seller" [47].
- "when the seller is facing strategic customers, she will autonomously tend to adopt a policy that protects the privacy of her customers" [223].

Many of the problems related to "the distrust of newcomers is an inherent social cost of easy identity changes" [192].

2.6.2 How can Economic Models and Technology Help?

Acquisti [2], argues that "economics can assist in the design process of mechanisms to solve the impasse when no party alone would have the incentive to perform certain actions", a premise that is achievable for example by parties sending "dummy" traffic to one another to increase anonymity within the system. Furthermore, in a more general sense, "socially-informed design of privacy technologies economics can be used to define what information should be shared, and what protected" [2].

These proposals can then be assisted by technology within the realms of law and governance. Samuelson [200] states that such an approach "should place constraints and liabilities on the side of the parties receiving private information, calibrating them in order to compensate the moral hazard and asymmetric information in the market of personal data, and combining them with information technology as a “commitment” device in the system".

This in turn creates an incentive and opens a market for third parties to create solutions that help to preserve anonymity off-line but make it possible to act privately on-line. Scoglio [203], adds that if privacy is a holistic concept, only a holistic approach can provide its adequate protection and define:

- economic tools to identify the areas of information to share and those to protect [2].
- law to signal the directions the market should thereby take [203].
- the technology to make those directions viable [203].

2.7 Prospect Theory

Prospect theory presents a detailed understanding pertaining to user decisions in risky situations. This notion fits our problem domain well and provides reasoning behind user actions that others may deem unusual or contradictory given the context and options available.

Kahneman [136] presents Prospect Theory (PT) as a critique to Expected Utility Theory (EUT) for situations of decision making under risk. He states that "choices amongst risky prospects exhibit several pervasive effects that are inconsistent with the basic tenets of utility theory". This approach proposes that people "underweight outcomes that are merely probable in comparison with outcomes that are obtained with certainty". The theory is intrinsically different to EUT with respect to the following areas:

- Value is assigned to gains and losses as opposed to final assets.
- Probabilities are replaced by decision weights.
- The value function is normally concave for gains, and typically convex for losses. It is mostly steeper for losses than gains.
- Decision weights are generally lower than the corresponding probabilities (except when faced with low probabilities).
- Overweighting of low probabilities may increase the likelihood of both insurance and gambling.

2.7.1 Definitions

Decision making under risk can be conceptualised as a choice between prospects and risks [136]. A prospect, defined as:

$$x_1, p_1; ...; x_n, p_n$$

is a contract that yields outcome x_i with probability p_i where:

$$p_1 + p_2 + ... + p_n = 1$$

To simplify notation, we omit null outcomes and use (x, p) to denote the prospect $(x, p;\ 0, 1 - p)$ that yields x with probability p and 0 with probability $1 - p$ [136]. The (riskless) prospect that yields x with certainty is denoted by (x).

Kahnman further specifies that the application of expected utility theory is based upon three tenets; expectation, asset integration, and risk aversion.

- Expectation: $U(x_1,p_1;...;x_n,p_n) = p_1u(x_1)+...+p_{nu(x_n)}$.
  That is, the overall utility of a prospect, denoted by U, is the expected utility of its outcomes.

- Asset Integration: $(x_1,p_1;...;x_n,p_n)$ is acceptable at asset position w if $U(w+x_1,p_1;...;w+x_n,p_n) > u(w)$.
  That is, a prospect is acceptable if the utility resulting from integrating the prospect with one's assets exceeds the utility of those assets alone. Therefore, the domain of the utility function is final states (which include one's asset position) rather than gains or losses.

- Risk Aversion: u is concave $(u^{''} < 0)$.
  A person is risk averse if they prefer the certain prospect (x) to any risky prospect with expected value x. In expected utility theory, risk aversion is equivalent to concavity of the utility function.

Certainty

The Certainty Effect (CE) [136] relates to the manner with which people overweight outcomes that are considered certain, relative to the outcomes which are merely probable. The CE uses demonstrable examples to prove this:

Problem 7: A: (6,000, .45), **B:** (3,000, 0.9).
N = 66, A = [14], B = [86]
Problem 8: C: (6,000, .001), **D:** (3,000, .002).
N = 66, C = [73], D = [27]
where (x, y) x = prize, y = probability; N = number of responses; [z] = percentage of N

Fig. 2.17 The Certainty Effect - A Worked Example [136]

From 2.17 we see a stark contrast between the participants' responses. **Problem 7** includes substantial winning odds and in this case the majority of participants chose answer B. In **Problem 8**, however, there is a possibility of winning although the probability of doing so is highly unlikely. In this instance, most participants chose the option with the higher odds and were less interested with the prize value. This difference highlights the more common aspects of risk that cannot be captured using the EUT. Kahneman [136] proposes the following empirical generalisation to the manner in which the substitution axiom is violated:

$$\text{If } (y, pq) \text{ is equivalent to } (x, p)$$
$$\text{then } (y, pqr) \text{ is preferred to } (x, pr), 0 < p, q, r < 1$$

2.7.2 Reflection Effect

The Reflection Effect examines the impact of negative rewards. It is important as it implies that risk aversion in the positive domain is accompanied by risk seeking in the negative domain [136]. To exemplify, let us refer to a different problem:

Problem 3: Positive Reward **A:** (4,000, .8), **B:** (3,000, 0).
N = 95, A = [20], B = [80]
Problem 3: Negative Reward **C:** (-4,000, .8), **D:** (-3,000, 0).
N = 95, C = [92], D = [8]
where (x, y) x = prize, y = probability; N = number of responses; [z] = percentage of N

Fig. 2.18 The Reflection Effect - A Worked Example [136]

Figure 2.18 details the change to risk seeking behaviour when choosing between negative prospects. Users are more willing to risk losing larger sums of money in preference of a sure loss of 3,000. Studies show that a "translation of outcomes produces a dramatic shift from risk aversion to risk seeking" [254].

Figure 2.18 details another important factor. Both the positive and negative prospects demonstrate that "outcomes which are obtained with certainty are outweighted relative to uncertain outcomes" [136], a finding that is inconsistent with EUT. In the positive domain, "the certainty effect contributes to a risk averse preference for a sure gain over a larger gain that is merely probable". Conversely, in the negative domain we observe that the same behaviour in fact leads to a risk seeking preference for a loss that is "merely probable over a smaller loss that is certain".

2.7.3 Probabilistic Insurance

Insurance, with respect to both large and small scale loss, is considered strong evidence for the concavity of the utility function when dealing with financial numeration. This belief stems from the fact that many people spend large sums of money on purchasing insurance policies that often exceed the expected actuarial cost. On closer inspection, however, it is noted that the relative attractiveness of various forms of insurance does not support the concave nature of the utility function with respect to money [136]. People often prefer insurance plans that offer "limited coverage with low or zero deductible over comparable policies that offer higher maximal coverage with higher deductibles - contrary to risk aversion".

Figure 2.19 depicts a scenario where one pays a certain cost in order to minimize the probability of a given event without eliminating it altogether. As is visible from the participant answers, this is not a popular choice. To summarise, $p/2$ is less desirable than $p = 0$.

Problem 9: Suppose you consider the possibility of insuring some property against damage, e.g. fire or theft. After examining the risks and the premium you find that you have no clear preference between the options of purchasing insurance or leaving the property uninsured.

It is then called to your attention that the insurance company offers a new program called *probabilistic insurance*. In this program you pay half of the regular premium. In case of damage, there is a 50 per cent chance that you pay the other half of the premium and the insurance company covers all the losses; and there is a 50 percent chance that you get back your insurance payment and suffer all the losses. For example, if an accident occurs on an odd day of the month, you pay the other half of the regular premium and your losses are covered; but if the accident occurs on an even day of the month, your insurance payment is refunded and your losses are not covered.

Recall that the premium for full coverage is such that you find this insurance barely worth its cost.

Under these circumstances, would you purchase probabilistic insurance:

Yes, No.
N = 95 [20], [80] respectively.

Fig. 2.19 Probabilistic Insurance - A worked Example [136]

In contrast to this finding, traditional EUT (with a concave u) would imply that probabilistic insurance is preferable to regular insurance [136]. This is an interesting yet perplexing finding because "probabilistic insurance appears intuitively riskier than regular insurance, which entirely eliminates the elements of risk". Therefore, it is reasonable to suggest that "the intuitive notion of risk is not adequately captured by the assumed concavity of the utility function for wealth" [136].

2.7.4 The Isolation Effect

The Isolation Effect examines how individuals choose between alternatives. Specifically, it focuses on the components of that choice and proposes that "people often disregard components that the alternatives share, and focus on the components that distinguish them" [232]. This method of choice resolution typically yields inconsistent preferences as a "pair of prospects can be decomposed into common and distinctive components in more than one way, and different decompositions sometimes lead to different preferences" [136]. For example:

Problem 10: Consider the following two-stage game. In the first stage, there is a probability of 0.75 to end the game without winning anything, and a probability of 0.25 to move into the second stage. If you reach the second stage you have a choice between:

(4,000, .80) and (3,000).

Your choice must therefore be made before the game starts, i.e., before the outcome of the first stage is known.

Fig. 2.20 The Isolation Effect - A worked Example [136]

In figure 2.20, the final outcome probabilities can be summarised as $(4,000,.2)$ and $(3,000,.25)$. From this study (N = 141), 78% chose $(3,000)$ and 22% chose $(4,000,.8)$. This is in direct conflict with **Problem 4** where:

Problem 4: A: (4,000, .2) **B:** (3,000, .25)

Fig. 2.21 The Isolation Effect - Problem 4 [136]

From this, we deduce that people ignored the first stage of the game (whose outcomes are shared by both prospects) and merely focused on the odds in the second part. To further examine this phenomena one can build a decision tree to more closely examine the mental process (This process forms part of the rationale behind Section 7). *Squares represent decision nodes, circles represent chance nodes.*

From analysis, we are able to visualise the different placement of the decision point between figure 2.22 and figure 2.23. More specifically, this different reflects a change to the decision prospects whereby figure 2.22 denotes a choice between two risky prospects and figure 2.23 requires a choice between and risky and riskless prospect. This is possible as there is now a "dependency between the prospects without changing either probabilities or outcomes" [136]. This is now apparent as the event ’not winning 3,000’ is included in the event ’not winning 4,000’ (in the standard formulation), but is separated into two independent events in the sequential formulation. Thus, "the outcome of winning 3,000 has a certainty advantage in the sequential formulation, which it does not have in the standard formulation" [136].

This is highly important as it details how preferences may be altered by different representations of probabilities and forms a mathematical basis for the theory of nudging (see section 2.3).

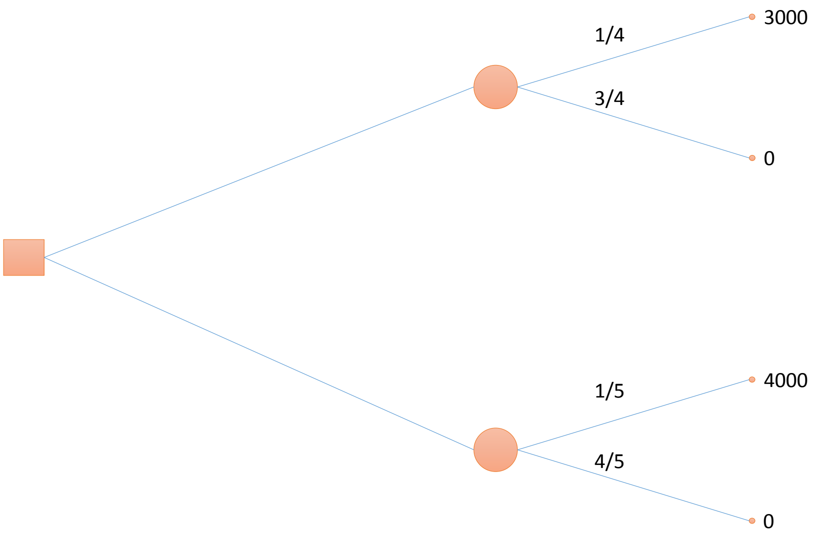

Fig. 2.22 Problem 4 as a decision tree (Standard formulation) [136]

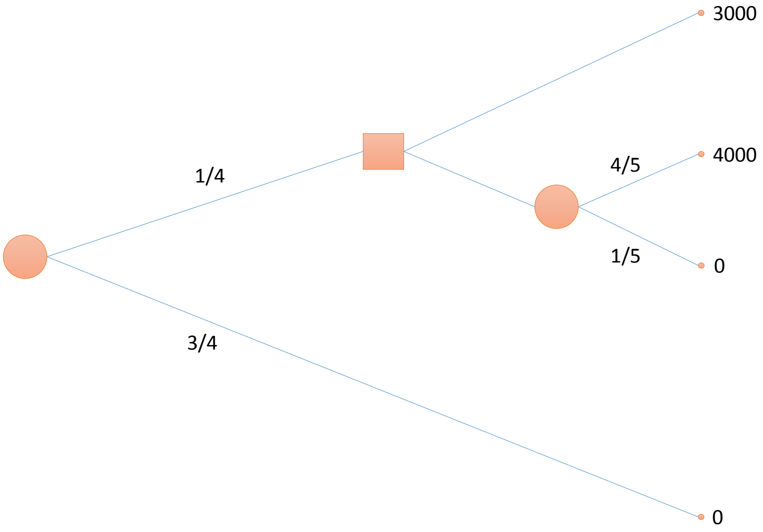

Fig. 2.23 Problem 10 as a decision tree (Sequential formulation) [136]

2.8 Trust

Trust is "the belief that someone is good and honest and will not harm you, or that something is safe and reliable" [176]. This definition holds true for Information Security and incorporates

Table 2.4 Summary of common set of trust characteristics [73].

Characteristic	Example
Trust is context aware	Entity A may trust B to download files but does not trust B to perform routing
Trust can be measured	Entity A has more trust in entity B than A's trust in C
Trust changes with time	The amount A trusts B may increase or decrease as interactions happen
Trust is socially aware	Entity A must trust entity C because C was presented by entity B, and A already trusts B
Trust may be directional	Entity A may trust B, but B may not trust A

both the user, "someone", as well as the physical and technological solutions, "something". Since 2010, there have been several influential papers that highlight the role of trust in users [182]. An important finding is the use of "trust to mitigate risk" [182] building on work conducted by the European Union in their multi-disciplinary, several-year project in Online Trust (iTrust)[1] [133] "documenting the many ways that trust can be created and broken" [199]. In addition, frameworks have been developed that aim to analyse the ways in which trust is built and maintained within computer applications [193].

Trustworthiness defines the "competence of an entity to act dependably, securely and reliably" [104]. It is not dissimilar from the concept of dependability which refers to the "ability to deliver a service that can be justifiably trusted" [23]. Trust and trustworthiness are not static and may change based on whether they adhere to specific criteria. The understanding of this has led to many organisations requiring management of "trust relations" [104]. This is highly relevant in dynamic environments where the understanding of disposition [67] (the level of tolerable uncertainty within an organisation), affects the dependency.

The following figure details the common set of trust characteristics [73]:

2.8.1 Trust in Information Security

Trust is an integral part of any **IS** policy and organisational structure. Organisations trust that employees will behave appropriately and in a trustworthy manner (often dictated by contractual obligations and policy), and employees expect that companies will act in a trustworthy manner ensuring their general well-being (securing personal details, honouring work contracts). The all-encompassing nature of trust within organisations can be visualised as:

With respect to trust, Albuquerque et al [74] argue that "when it comes to security, trust is zero or one", you either trust completely or not at all. Trust can usually be acquired

[1] Workshop papers available at: http://www.informatik.uni-trier.de/~wley/db/conf/itrust/itrust2006.html.

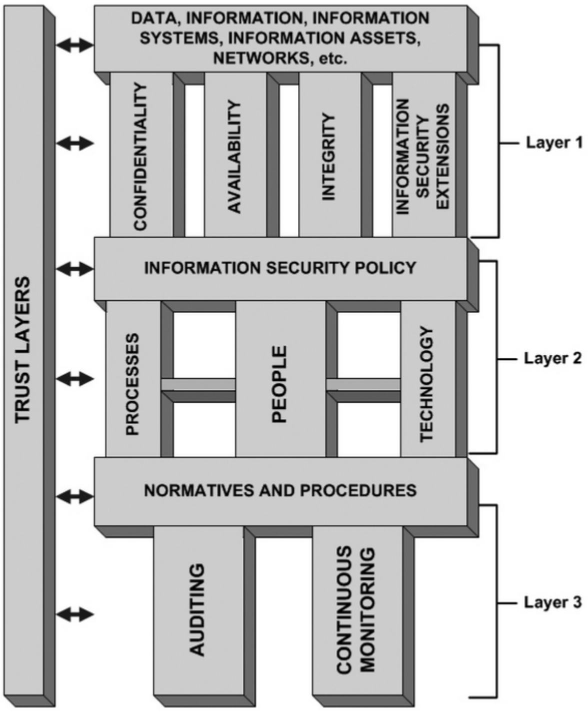
DATA, INFORMATION, INFORMATION SYSTEMS, INFORMATION ASSETS, NETWORKS, etc.
CONFIDENTIALITY
AVAILABILITY
INTEGRITY
INFORMATION SECURITY EXTENSIONS
Layer 1
INFORMATION SECURITY POLICY
PROCESSES
PEOPLE
TECHNOLOGY
Layer 2
NORMATIVES AND PROCEDURES
AUDITING
CONTINUOUS MONITORING
Layer 3
TRUST LAYERS

Fig. 2.24 Layered Trust Information Security Architecture [74]

through empirical observation, "by formal proof of the systems and the mechanisms involved and other techniques" [145]. Once all expectations are met, you may trust. The problem with this belief is that trust is "an expectation", "it is a probability that things will work and keep working as they are supposed to" [74]. Failing to do so results in the system being untrustworthy through its unpredictability. Further complications arise from the unpredictability of users in a number of scenarios stemming either from a reasoned decision to conform or defy policy, or involuntary non-compliance promoted through issues such as cognitive fatigue. These issues create an environment where trust is not necessarily binary and in fact is related to a number of components often out of the control of policy maker.

Security relies heavily on processes such as access control, authentication, non-repudiation. In reality, however, information practices should follow "trust, but verify" [60]. The layered approach in 2.24 allows one to see that "trust is connected as part of the security architecture" and may be addressed when required in "all layers and in all components that are part of the architecture" [74].

The IEEE Cybersecurity Initiative [204] placed trust as a key factor within their list of the top 10 security design flaws. They state that "data sent to an information system by untrusted clients or channels should be assumed to be compromised until proven otherwise". Ultimately, if the data cannot be verified it is inherently not trusted.

Chapter 3

Methodology for Decision Making

This chapter proposes a methodology for aiding the policy decision making process using data from empirical investigations, and continued interaction with the stakeholder. Specifically, we identify behaviours through CISO consultation and user interviews (our pilot study and the IRIDIUM study), alerting us to important behaviours and practices within our environment. This process forms the rationale for subsequent bespoke investigations (see chapters 4, 5, & 6) that aim to specifically identify how users form their decisions with respect to that behaviour, and how we can manipulate this process through behaviour interventions (predominantly nudges). This approach enables a more accurate understanding of the problem space as our articulation of critical components is improved through repeat investigations that are relevant to both our users and the problem space.

The methodology highlights the necessity to understand user behaviours, and to identify common technological processes that may pose a security risk within a mobile working environment. This chapter examines current trends in BYOD and consumerisation with respect to security, data mobility, and access. Within our assessment of the problem space, we adopt a case study that initially focuses on a university environment but subsequently expands on this to include an abstract overview of security practices in other domains. Successive chapters follow the steps highlighted within the methodology and help to improve our understanding of the problem space by investigating specific methods aimed towards populating and analysing these data sets through experimentation. Ultimately, this approach enables the policy maker to produce a more robust modelling environment to test future policies, and ultimately aids the policy decision making process.

Section 3.1 introduces our study and outlines our direction. Section 3.2 outlines a methodology for investigating the omissions in our understanding through repeat, small-scale, targeted studies and interaction with the CISO. Section 3.3 details a case study and uses a pilot study as part of the methodology providing the first iterative step of the process.

Section 3.4 discusses the results obtained within the pilot study. Section 3.4.1 defines the problem space and highlights the challenges we face with respect to policy design. Section 3.5 discusses the results with respect to future investigation. In section 3.6 we outline the remainder of the thesis and introduce the following chapters with respect to the key behaviours outlined within our pilot study and CISO interaction.

The IRIDIUM Project

Within the IRIDIUM project I was appointed as a researcher to fore mostly assess current policy documentation within the university. This required a university wide assessment of central policy along with investigation into faculty level, more specific policy related to the governance of bespoke practices. The necessity for such disparity and inclusion of additional bespoke policy is often tied to affiliation with industry and collaborative initiatives between the university and external parties that often required the dissemination and sharing of sensitive data.

Following this review I conducted a series of stakeholder interviews (33 in total of varying seniority) capturing common practices within specific departments (those defined as high value) to understand vulnerabilities and threats and whether current policies were indeed being adhered to. This process required assessment of current ’shadow security’ (a phenomenon where users often adopt their own security practices that are not necessarily in direct compliance of policy, but are often deemed ’sufficient enough’ to work securely whilst remaining productive), direct non-compliance and an assessment of where the stakeholder felt they were vulnerable. These observations and concerns were captured and reported back in order to understand our current threat landscape and determine where new policy could be designed to improve our security and minimize data loss.

3.1 Introduction

Universities have increasingly become targets for attackers who wish to obtain ‘desirable research’ and ‘student’s personal and financial details’, as documented in a January 2015 report on the Queen Mary University cyber-attack [106]. As such, it is necessary for academic institutions to evaluate their environment from an Information Security perspective and assess the effectiveness of their data management policies with respect to university IPR.

The objective of this chapter is to introduce and design a methodology for aiding the information security decision-making process. This involves determining stakeholder requirements and investigating user behaviours. In doing so, bespoke tools and studies are designed and subsequently conducted in order to understand current user behaviours and practices, and determine where potential threats may lie. We wish to improve our understanding of our problem space, and ultimately to empower the CISO to design more effective policies that are tailored to both the user and the organisation in an effort to maximize compliance and productivity.

Central to this methodology, which we call Model-Based Information Security Decision-Making methodology (or Model-Based methodology for short) is the development of a testing framework (discussed in chapter 7) to aid in the design process and subsequent testing pertaining to the possible consequences of new policy decisions. Special emphasis is placed on the understanding and representation of how employees behave within the current environment as we believe that users play a critical role within security. The presented Model-Based methodology is iterative, increasingly improving the expressiveness of human behaviour and providing a more precise quantification of environmental variables. Specifically, we investigate a set of activities and behaviours identified from out pilot study and CISO interaction (see chapters 4, 5, 6, 7) in order to test our approach and validate our iterative, small-scale, user-centric method. We believe this approach to be scalable, and that this work provides an assessment of its viability with respect to the problem definition.

We apply this methodology to a case study of practices within a university environment and aim to improve our understanding of data mobility and user security behaviours under the umbrella of Data Loss Prevention (DLP). To do this, we work closely with the IT team that governs the university's IT infrastructure. Within this initial chapter we focus on the use of USB memory storage and BYOD and discuss these mediums with respect to collaborative work. This is a conscious decision based on the needs of our partner, and also the popularity of the trend within the institution and other working environments (obtained through our pilot study and literature in chapter 2). Focusing on the devices most common to our environment adds feasibility to the study with respect to scope, as well as enabling finer granularity regarding their use, ultimately aiding our understanding of data transfer.

Throughout this process we maintain continual interaction with the CISO and support staff responsible for the policy decision making process. This is beneficial as it enables us to specifically target the requirements of the CISO and relate these to activities and behaviours witnessed through participant interviews. This approach benefits our methodology as we are able to base our analysis and subsequent recommendations on 'real-world' data that is obtained through focused investigation techniques. Specifically, we use the following empirically derived data sets:

- **The IRIDIUM Project:** The project [186] was funded by the Joint Information Systems Committee (JISC) and represented a collaboration between the University Research Office, the Digital Institute, the University Library, Information Systems &

Services and MEDEV, School of Medical Sciences Education Development. The aim of iridium is to make recommendations for a complete holistic plan and infrastructure for research data management in the University, making data generated by research at the University both available and discoverable with effective curation throughout the full data lifecycle in consultation with the researchers who produce it. It should be noted that:

  - central to the project is the formulation policy and appropriate support for its implementation.
  - information technology will likely play an important role in supporting compliance with newly developed policies, with a significant part of the project. investigating the assessment of this.

The principal project outputs are:

  - an institutional research data management policy (as required by the funding councils).
  - a costed business case for a sustainable institution-wide research data management infrastructure to support that policy.

- **CISO Interaction:** We maintain regular meetings with the CISO and support staff to document their main concerns. We subsequently encapsulate these via a process of problem formulation and identify the behaviours and practices we need to further investigate these areas through bespoke user studies.

- **Pilot Study:** The pilot study is part of our iterative methodology and aims to identify user behaviours and practices identified from the above CISO interaction. We include this process as we believe that understanding how users form their decisions and ultimately act upon these decisions is critical in understanding how to improve security and formulate appropriate policy.

3.1.1 Problem Domain

During initial meetings, the CISO has expressed concerns towards a lack of understanding of where vulnerabilities and potential exploits lie with respect to the protection of university IPR. We examine the policies currently in place and interview users to determine areas of further investigation to address this omission. Our aim is to provide a better understanding of the IPR within the university, assessing value (from a user's and organisation's perspective) user awareness and common practices. From this, we aim to identify areas for further investigation with the goal of developing model-based tools to aid in the policy decision-making process.

3.2 Methodology

The methodology aims to:

- Encapsulate the requirements of the CISO and identify high value assets/targets.
- Use this to formulate an initial investigation strategy.
- Conduct a Pilot Study based on this strategy to articulate our understanding of security practices.
- Extract data obtained from the pilot study to generate specific subsequent investigations.
- Understand the Human Decision Points (HDPs), data value and user practices.
- Begin to model the environment.
- Extend the modelling formalism if required.
- Determine which areas of our model are under/not represented.
- Repeat the investigation strategy to complete the omitted data sets.
- Iteratively refine the model to aid the decision-making process.

With reference to figure 3.1 it is necessary to discuss the individual steps.

Step 1 defines the initial contact with the CISO (ISS) and aims to formalise the problem. It is essential that the needs of the CISO and stakeholders are sufficiently understood and addressed so that we may aid the decision-making process. We determine ‘end goals’ (a collation of interests, preferences and requirements) that will be the basis of the work we conduct.

Step 2 conceptualises Step 1 and promotes the necessity for a case study. From our ‘end goals’ we understand some of the requirements that our Model-Based approach is required to address. We visualise additional steps towards fulfilling these goals as sub-processes to our model and begin to formulate data collection strategies.

Step 3 represents the pilot study and data collection phase. Within this process, we determine which variables we need, highlight which users we are going to investigate, and

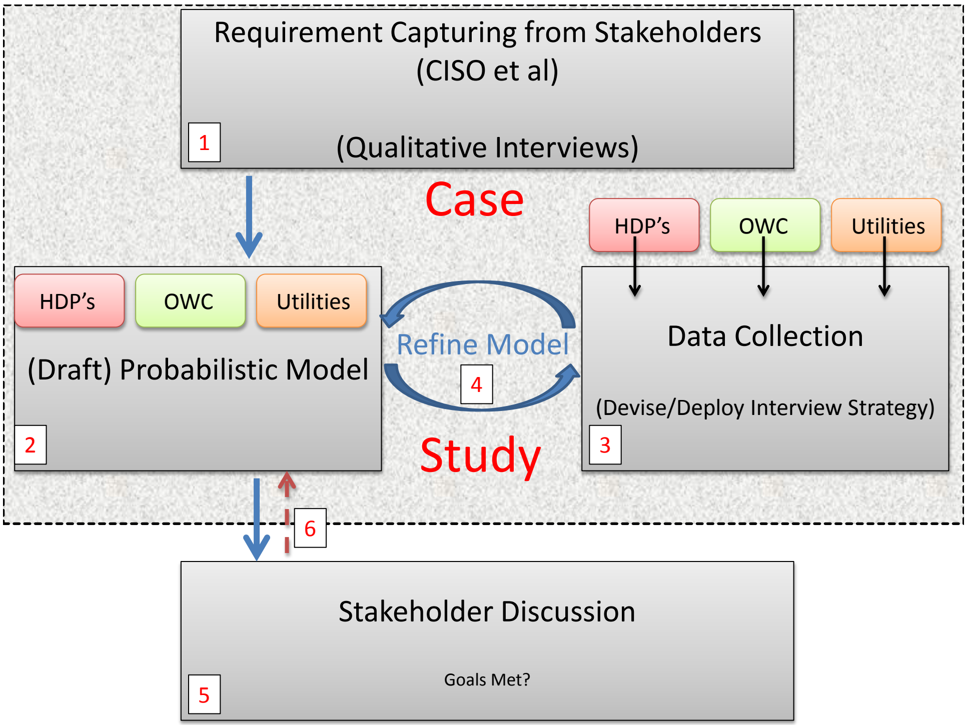

Fig. 3.1 Case Study in Relation to our Model-Based Methodology

formulate which questions we will ask. This is based on uncertainties and omissions from our understanding of the environment obtained through Step 1. As visible, these are an assessment of the Human Decision Points, and the user's Opportunity. Willingness and Capability to complete a task [89].

Step 4 denotes the iterative nature of our methodology and highlights the relationship between our modelling and data collection phases. Data collected from Step 3 is used to populate our model and determine where focused investigations pertaining to specific behaviours are required. This process may require further extension of the modelling formalism (see chapter 7). When completed, it is then necessary to determine whether to revisit Step 3 (obtain specific data regarding identified behaviours - chapters 4, 5, & 6), or move to Step 5. With 1 iteration, we argue that there is insufficient understanding to move to step 5 as it is not possible to accurately model the environment (a feature highlighted through the myriad of phenomena discussed in chapter 2). As such, we revisit Step 3, typically including a greater number of participants from a wider selection of appropriate user profiles in order to further our understanding of a specified behaviour or practice.

Step 5 illustrates a functional model that satisfies the goals highlighted from CISO interaction in Step 1. At this stage, our model can either be used to aid in policy decision making, or must be rejected forcing us to transition back to Step 6.

Step 6 is necessary in order to further our belief in an iterative strategy. As we discuss in chapter 2, threat environments and user behaviours change temporally as new technology influences our lives. This process enables a continued iterative process that should aid in the discovery of these trends and enable an adequate test-bed for future policy formulation.

3.3 Pilot Study in the School of Computing Science

The pilot study was conducted to expand our understanding of the problem space, and as a follow-up to the areas highlighted in our CISO interaction. This process represents our inclusion of users into the decision-making process, as we believe that understanding user behaviours plays a critical role in improving user compliance, and ultimately aids the successfulness of any policy implementation. The pilot study forms an important part of our methodology as it provides a quick, highly targeted method of obtaining data pertaining to user behaviours and practices within a given problem space.

Our study was piloted within the Computing Science Department using a stratified sampling method to reduce sampling bias. This involved placing participants into sub-populations that were homogeneous with respect to their job role and title. Although the sample was small, it was important to separate members in such a fashion as there are specific risks and attack vectors that are commonly associated with each population (risks that pertain to their given daily activities).

We believe that overall, participants operating in this environment will have a higher base knowledge of information security due to the nature of their research background and the working environment itself. We believe that this increased knowledge is beneficial as it will highlight more behaviours and practices that less technically skilled users may lack. This approach may also represent a 'worst-case' scenario as we would expect more informed and technically able users to adopt more secure computing behaviours. We this process documented in Fig 3.1, by interviewing people of various positions within the department categorized as either Lecturers, PhD Students, Technical Support Staff or Clerical. In order to justify this selection it is necessary to briefly describe each category in terms of policy.

Lecturers have access to student records, exam scripts, coursework and crucially their own research which often has ties with outside industry and other academic institutions as part of collaborative projects. This often has a value associated with it as research grants have tangible budgets that can often run into several million pound figures. There is also an expectation for return on investment from the research body and subsequent publication. Lecturers also mentor doctoral students further increasing their involvement into research and potentially sensitive information that may have contractual obligations or non-disclosure agreements.

PhD students are seen with a similar importance to Lecturers as research typically produces new IP which is funded by various institutions. This category is also responsible for marking students' work through demonstrating (a process where PhD students assist lecturers in practical classes). The way in which this work is stored and subsequently marked could have a significant risk in terms of data protection, identity and confidentiality.

Technical support staff are also included in the study as they are responsible for the maintenance of existing systems, and data backup of the department's servers. Part of their duties also include the procurement of software and hardware as well as its installation requiring access to many areas with multiple levels of physical access-based security.

Clerical staff typically comprise of members from the Human Resources department. Typical responsibilities of staff within this department include the management of student and staff finances (including personal details), purchasing of items using the Department's credit card, personal hardware requests and the restocking of office supplies, and the collation and storage of student's coursework and examination results.

The use of user profiling is beneficial as we enable the specific identification of practices and behaviours with respect user groups (a process made possible by an assessment of their seniority, access rights and contractual roles). We assume that lecturers and PhD students within this Department have a high preference for security and are making more security conscious decisions with respect to their data. We assume that they are taking adequate precautions to protect their data based on the risks they perceive. Countering this, we believe that clerical staff (specifically within the Human Resources Department) are less likely to have this technical awareness and as a result will have a lower knowledge of computer based security. This places them at a higher risk, but ultimately they have lower value data.

3.3.1 Investigation Strategy

In line with the methodology, it is important to produce a quick, small-scale assessment of our environment as the start of the iterative process (a sub-process of our talks with the CISO). To do this, we have produced a semi-structured interview as defined in Step 3 of Fig 3.1. With reference to Wengraf [250], we have produced a framework for interviewing individuals which aims to investigate users, their devices, and the data they handle. Barribal [154] discusses how to improve the credibility and reliability of data obtained from semi-structured interviews and this interview strategy has been conducted with respect to this. The questions have been specifically designed to investigate areas of concern highlighted from our CISO, and current trends and behaviours noted within our literature survey.

The interview questions follow a logical progression where the objective is to provide both qualitative and quantitative results as an assessment of an individual's current practices. This involves asking questions in a specific order so that answers are not inferred and so that users do not answer in a manner in which they feel obliged to. Both data types are important and can be utilized in different manners when assessing the problem space.

More specifically the questions are designed to assess current practices and behaviours within the target environment. The answers to these questions allow for the identification of security vulnerabilities that will later be investigated through our bespoke experiments that are aimed at assessing the validity of nudges as an intervention mechanism. We are not attempting to formally test a hypothesis, more encapsulate the behaviours and capture user practices and preferences.

The questions we asked are thus:

1. Which words and/or phrases do you associate with Data Security?
2. What security measures are in place where you work?
3. Please rank the above in order of importance.
4. What devices do you use to store university data?
5. What devices do you use to store personal data?
6. What devices do you use to transfer university data between locations?
7. What devices do you use to transfer personal data between locations?
8. How often do you transfer data? From where to where?
9. How do you carry your device during transit?
10. What types of data do you transfer?
11. What measures do you use to secure the data?
    (a) Have you considered other measures you do not use?
12. Briefly describe the implications if you lost university data.
    (a) Would you report it?
13. Briefly describe the implications if you lost your personal data.

    (a) Would you report it?

14. What can others do if they obtained university data you are responsible for?

15. What can others do if they obtained your personal data?

16. Do you ever transfer data without backups?

17. Do you create data remotely?

    (a) How long between backups?

18. What device(s) do you use to work remotely (off-site)?

19. Do you connect to non-university managed networks with this device?

20. Do people other than yourself have access to the device?

21. Do you submit error reports if they occur whilst on campus?

22. Do you submit error reports if they occur whilst off campus with the same device?

The questions follow a specific progression in an effort to gain honest responses where the individual does not question their own practices and thoughts. Open ended questions allow the user to comment freely on the subject at hand whilst requiring little prompting from the interviewer. As HDP's and OWC are critical to our understanding, this method allows for a detailed assessment of participant. From these questions, we wish to highlight and investigate specific behaviours and practices which are common throughout our participants in order to develop bespoke investigation strategies.

Initial questions (Q1-3) aim to assess the security of the given locale across a broad spectrum from the user's perspective. This enables assumptions to be made about a participant's awareness to current policies, and technical expertise in that domain or specific practice. Such questions aim to investigate whether the user is aware of their surroundings, and understand their preferences with respect to their perceived impact on security.

Questions 4-11 begin an in-depth analysis of data security and mobility with respect to data type (organisation vs personal). We gain insight into how users value data, and how this is reflected in their choice of device. Importantly, we also learn the types of data the individual transfers and the locations this entails, enabling the CISO to begin to quantify the risk the university faces upon a data leakage event. These questions can highlight whether users are sufficiently able (OWC) to select the correct device to use in a given scenario highlighting whether current policies, training and practice are adhered to or sufficiently understood for employees to operate securely.

Questions 12-15 aim to investigate whether users are aware of the impact that data leakage can have. This provides an assessment of personal and non-personal data and the role this has in the way users protect or select their device.

Questions 16-18 examines how users mitigate risk when transferring data and working remotely. We question which devices they choose in order to assess the penetration of BYOD in the workplace (rationale for chapter 4) and subsequently determine whether this medium is the preferred choice of access and storage with respect to university data.

Questions 19-21 further examines the role of BYOD and investigates multiple user devices. We ask how users behave when outside of the university network with respect to unknown and non-university networks (detailed further in chapter 4) and data transfer/access. From this we begin to understand how access and location are related with respect to specific practices and can better articulate our threat environment. Assessing multiple users furthers this as we investigate how sharing devices with non-university staff may impact data security. We ask specifically whether the device is used to access the internet outside of the university network to understand whether this is a legitimate attack vector (rationale for chapter 5). Furthermore, we examine whether users report errors if they occur in order to minimize the impact of data loss or breaches (rationale for chapter 5).

3.4 Results

The pilot study had n=35 participants and included the following breakdown: 3 Lecturers, 20 clerical staff, 10 PhD students, and 2 technical support staff. These numbers reflect the expertise of staff (we have a higher percentage of clerical staff as they are responsible for personal records and the transfer of financial documents) and the expertise of staff with the belief that less technical staff pose a higher risk to security. We include lecturers and PhD students for their involvement with current research, which may have industry data or sensitive material, as well as lecturers who may also have far more sensitive information. This breakdown is in line with recommendations from the CISO.

There were a number of significant results obtained from the pilot study that have helped to formulate subsequent investigations. Most surprising is the contrasting ways in which people of different profiles (Lecturers, PhD, Technical, HR) see and utilise security methods. This was witnessed in chapter 2 in several studies where seniority and a higher level of access often resulted in a far higher risk to the organisation (attackers will often expressly target such people). As such, we discovered that less technically skilled staff (clerical staff), were less capable of identifying which devices and methods of data transfer should be adopted when handling sensitive data (a reason why we chose to investigate non-experts in chapters 4, 5). Furthermore, there is a clear divide between physically-based accessed such as doors and swipe cards, and computer based solutions such as encryption and User Managed Access with respect to security mechanisms (a further example of the technical divide). This is an important observation as it suggests that current policies are ineffective with respect to security, or that staff are receiving insufficient training.

To further analyse our results, we will discuss them in order:

Question 1. 80% of the sample population mentioned IT based solutions such as access management (user authentication via user name password login). It is important to stress that this figure represented 100% of Lecturers, PhD students, or Technical staff, whilst only 65% of clerical staff. This suggests that there is a difference in technical awareness and ability amongst different employees. Of the remaining 35% of clerical staff, there was confusion as to what they actually considered security with many participants requiring significantly longer than their counterparts to answer.

Question 2. We witness an almost identical distribution with respect to IT-based solutions. The only difference was an increase to 68% (1 participant) of clerical staff who mentioned such mechanisms in the workplace. This is important and suggests that there is a disconnect amongst users. More specifically, it suggests that these people had either not actively thought about the security around them or were unaware due to a lack of sufficient training or prominence in the workplace.

Further analysis revealed that 73% described password based access as the most common and important form of security followed by 50% of participants mentioning physical access. This is an important find as whilst password based security may seem obvious, physical access as a form of security is often overlooked. The physical access response was most common from technical support staff, perhaps due to the nature of their profession; they have several locations throughout the department where expensive hardware and software are stored, and they also carry master keys that enable access throughout the department.

Perhaps most alarming is that 15% of the population (all clerical) were unable to identify any security measures at all.

With respect to question 3, user authentication followed by physical access to the building via smartcard were deemed to be the most important.

Question 4. From our sample, 100% of people used personal storage space on university provided computers (these are virtual drives located on university servers), and 75% of them used an additional personal Laptops or other storage mediums. As all participants used computers provided by the university, it is important to further analyse the use of personal devices. Of these, 24 used USB sticks and 3 used external hard drives; none of which were highlighted as encrypted or password protected. 12 participants highlighted they used tablets and mobile phones to transport data. This is an important observation and poses a significant security risk. We highlight this finding as significant and is our rationale for chapter 4.

Question 5. Personal data was less commonly stored in the workplace than university data with only 20 (57%) participants engaging in the activity. Of the respondents, all lecturers, PhD, and technical staff stored personal data but only 5 (25%) of clerical staff did. This is perhaps due to the working environment where clerical staff share their work space in an open plan environment (and perhaps do not wish to be shoulder-surfed) compared to the other participants who often have their own private spaces. The results could also depict an uncertainty in what constitutes as personal data. Of the devices chosen, 15 users chose USB sticks, personal laptops and mobile devices (combined response); and the remaining 5 used university or private email accounts via attachments.

Question 6. Of the population, 25 (71%) users stated that they moved electronic university data (others noted that they did not or that they simply logged into another computer and used their private server space). Of the 25, all lecturers, PhD students and technical staff were included as well as a further 10 clerical staff. The most common methods were the use of an unencrypted USB stick (15), personal laptop (15), mobile device (6), hard-drive (2) and email (12).

Question 7. The predominant method of moving personal data was by personal laptop or mobile device. Of the 20 who answered from question 5, 18 used either a personal laptop or mobile device with 2 using USB sticks. This indicates that there a great deal of mobile devices (including personal laptops) within the working environment.

Question 8. Daily (25) was the most common frequency of data transfer as users were either required to perform it to complete tasks, or would work from other locations. The most common location was to and from home (18) whilst collaborative meetings (3) and public locations (4) were also featured.

Question 9. 20 participants carry the media on their person (around neck, or in a pocket) if small enough (i.e. smart device, usb stick) and 5 store it in baggage or luggage elsewhere (typically laptops).

Question 10. The most common response was the transfer of ‘work’ data (21) which often included a mixture of personal data and university owned IPR (experimental results, papers etc.). This number featured all lecturers, PhD and technical staff, and 6 clerical staff. Of the 6 clerical staff, all transferred university data.

Question 11. Of the 21 respondents from question 10, 12 highlighted that authentication on their personal laptops was their method of securing the data. The remaining participants did not provide an answer although they may still have used the previous method.

Part (a). 12 participants highlighted specific hardware such as encrypted USB devices or external hard-drives. An additional 4 participants mentioned software based encryption that could then be placed on unencrypted media. The remaining 5 did not mention any other methods.

Question 12. All 35 participants answered this question. Of these, 12 indicated specifically that it may cause financial or reputation damage towards the university. 5 participants said that they would personally be affected as it would interrupt their work. The remaining 23 participants were unsure of the implications of losing university data but did mention that it would be damaging for both themselves and the university. With respect to (a), 8 participants (of which 2 were lecturers, 2 were PhD students, 3 were technical staff and 1 were clerical staff) were able to indicate that they should inform the support team or someone else at the university.

Question 13. All 35 participants answered this question. The most common response (12) was that time would be lost with respect to work (it is unclear whether there was an uncertainty as to who owned their work). 2 participants highlighted that personally identifiable information would be lost (photographs, bank details) and this may have future implications. With respect to (a), participants declared that their report would be context dependant (8 - if it was important, would not cause embarrassment) and if there was a genuine chance of recovery (2 - accidentally deleted emails or files on remote servers).

Question 14. Of the 25 respondents, the majority of answers were context specific and related to the role of the individual. All lecturers stated that either publications could be lost or stolen (2), the likelihood of future grant applications could be affected (1), and that other sensitive data could cause significant reputational and financial losses for the university (3 - specifically so in the medical department and the inclusion of corporate sponsors). All PhD students responded by reiterating publication loss (5), or mentioned facing penalties for the loss (6 - perhaps removed from the programme). Technical staff responded by mentioning the theft of volume licensing keys for software (1), gaining unauthorised access to user databases via stolen credentials (1), and further impersonation and associated access (2).

Question 15. Of the 25 respondents, 13 replied with concerns regarding unauthorised access to their personal files and the subsequent use of this (2 mentioned blackmail). 4 responded more specifically with unauthorised access to accounts as they stored access credentials in an unencrypted, non-passworded form on their device (3 of these were clerical staff). The remainder provided no specific response in relation to a particular issue, but eluded to it causing an inconvenience and were worried with respect to knowing that their data could be in the hands of somebody else.

Question 16. Of the 25 respondents who transfer data, 4 (clerical staff) noted that they did not have a back up of their data. Upon questioning, it was apparent that their device was the only copy of the document in question. It is unclear as to whether they were aware of the implications of doing so.

Question 17. Of 35 participant answers, 19 answered ‘yes’. All lecturers and PhD students said that they created data remotely either due to working from home or collaborating with partners elsewhere. 1 technical member of staff said that they also worked from home on certain days. 4 clerical staff also noted to working some evenings and weekends in order to complete their duties. With respect to (a), none of the respondents mentioned that they produced specific backups on dedicated devices but instead detailed that they had other consumerisation and BYOD based approaches. 5 PhD students mentioned that they used Dropbox for their backup and synchronisation of work whilst the remaining respondents used the storage medium on their device (be it personal or university owned).

Question 18. Of the 19 respondents, the most commonly used device was a personal laptop (19), whilst personal tablets featured (4).

Question 19. All participants connect to their home Wi-Fi network. An additional 12 of these also commented that they connected to public Wi-Fi’s with the device (Cafe’s 10 - public libraries 4 - and transport 6).

Question 20. Of the 19 respondents, 8 replied by stating that other users had access to the device (family members - 6, friends - 1 and partners - 6).

Question 21. Of the 35 respondents, 7 had stated that they suffered an error that required reporting. Of the 7, 3 participants engaged the official procedure of contacting support with respect to the issue. The remaining 4 either dismissed the error and continued (using alternative methods), or sought assistance from other colleagues without the issue being officially recorded. Of the 7 incidents, 3 were access issues related to their work allocated computer or email account whilst the remaining 4 were software failures, crashes, or malware instances.

Question 22. Of the 35 respondents, none had reported an issue that transpired off-campus even if it were relating to a university owned device. Whilst 3 participants acknowledged that an event had occurred, they did not wish to report it for fear of disciplinary action as they were unclear as to the rules governing device usage in such environments.

3.4.1 Challenges in the Policy Decision-Making Process

This section highlights the omissions in data sets and identifies the unknown behaviours within the study environment. We aim to define these unknown variables through our pilot study and determine their impact on our methodology, its development, and more specifically how it impacts our data collection strategy.

Current Policies

From meetings with the CISO and support staff, we know that it is current practice for all new and existing machines to be registered for use on the campus network (wireless devices are able to connect freely provided the user has an active campus account, i.e. they are a current member of staff or attending student). This network, however, is separate from the central network and requires additional verification through a VPN in order to be able to access remote storage and printing etc. This registration process records the devices Media Access Control (MAC) address, brand and model, along with its location and registered owner. Access to this database via the IRIDIUM project [186] has provided us with a rich dataset containing 34549 entries at present. This has enabled a quick assessment of the concentration of static devices and at finer granularity, the number of mobile devices (specifically laptops and mobile phones). Of the previously mentioned 34549 computers on campus, 1932 belong to the Computing Science Department alone. Critically, 620 are described as mobile devices and are typically in the possession of lecturers and PhD students in the form of laptops, tablets and mobile phones. This highlights the mobility of the data and justifies the concerns of ISS. The dataset has also provided an approximation with respect to the spatial distribution of machines that can be cross-referenced with device type and department in order to determine where possible high risk may lie (i.e. some departments are likely to hold more sensitive data than others). In summation:

- It is possible to highlight areas of high concentration - the most popular departments, libraries and other open study domains. Areas such as these potentially pose a higher risk to the network purely through the increased number of interactions and likelihood of collaboration (sharing of data and media devices).

- We can visualise when the device was last seen by the network. If it is a campus controlled device (it receives updates from a central server) we can assess whether virus definitions are up to date and whether critical operating system patches have been installed. It also highlights mobility and the frequency with which the device leaves the network (it is possible, however, that the device may simply be turned off).

- It is possible to categorise devices based on their owner with the belief that higher levels of seniority grant higher levels of access, thus posing a more significant threat if successfully attacked.

Having access to such a rich dataset has also allowed for a rapid deployment of our pilot study (utilising findings highlighted in Chapter 2 and observations from the above). We have been able to identify areas of significance related to the CISO's concerns, as well as those areas of potential high-risk as denoted by the analysis of the dataset.

Issues With the Current Policy and Environment

To devise a new strategy and design more effective policies we must first identify common threats with respect to current practices and user behaviours and understand the relationships between these and our environmental variables. To do so, we must highlight these issues with example.

The first issue is the unknown number of non-registered, mobile/static machines within campus under ISS governance (these are typically devices that are personal or bought outside of ISS jurisdiction through project grants). Without this knowledge, it is almost impossible to understand the scope of the problem as we do not know the percentage of devices that are currently under management. This also presents yet another problem. Even if this value were to be known, it is highly unlikely that there is a uniform data value distribution (users may prefer to store more sensitive data on their own personal machines) and also it is very difficult to ascertain the level of security (patches, virus definitions etc) on the device, a problem noted in Morrow [166]. Factors such as the machine's user, their practices, area of study and access (a property of location) further add to this complexity.

Of particular interest is the issue of value and what exactly this metric reflects. Moreover, how it impacts user behaviours and their decision-making process. This typically falls into personal value (that which individuals place on their data) and the organisation's value (what the data is worth to the organisation - often a calculation of financial, legal, reputational and time based approximations). For this study we concentrate on the latter (using a rating system similar to the SANS Institute [201]) but understand the importance of the prior with respect to behavioural influence. Being aware of this phenomena is important with respect to our pilot study and user interviews. To exemplify the issue, consider the following:

- **A student's essay:**

  User Perspective: The student will most likely place a high value on their essay as it has a profound impact on their progression and success during studies. It takes time to complete (and ultimately lose this if the essay is lost) and is a personal reflection of their ability and in some cases beliefs.

Organisation's Perspective: From a university perspective there is little value placed on the asset. It is the student's responsibility to complete the work by a specified deadline (often detailed within university regulations) and ultimately there are no repercussions for the university if the student fails to do so (one could argue that it may affect the overall pass-rate).

- **Data Loss from a Corporate Contract:**

  User Perspective: Whilst the user may be responsible contractually for the security of the data in some circumstances (failure to do so may result in the loss of their job, or legal action if it is determined to have been malicious), there are numerous cases with which the user may be exempt from prosecution or fault. In this example, the user is not personally attached to the data and therefore they may regard it with having a lower value. They may feel that the protection adopted by the organisation is proportional to the value they place on it and ultimately it is not their responsibility.

  Organisation's Perspective: Contracts within organisations can be highly sensitive especially when dealing with classified information (medical research, defence etc). Such contracts are often influenced and governed explicitly under law and contractual obligations with heavy financial and reputational losses related to breaches of contract. From this perspective, the value placed from the organisation will most likely be several orders of magnitude higher than that of the user even if the user is aware that they need to protect it.

Without being able to quantify this value, it is difficult to assess the appropriateness of current security practices as asset value should be mirrored by the level of security. In a non-uniform environment (uneven data value distribution) this will require a detailed approximation of asset values across a broad area. For this, we will need to determine how the university categorizes data types, determine the data type users transfer as well as the locations and devices involved during the process. It is then possible to approximate the data value based on the findings detailed in chapter 2. Doing so ultimately aids the model building process as we are able to express these values within our variables and bounds to provide visual assistance to decision maker.

In an effort to address the above issue and provide basic protection for users who request it, ISS have initially responded to data security concerns by piloting an encrypted USB stick programme granting a 'secret' (an industry standard level of encryption suitable for secret documentation [10]) level of security (encryption based). The introduction, however, does not follow a controlled implementation and has been conducted irrespective of any statistical data to backup the claim of increased security gains. This programme is small-scale and has had less than 50 adopted users at the time of questioning. Importantly, there are no tools in place to monitor the effectiveness of this strategy or analyse its impact.

3.4.2 Human Behaviour

Human behaviour with respect to policy design is critical to our methodology. We wish to develop a modelling environment that depicts specifically targeted user actions to be tested with a prototype policy. The actions we investigate are chosen from our pilot study participant answers and are further examined in separate chapters to fully highlight their intricacies and understand the behaviours that govern them. This requires bespoke experimental design and provides a scalable framework with which to test in order to populate our model variables. Specifically, we investigate how to change these behaviours (a process of understanding user preferences through MINDSPACE), and how users form decisions with reference to security decisions. From this increased expressiveness of user behaviour, we are able to build more accurate models that reflect the target environment and enable to the policy decision-maker to design, implement and test new policies.

With any security decision, it is important to note that there are often two seemingly conflicting objectives. Firstly, users do not want to be inconvenienced (lose productivity) by ‘overly aggressive’ encryption or security policies, and secondly that organisations wish to protect their data and assets (but need to balance risk aversion). To remedy this, we need an holistic approach that involves users’ behaviours in the policy decision-making process.

If poorly implemented, an increase in security often hampers productivity through additional user authentication procedures. These extra processes impact the time it takes to complete a task either through additional verification processes, an increase in the time it takes to be authenticated, or a combination of the two. Introducing such policies modify a user’s typical behaviour presenting a somewhat unnatural activity pattern from perceived habitual processes. Deviations from expected procedures create additional HDPs that in turn have an associated thought processing time and ultimately affect the decision outcome (we are often causing the user to turn from system 1 to system 2). Moreover, from chapter 2 we are aware that users have a finite cognitive budget, and once exhausted, the willingness to comply with procedures greatly diminishes. Increasing the number of additional steps a user must navigate in order to complete a task consumes this budget. In chapter 4 we aim to combat this problem by employing nudges to improve the security decisions that users make without inconveniencing them with additional processes or authentication measures.

User profiling plays an important role within our pilot study and participant selection requires grouping prospective participants in relation to common practices and procedures they conduct (often a result of occupational contract). This will typically relate to the data that they handle and the department they belong to with technical knowledge and security awareness being a crucial factor.

3.4.3 USB Sticks and the Role of BYOD

The methods in which users utilize hardware and software (specifically USB Sticks and mobile devices) in order to secure and transfer their data is a complex process. From a Human Influenced Task Orientated Process (HITOP) [89] perspective we are able to represent these processes using a subset of variables which we obtain through experimentation and subsequently implement into our modelling environment.

Within the field of HCI (detailed in chapter 2) it is important to understand the complexities and relationships that are present between user and device, and the method in which policy can influence this. As discussed in 'Human-Computer Interaction' [81], HCI is often regarded as the junction between Computer Science, Behavioural Science, and Design. How we represent and populate such data sets is critical to the validity of our work and the process in which we design successful policy. This is useful in our example when we try to understand how users utilize their chosen hardware for data migration.

Opportunity, Willingness and Capability (OWC) from a HITOP perspective is necessary in the understanding of the present environment. It is an assessment of the user's ability to complete a given task [89]. We use an assessment of this in our Data Collection Strategy.

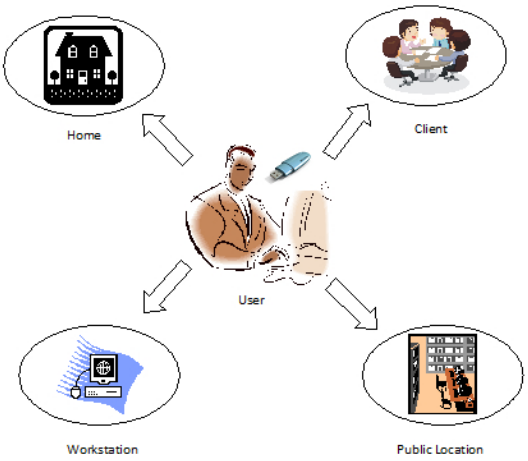

Fig. 3.2 Access locations and Data Mobility

Fig. 3.2 details the complex nature of the study environment with respect to mobility. There are a multitude of access locations each with an associated risk. The manner in which the USB stick or mobile device is transported and subsequently accessed has a profound impact upon data security. For instance, public locations are heavily susceptible to malicious behaviour as there is no record of who has previously used it, and what they were doing (has the network or machine been compromised?). There is also more importantly no control over how that environment is maintained. From an organisational standpoint, the machine is not trustworthy as we cannot trust the environment nor verify the integrity of the device. Governing access from devices outside of the university's control is a difficult issue.

The overarching problem with a standard unsecured USB stick from a DLP perspective is paradoxically its main advantage. Once data has been placed on the USB stick it can be accessed on any device with a USB interface provided it is unencrypted. This provides quick, simple access to anyone who has acquired the USB stick and presents a significant opportunity for data loss and misuse in the hands of an attacker. This problem is exacerbated for unsecured mobile devices that run their own operating system as they can be accessed on-site without the need for additional hardware. Moreover, these devices often boast networking hardware enabling the rapid transfer of data. Such devices often store saved user credentials and personally identifiable documents that often contain private and perhaps confidential information. One may feel that enforcing encrypted USB stick usage or making password protection mandatory for mobile devices would help to remedy this but it is important to consider the impact such a strategy would have on the workplace. Our methodology is aimed towards investigating such user behaviours and subsequently providing a modelling test-bed of our environment to provide a valid method of testing such a policy.

3.5 Discussion and Outline for Future Study

The pilot study has enabled a rapid assessment of our problem space. This forms the preliminary stage in our iterative methodology as depicted in figure 3.1. Prior to this investigation, the CISO had a limited understanding and restricted visibility pertaining to the current environment and the practices of users within the environment. Through conducting this investigation, we have enabled a more thorough, scientific evaluation of the risks and threats within our problem space. From our results, we are now able to target specific behaviours and formulate bespoke investigations in order to increase our understanding of specific threats and where vulnerabilities may be exploited.

From our results, several important factors have been highlighted:

1. User Expertise: There is a stark divide between users in relation to their knowledge of IS and risk aversion. Lecturers, PhD students and technical staff were more risk averse than clerical staff who were typically unaware of current procedures and often lacked knowledge pertaining to data security (Q1-7). This divide is further visualised through an assessment of the types of data that were being handled and the method in which this took place. Often, their choice was not clear, nor could it be rationally justified through further dialogue.

2. Data Mobility: Users actively transfer data between locations on a regular basis. Subsequent access to this data is typically off campus and is done so whilst connected to unknown networks. All four categories of users partake in this practice, and it includes a mixture of many different data types with differing values and consequences regarding loss. The use of multiple device types both personally and university managed furthers the risk the university faces with respect to security.

3. Devices: Many users employ BYOD in their daily practices. Personal laptops are seen as the most common form of device for data transfer and off campus data creation, whilst USB sticks are seen as the most popular transfer media. Of note, however, is the use of smart devices such as tablets and mobile phones which some users adopt and are typically less secure (for example, many users adopt swipe to unlock to hasten a common task). As previously discussed, these devices typically hold more personally identifiable information and often store user credentials.

4. Device Access: Numerous users share their devices with other people. Whilst it is understandable to trust family members not to maliciously steal or delete/alter data (irrespective of whether policy allows sharing in this fashion), it is possible for these users to unintentionally install malware or other malicious software through unsafe internet browsing and other activities. This poses a significant risk for the user as they may be unaware that their device is infected, and indeed infect other devices on the campus network.

5. Method of Transport: We discovered that the majority of users transport their chosen devices on their person. This is most likely due to the popularity of mobile devices and USB sticks and their attributed portability. However, some users still chose to transport laptops and leave them in communal baggage areas providing an opportunity for thieves and targeted attacks.

6. Data Loss: Only one third of participants were aware of the implications that data loss may pose for themselves and the university. This is important as it suggests that

users do not actively match their level of security, represented by device choice, to the sensitivity of the data they are transferring. This is furthered by the discovery that many of the participants were unaware as to the implications of what losing this data may cause.

Consultation with the CISO regarding these results and the above findings has highlighted numerous areas for further in-depth investigation. These areas reflect the concerns of the CISO and the necessity to adopt policies to govern these particular practices based on user assessment. With respect to our methodology, the above numerical list has demonstrated the benefits of a rapid, targeted user interview strategy that has identified numerous relationships and trends with which to investigate. For the remainder of this thesis we identify the top three behaviours and practices as highlighted by our CISO in order to demonstrate and test our methodology. This decision aids in scoping the project and provides sufficient example to demonstrate our aims and objectives. As stated, we believe this process to be a proof of concept and one that is scalable. As such, it would be possible to investigate numerous identified behaviours simultaneously with sufficient resources.

The following sub-headings denote the behaviours we wish to examine. These behaviours relate to the following chapters within the thesis where we detail bespoke experimental design, data collection and a more thorough analysis of these specific behaviours.

3.5.1 Connecting to Unknown Networks Using BYOD

Connecting to an unknown, unsecured network presents a myriad of risks that increase the likelihood of data leakage. Misconfiguration of device settings such as auto-network enrolment, sharing of shared folders along with using non-encrypted web tasks and connections increases the likelihood of attacks through avenues such as man-in-the-middle. We have chosen to examine this particular practice as it presents a genuine threat to the users within our pilot study. In this investigation we aim to understand how users connect to networks, what preferences affect their selection process, and understand their rationale for a given choice. We wish to simulate a real-world scenario where we promote a choice architecture that utilises nudges in order to steer users towards selecting a more secure wireless network. If successful, our behavioural interventions should reduce the likelihood of data leakage and provide a strong evidence base for policy design.

The experiment is detailed within chapter 4.

3.5.2 Online Behaviour and Privacy

Secure behaviours when online are a necessity in today's cyber environment. The annual cost of cyber-crime is approximately \$ 100 billion [101] and rising at an estimated 17% annually. It is imperative that organisations take the necessary precautions to protect themselves from this avenue of attack. One such method is to understand how to protect users, or perhaps more importantly, how to enable users to help protect themselves. One such area where this is highly relevant is in online privacy. Users value their privacy, but do not readily protect it - a phenomenon known as the privacy paradox [210]. As such, we must understand how users form their privacy decisions by examining their decision making process.

The experiment is detailed within chapter 5.

3.5.3 Error Reporting

The importance of error reporting cannot be underestimated. It is a critical component of error prevention as the mantra 'we learn from our mistakes' dictates. For security systems, it can be viewed as an integral part of an iterative process. By definition, no system is perfect, and thus when errors do occur it is important to understand what happened and how it occurred in order to emplace barriers against that particular event. This process relies heavily on reporting and often involves user interaction with the submission of logging events of system errors. If errors are not reported, it can leave a system vulnerable to attack through the same vector. As we witnessed within our pilot study, almost a quarter of users had experienced an error where they needed assistance, however, less than ten percent actually reported them to the proper authorities.

In response to this, we wish to understand how users behave with respect to error reporting. We wish to investigate what decisions ultimately determine whether or not an individual will report a given error and devise strategies for influencing this. If we can increase the number of reports that are submitted, we will reduce the likelihood of preventing future events through mitigation strategies and countermeasures. Ensuring reports are easy for the user to conduct will be an important process as we do not wish to increase their cognitive load or diminish their compliance budget.

The experiment is detailed within chapter 6.

3.6 Chapter Contribution

This chapter has demonstrated that an iterative methodology is an effective approach towards aiding the design process of DLP policies. We have shown the value of targeting user behaviours to develop bespoke investigative methods that aim to identify how users make their choices and to ultimately gain a more detailed insight of our problem space. Through these investigations we improve the expressiveness of our modelling process through finer articulation of the processes and sub-process that are present. Furthermore, we have demonstrated the importance of conducting a pilot study in order to begin this iterative process. Feedback from our stakeholders has validated our findings and examining these behaviours has enabled subsequent experimental design witnessed in the following chapters.

Chapter 4

Nudging Towards Security

The previous chapter detailed the methodological approach towards aiding the policy decision making process. Within, we emphasise the necessity to investigate specific behaviours and practices within the current environment that reflect the needs of our partner in managing such an environment. This chapter represents one such area of enquiry and aims to determine how users connect to networks to access and transfer data remotely.

This chapter details experimentation with respect to nudging users towards selecting a more secure wireless network on their mobile device. As discussed throughout chapter 2, users pose one of the most significant risks to organisational information security and the trend in BYOD is set to continue. By understanding how users select networks, we can begin to build intervention methods that ultimately influence user selection with the goal of mitigating data loss and exposure through understanding which networks pose the most significant threat. This knowledge ultimately aids the policy decision making process as we empower the CISO to design more appropriate policies that aim to operate holistically with users whilst preserving security.

The remainder of the chapter is as follows: Section 4.2 provides an overview of the problem space and the formulation of an investigative aim. Section 4.3 details the development of a prototype application for investigating the issues detailed within section 4.2. Section 4.4 evaluates the results from experimentation and presents an analysis of the adopted nudges. Section 4.5 discusses the significance of the findings with respect to aiding the policy decision making process and highlights limitations within the research along with future experimentation.

This chapter includes results and experimental design detailed in Turland et al [231] and includes collaborative work declared within the 'Publications'.

4.1 Wi-Fi Experiment

This work aids our understanding of the problem space through the investigation of user behaviour with respect to network connection. As previously highlighted in chapter 3, the CISO has little knowledge pertaining to the amount of information transferred by users, and the methods in which this is done. By assessing how users connect, and importantly what influences their decision making process, we can design policies that aim to ‘nudge’ or guide users towards more secure connections with the goal of minimizing data loss.

4.1.1 Chapter Contributions

Within this chapter a prototype application is presented which promotes the choice of secure wireless network options, specifically when users are not familiar with the wireless networks available. The application is developed based on behavioural theory, choice architecture and good practices informed by HCI design. The application includes several options to ‘nudge’ users towards selecting secure public wireless networks. Explicitly, this chapter outlines the development and the results of an evaluation of some of the potential application nudges (specifically, presentation order and colour coding of the wireless networks). In summary, colour coding was found to be a powerful influencer, less so with the order in which we listed the Wi-Fi networks, although the colour × order combination was the most effective.

4.2 Introduction

Organisations rely on their staff to make numerous different security decisions when completing different tasks, whilst on the premises and on the road. An example of such a decision is the selection of which public Wi-Fi to connect to if working remotely from the office. Over one billion workers are believed to work remotely [64] - over a third of the total worldwide workforce. Unsecure wireless network selection becomes problematic as more employees use mobile devices (possibly their own) to access various networks outside their workplace and transmit potentially sensitive information. Unless employees have local, trusted wireless networks or use 3G/4G data communications, they will often utilise locally available, public options. When employees face time pressure, they may be tempted to make hasty decisions that allow them to access the internet via unknown hotspots. This behaviour poses a significant risk to the security of the device and its data, as these networks provide many opportunities for cyber-attacks, including spoofing and man-in-the-middle [45].

Exploiting and influencing wireless network selection is possible in numerous fashions, some as simple as changing the network name [91]. Aiding users to select secure, appropriate networks for their tasks, be it on personal or company owned devices is imperative to maintaining security. Wireless network selection is a potential area in which to encourage a more security-driven decision. Selecting a secure and trustworthy Wi-Fi connection is one of the top 10 security behaviours promoted on sites such as www.staysafeonline.org. The question is how can we ensure that people make a secure decision whilst still permitting choice and preserving productivity?

Through understanding how users connect and ultimately share/access data we are able to begin to more accurately articulate and identify where threats lie. This process details an investigative method which compliments the overall methodology highlighted in chapter 3. The work detailed here aims to understand how users choose networks and investigates the effectiveness of nudging towards more secure alternatives as part of improving the policy decision making process. The development of tools provides an opportunity to deploy and test an intervention method in a real-world environment, a key limitation of simulation-based study.

4.2.1 Designing for Persuasive Security

When using a mobile device and connecting to the internet, people rarely think about security, yet the choices they make have implications for the security of their device. Leaving such decisions to chance and simply 'trusting' employees to do the right thing may not be the best course of action. We need to understand how we can influence this behaviour to be more security orientated and develop more effective policies.

In recent years, we have seen the popularisation of Thaler and Sunstein's [228] approach to 'nudging' behaviour towards a particular outcome. Nudging is based on choice architecture and promotes the idea that the manner in which a choice is presented will affect the decision outcome. Several studies have considered nudging as a means to affect behavioural change in the area of privacy. For example, behavioural nudging has been considered in terms of encouraging less information disclosure via social media [245], on mobile devices ([26]; [55]) and to improve general privacy behaviours ([3], see also [110]).

Nudging is a logical approach to investigate within this problem space. As documented in section 2.3 the application and merits of such an approach are witnessed through numerous examples. As we are investigating an undocumented environment, nudging can be seen as a bridge between enforcement and open access. This is beneficial as the introduction of new systems that are considerably different than current systems often meet the greatest resistance or adaptation time (particularity so if perceived as negative), and subsequently have a lower adoption rate as discussed in 2.2.2. The 'art' of this approach is to present users with the complete set of choices, but fashion the most desirable choices in such a manner that it makes choosing sub-optimal choices difficult to justify.

The principles of a good choice architecture ([228]) can be attributed to how different aspects of security options are depicted. For example, in terms of selecting wireless network options, we must consider [231]:

1. Incentives: What rewards (or disincentives) are in place to encourage the desired secure choice? For instance, does the most secure wireless connection require payment?

2. Understanding mappings: Can people understand how to use the system properly to achieve their goals? Is it possible to map the choice to a desired outcome? Some choices are simple and vary only on one factor, others vary on multiple factors simultaneously and we have to decide which is of greatest importance.

3. Defaults: If for a given choice, a default is set, many people will simply follow that path. Is the default the secure choice?

4. Give feedback: Behaviour will improve if feedback is provided about the successfulness of the behaviour.

5. Expect errors: Users make mistakes, so how can we reduce the likelihood that an unsecure Wi-Fi connection will be chosen?

6. Structure complex choices: Help people to make a choice.

These rules are not new to HCI studies. Norman summarised several approaches in ‘The Design of Everyday Things’ [171] where he discusses the connection between design and the user; specifically, how to optimize this experience and make it more pleasurable for the user. Similarly, B.J.Fogg and ‘Persuasive Technology’ [93] focused extensively on the role of technology in this process. Persuasion design embeds various forms of influence and ‘choice architectures’ in products and services to maximize the likelihood of positive behaviour change and experience. This work has been further explored by Lockton et al [153] in ‘Design With Intent’. These authors utilised choice architecture within the context of ‘Design with Intent’ - design intended to influence user behaviour.

4.2.2 Research Aims and Scenario

We employ several of these good principles of choice architecture in the development of a new application in order to support decision-making in situations of uncertainty. Our aim was to investigate security-related decision-making via design in a specific setting: The user has no access to a whitelist of wireless networks, their device has no mobile data access to the internet, and they are under time pressure to complete a task involving submitting potentially sensitive information from a personal mobile device. This task was selected for several reasons in full knowledge that an increasing number of mobile devices have inbuilt 3G or 4G capabilities. Importantly, however, there are several physical issues that can render this technology unusable. For instance, network coverage may not be present in all areas, especially when travelling abroad, and tariffs for overseas internet access tend to be significantly higher than national rates making them financially non-viable. Data caps may also influence wireless network behaviour as exceeding tariff quotas often causes significant financial burden.

This approach is of merit for the evaluation of current Wi-Fi behaviours and the documentation of an application constructed using behavioural insight and user-centric design. At present, mobile devices tend to use different parameters for displaying available wireless networks even on the same mobile operating system (proprietary 'skins' often create visual discrepancies). For example, at present the default Android behaviour ('non-skinned') lists all available networks in alphabetical order with a padlock denoting those which require a password. This presentation of options may influence decision-making in a number of ways. Either individuals select the first available choice, they search for a trusted wireless (one they recognise by name) or they search for the strongest signal. We investigate which options users select and assess the potential to change their decision-making process.

4.3 Application Development

To develop this application we consulted and followed guidelines presented in the SCENE process ([63]). In summary, we identify target behaviour, i.e. selecting the most secure Wi-Fi network available, explore the underlying mechanisms that influence that behaviour and design the application to utilise these factors and nudge the target behaviour.

An important part of the process is to understand the different ways in which behaviour can be influenced. This allows for future policy development to be sensitive to user actions and behaviours.. A critical component of this is the MINDSPACE framework summarized by Dolan et al ([83]). This framework summarises a number of the 'influencing factors' that have been shown to affect behaviour through years of research in the behavioural sciences ([1]; [161]). Many of the guidelines and observations within feature prominently in the design and implementation of the application.

MINDSPACE [83] can be adopted to explore design solutions in different ways:

- **Enhance:** MINDSPACE can help designers to understand how current designs are influencing behaviour and how they could be improved. A pre-cursor to effective application design.

- **Introduce:** There may be some of the elements from MINDSPACE that are not currently used within the design and there may be space for an innovative use of some of them.

- **Reassess:** The designer needs to understand the ways in which the design may be influencing the behaviour of its users unintentionally. It is quite possible that the design produces unintended and possibly unwanted behaviour.

Using the MINDSPACE framework, we generated a number of potential design possibilities that considered the following 'influencing factors' in relation to wireless network selection (Table 4.1):

- the extent to which messages are trusted due to the messenger is (e.g., messages from company manager rather than network provider).

- the incentives to stay on a secure network and drop an insecure network.

- the behavioural norms or common experiences of the user's important social groups (e.g. infection rates of users of the same network).

- changing the default options (e.g., by listing the most secure networks at the top of the available networks list).

- changing the salience of secure vs. unsecure options (via highlighting company-trusted networks in green, secure networks in yellow and unsecure networks in red) and

- the use of emotive (affect-based) interventions (e.g., voice concern over selecting flagged options).

Many of the above scenarios can be viewed as the testing phase of a prototype policy implementation. Through investigating users (their behaviours and practices) and the direct manipulation of security mechanisms, the CISO is able to test proposed impacts on the user and their environment. This approach enables an optimization of the policy design process through a rapid deployment and small-scale testing platform with the advantage of understanding impact analysis at an early stage.

Table 4.1 Definition of MINDSPACE Influencers

Influencer	**Definition**
Messenger	we are heavily influenced by who communicates information
Incentives	our responses to incentives are shaped by predictable mental shortcuts, such as strongly avoiding losses, optimism bias and perception of risk.
Norms (Social)	we are strongly influenced by what others do
Defaults	we ‘go with the flow’ of pre-set options and take the lazy way to decisions and behaviours, habits become pre-set behaviours.
Salience	our attention is drawn to what is novel and seems relevant to us
Priming	our acts are often influenced by subconscious cues
Affect (Emotion)	our emotional associations can powerfully shape our actions
Commitment	we seek to be consistent with our public promises, and reciprocate acts
Ego	we act in ways that make us feel better about ourselves

4.3.1 Problem Identification and Applicability

Encrypted Wi-Fi connections on trusted networks, and Virtual Private Networks (VPN) over any Wi-Fi offer a reasonably secure method with which to connect. It is important to note, however, that they are susceptible to attack through various hardware and software based products such as WiFi Pineapple and Dumpper that exploit weaknesses in the implementation of popular protocols by certain vendors. These attacks however, typically require large periods of time as packet analysis requires large volumes of traffic. For this work, we will ignore these attacks and focus on open networks such as those commonly found within many public environments. There are numerous situations where it is not possible or practical to use encrypted connections and instead unsecured networks are necessary.

Public Wi-Fi’s often require a subscription that may dissuade users from choosing it, thus increasing the likelihood of selecting other free, open access points which in-turn pose a multitude of security risks. Furthermore, many Small-Medium Enterprises (SME’s) often lack the finances to implement and maintain a VPN for their employees. This, coupled with the ever increasing need to ‘work on the go’ produces an environment where it is imperative that you are able to work securely. Consumerisation further exacerbates this problem as not only is the device multi-use (often changing the way users interact and secure their devices) it often contains a mixture of both company and personal data on third party servers (often in an unknown locale).

Ferreira et al ([91]) states that it is relatively simple to influence a user's choice when selecting a network. This provides a perfect opportunity for an attacker who wishes to gain access to a device or sensitive data as the user is placed in an unknown environment with unknown threat vectors. A relatively simple example of such an exploitation is the changing of the broadcast ID (SSID) or more fundamentally the changing of the MAC address, (known as MAC spoofing). Employing such a strategy exploits the user's familiarity and habituation to certain networks and working environments. Upon erroneously connecting to a spoofed network, users may believe that they are indeed safe as they trust the provider and are more likely to undertake confidential activities (i.e. checking one's bank account, sharing files). If indeed this is a man-in-the-middle attack, the user is at serious risk of losing data or credentials (some of which may not be their property).

It is therefore essential that we have some form of intervention method at the point where the user chooses a Wi-Fi network. We must influence the user to choose a more appropriate network for the type of activity they wish to undertake. This mentality is the rationale behind our application which aims to impact the 'thinking fast' system 1, habitual process of wireless selection (see Bravo-Lillo et al [42]). It is important to stress, however, that we do not wish to hamper or increase the time in which it takes to connect to a network; simply that we wish to empower the user towards making a more appropriate, informed choice with respect to security for the task they are undertaking.

4.3.2 The Prototype Application

The Android mobile operating system was chosen as the platform for our application as it was the most popular mobile operating system in the third quarter of 2013 with an 81% market share ([40]). The interventions designed and tested on this platform are not operating system specific and can be adapted to suit other devices. One attractive feature of Android, however, is the openness of the environment and source code enabling one to have total control over the specific device. Such an approach differs from competitors like iOS [19] where many alterations of source code, or changes in device behaviours are prohibited. Whilst such an openness may create security vulnerabilities, we are not concerned with this in the current investigation and only adopt default calls and methods.

The prototype application provides the following additional features:

- Reordering the network list so that the most 'secure' networks are first (appropriate default). Secure is defined as the most desirable to the CISO.

- Applying a colour scheme (appealing to user's social norms). The most secure choice is coloured green and the least secure red (explicitly defined later in section 4.3.3).

- Implementing an approved wireless network list that is editable by the user/CISO (Chief Information Security Officer), marking reliable and secure networks as ‘trusted’. Approved networks are placed at the top of the network list.

- Showing all named wireless networks at the scan location (Android currently combines SSIDs and does not display duplicates making it difficult to verify authenticity and avoid spoofing).

- Additional information buttons presented alongside scan results to inform the user about potential risks pertaining to their actions (messages to educate and empower the user if they are unsure).

- Pop-ups and notifications (providing potential audio, visual and haptic feedback) alerting the user when connecting to an unsecured network (expect errors, provide feedback) and discouraging them from persisting.

- No additional confirmation steps when connecting to a network. Does not hamper productivity or deviate from familiar behaviours. Instead, provides additional contextual information with which to make a more informed decision (expect errors, provide feedback).

The application is written within the Official Android IDE (Eclipse) and conforms to API 8 with a target API 18. This ensures that the application is compatible with older Android devices but can take full advantage of newer features found in more recent releases. Specifically, we take advantage of numerous new notification features (audible, haptic and visual feedback).

The application in its current iteration consists of two activities. An activity represents ‘a single screen with a user interface’ ([102]). The rationale for this design is that we wish to present the user with a familiar splash screen, in our case the organisation’s logo and additional relevant contextual information first(perhaps a contract or code of conduct). This approach adopts the Messenger element within MINDSPACE (‘we are heavily influenced by who communicates information to us’), reminding the user of either expected behaviour or perhaps more importantly, contractually obligated behaviour. Figure 4.1 denotes a possible implementation of this proposal (this is untested at present) and figure 4.4 represents a possible popup aimed at informing the user.

Figure 4.1 presents a design view of activity 1 from within the Eclipse workspace. This rudimentary design features a company logo, an information button for assisting users, the ability to add and delete trusted SSID’s (automatically given preference in the Wi-Fi list in activity 2) and the button to perform a scan.

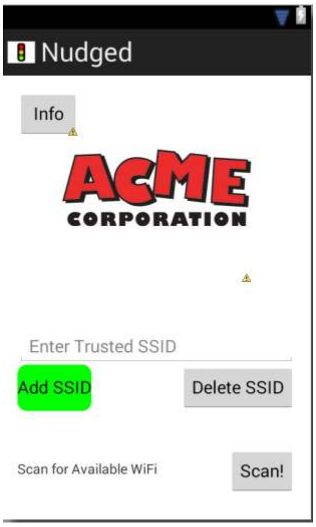

Fig. 4.1 Activity 1: Potential Mockup (Untested)

The second activity is called when the user presses the scan button. This displays the wireless scan result at that given instance. Results shown follow our colouring, ordering and contextual format as described previously. Results are not auto refreshed but can be manually rescanned by pressing the ‘rescan’ button. This is a conscious design choice which allows the user to focus on a static list rather than one which dynamically updates. Dynamic lists (android default is 5s refresh) increase the risk of selecting the wrong, potentially unsecured network through user error creating security risks and additional user input as they must re-select a Wi-Fi. A comparison of the before and after effects in a real-world example of our interventions can be witnessed in figures 4.2 and 4.3 (discrepancies in the number of networks and signal strength are the result of natural variation at the time of capture).

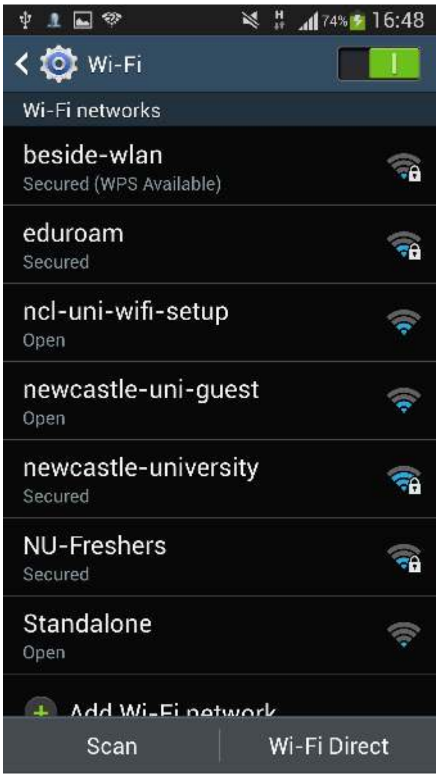

74% 16:48

Wi-Fi

Wi-Fi networks

beside-wlan
Secured (WPS Available)

eduroam
Secured

ncl-uni-wifi-setup
Open

newcastle-uni-guest
Open

newcastle-university
Secured

NU-Freshers
Secured

Standalone
Open

Add Wi-Fi network

Scan | Wi-Fi Direct

Fig. 4.2 Default Android Wi-Fi List without Interventions

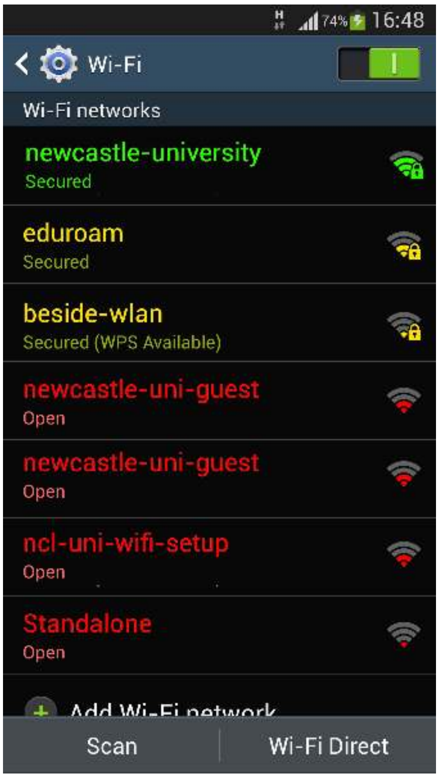

74% 16:48

Wi-Fi

Wi-Fi networks

newcastle-university
Secured

eduroam
Secured

beside-wlan
Secured (WPS Available)

newcastle-uni-guest
Open

newcastle-uni-guest
Open

ncl-uni-wifi-setup
Open

Standalone
Open

Add Wi-Fi network

Scan | Wi-Fi Direct

Fig. 4.3 Default Android Wi-Fi List with Interventions

The second activity also utilises popups (4.4) that are generated on user interaction. The ‘help’ button displays useful information relating to which network is appropriate for the task wanting to be performed. This empowers the user to make their own decision and utilises incentives (‘our responses to incentives are shaped by predictable mental shortcuts such as strongly avoiding losses’) from the MINDSPACE framework. The popup itself builds upon findings supported by Raja et al ([190]) with the aim of providing warnings that are ‘understandable to users’, rather than the often convoluted and technically based solutions that many software programmes implement. Example tasks and solutions are displayed in a brief pictorial form.

4.3.3 Ordering Networks

Upon scanning, results are ordered using an algorithm that lists the Wi-Fi’s in order of their whitelist approval (managed by the CISO) and security (encryption protocol strength, 802.1x > WEP etc). We define whitelists as a set of preferred networks based on these criteria. However, this does not include networks that the user has previously visited (so the whitelist is not a visitation list). The ordering is such:

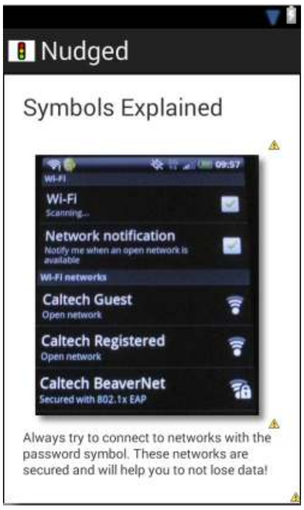

Nudged

Symbols Explained

Wi-Fi ☑
Scanning...

Network notification ☑
Notify me when an open network is available

Wi-Fi networks

Caltech Guest
Open network

Caltech Registered
Open network

Caltech BeaverNet
Secured with 802.1x EAP

Always try to connect to networks with the password symbol. These networks are secured and will help you to not lose data!

Fig. 4.4 Pop-Up: Potential Mockup (Untested)

1. **Whitelist** (Green) (protocol strength) (high to low signal strength)
2. **Secured** (Amber) (protocol strength) (high to low signal strength)
3. **Open** (Red) (protocol strength) (high to low signal strength)

Firstly, the algorithm ensures that company approved whitelisted Wi-Fis are always presented at the top of the list and coloured green (the most commonly chosen from our results - see 4.4). This implies that a user is most likely to choose a network that is trusted by their organisation and offers a ‘secure’ encrypted connection (such as WEP, WPA-PSK, 802.1x). Importantly, it ensures that if a data leak were to occur, the employee was operating within company policy. The term ‘trusted’ is added to the network metadata to inform the user that the network is known to the device (i.e. in the whitelist and managed by the user’s organisation). This builds upon the standard use of the word secured that is default Android behaviour and expected by the user.

Secondly, ‘Secured’ networks are displayed in amber and below the whitelist. These are networks which encrypt traffic (802.1x, WPA etc) but where the network owner is unknown. To reiterate, we do not want to limit choice, we wish to nudge users into making a more informed and appropriate choice. In our definition, ‘Secured’ networks represent Wi-Fi’s that are not trusted (not in the whitelist) but do offer an encrypted communication. Typical occurrences include home/public hotspots where the network key is shared. Whilst this does not protect the user from ‘man in the middle’ attacks (even though each client derives its own session key) it does provide the user with a choice if there are no whitelisted Wi-Fis available, or the connection of a trusted Wi-FI is unstable. The use of such networks will be familiar with many people in public locations and requires the user to assess their threat environment.

Lastly, ‘Open’ networks are displayed in red and presented at the bottom of the list. These represent Wi-Fi connections that offer no form of encryption, require no login credentials nor offer vendor identification. Anyone is able to freely connect and use the service. As such, this is by far the least secure network type for data transfer and sensitive activities and thus the least desirable. Fig 4.5 represents a test slide presented to users with our interventions included (network names are obscured to remove user familiarity).

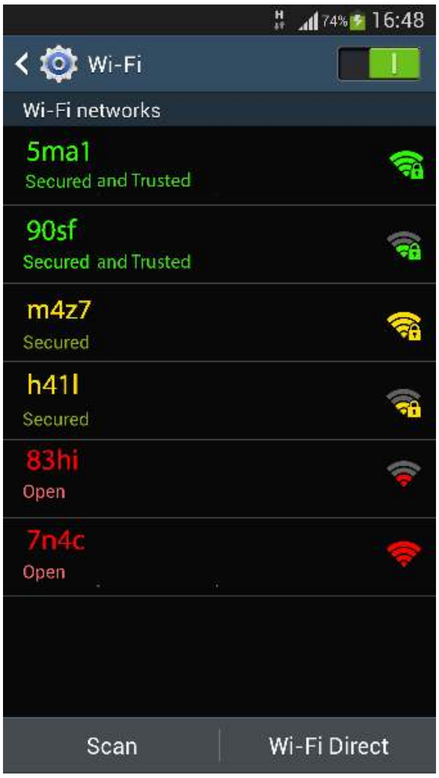

Wi-Fi networks

5ma1
Secured and Trusted

90sf
Secured and Trusted

m4z7
Secured

h41l
Secured

83hi
Open

7n4c
Open

Fig. 4.5 The Default Android Wi-Fi List with Colouring and Ordering

Currently, default Android behaviour combines SSID names into a single access point (see Fig 4.2) regardless of whether the access point is physically different. This is potentially highly problematic and something that is easily exploitable as highlighted by Cassola et al ([52]). This application has the ability to display multiple networks with the same SSID (see Fig 4.3) or MAC address by displaying all results. By listing all access points with the same name we raise the user’s awareness by indicating that there may be more AP’s than should be expected, a potential security risk. For instance, if a user expects to find one access point but instead finds two in their local café, they may hesitate, think, and ultimately decide not to connect to it, thus avoiding connecting to a malicious network. In larger networks, several AP’s with the same name may be displayed, but listing these presents the user with greater choice and increased control over network selection (for instance they can connect to the network with the strongest signal for optimal productivity). The use of a whitelist can help to reduce the number of repeat network SSIDs by actively adding MAC addresses to the allowed connections.

The whitelist itself is a top level filter that parses the scan results subject to specified criteria (essentially, but not restricted to, trusted MAC addresses). The whitelist is operated and maintained by the company CISO and uploaded to the user's handset at regular intervals at a trusted location or network (possibly on the premises). This allows the whitelist to be modified and importantly kept up-to-date. It also removes user inconvenience and further limits user error. The specific deployment of the whitelist and the security concerns related to this are not covered.

Selecting a Network

Selection of a Wi-Fi remains synonymous with the default Android method (scan -> select) for whitelisted networks. To connect to a whitelisted Wi-Fi, a user simply presses the appropriate list field and connects, entering any required details when prompted. This is the most convenient method for the user as it compliments habituation and is the desired 'nudged' behaviour from the CISO's perspective. Upon selection, the default connection method will be called and the standard Wi-Fi symbol will be displayed in the notification area signalling a successful connection. It is important to stress that with this example (desired behaviour) there are no additional steps either during connection or once successfully connected, and the phone 'behaves' in the default Android fashion.

Connecting to a Wi-Fi that is outside of the whitelist prompts the user with additional steps and notification frequencies. These increase as you move from amber to red. These interruptions specifically aim to hamper productivity when not behaving securely and promote the adoption of whitelist network selection by reducing interruption frequency. To exemplify, consider a connection to a red (Open) network. Here, the user adopts the typical behaviour of scanning and subsequently selecting a Wi-Fi name. However, instead of automatically connecting or being prompted with credentials, the user is notified that the network may not be secure (Fig 4.6). This additional step breaks the user's habituation and introduces an unknown process which requires additional cognitive processes to bypass. By alerting the user in such a fashion we empower the user to validate their network choice and ask them to question their selection.

This approach is similar to Maurer, De Luca and Kempe ([159]), where prompts and warnings 'appear in-context' to user actions. This heightens their impact as habituation dictates that 'passive indicators are mostly overlooked'. By interrupting typical procedures we are impacting user behaviour and increasing the likelihood that notifications will be

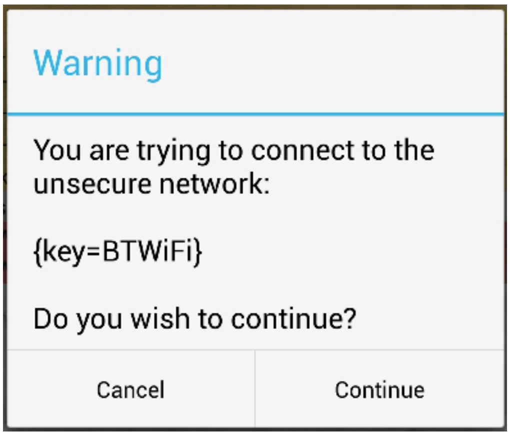

Warning

You are trying to connect to the unsecure network:

{key=BTWiFi}

Do you wish to continue?

Fig. 4.6 Warning Dialogue Presented to Users When Trying to Connect to an Unsecured Network

observed. Upon selecting cancel, the user is returned to the previous Wi-Fi list with no additional prompting (see Figure 4.6 and 4.7) and allowed to select another network.

In contrast, if the user wishes to continue with this unsecured network, they are provided with a three pronged notification (audio, haptic, and visual see Fig 4.7). To implement this, we call the default Android Notification Manager (Android, 2003). This provides customisable text informing the user that they are connected to an unapproved network (this could include company policy to further stress the implications to the user). It also makes the device sound the default warning tone, flashes the LED, and vibrates the device (if applicable). These processes may occur at regular intervals (or not at all, defined by the policy), and can be seen as a means of limiting productivity for an undesired network. Whilst this approach may seem intrusive, it may be warranted if the user is accessing sensitive data over a less secure/unknown network and placing the organisation at risk. The immediateness and method of the notification may be sufficient to dissuade users from continuing on the unsecured network.

This implementation provides a clear, simple yet informative user interface that enables the user to quickly make informed decisions as to which Wi-Fi to choose. The display builds on defaults, familiarity and habituation, as well as common Android methods and behaviours (the way in which events are presented) and format (takes advantage of newer Android draw features) but adds additional layers of usability that enable a safer, more secure environment. This framework is scalable between devices and operating systems and requires limited interaction from the user.

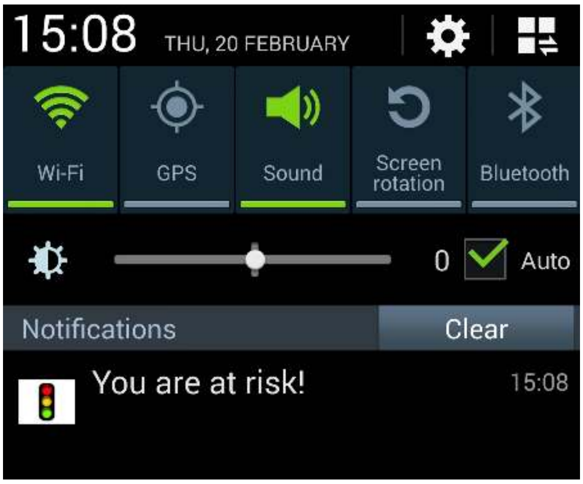

You are at risk! 15:08

Fig. 4.7 Notification Symbol in the Android Status Bar Upon Selecting an Unsecured Network

Stage-wise Development of the Application

As with any prototype application it is important to test the validity of certain features before implementation. As our interventions are designed and influenced from recommendations within relevant literature, it is necessary to ascertain whether the findings within hold true for our explicit purpose and execution.

An important element at this stage was to consider the need to evaluate each possible nudge, or a combination of nudges. This would allow for an independent assessment of how effective each nudge was, the outcomes of such a nudge on behaviour, and any new or unexpected behaviours that may result as a combination of the two. As such, the remainder of this chapter discusses how the first two nudges were employed within the new application; chiefly colour and ordering. This approach combines salience (prominence of certain options listed at the top) and affect (the impact of using red->green colours and the associated emotions and meanings). Support for this approach is seen in Choe ([55]) where visual framing was adopted to nudge individuals away from privacy-invasive applications.

Materials to Test First-Stage Nudges

The new application is designed to change user choices and defaults by moving individuals away from the Android default of alphabetical ordering to one where network options were ordered by security or organisational preference (see 4.3.2). In addition, we introduced a colour coding scheme where whitelisted networks were green, secured networks were amber and the least secured (open) were red. The new application subsequently includes two nudges: first, one involving ordering (thus eliminating the bias of picking the first and most convenient network regardless of security) and second, colour (informing and capturing the attention of individuals when selecting wireless networks). Similar visual feedback and framing of information techniques have been employed in other studies (see [55], [236]).

In order to be able to examine the potential effectiveness of these nudges in combination, and to compare them to more standard network displays (without these nudges), we used five different screenshots of wireless network lists. We eliminated familiarity and signal strength effects by creating random network names and presenting network options with the same number of bars (indicating signal strength) which were either half or full strength (2/4) for the same number of options on each screenshot (Fig 4.3 & Table 4.2).

	Colour nudge	
Security	**Colour**	**White**
Order present. (most secure on top)	OCp	OWp
Random present. (no order)	RCp	RWp
No Padlock (Order)	OCn	–

Table 4.2 Overview of screen shot characteristics including nudges (security ordering and colour)

4.4 Evaluation

We evaluated the effectiveness of our nudges with 138 participants (university non-computing students and staff) who were familiar with using wireless networks on a university campus but were not technical experts. As we were sampling non-experts, we also included a fifth screenshot that featured order and colour but no padlock (OCn), thus representing an alternative to the same option with a padlock (OCp). We include an overview of the four out of five screen shots in Table 4.3 above.

All participants were given the following scenario: ‘You have an hour to complete and submit some urgent work and have decided to go to a public cafè to connect to the internet’. In this setting, subjects are presented with various network options relating to the above declared screenshots. Participants were then asked to pick which network (out of six listed) on each screen shot (totalling 5) they would select. Furthermore, participants were asked to order the 6 network names for each screenshot (participants took 10s on average to make a choice). Network names were randomly generated and six characters in length to avoid familiarity [91] and thus potential bias.

Participants were further asked to discuss why they had picked specific networks in order to examine which features were effective and how the presence or absence of the padlock symbol influenced their decisions. All screenshots for all participants were randomly

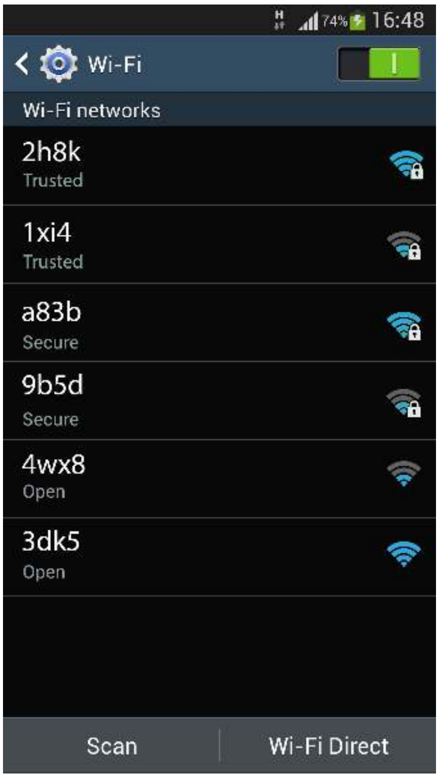

2h8k
Trusted

1xi4
Trusted

a83b
Secure

9b5d
Secure

4wx8
Open

3dk5
Open

Fig. 4.8 Nudged Application Screen Including Security Ordering Without Colouring

presented to reduce ordering effects. Figure 4.8 provides an example of the OWp condition for reference where all secured networks are listed at the top, while the less secured options are listed at the bottom (ordered, no colour labelling, with padlock) (see Table 4.2).

4.4.1 Evaluation Results

Our evaluation showed that ‘trusted’ also implied ‘secure’ for almost all participants, which led us to cluster secure and trusted choices into one group. When computing Chi-square for all five screenshots simultaneously, we observed a significant difference across the response options (p<.001). Following this, we wished to examine specific features by comparing the different screenshots against one another. We tested using a Chi-square test whether or not the results were statistically significant when comparing the frequencies for open vs. secure/trusted network selection across the screenshots that varied with respect to a specific feature (e.g., order or colour). We found that nudging by order alone was not significant but that colour was (leading to the selection of secure and trusted network options (p=.002)). When colour and order were combined, 70.3% of participants selected secure options - a significant improvement on the default condition (p<0.001). An overview of the preliminary results is provided in Table 4.3.

	Participant Choices	
Screenshots	**Open**	**Secure / Trusted**
RWp - networks not ordered by security, white labels (default Android)	80 (58%)	58 (42%)
OWp - networks ordered, white labels	71 (51.4%)	67 (48.6%)
RCp - networks not ordered, coloured	49 (35.5%)	89 (64.5%)
OCp - networks ordered, coloured	41 (29.7%)	97 (70.3%)
OCn - networks ordered, coloured, no padlock)	1 (.7%)	137 (99.3%)

Table 4.3 Frequencies Observed (N=138)

Further improvements were noticed when we compared the coloured and ordered results to the final and fifth screen shot featuring no padlocks. In the absence of a padlock (OCn), users were more likely to select secured (amber) options, which suggests that a key component of the decision-making process involved an assessment of the symbol itself. Open response options informed us that this effect may have been due to the fact that some users interpreted the padlock as a barrier (‘locked out’) rather than ‘security’. It is necessary to stress that this does not imply that OCn is a better design, but that user’s perceptions play a role in network selection. Figure 4.9 shows this more clearly. The vertical axis lists the frequencies. All options with C included colour, all options with O included ordered networks (see abbreviations in Table 4.1).

Open responses also suggest that some users look primarily for familiar names, a common reason why SSID and MAC spoofing is a successful venture. What is also apparent in these results is that the colour nudge in isolation produced a significant change in behaviour from the default, where as the order nudge by itself did not. As stated, however, both colour and order operated most effectively in combination.

4.5 Discussion

This work makes several important contributions to the field of HCI and proves the legitimacy of such an approach. Firstly, it outlines an application development process that integrates behavioural theory into the creation of a security intervention. Whilst the software development process was specific to this application, it may present an effective example of how interdisciplinary design considerations can support the development of application-based interventions. Secondly, the intervention achieves this without changing the default behaviour

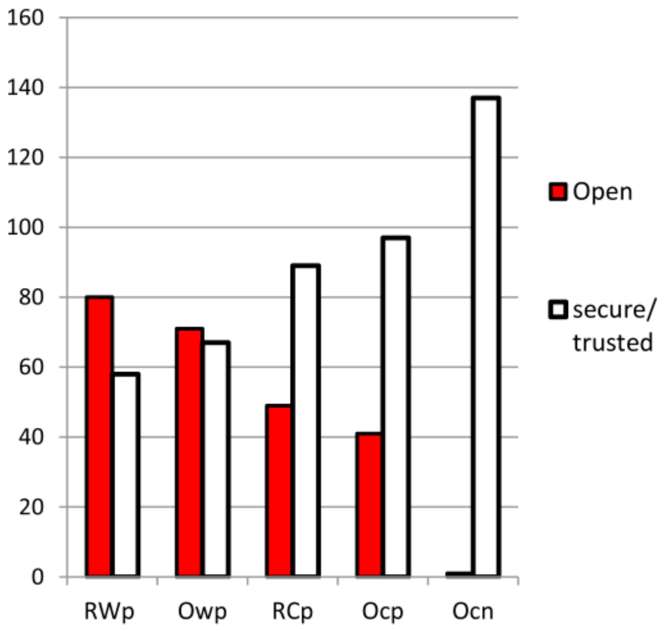

Fig. 4.9 Nudged Application Screen Including Security Ordering

(actual available options, i.e. it does not create a more restrictive environment nor does it introduce additional steps where security inconveniences the user). Instead, the application modifies and enhances the existing choice architecture to empower the user to make a more informed decision, providing the same set of choices as Android provides but changing the default presentation of those choices. Complimenting existing behaviours whilst also increasing security is an effective method for intervention.

Another important aspect of our work concerns the particular method in which it was framed. In essence we wish to preserve the 'your' in BYOD. We are not concerned with how people actually use their mobile devices in private or for work on an average day in familiar surroundings. Instead, our focus was to help participants make securer decisions in situations of increased uncertainty and unfamiliarity whilst under time pressure and accessing or transferring data that was not their property. Specifically, where users are forced to use unknown networks without the aid of the whitelist or the use of a VPN. This is also a situation where the user is potentially using unknown, local equipment that may not be secured effectively.

With respect to the policy decision making process, it is necessary to refer to graph 4.9 and the previous declaration regarding the padlock symbol. This figure details important findings and highlights the importance of tailoring interventions and policy implementation towards user perceptions. We must design policies that are simple to understand and compliment the user's method of thinking. Many ineffective policies fail not due to their purpose or aims, but the way in which they are deployed within the environment. Without fully understanding issues such as user interpretation and behavioural practices one cannot deploy effective policy. Critical failing in this area leads to increased circumvention of policies and outright infringement.

4.5.1 Information Security and BYOD Issues

From an information security perspective there are a number of alternative methods that need to be considered with respect to network choice and to further validate this approach. In particular, in the event that network selection is strictly controlled by the organisation (the user does not have a choice). Whilst this may be the case for company owned devices (and is perfectly reasonable), it is an unfair assumption that users are willing to conform to such an approach when using BYOD as issues of device ownership often compliment expected behaviour. This work predominantly considers the trade-off (security vs. productivity decisions) associated with this and nudges towards a trade-off optimum, instead of a security-only optimum. This creates interesting challenges in determining the trade-off and the optimal choice in which one tries to nudge towards.

One of the major challenges facing companies that embark on a BYOD strategy is the ability to understand, and more importantly influence user behaviour (see also Yevseyeva et al [258]). The prototype presented in this paper, and the subsequent evaluation results highlight the effectiveness of the nudges (colour and presentation order) and suggest that this application can remedy some of the potential security issues that arise in relation to wireless use with the adoption of BYOD. This is a critical component of the policy decision making process and further reiterates the importance of understanding user behaviours within the target environment.

Perceived familiarity with network names [91] may also create potential security vulnerabilities as users may not scan such network names as carefully, relying on recognition and habituation to select networks (also increasing the likelihood of spoofing). Our interventions combat this by providing an alternative method of ordering and displaying of network lists.

4.5.2 Limitations

Some, possibly critical, points must be raised with respect to the validity of our findings. These include usability concerns, context / applicability, and methodological / design issues.

Firstly, aspects concerning the usability issues related to coloured screens must be addressed. Whilst the design elements related to this are a core component of effective HCI design, and an integral part of our interventions, it is important to note that they are not effective for everyone and a solution must be offered. For example, it is possible to accommodate visual impairments such as colour blindness (a condition that affects 7% of people in some form) in numerous methods. There are also tools which test whether or not colours are distinct enough for people who are colour-blind (usually via contrast functions or other visual distinctions). Jefferson and Harvey ([126]) detail the importance and methods with which you can begin to reduce the impact of such factors by manipulating the way in which images and texts are presented. Their findings portray an effective method with which to alleviate some of the issues revolving visual impairment. Ensuring that policies are fair for all users, and importantly are equally able to be adhered to is another critical component of the policy decision making process. Failure to include such user groups may lead them to face heightened risk and ultimately higher consequences to both them and their organisation.

Secondly, with respect to the applicability of the application, we readily acknowledge the trend that many phones are now equipped with 3G or even 4G networking options. Nevertheless, as previously outlined in our introduction, there are numerous situations where convenience, costs and unfamiliarity in a new setting leads to Wi-Fi usage being the preferred connection method. As a result, this remains an issue for companies especially considering the fact that their employees are using BYOD or employer-issued mobile devices. As mentioned earlier, the trend in BYOD is increasing a growing rate and the need to adapt the policy decision making process towards tackling this issue in a timely fashion is becoming evermore pressing.

Lastly, it is necessary to reflect upon the experimental design of our investigation as it is not clear if our participants would select the same network choice if the passwords had been made available to them (this information was not included in the experiment as many situations involving public Wi-Fi frequently feature similar uncertainty). It is possible that making the explicit statement of the password may have led a greater number of users choosing the whitelisted results. Understanding this will be critical to further improving the decision making process for policy creation.

There are methodological decisions that may also warrant discussion. We decided to select six options per screen in the evaluation in order to control for list length, as the number of perceived options may influence decisions. This may have an effect on the user's decision making process because if you only have two choices you are somewhat forcing the decision compared to having several options where there are a greater number of variables to consider. At the same time, we did not want to present too many options that required scrolling or lengthy evaluation times. Further work is required to see if the effects still hold true if scrolling and larger lists are introduced.

Furthermore, the urgent task was not specified which may cause some participants to behave differently to how they would in situations of urgency. The sensitivity of the data was also undisclosed and its data type undefined. This ultimately may have influenced participants' responses.

As a smaller follow-up to this work, we also added a sixth screenshot featuring the default setting but eliminating the padlock so that we could compare the effect of padlock presentation on network choice in the default as well as nudged conditions. This further confirmed the role of the padlock as an important feature (for participants); however, more data is required to determine how relevant this feature is for different user groups.

Lastly, some research further suggests that habituation can possibly reduce the effectiveness of certain nudges such as messages warning users about posting certain content, e.g. Wang et al ([245]). However, that is a potential issue for all interventions and not something that is specific to our implementation.

4.5.3 Future Work

Future work will investigate more intricately the role of other interventions (see Figure 4.6 and 4.7), to educate the user and empower them to make a more informed decision that does not rely on Android defaults. This will be tested using the prototype application on a mobile device to improve ecological validity. Information will include explicit location references (e.g., expected networks within in the user's given location) and relate to real-world, possibly personal settings. In addition, the statement to users would provide data type and data value details to the users to allow them to make a more appropriate choice based on a greater number of variables.

To test the whitelist algorithms' dynamism with respect to variable data type and data values (e.g., the network order can change based on user input where a user selects which data type they have and what they are trying to accomplish). This will require testing a robust and clearly understandable method with which users can successfully assess their data sensitivity.

Furthermore, we wish to implement a monitoring solution that will collect anonymous data relating to the user's actions in a field experiment in order to evaluate the intervention in a real-world, real-time environment where users are making spontaneous decisions. A field study of this nature would also pave the way for data collection that may help organisations to determine the type of networks being accessed and thus inform organisational policies (which includes the whitelist). This may also enable organisations to track behavioural trends when users are presented with two identical networks (one being possibly spoofed).

Another challenge for future research is to conduct a more complex evaluation of the proposed intervention that considers security as one of only a plethora of trade-offs (security vs. performance/time required for connecting, convenience, or for particularly sensitive tasks such as choosing a Wi-Fi when conducting online banking). We do not at present know how the nudges play out in such circumstances.

4.6 Chapter Contributions

This chapter has highlighted the method in which a joint disciplinary approach can begin to manage and understand a user-centric information security problem. The methodology highlighted here presents a small-scale, rapidly deployable approach that aims to better understand user behaviours with respect to network security and data transfer. Such an approach provides an effective method in which to test-bed new policies in a scalable and controlled fashion with a statistically significant evaluation of the impact on the user base before adopting a more widespread deployment.

Through better understanding the users in our system from a user's perspective, and focusing on addressing specific issues that are highly relevant to the problem outline we are able to more effectively solve these issues. In relation to this work, we are able to understand how to work with users and compliment their current behaviours whilst also improving security. This is beneficial as it does not inconvenience users, reducing their productivity, and ultimately reduces the chance of non-compliance.

Chapter 5

Understanding Security Behaviours

The previous chapter documented the design and testing of an experimental nudge that aimed to modify a user's behaviour with respect to their wireless network preference. The following chapter documents an investigation conducted in Jeske & Turland et al [128] that examines how users behave with respect to online practices and privacy, specifically the role of cookie acceptance. We examine the role of cognitive factors and social nudges to influence the acceptance rate of such cookies. Understanding these behaviours is important to our methodology and were highlighted through our pilot study and subsequent interaction with the CISO.

The chapter is organised as such: Section 5.1.1 introduces our investigation and problem domain. Section 5.2 identifies our aims and scenario with reference to our methods and investigative procedure, section 5.5 highlights our results and subsequent analysis of these, section 5.7 concludes our findings with reference to literature, limitations and future work.

5.1 Cookie Study

Within this chapter we examine the effect of framing instructions and social norm nudges with respect to the acceptance of cookies. Through this, we understand how personality influences such responses through an assessment of risk-taking, impulsiveness, sociability, and the users willingness to self-disclose information about oneself. Through the experimental design process highlighted here, we gain a more detailed understanding of how users make privacy decisions. This in turn, enables a more expressive modelling environment where we can more accurately represent the users within our system.

5.1.1 Introduction

The *privacy paradox* declares that users value their privacy, but do not readily protect it [210]. This is often due to a combination of factors such as: being unaware as to the numerous methods in which they leave themselves vulnerable online, their inability to protect themselves, the effort required to preserve such privacy, or they may readily accept a trade in privacy for improved productivity and personalisation (sharing details to improve shopping experiences etc.).

From an organisational standpoint, privacy is an important issue. To incorporate this, organisations must develop tools and solutions that employ privacy preservation. The HCI development cycle must adopt a *privacy by design* approach [43] to ensure that privacy is a core component of the final design outcome. Alternatively, it is possible to augment final systems (systems that are already implemented) through the addition of new functionality, forcing mechanisms and also choice architectures that aim to nudge users towards desired behaviours.

Nudging in this fashion has been extensively documented in the *economics of privacy* [8] where users' privacy decisions change with respect to time as social and economic benefits present themselves. The implication of this means that users are susceptible to framing effects and personality variables where risks can be presented as more or less salient. In essence, "what people decide their data is worth depends critically on the context in which they are asked - specifically, how the problem is framed" [8].

With this knowledge, we examine some of the behavioural and contextual influencers that affect the privacy decision making process with specific reference to the acceptance or rejection of cookies. We wish to 'nudge' user choice through presentation effects to determine which interventions are most effective. In doing so we examine:

- Personality effects: The influence of different behavioural traits.
    - Impulsiveness - Making decisions without having all the facts at hand [78].
    - Risk taking: Those who take risks are less concerned with the consequences; those who are more risk-averse often regret taking risks [263]. Risk-takers often make decisions in the moment and are less influenced by emotions.
    - Willingness to self-disclose: A greater willingness to self-disclose (i.e. through social media) may expose individuals to privacy and security risks as they do not fully consider the consequences [4].
    - Sociability: a form of extroversion complimenting a greater interest in relationship management and information sharing. Sociable individuals are often more vulnerable to social influencing [146].

- Cognitive framing effects: The impact of the current task on the decision making process. It does not include the detailing of new information and instead relies on the way existing information can be manipulated in order to influence change (e.g. showing the benefits and risks associated with a particular choice [233]).
- Social framing: The role of social norms in our decision making process.
  - Social conformity: "if we engage in behaviours of which other approve, others will approve of us, too" [56]. An increase in the likelihood to follow social norms.
  - Adaptation of behaviour: Users may adapt their behaviour to match that of those they care about or trust [72]. People try to behave in a manner they believe to be appropriate to their social setting [77].

5.1.2 The Role of a Choice Architecture

From chapters 2 and 4 we have understood the importance of a choice architecture and demonstrated its application. For privacy, a similar strategy may be applied that recognises that users have the freedom to readily trade their privacy for a range of benefits they see fit [8] even if this is seen to hamper productivity in cases such as reading additional terms and conditions. Such a case is witnessed with respect to the following UK government report:

> ‘One key aspect of the use of social media data is the tension that exists between the generation of data by individuals and the use of that data by organizations. We have not been convinced that the users of social media platforms are fully aware of how their data might be used and what redress they may, or may not have if they disagree with how an organization exploits that data. This is exacerbated by our finding that terms and conditions contracts are simply too long and complex for any reasonable person to make any real sense of’ [175].

Within this study we examine how the adoption of a choice architecture can influence cookie acceptance or rejection. Through adopting an *economics of privacy* style approach, this enables an assessment of known benefits (e.g. personalisation) as well as known costs (targeted advertising and undesirable profiling). We acknowledge that users often default to habitual behaviours with respect to accepting cookies, often unaware of the implications of their actions. As such, users typically do not read the associated text which is presented alongside the choice and ultimately have little idea who has access to or can use the information that they are agreeing to share. This phenomena is known as the *control paradox* [7] and is often exacerbated by the user’s lack of understanding as to the means that exist to protect one’s privacy [5].

5.1.3 Cookies

A cookie is a small file that is designed to ‘hold a modest amount of data specific to a particular client and website’. This cookie is subsequently accessed by the client computer and web server in order to negotiate authentication on a particular website and allow it to create tailored content specific to the verified user. Common applications of such a process include ‘remember me’ logins and a plethora of retail and social media based websites.

Cookies are promoted as being necessary to enhance the user experience by enabling the streamlining of access (automatic login and authentication), personalisation, and targeted advertising. The issue of privacy has long been discussed with reference to cookie acceptance as there is an unclear notion of anonymity that pertains to their acceptance. This is especially evident with the ability to track user behaviours and browsing patterns through the monitoring of these files [177]. Therefore, although cookies cannot contain malware or viruses, it may be possible to intercept these cookies and use them for malicious purposes such as masquerading as somebody else. A man-in-the-middle attack on an open, unsecured network (see chapter 4) would present an ideal opportunity for such an attack provided any encryption methods are successfully decrypted. As cookies provide user identification and an associated token, it is possible to gain access to specific accounts. Moreover, it is important to stress that the security of the cookie is also a component of the target web server, the users computer as well as the browser itself, domains which the user may no influence or governance over.

The scale of adoption and the continued concerns regarding the use of cookies led the EU to introduce a directive in 2002 that required online providers to seek consent from users with respect to cookie usage. This later became law in 2011. Importantly, what started off as a somewhat explicit consent has been seen to shift to a more implicit choice where users are offered little in the way of actual choice regarding the manner in which their privacy is handled. The problem with this is that users become habituated to the processes of rapidly dismissing prompts in an effort to minimize productivity loss through the addition of extra user input requests. This can have knock-on effects with respect to additional security behaviours such as error-reporting (chapter 6) and the dismissal of software updates through ‘update later’ functionality [264].

5.2 Research Aims and Scenario

In this experiment, we wish to examine the following research questions:

- To what extent can contextual cues pertaining to the privacy and security implications of a task influence the decision-making process?

- Can social framing effects influence this choice?
- Do personality factors affect how users are likely to respond in a privacy context, and how does the implementation of a nudge within the choice architecture impact this?

To address this, we examine the following hypotheses:

1. Cognitive framing: Participants who receive instructions that make reference to security concerns are less likely to accept cookies than participants who receive no such reference to security [128].
2. Social framing: In comparison to a base rate (control) which is likely to lean towards acceptance of cookies, participants' responses will be affected by a social nudge, such that a low social norms reference will effectively nudge participants away from accepting cookies [128].
3. Personality: Cookies are more likely to be accepted when individuals are more impulsive and willing to share information, when they are greater risk takers and more sociable [128].

5.2.1 Study Design and Task

The study involved six experimental conditions in a 2*3, independent groups factorial design with 2 cognitive framing conditions (security vs convenience) and 3 social framing conditions (low, high and social norm).

With respect to our security frame, we presented participants with the following statement "We would like you to tell us what security/risk factors you consider when utilizing such websites" [128].

For our non-security frame we asked "We would like you to tell us what factors (e.g., convenience) you consider when deciding to go shopping using different sites" [128].

For social framing we augmented the text in the dialogue box. The general text in all cases was identical and read as follows: “Our use of cookies. This cookie stores basic user information on your computer, potentially improving the browsing experience and helping us deliver more relevant information to you" [128]. The control condition had no further text, whereas the other two conditions made reference to a minority and majority social norm as follows: "37% (74%) of MTURKers like yourself have used this option". All cookie dialogues concluded with the following: "Do you want to use this option? Accept/Don't accept". The allocation of each cookie was randomised from the possible 6, and an example

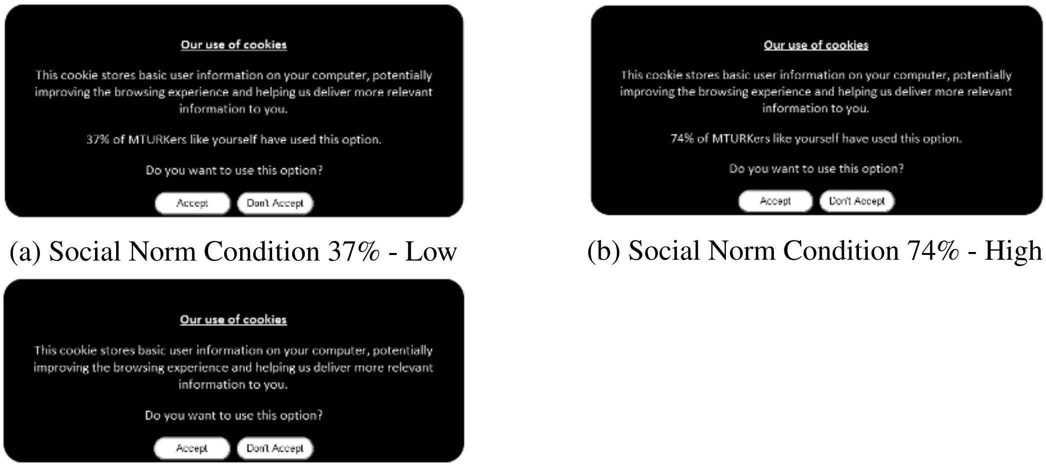

Our use of cookies

This cookie stores basic user information on your computer, potentially improving the browsing experience and helping us deliver more relevant information to you.

37% of MTURKers like yourself have used this option.

Do you want to use this option?

(a) Social Norm Condition 37% - Low

Our use of cookies

This cookie stores basic user information on your computer, potentially improving the browsing experience and helping us deliver more relevant information to you.

74% of MTURKers like yourself have used this option.

Do you want to use this option?

(b) Social Norm Condition 74% - High

Our use of cookies

This cookie stores basic user information on your computer, potentially improving the browsing experience and helping us deliver more relevant information to you.

Do you want to use this option?

(c) Social Norm Condition - Social Norm

Fig. 5.1 Example of Cookie Conditions - Security & Convenience [128]

of each is included in figure 5.1. Please note, these cookies were used in both cognitive framing conditions (security and convenience).

The descriptive social norms (37% and 74%) are in acknowledgement of Glynn et al [100] who used the same figures in a study of similar context and found these values to effectively demonstrate significant differences in responses. Similarly, group references have been shown to be effective in behaviour change in Terry & Hogg [227], providing the rationale for our choice of phrase "mTurkers like yourself". As such, direct referencing in this fashion is expected to increase potential conformity to social norms.

The task itself required participants to rate the trustworthiness of four different accommodation booking websites. Participants were instructed as follows: "We will present four different screen shots to you. Please have a look. Below each screen you will find five questions on whether or not you consider these trustworthy and would consider using them to book your next trip" [128]. Of the five questions, three related to trust, one investigated familiarity with the site, and the final question related to the probability of using the site in the future. Figure 5.2 provides an example of one of the web pages and subsequent questions.

Of the websites selected for our study, two were of high familiarity, and two were of low familiarity as reported in Skift [207]. These were 1) Booking.com, 2) Tripadvisor.com, 3) Bookingbuddy.com, and 4) Airbnb.com respectively. With respect to figure 5.2, the screen shot featured the home page of the web site and it was specifically designed to fill the screen to emulate visiting the official website.

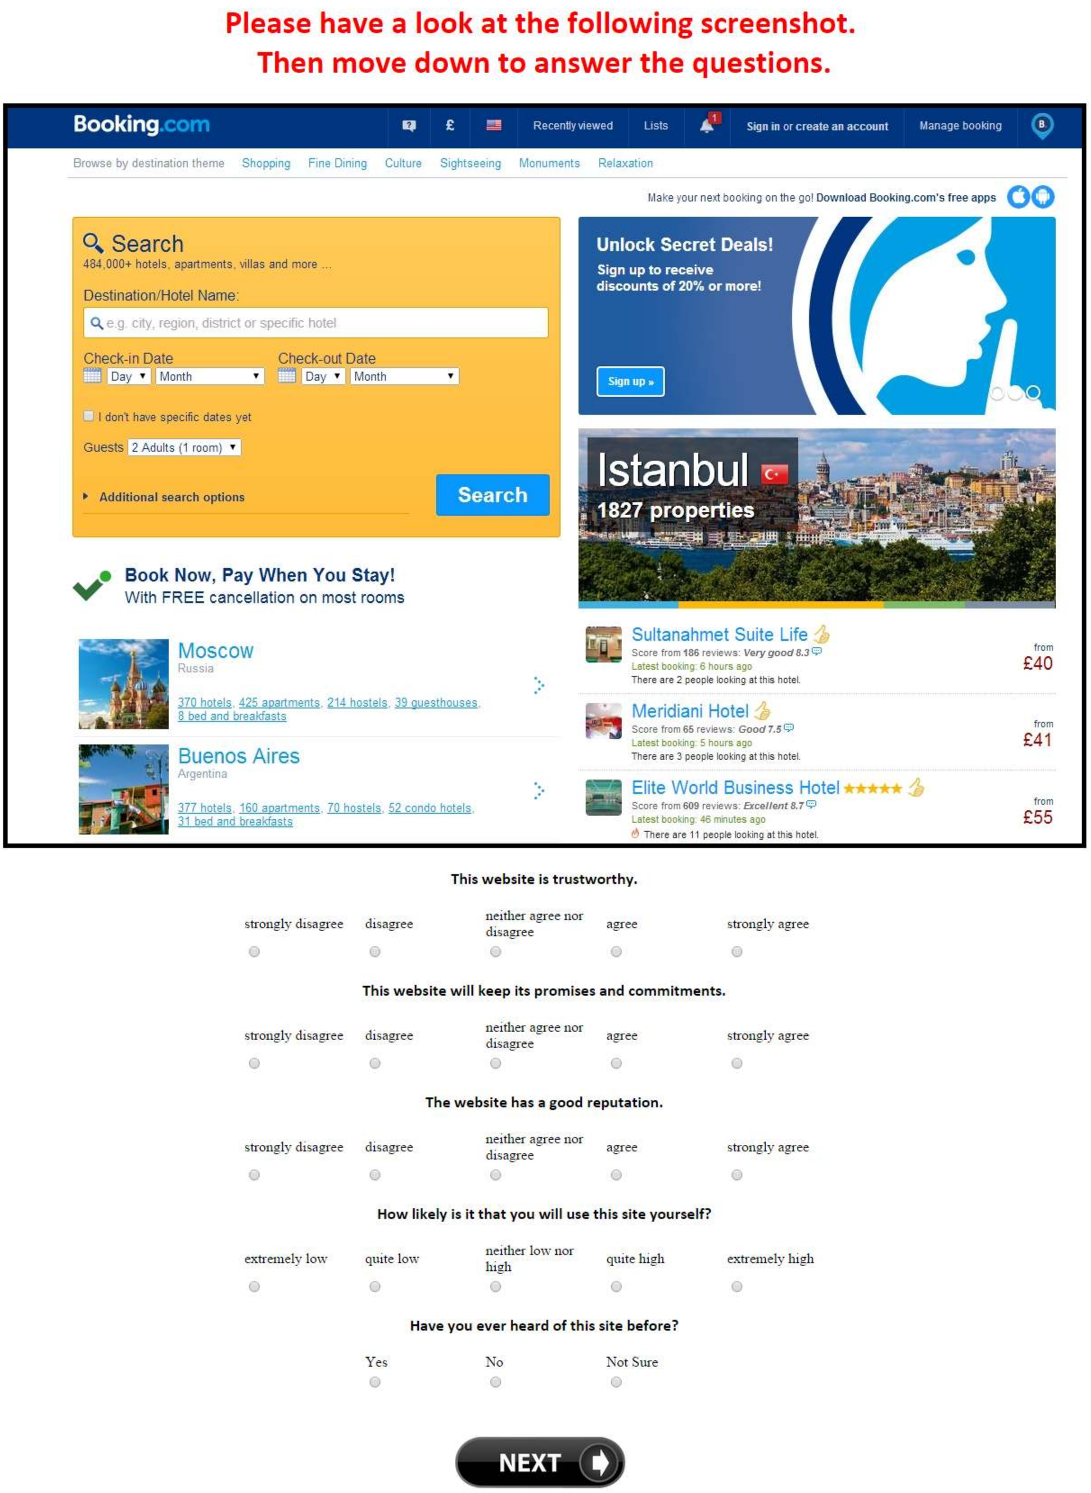

Please have a look at the following screenshot.
Then move down to answer the questions.

Make your next booking on the go! Download Booking.com's free apps

Search
484,000+ hotels, apartments, villas and more ...

Destination/Hotel Name:
e.g. city, region, district or specific hotel

Check-in Date: Day ▾ Month ▾
Check-out Date: Day ▾ Month ▾

□ I don't have specific dates yet

Guests 2 Adults (1 room) ▾

Additional search options

Unlock Secret Deals!
Sign up to receive discounts of 20% or more!

Istanbul
1827 properties

Book Now, Pay When You Stay!
With FREE cancellation on most rooms

Moscow
Russia
370 hotels, 425 apartments, 214 hostels, 39 guesthouses, 8 bed and breakfasts

Buenos Aires
Argentina
377 hotels, 160 apartments, 70 hostels, 52 condo hotels, 31 bed and breakfasts

Sultanahmet Suite Life
Score from 186 reviews: Very good 8.3
Latest booking: 6 hours ago
There are 2 people looking at this hotel.
from £40

Meridiani Hotel
Score from 65 reviews: Good 7.5
Latest booking: 5 hours ago
There are 3 people looking at this hotel.
from £41

Elite World Business Hotel ★★★★★
Score from 609 reviews: Excellent 8.7
Latest booking: 46 minutes ago
There are 11 people looking at this hotel.
from £55

This website is trustworthy.

strongly disagree	disagree	neither agree nor disagree	agree	strongly agree
○	○	○	○	○

This website will keep its promises and commitments.

strongly disagree	disagree	neither agree nor disagree	agree	strongly agree
○	○	○	○	○

The website has a good reputation.

strongly disagree	disagree	neither agree nor disagree	agree	strongly agree
○	○	○	○	○

How likely is it that you will use this site yourself?

extremely low	quite low	neither low nor high	quite high	extremely high
○	○	○	○	○

Have you ever heard of this site before?

Yes	No	Not Sure
○	○	○

Fig. 5.2 Example of the Website View and Questions

5.3 Method

For this study we the Mechanical Turk Platform (mTurk). The rationale for this decision is two-fold. Firstly, many of the users of this site are familiar with online browsing and cookie acceptance procedures, and secondly, these participants are predominantly driven by financial gains and optimal task completion times (a process similar to individuals working online in other tasks). As such, users will be driven by their primary goal of task completion which provides an ideal opportunity to explore how a choice architecture can be applied with respect to privacy and security nudges.

The experiment featured a total of 309 participants of which 19 were removed for repeat participation, errors in data or omission of data fields. The average age of participants was 35 years (MN=35.30, SD=11.96), with an age range of 18-71 years. Gender representation was almost equal with females (53.1%, n=154) and males (45.9%, n=133). Three participants did not wish to disclose their gender.

5.3.1 Participant Procedure

All participants accessed the study through the mTurk portal and were given $1 in payment for their time (a figure determined to be in-line with other studies of similar duration). After completing a consent form, the users were taken to the study site and randomly allocated one of the six conditions (1 of 2 conditions, and 1 of 3 cookies). Regardless of their response to the cookie acceptance, all participants viewed and completed the same content. First, they were asked to rate the trustworthiness, familiarity and intention to use the current site and secondly, participants were were presented with a questionnaire to asses their compulsivity, self-disclosure, risk-taking and sociability. Formal control questions and demographics were presented last along with a statement detailing that no actual cookie was downloaded by the user (we use Session Variables in the browser to call mySQL functions).

5.4 Results

The results section is formatted thus; first we outline the task results (the rating of trustworthiness) and secondly, the items and scales detailed in the follow-up questionnaire. The responses to all items have been 'combined into one score for each scale and the composite subsequently mean-centred, creating a scale measure with the same range as the original items' [128].

- Cookie reporting: More participants accepted the cookie than rejected it - 201 (69.3%) to 89 respectively.
- Trust/task-related questions: Trust was measured using three questions adapted from Lynch et al [155] to assess the trustworthiness of the websites, an example of which is "This website is trustworthy" with answers ranging from "Strongly agree to Strongly disagree" (see figure 5.2). The reliability and scale statistics for each of the four tasks are reported separately. The reliability was generally high (above 0.7 [172]): Task 1: α=.90, M=3.65, SD=.71 booking.com; Task 2: α=.86, M=4.26, SD=.56 Tripadvisor.com; Task 3: α=.86, M=3.33, SD=.61 Bookingbuddy.com, Task 4: α=.89, M=3.40, SD=.73 Airbnb.com.

Following our trustworthiness questions, we asked “How likely is it that you will use this site yourself?” with responses ranging from "(1) “extremely low” to (5) “extremely high.” and “Have you ever heard of this site before?” with the available options "Yes, (2) No, (3) Not sure" (see fig 5.2). A detailed discussion of the results from these questions follows in the next section.

The following questionnaire (see fig 5.3) included several measures to assess the users impulsiveness, risk-taking, self-disclosure, sociability, and demographics. Upon completion, participants were able to progress to the debrief and final page of the survey. The instructions of the follow-up stated: “Thank you for completing the task. We would hereby like to ask you a few questions about how you generally make decisions and share information about yourself.”

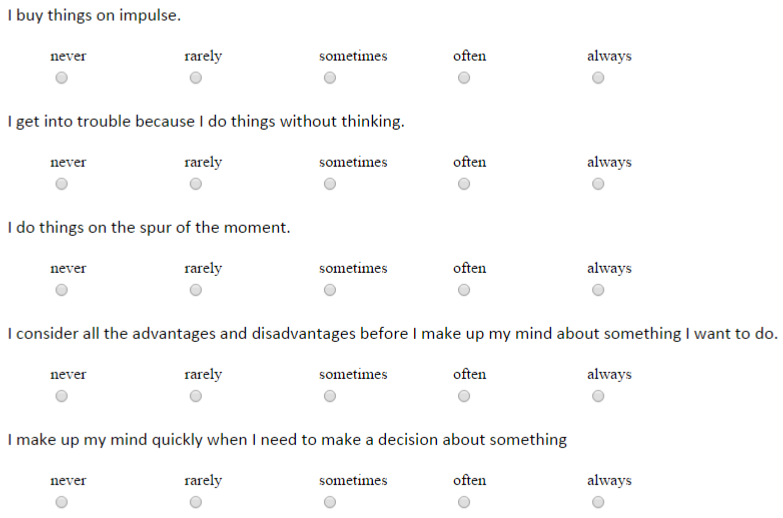

I buy things on impulse.

never	rarely	sometimes	often	always
○	○	○	○	○

I get into trouble because I do things without thinking.

never	rarely	sometimes	often	always
○	○	○	○	○

I do things on the spur of the moment.

never	rarely	sometimes	often	always
○	○	○	○	○

I consider all the advantages and disadvantages before I make up my mind about something I want to do.

never	rarely	sometimes	often	always
○	○	○	○	○

I make up my mind quickly when I need to make a decision about something

never	rarely	sometimes	often	always
○	○	○	○	○

Fig. 5.3 An Example of Part of the Website Questionnaire

For reference, we will take each in turn:

1. Impulsiveness: This was measured using a five-step sub-scale as witnessed in Eysenck et al [90]. Instructions asked participants: "Please tell us how you tend to go about making decisions on a day-to-day basis." We changed the questions into personal statements. An example item is: "I buy things on impulse." Each item included with five response options that captured the frequency with which participants engaged in the behaviours, ranging from (1) "never" to (5) "always." The fourth item was reverse-scored while the fifth items was excluded (α=.74, M=2.27, SD=.58) [128].

2. Risk-taking: We measure risk using a scale adopted by Dahlbaeck [68]. An example item is "I think that I am often less cautious than people in general." The answering options ranged from (1) "strongly disagree" to (5) "strongly agree." The third item was reverse-scored (α=.62, M=2.31, SD=.73) [128].

3. Self-disclosure: Measured using four items from the International Personality Item Pool (IPIP) and Wheeless [251]. The instructions were as follows: "Please tell us to what extent you share personal information about yourself on a day-to-day basis." An example item is: "I share and express my private thoughts to others." The answering options ranged from (1) "never" to (5) "always" (α=.77, M=2.93, SD=.60) [251].

4. Sociability: Measured using four items from Zuckerman-Kuhlman-Aluja Personality Questionnaire [16] and retaining the original instructions. An example item is: "I have a rich social life." The response options ranged from (1) "strongly disagree" to (4) "agree strongly" (α=.85, M=2.47, SD=.73) [128].

5. Control question and demographics: These included: "Do you normally accept cookies on websites?" Participants could select either "Yes" or "No". Demographics were also collected, such as age (including an option "prefer not to say") and gender (including an option "prefer not to say") in order to describe the sample in later analyses [128].

5.5 Discussion

The discussion is separated into the following respective sections: Firstly, we examine the data relating to cookie acceptance, and cognitive and social framing; Secondly, we detail the impact of personality effect, before briefly considering some of the incidental findings around personality and trust in the accommodation booking task.

5.5.1 Cookie Acceptance and Cognitive Framing

We predicted that participants who received instructions framed with security concerns would be less likely to accept the cookie than participants who received instructions framed with no reference to security (see Table 5.1). In terms of the number of participants in each of the experimental conditions, 134 were in the non-security domain and 156 participants received the security-related instructions. Chi square analysis showed that the cognitive framing manipulation did not influence the likelihood with which participants would choose to accept or reject the cookie (x^2(1)=2.446, p=.118), i.e. the security framing hypothesis was not supported [128].

	Cognitive Framing Conditions		
Cookie	Non-Security Frame	Security Frame	n
No (rejected)	35	54	89
Yes (accepted)	99	102	201

Table 5.1 Cookie Responses to Framing Instructions [128]

5.5.2 Cookie Acceptance and Social Framing

We predicted that cookie acceptance would be responsive to social framing effects believing that people will be less likely to accept the cookie if they believed that similarly-minded people were also rejecting it. Our results support these predictions. A Chi-square analysis revealed a statistically significant effect across the three conditions (x^2(2)=22.153, p<.001). The Cramer's V statistics result (Cramer's V=.276) also indicates that there is a moderately strong relationship between cookie framing and subsequent acceptance (range: .25-.30). Cookie acceptance across the three social conditions tends to increase, as shown in Table 5.2.

	Social Framing Conditions			
Cookie	Control Conditions	Minority (37%)	Majority (74%)	n
No (rejected)	19	51	19	89
Yes (accepted)	69	57	75	201

Table 5.2 Cookie Responses Based on Social Group References [128]

Additional pairing analysis (37% and 74%) was conducted on the results within table 5.2 highlighting that the rejection rate was higher when participants believed that their peers had higher rejection rates (x^2(1)=16.191, p<.001; Phi=.283). This suggests that the minority accepts frame moved participants away from cookie acceptance. Importantly, however, there was no statistical difference between the control group and the majority accept group (x^2(1)=.052, p=.819). This finding is not unexpected as default settings typically require participants to accept cookies.

We depict the ratio of observed and expected counts across all three group conditions in figure 5.4. If observed and expected values were equal, all bars would lie between the 1-2 range. The truth, however, is that there are two separate trends. Several bars fall below the 1 value reflecting a trend where the observed count was significantly lower than expected. In contrast, those values over 1 represent conditions where the count was significantly higher than expected. This leads us to conclude that cookies were rejected more often when presented with the smaller percentage (37%). Cookie rejection rates were lower for both control condition and the majority condition (74%) so these did not visually differ.

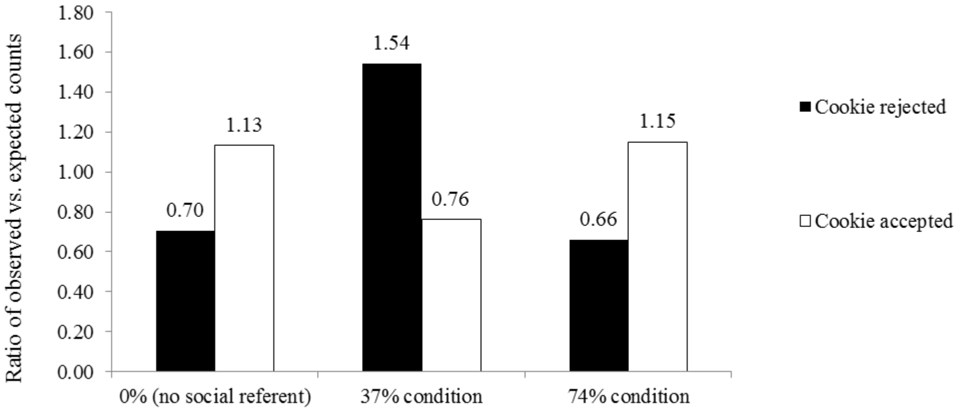

Fig. 5.4 Ratio of Observed vs Expected Counts

5.5.3 Cookie Acceptance and Personality Effects

We proposed that cookies were more likely to be accepted than rejected for individuals with high impulsivity, sociability and willingness to share information. This hypothesis was examined using an analysis of covariance technique with control for age and gender. We visualise this in table 5.3.

Cookie	Impulsiveness M (SD)	Risk-taking M (SD)
No (rejected)	2.13 (.51)	2.09 (.69)
Yes (accepted)	2.32 (.59)	2.41 (.73)

Table 5.3 Cookie Responses as a Function of Impulsiveness and Risk-Taking [128]

[a] The cookie was rejected by 89 participants, whilst 201 accepted it.

As hypothesised, cookie acceptance or rejection was affected by impulsiveness (F(1,288)=6.353, p=.012, partial n^2=.02) and risk-taking (F(1,288)=11.660, p=.001, partial n^2=.04). Furthermore, those who accepted the cookie also reported higher levels of impulsiveness and risk-taking. The figures also suggest a moderate effect from personality (.01 and .06). Analysis of cookie acceptance and rejection based on different norms produced no further significant results. Additionally, no significant observations were noted in relation to either willingness to disclose information (F(1,288)=.032, p=.857) or sociability (F(1,288)=.322, p=.571) and as a result, only partial support was obtained the our personality hypothesis (chiefly, risk-taking and impulsiveness).

5.5.4 The Booking Task

Numerous analyses focused on the trust scores for the online booking sites, with four trust measures (one for each task) being positively correlated with each other as expected (p<.001). Trust reported for each respective website was positively correlated with the intended use of the first (r=.723, p<.001; n=274), second (r=.594, p<.001; n=289), third booking site (r=.663, p<.001; n=289), and the fourth booking site (r=.594, p<.001; n=289). The correlation matrix illustrating the relationship between personality and trust in the online booking sites is provided in Table 5.4 [128].

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
(1) Trust_t1	1								
(2) Trust_t2	.291**	1							
(3) Trust_t3	.433**	.210**	1						
(4) Trust_t4	.232**	.270**	.279**	1					
(5) Impulsivity	.011	-.006	-.52	.28	1				
(6) Risk-taking	.10	-.34	-.32	.185**	.589**	1			
(7) Self-disclosure	.060	.073	.129*	.192**	.106^{t}	.114^{t}	1		
(8) Sociable	.054	.083	.056	.181**	.060	.193**	.264**	1	
(9) Age	-.094	-.136*	-.152**	-.273**	-.169**	-.177**	-.122*	-.81	1

Table 5.4 Correlations Between Trust Responses in Relation to Hotel Booking Sites, Age, Personality [128]

[a]** p<.01, * p<.05. ^{t} indicates marginally significant correlation (p<.10)

We acknowledge that the tendency to accept or reject cookies was not associated with differences in trust across all tasks (F(1,285)=.917, p=.339), or for any individual task. No significant differences emerged in relation to cookie responses task 1 (F(1,288)=.162, p=.687), task 2 (F(1,288)=.006, p=.936), task 3 (F(1,288)=.257, p=.613), or task 4 (F(1,288)=.588, p=.444).

Following this, we wished to investigate whether our cognitive framing manipulation also affected website trust values. There was no evidence obtained that suggested that security or non-security framing was associated with different levels of trust across all four tasks (F(1,285)=.002, p=.969), nor was there any significant effect for any of the individual booking tasks: task 1 (F(1,271)=.434, p=.511), task 2 (F(1,286)=.023, p=.880), task 3 (F(1,286)=.546, p=.460), or task 4 (F(1,287)=.051, p=.821). To summarise, there was no evidence to support the claim that the security vs. non-security nudge affected trust for each of the websites.

Our final analysis investigated website familiarity and was examined with respect to trust (see also Table 5.5). We observed that trust values were higher for the two high familiarity sites than for the two low familiarity sites. Intended use was also higher when the booking site was declared more familiar as suggested by our online site ranking source [207].

		High Familiarity Condition		Low Familiarity Condition	
Site Reports		Task 1	Task 2	Task 3	Task 4
Trustworthiness		3.64 (.71)	4.26 (.56)	3.33 (.61)	3.40 (.73)
Intended Use		2.92 (1.06)	3.90 (.93)	2.67 (1.02)	2.62 (1.12)
Familiarity	Yes	126 (46%)	255 (90.7%)	39 (13.6%)	89 (31%)
	No	130 (47.4%)	26 (9.3%)	255 (78.4%)	189 (65.9%)
	Not Sure	18 (6.6%)	0 (0%)	23 (8.0%)	9 (3.1%)

Table 5.5 Responses to Websites With Respect to Familiarity Condition [127]

[a] The first two rows list the mean and standard deviation in brackets. Task 1: booking.com; Task 2: Tripadvisor.com; Task 3: Bookingbuddy.com, Task 4: Airbnb.com

5.6 Summary of Findings

Within this chapter we have investigated the impact of social and cognitive framing, as well as personality effects, on cookie acceptance. Throughout this process, we have examined cognitive effects through the manner in which the task is introduced, as well as social effects through the implementation of high, low norms in relation to a control that provided no social information. In doing so, we have also obtained several personality characteristics such as impulsivity, sociability, self-disclosure and risk-taking.

By implementing our cookie discretely into the online survey, it has been possible to observe actual rather than intended behaviours. This is important as intentions in security do not necessarily translate into behaviours within security [65]. It is necessary to discuss these findings in relation to our hypotheses.

5.6.1 Hypothesis 1 - Cognitive Framing

Even with our holistic approach to cookie inclusion, our first hypothesis, with respect to cognitive framing, was not supported. We proposed that when individuals are made explicitly aware of the cookie with respect to privacy concerns, the adoption rate would fall as users would question the trustworthiness of the website. There are a number of reasons why this may not have occurred.

1. Users did not fully understand the privacy and security concerns related to cookie acceptance [4], nor did we highlight any specific negative outcomes of accepting it [262].

2. As proposed by Levin et al [149], the impact of persuasive messages depends on what kind of consequences the user may expect when they do not perform an action. As our task did not actually require users to purchase anything, they may not have interpreted the cookie as a threat to their security or privacy. Furthering this, our study was highlighted to be part of a university experiment and therefore users may have placed trust within our actions not to deceive them.

5.6.2 Hypothesis 2 - Default Control Behaviour

Our second hypothesis predicted that cookie acceptance is a normative behaviour, i.e., more people would accept the cookie than reject it. As previously highlighted, EU law now dictates that websites must explicitly disclose their use of cookies which has led to a change in the default behaviour towards acceptance. A possible impact of this can be observed in our results where the social nudge aimed towards increasing cookie acceptance had no measurable impact. Das et al [72] predicted, however, that social references can in fact shift cookie acceptance and rejection rates introducing some fuzziness with respect to their impact. On the basis of our research, however, we can confirm the hypothesis.

We can also confirm this conversely when a social nudge was introduced to highlight that others were less likely to accept the cookie. Such a statement led to a higher rejection rate in our participants. This was therefore in line with previous studies such as Stok et al [216] which stated that social identity is a salient feature and can affect attitudes within computer-mediated communication.

5.6.3 Hypothesis 3 - Personality & Decision-Making

Our third hypothesis reflected on the role of personality within the decision-making process stating that users would be more likely to accept if they were 1) more impulsive, 2) were greater risk-takers, 3) more willing to share information and 4) more sociable. This was partially supported for two of the four traits where we determined that the more impulsive individuals were more likely to accept cookies, which also fits with findings that participants are less likely to deliberate on their options (see Halpern [108]). With respect to our implementation, the manner in which the cookie appeared unexpectedly when individuals wanted to start may also have tapped into the distractibility aspect of higher impulsiveness (see Stanford et al. [212]).

Contrary to our predictions, there was no evidence to suggest that cookie acceptance was related to either self-disclosure or greater extroversion (sociability). It is possible that the association between self-disclosure and privacy concern is only found when the primary task is made to more explicitly reflect the sharing of sensitive information. This suggests that context-dependency plays an important role, as different environmental cues may influence the disclosure of private information (seen in Acquisti [8]; Joinson et al [132]). Moreover, the fact that the cookie was presented as part of a task in a separate section of the survey, rather than the questionnaire on self-disclosure, may have reduced participants' recollection of the cookie. This could mean that the observed results may be due to the online setting of the study and the fact that self-disclosure and information were not made salient [128].

Our results compliment similar findings within Nurse et al [174] where it is stated that is imperative to design systems and communications in a way that caters towards individual strengths. Personality traits may subsequently shape these strengths and weaknesses. From this, along with our results, we can conclude that impulsiveness and risk-taking may present potentially unfavourable traits in security-related decision making.

5.6.4 Limitations

With respect to our findings, we highlight the following limitations within the study:

- Nudges: it is difficult to assess the longevity of our nudge interventions as they impact the system 1 thinking process (see Mols et al [164]).

- Sensitivity to risk: Those who accepted the cookie also reported higher risk-taking tendencies. It is therefore possible that our sample was largely insensitive to security risks (see Druckman [86]).

- Group identification: It is difficult to ascertain group affiliation and linkages between mTurkers. This may have caused us to overlook some of the greater persuasive effects.

5.6.5 Practical Recommendations

This chapter documents the continued development and application of nudges within a choice architecture in the field of security and privacy decision-making.

Our results support several principles and recommendations, chiefly:

- Our results support that Nudges are most effective when placed just before the moment of decision-making.
- If a nudge is to be based on social norms, they should be credible and based on empirical evidence (Stok [216]). Furthermore, these nudges should be clear with respect to their direction of intended behaviour change (Das [72]).
- The most effective choice architecture will consider personality characteristics (e.g., impulsivity and risk-taking) as these may influence the effectiveness of nudges.
- Numerous theoretical frameworks exist which may, in combination with evidence-based research, provide useful starting points for future nudge-based interventions [128].

5.7 Conclusion

The current research contributes to current knowledge by building on the work by Das and colleagues [2014], but furthermore expanding on this work by considering the role of personality. The effectiveness of social proof in this setting also builds on past evidence [see Glynn et al., 2009; Stok et al., 2012] but in a different context, which further cements the influence of others on decision-making.

5.7.1 Future Work

The listed contributions would benefit from repeat experimentation through increased sample size and closer examination of specific areas. Such an approach would enable the following:

- The investigation of affect and risk communication (detailed in Visschers et al [241]) which could provide further avenues of investigation in relation to IS.
- An examination of other individual characteristics that were not under investigation within this study such as agreeableness [150], gender [86], conscientiousness, and risk preferences.

- As discussed in Levin [150], further research into framing effects and the role of individual differences on decision-making may help us to better understand under what circumstances and with whom framing is effective. This is also relevant because it would allow practitioners to frame decisions to maximize impact.

5.7.2 Chapter Contributions

This chapter has contributed to our understanding of nudge applicability in privacy and security. Specifically, in the realms of online security and privacy with respect to the acceptance and rejection of cookies, a common procedure in day-to-day browsing.

The findings presented here not only support similar applications and results in other studies, but also expand them by applying a nudging framework to cookie acceptance. Our methodology benefits from being able to more accurately articulate this process and ultimately represent the problem space. This empowers the policy-maker to form a richer expression of our environment and the users that populate it, allowing for a more encapsulated approach towards policy creation and implementation.

Chapter 6

Understanding Security Behaviours - Part II

The previous chapter documented the testing of a nudge on privacy and security using two framing conditions. The following chapter highlights work conducted in Jeske et al [127] which aims to understand user behaviours with respect to error-reporting; a behaviour highlighted as 'of significant importance' from interaction with the CISO and results obtained from our pilot study. We examine how users form their decision-making process with respect to error submission reports. In doing so, we develop and test an intervention method based on social norms and nudges with respect to framing effects to investigate how we can modify the report and non-report rate. The application of this work will aid the understanding of how to build policies that increase the likelihood of error-reporting, ultimately improving security through the subsequent development of mitigation strategies.

Section 6 documents work conducted in Jeske et al [127] and aims to investigate the error-reporting process

6.1 Error-Reporting Study

User feedback is an important component of both security and software development, enabling the detection and response to current threats. Error reporting is an important feedback mechanism, but many users do not respond to error reporting requests; a trait acknowledge within our pilot study (see chapter 3). The following chapter presents an investigation aimed at improving the rate of error reporting by presenting 126 participants with different types of messages as part of an online study. Within our study, we asked users to rate four booking sites in terms of the sites' trustworthiness. The displayed messages differed in terms of how the issue was described (cognitive framing: messages triggered due to technical or security), and the benefits associated with the reporting of the error (benefit with respect to the user, and user to other user). The inclusion of the benefit statement is based on the literature of social loafing which suggests that redundancy may affect performance when individuals perceive themselves to be one of many attempting the same task. Our results revealed that the cognitive framing had a significant effect on whether or not participants reported an error intentionally when it presented itself. Error reporting was found to be lower when the message suggested a security issue rather than a technical problem. Furthermore, participants were more likely to report the error if doing so implied some self-benefit. Without reference to such benefits, error reporting was significantly lower. Our findings reveal that error framing has a significant effect with respect to the likelihood of reporting (security and technical). Additionally, the benefit statement may have made error reporting more relevant and more personally meaningful for the individual users, an important effect that is known to reduce social loafing in group settings.

6.2 Introduction

Anticipated and actual user behaviours often diverge, particularly in regards to their responsiveness to different system requests and software. Several theories explore why this might be the case. One such theory is based on the concept of social loafing (Karau [138]) that stipulates: when in the presence of many other users, the individual user may not react to a request that many other users will have already responded to. One explanation for this inactivity focuses on the perceptions of redundancy of effort, where any action from the individual user may be interpreted as a duplicate of inputs generated by others (particularly prominent in a collective working environment - see Karau [139]). This impacts the user as the apparent significance of any contributions that they may make is reduced, further lowering their motivation to engage with subsequent requests.

Error messages are a frequent occurrence when individuals interact with many forms of technology, often requiring users to submit report errors on any issues they encounter. These requests typically require users to decide whether or not they wish to respond to these requests. Despite their unpopularity, legitimate error requests serve a variety of important functions. For example, these reports enable programmers to identify and formulate solutions towards specific problems [21], [117], [112]) by using patches, improving design and quality of products to name but a few [21]. They also represent an important feedback mechanism between organizations and their users enabling awareness of potential errors that may affect their operations [265]. This is of particular significance to our CISO through continued interaction.

We must therefore determine how to most effectively assess this user behaviour utilising existing theories and frameworks. The following section details numerous reasons as to why error requests may not be addressed. In addition, we outline several suggestions on how to reduce this redundancy, of which some are subsequently investigated to evaluate their utility.

6.2.1 Understanding User Reactions to Error Requests

The presentation of error messages with respect to HCI is often met with an habitual response. This typically manifests with a feeling of annoyance from the user and the subsequent dismissal and rejection of the notification. This reaction is often due to a variety of reasons. Firstly, error reporting requests share characteristics with other prompts that are often unanticipated, potentially disruptive and, on occasion, indicators of an infection (e.g., updates and cookie requests - see chapter 5). Secondly, it is often unclear how doing so benefits the user and whether or not the error requests are generic or specific. Thirdly, the degree of effort required to report errors is often uncertain and undocumented within the message. Fourthly, many system errors are encountered repeatedly within day-to-day tasks, reducing the sense that a report will be useful or lead to some personal or overall benefit and thus contributing to non-reporting [117]. These factors are often exacerbated when the semantics of system settings and messages are unclear to the non-technical user, leading to misunderstandings and possible confusion [160]. Moreover, the standard user is typically unaware of security threats and vulnerabilities [95].

Some errors may be triggered by new software and the implementation of new processes. Having little prior knowledge pertaining to how a user should respond to errors may also influence an individuals' response to error requests. In this context, many users prefer to avoid or reject information when they have this option [152], [115]. This may also be the case when engaging with new information that may also disrupt the processes, having a particularly negative effect on subsequent but primary task performances [87]. These negative effects may differ depending on who is disrupted and during what kind of tasks (e.g., for those with low working memory capacity [85]) and the amount of cognitive load that users experience at the time [173].

Other issues related to error reporting include whether or not an individual believes that their actions or non-actions are identifiable [111]. Specifically, some individuals may not wish to report an error because they are unclear as to what information they are sending (e.g., the data they send may identify them). Furthermore, anonymity may lead individuals to not report errors, since they believe their inaction is not likely to lead to any negative consequences. These findings explain why error reporting has been linked to threat (mis)perception, security beliefs, perceptions of risky behaviour, and avoidance of security interventions [118]. Concerns about what is reported may be exacerbated when large corporations publicly acknowledge that such personally identifiable information (e.g.,Microsoft) is included in error reports, and this information is subsequently shared with their partners [224]. This anxiety may be reinforced by the difficulty that users have in assessing the credibility of different messages, which may then reduce the acceptability of the message and, over time, lead to a reduction in the error reporting rate. Some support for this is found in Workman et al. [257] who observed that perceived vulnerability, perceived response efficacy and response-cost benefit all predict omissive behaviours.

6.2.2 Redundancy of Effort and Perceived Personal Relevance

With respect to social loafing, there are a number studies from social psychological literature that may help to understand why the under-reporting of errors is so prevalent. Some research suggests that task characteristics may play a role in terms of whether or not individuals will be motivated to contribute and complete a task. Such task characteristics include the 'attractiveness' of the task [261] providing some insight into why many organizations invest in incentives to encourage employees to report errors. Another interesting factor is that complex tasks are also less likely to lead to personal engagement [124] as the unknown requirements involved in error reporting may dissuade users from doing so [194]. Thirdly, error requests do not represent creative tasks that are typically preferred with IT users [29]. Lastly, it is often unclear to the user how meaningful it is to report errors, a phenomenon similar to the redundancy effect noted earlier.

It is possible that there may be no immediate solution available to addressing redundancy concerns and the characteristics of error requests within HCI. While providing incentives for reporting errors in an organizational setting may be an option, incentives in other settings may not be as easily implementable or effective in persuading users to respond to error requests. Several generic suggestions, however, can be found within social loafing literature. For instance, error requests could be timed better (e.g., when these requests are not part of a critical malfunction) to reduce the likelihood that users will ignore these request in order to proceed with a primary task. This is in line with calls for designs that do not reduce performance [173]. While it may not always be feasible to make error reporting "attractive", it ought to be in the interest of developers and users to make them less difficult to comply with.

Another option is to consider the extent to which error reporting is due to perceived redundancy. Research has shown that information processing is motivated by the extent to which the information has personal relevance [181]. Thus, providing a rationale with respect to the error request context may alleviate the perceived redundancy, and in-turn, highlight the benefits associated with reporting errors [173]. This could be achieved by either personalizing the message or outlining the repercussions of leaving a potential issue unresolved over time. Personal relevance may further promote greater involvement, which has been shown to lead to more systematic information processing [54]. By increasing personal relevance, it may also be possible to reduce habitual responses that lead to participants to ignoring system-generated messages [42], [239].

These findings suggest there is a need to raise the perceived value of the task and individual input with respect to error reporting. From this, we need to determine how the expected value of error reporting can be influenced using evidence from social loafing in both the context of the user and the perceived value to other users [205].

6.2.3 Research Gap and Hypotheses

Our goal is to reduce the potential redundancy of effort and increase personal relevancy by clarifying the benefits of error reporting. Likewise, this chapter is a response to past calls for more research in this area [265], particularly research that addresses the lack of theories that could be used to explain why and how individuals may respond to error requests. Importantly, this necessity for further research was also indicated from our pilot study (chapter 3) and discussions with the CISO. The focus of this study is therefore to reduce the perceived redundancy and increase the meaningfulness of error reporting whilst identifying the role of error message type. A critical part of this is to understand how error framing, e.g., as a technical or security-related issue may trigger different responses by individuals.

Framing Hypothesis: 1 & 2

This study examined three hypotheses.

1. Firstly, we predict that the frequency of error reporting is higher for a problem that is framed as a security than a technical issue (cognitive framing effect).

2. Secondly, we predict that the frequency of error reporting is higher when the reporting has an implied benefit to the user. When users perceive a benefit to themselves or to others, they are expected to be more likely to report errors.

Individual Difference Hypothesis: 3

In addition to personal relevance, individual differences also affect technology interactions. Past evidence suggests that error-message interruptions can negatively impact the performance of users depending on their memory capacity and use of memory strategies [85]. Other traits such as perceived self-efficacy may also impact how often, regularly or for what reasons users may report errors they encounter [113]. In addition, impulsivity has been linked to greater distractibility, less deliberate and attentive decision-making as well as selective information processing [212].

3. We consider the role of individual differences in error reporting and predict it to be higher amongst individuals that rank highly on impulsiveness and risk-taking, low on privacy concern, and higher on expectations of reciprocity [59].

Privacy concern were included as they may capture fears pertaining to the type of information that is transmitted with an error, while reciprocity has been hypothesized to play a role in shaping contributions in settings such as computer-mediated discussion forums [247]. It is also expected that perceptions of security vulnerability may be associated with higher error reporting.

Underlying Motives for Error Reporting

In the final part of our results, we also consider the role of generic error responding in an exploratory fashion, exploring the extent to which these behaviours can be attributed to particular motives that either encourage or discourage error reporting.

6.3 Method

6.3.1 Study design and tasks

The primary task for users was to rate each individual sit; whilst the main focus for researchers was the secondary task which had been embedded into the error-reporting (first task). It is necessary to discuss this in-turn.

- Site rating task (primary): focused on how trustworthy participants would rate four different accommodation booking websites. These had been selected from a list of popular websites ([207], the same as chapter 5) and included (1) Booking.com (2) Tripadvisor.com, (3) Bookingbuddy.com and (4) Airbnb.com. All four screenshots presented the home pages of each of these sites (see 5.2). The instructions to participants informed them that they will be presented with four screenshots of websites and will be required to rate how trustworthy they find these sites. Following this assessment, they would additionally be requested to report how familiar they were with these sites and how likely they would be to use them in the future.

- Error reporting task (secondary): All participants were presented with a dialogue box that stated 'We noted a problem on this page'. This was followed by either a security or technical message frame. The security message frame informed participants that: 'This problem may indicate a security issue.' The technical message frame was almost identical, only one word was changed: instead of mentioning it as a 'security' issue, this frame referred to a 'technical' issue. The rest of the message differed depending on what kind of benefit condition participants were allocated to. In the benefit-to-self condition, the message stated 'Problem reporting will help us identify the source of the problem and protect you'. In the benefit-to-others condition, the message reported 'Problem reporting will help us identify the source of the problem and protect you'. No such message was presented to participants in the control condition. Each error message ended with a question: 'Do you want to report the problem? Report/Don't report'.

Overall, the error reporting tasks involved six experimental conditions based on a factorial design (2*3). This included two cognitive framing conditions (security vs technical message frame) and 3 implied benefit frames (implied benefit to self or others; control condition without benefit statement). All participants were placed in one of these conditions, but all would be asked to complete the two tasks. Allocation to the six conditions was randomized and the error message was presented as part of the first task. The message appeared after the third screenshot, just before the fourth (always the same screenshot for tripadvisor.com).

6.3.2 Participants

Recruitment involved a pool of university students situated from different departments (social and natural sciences). Information regarding the task was circulated via email and a dedicated university online recruitment portal. Students were highlighted as a relevant sample in this study as they tend to use many different online services throughout daily activities. In addition, they often rely on and make assumptions with respect to the security of the university infrastructure and their publicly available open-access services [140]. This may also make this group of students less informed with regards to the role and utility of error reporting. As an incentive for participation, all participants could earn research credits for their respective degree programmes. Participants were recruited between December 2014 and

March 2015 totalling 147 by the end of this date. Studies that had not led to full completion of the survey were excluded (n=19). This reduced the final dataset to n=126. Participants were 18 to 36 years old (M=20.15, SD=2.79, n=125). The sample composed of 84% female students (n=105, two missing values) and 16% males.

6.3.3 Procedure

Once participants had given their consent to participate they were presented with instructions to rate the trustworthiness of four online travel sites. Each of the four screen shots were presented separately with the first 3 being randomised. On the fourth screenshot, they encountered an error message that gave them the choice to report or not report the error. This was followed by a series of questions about their online use of travel sites, familiarity and review activity of such sites. This was followed by a number of short personality questionnaires. The questionnaire ended with demographics and the debrief statement about the study.

6.3.4 Measures

The next section describes the different experimental conditions that were employed, the screen-specific measures, and the follow-up questionnaire.

Screen measures: All four screenshots were presented individually. Each screen had to be rated on trust, potential use of this site, and familiarity with site.

1. Trust: each of the sites were assessed using three questions from Lynch et al. ([155], also used by Belanger et al [34]), an example of which is: ‘This website has a good reputation’. The five response options ranged from (1) “strongly disagree” to (5) “strongly agree.” The reliability for each site-specific trust measure was appropriate: Trust 1 (M=3.73, SD=.69, α=.89); Trust 2 (M=3.32, SD=.69, α=.82); Trust 3 (M=3.26, SD=.62, α=.819), Trust 4 (M=4.42, SD=.64, α=.89). All trust measures collected for each of the four booking sites correlated to differences between one another. As such, there was no evidence that trustworthiness was linked to any personality traits.

2. Use of and familiarity with the site: three questions on trustworthiness were followed by two items measuring the user’s intent to use the site. The first item asked participants ‘How likely is it that you will use this site yourself?’ Response options ranged from (1) “extremely low” to (5) “extremely high.” The second item asked: ‘Have you ever heard of this site before?’ Response options included: Yes, (2) No, (3) Not sure.

Follow-up questionnaire: The follow-up questionnaire included numerous questions with respect to the participants' engagement on the various travel sites, general error reporting, short personality questionnaires, demographics and the debrief information. Unless otherwise specified, all items in the personality measures were combined and divided to create a mean-based composite that had the same response range as the original items. Skew and kurtosis of the new scale composites were largely unremarkable and are hence not listed.

1. Review activity on sites: Participants were asked if they mainly read travel sites and/or reviews or were typically contributors (reviewers) to travel sites. The answering options included (1) "Yes" and (2) "No". In our sample, none of the participants reported reading travel sites. However, 103 participants stated that they reviewed travel sites.

2. General reporting tendency: All participants were prompted with respect to their general error reporting tendency in the follow-up (two screens after the presentation of the error message). They were asked: 'When you are notified about a problem, would you normally report it if prompted with a message?' The answering options included (1) "Yes" (n=43) and (2) "No" (n=82; 1 missing value). The next question was: 'When other people get notified, do you think they would report the problem'. The answering options included (1) "Yes" (n=34) and (2) "No" (n=90; 2 missing values). Each of these questions were followed up by the appropriate open question: 'Please tell us why you do (not)' and 'Please tell us why you think that others would (not) report problem'.

3. Impulsiveness: This was measured using five adapted items from the impulsiveness subscale by Eysenck et al. [90]. An example item is: 'I do things on the spur of the moment'. We presented each item with five response options that captured the frequency with which they engage in these behaviours and deliberations, ranging from (1) "never" to (5) "always." The fourth item was reverse-scored (M=2.76, SD=.64, α=.76).

4. Risk-taking: This was assessed using three items from Dahlbaeck [68]. An example item is 'I take chances in various situations'. The answering options ranged from (1) "strongly disagree" to (5) "strongly agree." The third item was reverse-scored. Only two (r=.538, p<.001) out of three items were used to create a composite as reliability for the three item was low (M=2.87, SD=.84, α=.70).

5. Privacy concern: This was measured using five items. The first three items came from a sub-scale by Dinev & Hart [80] which measured concern about abuse of personal data. The fourth item was proposed by the authors and the fifth item was taken from the perceived severity scale used in Workman et al [257] and Ifinedo [120]. An example

of the fifth item is: ‘Having my confidential information accessed by someone without my consent or knowledge is a serious concern for me’. The response options ranged from (1) “strongly disagree” to (4) “strongly agree” (M=3.01, SD=.58, α=.85).

6. Security vulnerability: This was measured using four items copied from the perceived vulnerability sub-scale used by Workman et al [257] and Ifinedo [120]. An example item is ‘The likelihood that my information and data is vulnerable to security breaches’. The response options ranged from (1) “extremely low” to (5) “extremely high” (M=3.01, SD=.77, α=.83).

7. Reciprocity: Reciprocity was measured using three positive reciprocity items by Perugini [180]. An example item is: ‘If someone does a favour for me, I am ready to return it’. The responses ranged from (1) “very untrue” to (7) “very true” (M=3.47, SD=2.04, α=.87).

8. Sociability: Sociability was assessed using four items from Zuckerman-Kuhlman-Aluja Personality Questionnaire [16]. An example item is: ‘I have a rich social life’. The response options ranged from (1) “strongly disagree” to (4) “agree strongly” (M=2.87, SD=.698, α=.73).

Participants were also asked to report their age (including an option “prefer not to say”) and their gender (including an option “prefer not to say”) in order to describe our sample in later analyses.

6.4 Results

6.4.1 Allocation to Cognitive and Benefit Framing Conditions

With respect to framing, 65 participants received the technical framing and 61 participants received the security framing. In terms of the implied benefit conditions, 41 participants were in the benefit-to-self condition, 37 were in the benefit-to-others condition. The control condition included 48 participants.

6.4.2 Error Reporting Across All Conditions

The main output variable of interest was whether or not participants chose to report the problem when prompted with the error message, with respect to the benefit and framing condition they were allocated to. Seventy-three participants (57.9%) did not report the error. However, 53 individuals did (42.1%).

		Error Reporting (n=125)					
		Don't Report (n=73)			Report (n=53)		
	Implied Benefit	to self (n=10)	to other (n=22)	none (n=41)	to self (n=31)	to other (n=15)	none (n=7)
Cognitive framing	Technical (n=65)	0(0%)	11 (16.9%)	19 (29.2%)	24 (36.9%)	6 (9.2%)	5 (7.7%)
	Security (n=60)	10 (15%)	11 (18.3%)	22 (36.7%)	7 (11.7%)	9 (15.0%)	2 (3.3%)

Table 6.1 Error Reporting Across All Conditions (Framing & Implied Benefit)

6.4.3 Cognitive Framing Effect on Error Reporting (Hypothesis 1)

Security Vs Technical

The x^2 statistic was used to examine whether being placed in specific conditions (framing and benefit) increased or decreased error reporting.

The first hypothesis predicted that the frequency with which individuals report the error would be higher when the message framed the problem as a security issue, rather than a technical issue. A significant difference was observed in terms of error reporting (x^2(1)=7.649, p=.006). Error reporting was higher when the problem was framed as a technical issue (obs/exp=35/27.3) rather than security issue (obs/exp. 18/25.7). The Phi statistics (Phi=-.246) also indicate that there is a moderately strong relationship between framing and error responses. Error reporting was lower when the problem message suggested a security issue. Hypothesis 1 was therefore not supported (see Table 6.2).

Condition	Counts	Don't Report	Report
Technical	Observed	30 (46.3%)	35 (53.8%)
	Expected	37.7	27.3
Security	Observed	43 (70.5%)	18 (29.5%)
	Expected	35.3	25.7

Table 6.2 Frequency of Error Reporting for Both Framing Conditions

6.4.4 Benefit Framing Effect on Error Reporting (Hypothesis 2)

Implied Vs None

The second hypothesis predicted that the frequency with which individuals report a problem is higher when problem reporting has a benefit to self (benefit effect). A significant differences was observed (x^2(2)=33.842, p<.001). Error reporting was higher when a benefit was implied, particularly a benefit to self (obs/exp. 31/17.2). In the absence of such a statement, error reporting was much lower (obs/exp. 7/20.2; see Table 3). The Cramer's V statistics (Cramer's V=.518) also indicate that there is a strong relationship between message contents and error responses. This provides support for hypothesis 2.

Condition	Counts	Don't Report	Report
Benefit to Self	Observed	10 (24.4%)	31 (75.6%)
	Expected	23.8	17.2
Benefit to Other	Observed	22 (59.5%)	15 (40.5%)
	Expected	21.4	15.6
Control	Observed	41 (85.4%)	7 (14.5%)
	Expected	27.8	20.2

Table 6.3 Error Reporting Across Different Benefit Conditions

6.4.5 Individual Differences and Error Reporting (Hypothesis 3)

Several personality measures were correlated with each other (see Table 6.2). Impulsivity correlated with risk-taking (r=.611, p<.01), sociability (r=.181, p<.05) and norms of reciprocity (r=.187, p<.05). Sociability also correlated with risk-taking (r=.294, p<.01) and norms of reciprocity (r=.184, p<.05). Privacy concerns correlated with perceived security vulnerability (r=.424, p<.01).

Hypothesis 3 proposed that error reporting was lower amongst individuals who score higher on impulsiveness and risk-taking. ANCOVA was used to examine this group difference in error reporting, also controlling for potential covariates such as gender, age, and general problem reporting tendency of participants (q3) and their perception of other people's error reporting (q4). No significant results were obtained for impulsivity (F(1,117)=.161, p=.689) or risk-taking (F(1,117)=.030, p=.863). Further analyses with privacy concern (F(1,117)=.727, p=.395), security vulnerability (F(1,117)=1.143, p=.287), sociability (F(1,117)=.146, p=.703) and reciprocity (F(1,117)=.621, p=.432) rendered similar non-significant results. Hypothesis 3 was therefore not supported.

6.4.6 Qualitative Analysis of General Error Reporting (Exploratory Analysis)

Of the 126 participants, 43 said they would report errors in general whilst 82 said they would not (1 missing response). In order to explore this further, it was necessary to examine the qualitative comments participants made with respect to what motivates their error reporting practices. Comments were available from 121 of 126 participants. These comments were collected to explore why participants would and would not typically report errors. Several different themes emerged.

The reporting of errors may have been influenced by a variety of factors. The reporting appears to be subject to the extent to which participants sensed a potential threat (fear appraisal), or lack thereof. The first group of participants actually felt that the error message represented a threat (as an indicator of a Trojan or virus). Participants reported, for example, that they do not report errors because: "I always feel like the message is a virus rather than an actual warning;" "[it] could be a virus;" and "in case it's a scam or a virus". The second group of participants did not perceive a threat and hence decided not to report an error citing reasons such as; "I have anti-virus software." The third group of participants accepted the error message as they determined that it informed them about a legitimate threat, thus complying with the request to report the errors.

Another factor was the perceived efficacy of responding to the threat (including the benefit of response). Those who did not report the error message suggested that this would have negligible or no significant benefits to one's self or others. Several commented that it would be "time consuming". A few additional quotes from participants who do not report errors reflect their logic for not doing so: "don't really think that it is important, but of a time effort". Not reporting errors may also be linked to uncertainty. For example, two participants reported that they were "not sure how it works" and "because I don't know what it means or what it is". In addition, some participants believed action would not lead to an improvement ("when reporting incidents in the past nothing has happened"). Another group of participants did not comply because they expressed a lack of information. That is, they may not have been knowledgeable enough to know what is required of them or where the information would end up. This is shown in citations such as "because I don't know what it means or where it goes". The third group of participants recognized the importance of error reporting not just for themselves, but others, leading them to comply with such error request generally: "To hopefully draw attention to the problem and ensure it is more likely to be fixed"; "to bring the problem to the attention of the website administrator so they can sort it out faster"; "to try and stop it from happening from again"; "because it may improve future services"; and to "improve site."

A third factor concerns the potential costs associated with error reporting (e.g., in terms of productivity costs incurred due to error reporting). Individuals who did not report errors were particularly attentive to the potential costs associated with reporting errors (including the time and effort involved). For example, one participant stated "makes it go away quicker if I say no". In addition, participants reported that "I just want to continue doing what I was previously and did not want to report an error because of the potential for disruption that may result in terms of “time and redirection". Participants who decided to comply with the error requests did not comment on the immediate cost to themselves, only recognizing that the error “needed telling” and that reporting it will ensure “it can be fixed” and “to ensure it is solved.”

6.5 Discussion

Perceived redundancy and the lack of personal relevance of individual contributions are important work characteristics known to increase the social loafing of individuals in larger groups [138], [139]. Applying this knowledge to the user on a far grander scale, the current study aimed to use some of these social psychological findings to promote error reporting by avoiding perceptions of redundancy that many users experience and subsequently report when encountering such error requests. The next session summarizes our results and relates these findings to the existing literature.

Our first hypothesis examined whether or not the cognitive framing of an error message influenced error reporting in the current study. In contrast to our predictions, error reporting was significantly higher when participants were presented with a technical, but not security, framed error message. The difference may be attributed to several different factors. The first explanation relates to concerns that participants may have had when the error is security-related. Participants may be more likely to opt against reporting such an error as the label ‘security’ may imply potentially punitive or other unknown serious consequences [137]. This may evoke protection motivation that discourages compliance with such error requests [113].

In addition, a security framing may raise concerns about perceived control over what is implied. If individuals feel they are not in control over the consequences, they may not engage in error reporting behaviour [15]. In essence, the issue may be perceived as less manageable by the users themselves. This may explain why these participants were more likely to decide for inaction (not reporting the error). A third and related explanation concerns the lack of knowledge and experience as a motivation for inaction. Whether or not users understand and know what to do can also influence their interactions with a system [160]. This may discourage error reporting.

The second hypothesis addressed the lack of meaningfulness. The hypothesis proposed that when the message implies some purpose (specifically, the benefit of reporting), participants would also feel inclined to report the error message presented in our study. This hypothesis was supported. Error reporting frequencies were higher when the message implied a benefit to self, followed by a benefit to others. Error reporting frequencies were lowest when no information was provided. Since this control condition also represents the default in the majority of messages users will encounter, the results provide evidence for the role of meaningfulness in what a user is expected to do. By reducing redundancy of effort, the error reporting frequency was significantly increased. When considering both factors (framing and benefit) in combination, the results further indicated that both labelling of messages and the type of implied benefit could increase error reporting. Error reporting was highest when the error was technical as well as implied as a benefit to self. These results suggest that the combination of factors led to highest reporting rate; in line with protection motivation [113]. The findings also speak to the importance of providing information to the user about why certain actions are required rather than relying on uninformative requests, similar to the messages presented to the control groups in our study [173].

The third hypotheses considered the possibility that error reporting was a function of specific personality traits. In this study, the main traits of interest were risk-taking and impulsiveness. The suggestion for testing these was based on the fact that the continued rejection of errors may indicate (a) poor self-efficacy to deal with such challenges, or (b) an attitude towards ignoring risks and not considering the long-term consequences of ignoring such errors. There was no significant evidence that error reporting in any of the conditions (control or implied benefit) were impacted. Further analysis with other traits (such as sociability, privacy concern, security vulnerability, sociability or reciprocity) played no significant role in error reporting.

Exploratory analysis focused on the extent to which generic error reporting may indicate specific motives that encourage or discourage error reporting. This analysis involved the review of open response options and the coding of all reported rationales. The main themes that emerged included fear appraisals, response efficacy and productivity costs. These findings link to existing literature within the field.

6.5.1 Practical implications

The problem of low error reporting is not just specific to HCI, it is also a concern in other domains. For example, a wrong delivery will often lead to the customer not obtaining his or her product. However, many customers are reluctant to complain about delivery service [265], particularly when they have an existing relationship with the provider that they wish to maintain in the future.

The current findings may have several practical implications. First, users will not blankly respond to system messages, an observation that may require some revision of the guidelines outlined by Nurse et al. [173] on how to support usable cybersecurity. These authors proposed that cybersecurity functionality should accommodate all types of users. At present, many error and system messages are standardized, without taking into account the knowledge or availability of the user to report an error. Another important guideline related specifically to 'error prevention, handling and recovery/Undo'. In order to accommodate all users more effectively and support their decision-making, it may be important to reconsider what information is and isn't shared with novice users to support them more effectively. The current study suggests that the content of the error message can play a critical role in the reporting likelihood. This is in line with work that explored how risk salience on interfaces can help increase the attention of users to the message [42]. Further research may wish to explore if providing more information (and hence increasing the meaningfulness of error reports) will effectively increase reporting amongst novice and expert users in a similar fashion as observed in the current study.

Second, concerns about the information that is transmitted with an error message may be an important barrier when practitioners try to increase error reporting in various settings. This is in line with current evidence that states: when a system also requires identifying information, users of this system are often reluctant to use it to report errors [141]. In order to overcome this apprehension, IT professionals may wish to encourage in-person error reporting (e.g., via the phone or in person). This gives employees the option to report an issue without necessarily responding to a system message themselves. This approach may also make IT departments aware of any issues that may indicated malicious or unauthorized pop-ups due to some virus or Trojan on the computer of the user.

Third, error reporting represents a situation with minimal or unknown 'return on investment', interpreted as the time the users spends on reporting errors. This situation may only be reversed when the users also receives some feedback; or systems and programs at least acknowledge that errors have been detected, reported and subsequently fixed (e.g., when updates are triggered). This suggestion is based on work from Holden and Karsh [117]. They argued that error reporting will only be perceived as useful when the data is also used in system improvement and the reporters are made aware that it was their feedback that led to these improvements. In addition, habitual responding may be reduced if error messages are less uniformly and frequently presented, but occur at specific intervals.

6.5.2 Limitations and suggestions for future research

Several small issues arise, some of which may also be addressed in future research. It is our assertion that the error message was perceived as legitimate. However, this was not assessed in the follow-up. In addition, it is not clear as to why the personality characteristics did not relate to error reporting. It is possible that decision-making 'in the moment' (as in the case of error reporting) may not reflect personality but instead situational demands. These findings point to several future research avenues that may build on the current study but also address some of the remaining research gaps.

Firstly, future research may wish to consider how apprehension about what information is shared influences error reporting. For example, the use of punitive responses to error reporting may essentially thwart voluntary reporting [117]. While no evidence was obtained in this study that privacy concerns and security vulnerabilities related to error reporting, these variables may not capture other variables. Research by Maner and Gerend [158] may provide an explanation for this. They propose that motivational orientation associated with approach and avoidance may also influence individual judgements of whether or not a decision will lead to positive and negative outcomes. While the analysis of the qualitative responses in the current study provides some starting points in this direction, future research may wish to consider avoidance and approach orientation in relation to error requests as well.

Secondly, how are errors managed across different disciplines and occupations? Are they anonymous to engender trust into the reporting process [30]? What are the consequences? Finally, to what extent does the perceived severity of a problem feature in the decision to report an error [137]? For example, it is not clear how occupational practices shape tolerance for errors and influence error reporting practices. For instance, different professions have error management processes that may also shape how individuals view and respond to error messages [240]. Certain professions, such as those in the medical realm, have both voluntary and mandatory reporting schemes [137], [117]. This means there may also be lessons to be learned from other disciplines. By extension, what kinds of person versus system reporting systems are most effective and supported by users? Is it easier to avoid error reporting on a system than in everyday interactions (such as staff meetings that include system or patient reports)?

Lastly, what kinds of errors are encountered, how often and by whom? The study's results propose that how error requests are labelled (in this case, either as technical and security issues) plays a key role in shaping reporting. The group difference was explained in relation to the potential lack of control that participants may have felt when facing a 'security issue'. Future research may wish to examine the role of stating the benefits to others or self, and address the response efficacy and costs associated with error reporting.

6.6 Chapter Contributions

This chapter represents further investigation with respect to user behaviours and nudge interventions, specifically regarding the reporting of errors. This behaviour was highlighted as an area of significant interest and importance from both our pilot study and CISO interactions.

Within, we have conducted a review of the relevant literatures and utilised these findings to produce and conduct a bespoke investigation to highlight the behaviours that affect error-reporting likelihood; specifically, the role of social loafing and social norms.

The outcome of our work has highlighted important behaviour characteristics of users which has enabled a more detailed assessment of security behaviours and practices. Several of our results compliment the findings from our literature review aiding both the validity of our work, and extending the applicability of previous studies within security behaviour analysis. Furthermore, we have highlighted several important behaviour traits that are specific to the realms of error-reporting, providing valuable insight with respect to the modification of user behaviours through nudge interventions. From a CISO perspective, we are ultimately more aware of the users within our environment with respect to current behaviours and practices, and are able to identify how to improve the likelihood of error-reporting. This helps to promote a more secure environment where users are actively contributing to the process, and are fully aware of why they are doing it with respect to the greater benefits it provides to both themselves and others. Further benefits are witnessed with respect to our methodology and modelling approach as we are able to better articulate our problem space with respect to user behaviours and actions. With this knowledge, we are able to create and deploy tested interventions that shift users towards being compliant, and thus more likely, to report errors.

Chapter 7

Business Process Modelling Notation

The previous chapter detailed the investigation of two user behaviours; error reporting and cookie acceptance ([127], [128]). Both of these findings, along with chapter 4 ([231]), present methodologies for the testing of user behaviour interventions and allow for a better understanding of the problem space and the users who operate within. Improving our understanding of user behaviours and practices is an important process of aiding the policy decision making process.

This chapter introduces the Business Process Modelling Notation 2.0 (BPMN) formalism and discusses its applicability with respect to our problem space. This chapter represents the modelling phase of our model-based methodology as depicted in figure 3.1.

Through the utilisation of BPMN and the knowledge obtained in previous chapters pertaining to the increased understanding of user behaviours, we are able to more accurately model our threat environment and determine potential risk areas related to user practices. BPMN allows for an abstract modelling of this process in an effort to provide the visual representation of tasks with the goal of optimising throughput and maximizing performance (the efficiency of a user to complete a given task). Within, we utilise BPMN to focus on a case study related to thrombolysis (see Nesbitt & Turland [69]) and the subsequent optimisation of patient throughput (specifically reducing the waiting time of patients that require urgent medical attention). We use the BPMN formalism as a decision support tool and through this process we identify, propose and implement six new notations to the BPMN 2.0 specification. These notations are then applied to the policy decision making process with respect to our methodology in chapter 3 and build upon the knowledge obtained through previous chapters relating to better understanding user behaviours. Ultimately, through improving the expressiveness of the formalism via new notations we aid the modeller within decision making process.

The remainder of this chapter is as follows: section 7.1 introduces the BPMN formalism and discusses its purpose and goals. Section 7.2 documents other work regarding extending the formalism. Section 7.3 introduces the case study and its applicability. Section 7.4 discusses the proposed additions to the formalism with respect to the case study. Section 7.5 discusses the impact, limitations and future work.

This chapter includes discussions and experimental design detailed in Nesbitt & Turland [69] and includes collaborative work declared within the 'Publications'.

7.0.1 Chapter Contributions

This chapter adds to the BPMN 2.0 formalism by proposing further extensions in the form of new notations. These notations represent new model nodes and transitional visualisations that aid with the shared decision making process under the umbrella of enhanced tool support. The additions specifically aid in the expressiveness of the formalism by introducing further visualisation elements that provide further contextual information to the decision maker.

7.1 Introduction

The Decision Making Process is complex and features a high degree of uncertainty. This complexity is furthered when the decision making scenario involves human actors who perform unknown behaviours, practices and processes. As discussed within chapter 2, human actors make decisions based on a plethora of variables that are further masked by influence and interpretation (the method in which they may be presented) as well as personal preference. In shared decision making the effect is multiplied as additional parties have an impact on the overall decision output. The challenge therefore, is to attribute these variables to a modelling environment to effectively reflect the intricacy of the problem space.

BPMN 2.0 provides a robust environment in which to model such an environment. Using a graphical formalism such as this to provide a visual representation is highly beneficial to the overall decision making process. The adoption allows for a party to design and coordinate the sequence of processes and messages that flow between participants (and interdependent systems) of different activities [252].

Decision making activities are represented within BPMN via the use of 'gateways', and are used to define the node traversal within the model following a successful decision outcome. As such, a gateway can be interpreted as a decision point with the outcome of the decision impacting upon the next successive node. This symbiotic definition promotes the need for the collection of ethnographic data relating to the users and their environment (methods for obtaining such information can be seen in chapters 4, 5). This is complimented by a 'rich set of flow and connecting objects' which enables the modeller to create bespoke, user-centric representations of the problem space [191]. These factors combine to create an effective method for aggregating processes where macro-level observations are required (typically large organisations or populations such as a university or hospital).

Our extensions to the formalism are expressed through the Unified Modelling Language (UML) in the form of new node symbols similar to Stroppi et al [217]. Specifically, we utilise XML to define the problem space both graphically and programmatically and attribute contextual information (both human and environmental) in this fashion. Once defined, we apply these to real-world scenarios highlighting the benefits of our proposals. In addition, we evaluate how technology from a Computing Science approach can use contextual information captured within our model to aid the decision making policy. This involves an assessment of devices (hardware) and their applicability for a given scenario. We have coined this extension BPMN for decision making, or BPMNdm.

The need for such an approach is justified and supported by both literature and empirical industrial observations. In the medical profession alone, the use of decision making tools is commonplace and can be seen through software installed on computers, manually via whiteboards, and increasingly on mobile devices such as tablets and smartphones (health applications etc). Each platform has attributes which make it more suited to a particular case. We aim to identify these strength and weaknesses and discuss how they can be adopted for specific scenarios.

The goal of this work is to extend the BPMN formalism and to enable BPMNdm the flexibility and scalability to more accurately define and support the decision making process. Given the inter-disciplinary nature of the problem space and the collaborative decision making process this extension provides an effective method of increasing expressiveness and encapsulating the multitude of aspects that are necessary in formulating a solution.

7.2 Related Work

Numerous studies adopt the BPMN formalism because of its adaptability and flexibility in a wide range of scenarios. Specifically, the ability to extend the existing set of notations is appealing as it allows bespoke expressions of unique problem spaces. It enables designers to compliment the existing notation set with additional nodes that suit specific domains as detailed in [41], [96], and [62], making BPMN a natural selection for our problem space.

Modelling the decision making space is not a new phenomenon [35], however, the need for Computing Scientists to adopt strategies that fully understand the user requirements and issues related to non-compliance is comparatively recent (see chapter 2) and moreover, necessary with respect to policy decision making (see chapter 3).

The decision making process has been described as a cyclical loop [116] where the aims of a particular venture enter into an iterative state.

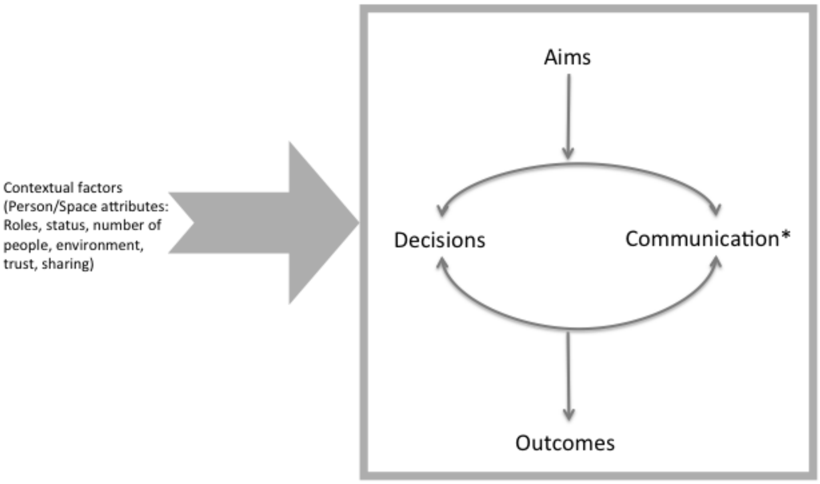
Contextual factors (Person/Space attributes: Roles, status, number of people, environment, trust, sharing)
Aims
Decisions
Communication*
Outcomes

*Processes: Analysis, combination of information, generation, elaboration and evaluation of ideas.

Fig. 7.1 Diagram of decision-making activities [69]

Figure 7.1 highlights the transitions between ‘decisions’ and ‘communication’. Within this, we define decisions as agreements or disagreements with specific points amongst party members. Communication is argued to be far more complex [116] involving mental processes such as analysis, information generation and exchange, incorporation of contextual factors related to the meeting environment and task, visibility and availability of information, and the roles of particular individuals within the decision making space.

This assessment and conceptualisation of the problem space aids greatly towards addressing many of the psychological phenomena outlined in Groupthink [125]. As highlighted, the main concepts of Groupthink which promote failures in decision making are:

- Decision is limited to a narrow range of alternatives.
- Groups do not survey the objectives that need to fulfilled.
- Groups fail to evaluate the non-obvious risks and issues with their choice.
- Groups do not re-evaluate approaches that were previously rejected.
- The group does not consult experts in the field.
- Selective bias - decisions are made based on external influences.

- Members fail to sufficiently discuss how to overcome bureaucratic blocks and other forms of institutional inertia.

This knowledge can be applied to the method in which we govern tool support, specifically, how we utilise technology in our decision making process. We need to support and adopt technologies that combat Groupthink and actively aid the group decision making process [39], [76].

Other examples of BPMN extension include Herbert & Sharpe [114], where the formalism was applied to multiple decision criteria such as cost and technical quality. Their work developed a limited formalised variant of BPMN that extended support towards simple probabilistic branching and rewards. Planning Support Systems (PSS) have also witnessed a similar approach in Campagna et al [49] where the authors claim that it is simple to adapt the BPMN formalism to aid with designing PSS. Bahrani et al [25] similarly adopt BPMN in their work focusing on short-term operational decision making in healthcare, reporting that the model aids the maker with more accurate and timely decisions.

7.3 Case Study and Worked Example

The following four activities provide real-life decision making scenarios that help to exemplify the benefits of our notation changes. The principle example in this work is the treatment of thrombolytic stroke patients as this process involves high-risk, multi-person decision making which highlights effectively the applicability of our contributions.

Three other smaller-scale examples are also documented utilising some of the proposed notations.

Thrombolytic Stroke Treatment

In a clinical environment, shared decision making is increasingly considered good practice. In this context we define shared decision making as the interaction between medical practitioner and patient (the sharing of medical information and even particular faiths, beliefs or practices related to treatment). In essence, the medical staff will provide the medical knowledge for the process, and the patient will give their preferences. This new approach to care replaces the traditional more paternalistic role of the clinician and patient where many of the decisions were solely conducted without the patient's input [88].

The following scenario is taken verbatim from Nesbitt & Turland [69]:

'John (67 years old) was admitted to hospital with a suspected stroke. His wife, Sheila attended with him. During the initial assessment phase, the A+E consultant asked Sheila the approximate time at which John began to exhibit the signs of a stroke. John was asked if had taken either Clopodogrel or Aspirin recently and if he was receiving any hypertensive treatment. Following these questions, the consultant looks at the patient's medical records for indications of a previous cerebravascular event or a history of diabetes. At this stage, the consultant can ascertain whether the patient is a candidate for thrombolysis. In the meantime, an A+E nurse has recorded John's systolic blood pressure, his blood glucose value, approximate weight and has performed a National Institutes of Health Stroke Scale (NIHSS) assessment on him. Finally, John has a CT scan to look for evidence of infarctions in his brain. At this point, John's indicators meet the licensing criteria for thrombolysis. Since his stroke is moderate, the treatment has the greatest net benefit compared to non-treatment. With this in mind, the consultant asks Sheila for consent for treatment, explaining what her husband has and discussing thrombolytic treatment. The consultant states that out of 100 patients like John, approximately 33 extra patients would make a successful recovery with thrombolysis than without it. However, there is a small risk of a brain haemorrhage that could result in death or serious injury. After some discussion, Sheila asks what the consultant would recommend; the consultant recommends treatment. Sheila signs the consent form and the drug is administered'.

With this scenario the BPMN formalism can aid considerably in modelling the problem space and mapping the decision points. For example, whether or not to proceed with further decision making if the patient does not fit within the licensing criteria. With any tool, however, design is critical (see HCI in chapter 2.4, 4).

Security Policy

Security policies are often governed by fixed budgeting and resource allocation. Relating specifically to a university environment and the transferral of data (see chapter 3), our notations enable a better understanding of these finite variables through graphical and programmatical representation. The impact of this is to suggest new tools and technologies that may or may not be available at a given instance related to a particular resource (for example, with cost - less expensive options are available once a certain budget constraint is exceeded, with time - more efficient processes are highlighted for optimisation purposes).

Security Policy - with Tools

The use of tools in this scenario would allow for automatic allocation of node traversal based on set criteria and bounds within the model.

Cardiovascular Disease Risk Reduction

'A GP is using their practice's EMIS (Egton Medical Information Systems) computer system (a database that stores a practice's patient records) to check patient records. In this scenario, EMIS will flag patients who are at high risk of cardiovascular disease (such as stroke or heart attack). The GP will invite each patient flagged at risk for a discussion on reducing their risk. This is achieved using a tool called CVdecide that displays risk presentations in the form of a Pictograph, showing risk of disease based on such factors as the patient's weight and smoking status. Based on these factors, both GP and Patient will discuss a plan for reducing the patient's cardiovascular risk' [69].

With the above scenario, a consent point could be added to the BPMN model indicating the goals that have been agreed by both parties.

The following table (7.1) summaries these scenarios and relates specifically to the application of the new notations detailing their salient features.

Feature	CVD	MBA	Security	Stroke	Comparable BPMN construct
Cost [†]			*		*none*
Time		*	*	*	Time start event only
Multi-person	*	*	*	*	Swimlanes
Confidentiality	*		*	*	*none*
High risk			*	*	*none*
Information capture	*	*	*	*	Data object
Private	*	*	*	*	*none*
Defined goals	*	*	*	*	*none*
Output information	*	*	*	*	Data object
Share documents		*	*		*none*
Workflow data	*	*	*	*	*none*
Scope data	*	*	*	*	*none*

[†] In the context of the scenarios, cost is defined as a scalar values that is given to a financial, effort or other quantifiable variable.

Table 7.1 Summary of salient decision making activities in each of the examples [69]

The table documents important similarities between the applications, notably that all scenarios generate their own data (risk presentation, policy document etc.) and can defined on an aggregate level by the existing data object constructs. In addition, the actors from each scenario conform to the swim-lane constructs already defined within the BPMN specification along with information flows denoted by arrows.

Workflow data is defined as ‘data that is accessible to the entire group’ [69]. Scope data can be considered as ‘any kind of data that is private to an individual or a subset of individuals’ [69]. The existing BPMN formalism includes document elements that can be adapted to encapsulate these details.

The remaining contributions have no similar definitions within the existing BPMN formalism. The next section details their features and implementation.

7.4 Extending the Formalism

When modelling a decision making process we must first consider how we are going to capture and represent a plethora of environmental and contextual data sets, and second determine the constructs that define how the model will simulate real-life decisions. These requirements apply to both the environment in which the scenario is based, as well as each individual task and sub-task within the model. To enable this we utilise a BPMN management suite (jBPMN) that enables us to ‘store information by applying the principles of composite types as seen in programming languages by providing data structures to all of the BPMN objects’ [69].

By adopting jBPMN and providing the definitions to our notation extensions, we have enabled developers the ability to ‘freely extend a given BPMN object’ [69] allowing data storage within the object. Using this platform, modellers will be able to store and access ‘strongly typed variables, such as Ints, Strings, Booleans’ simply through basic BPMN XML and Java class instances allowing on-the-fly manipulation and declaration [69].

A pseudo code representation of this process and connections can be visualised in table 7.2.

Example code	Description
`{Object Name}.all`	Returns all items in the named object
`{Object Name}.age`	Returns the item “Age” from the named object
`{Object Name}[0]`	Returns an item located at the named object’s array index 0.
`{Object Name}.size`	Returns the number of items a named object stores.
`{Object Name}.has(item)`	Returns TRUE if an object has an item that corresponds to the parameter, returns FALSE otherwise.

Table 7.2 Summary of BPMNdm Object Notation [69]

To exemplify table 7.2 let us consider figure 7.2.

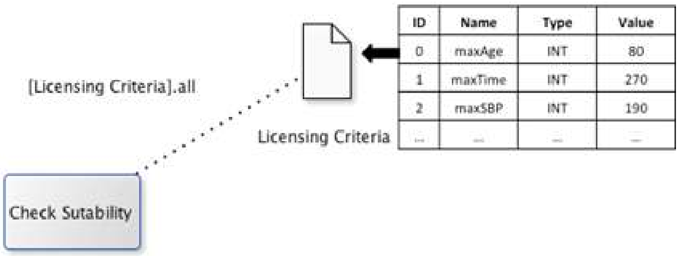

ID	Name	Type	Value
0	maxAge	INT	80
1	maxTime	INT	270
2	maxSBP	INT	190
...	...	...	...

Fig. 7.2 Illustration of License Criteria document and Check sustainability task [69]

Observing figure 7.2 we see that the data object ‘Licensing Criteria’ contains an array of integers that relate specifically to the maximum permitted patient statistics for thrombolytic treatment (age, time, etc.). The blue box labelled ‘Check Suitability’ denotes a condition that checks whether the patient’s values are lower than that of the specified Licensing Criteria. Individual values can be checked by providing the name of the object in square brackets [example] and performing an array search similar to many other programming languages. As seen in figure 7.2 explicitly, the .all method can be called to check all criteria within the array.

Traditional BPMN gateways are preserved in our notation change, but receive enhanced features. We extend the expressiveness of the gateway to incorporate common programming statements such as IF/ELSE/ELSE-IF (Fig 7.3). This allows us to specify triggers that determine model traversal and defines explicit pathways.

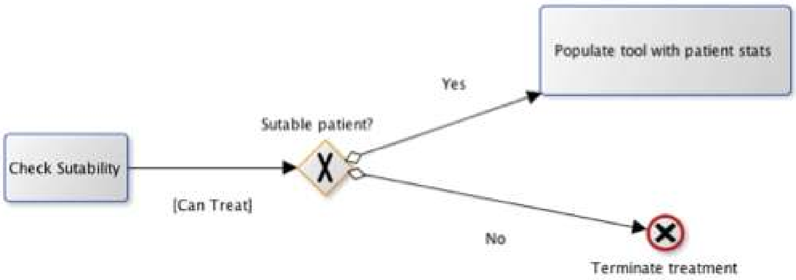

IF[Can Treat] == Yes	Return Yes
ELSE	Return No

Fig. 7.3 Gateways and Underlying Statements [69]

With specific reference to figure 7.3, we can see that the ‘Check Suitability’ process returns a boolean object. The boolean value is returned subject to a checking condition and the subsequent path (whether or not the patient is suitable) is determined. This process uses a simple XML schema and is easily adjustable by the modeller allowing for specific scenarios to be depicted within the node. This allows for extensions and further additions to be made as the modeller becomes more informed with respect to the target environment and/or users. The iterative nature of our overall methodology (see chapter 3) compliments this approach.

7.4.1 Proposed Notation Changes

The following sub-sections denote the proposed notation changes with respect to their purpose and scenario applicability (see table 7.1). Of the six extensions listed in order, Daniel Nesbitt [69] is responsible for Workflow Data, Scope Data, and Decision Points (listed first), I claim no ownership. They are defined with respect to their sub-heading and are included for the overall validity and expressiveness they provide. All other notations and subsequent evaluations are my own work.

Within our proposal, ‘Workflow Data’ and ‘Scope Data’ form the basis of our data representation and data flow. For decision-making , this relates to both the physical (i.e. a hand written piece of information) and the virtual/mental (anything that is stored and visualised electronically) data. The summation of this can be seen in table 7.3.

Workflow data

As stated, this extension definition is the work of Daniel Nesbitt [69]. It is included to aid the overall validity of our approach and is taken verbatim from Nesbitt & Turland [69]. The sub-heading ‘Information Security Applicability’ is my own work.

Proposed BPMNdm construct	Description	Implements missing feature(s) (see Table 7.1)
Cost Swimlane	Modified Swimlane that represent increasing, or decreasing cost on a scale with child nodes	Cost
Time Swimlane	Modified swimlane that represent increasing, or decreasing time on a scale with child nodes	Time
Scoped and Workflow data objects	Scoped Data objects are owned by a actor, and not visible to other actors unless shared. Workflow data objects are accessible to all users	Confidentiality, Private data, Sharing documents
Weighted Paths	Weighted paths represent the predetermined preference of a outgoing edge.	Defined goals

Table 7.3 Summary of BPMNdm Formalism [69]

Russel et al. define workflow data as: ‘Data elements that are supported with and are accessible to all components in each and every case of the workflow’ [197]. In decision making contexts, workflow data can be considered public data that is accessible to many/all in the decision making environment. Examples of this include documents that are shared between participants or files on a server that are accessible to all users. BPMN does not support the representation of workflow data [238]. In the context of decision making, using the BPNM data object, workflow data could be represented by data objects that are not explicitly connected to any other object and are outside of any swim lane or similar construct.

Figure 7.4 illustrates the concept of workflow data (Risk Presentation) seen alongside Scope data (Pateint stats/history and Persons characteristics). The definition of workflow data follows Russel et al's definition with all given tasks and processes having access to the object.

Information Security Applicability

In a typical IS setting, the ‘workflow data’ extension would represent data that is accessible globally. For instance, this may include the company e-mail directory that houses all employee e-mail addresses. On a smaller scale, this could be visualised in collaborative work such as a data repository that is available to all parties (e.g. shared Dropbox). The mapping of such data is important in reference to its availability. As more users have access to the document or repository, a more closely monitored access management system is required. This will no doubt scale as data sensitivity and the number of users who have access increases.

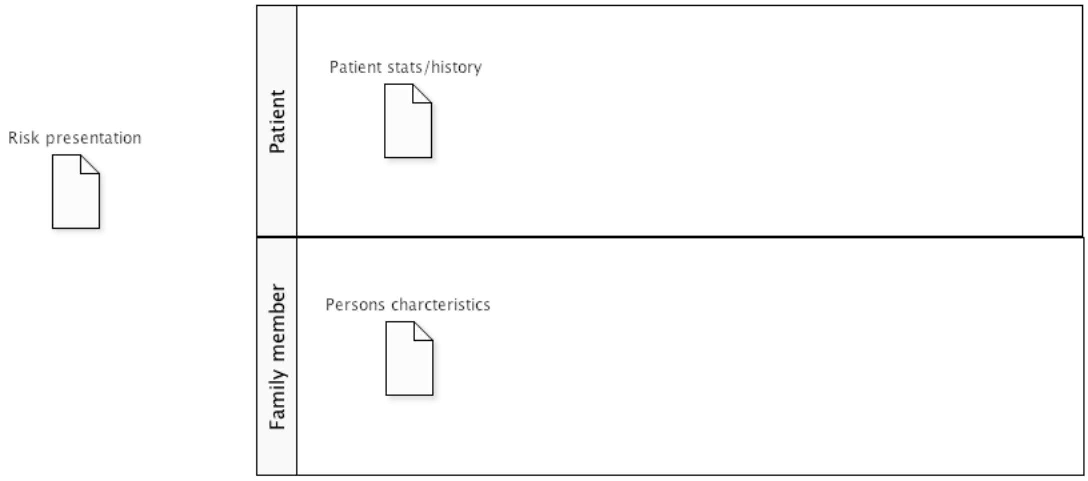

Fig. 7.4 Workflow data object (Risk presentation) and Scope data objects (Patient stats./history, Person's characteristics) [69]

The main advantage of introducing this expression is the ability to monitor assets. Through an assessment of risk and data value (detailed in chapter 2) one can attribute value to the asset as well as its perceived risk. Visualising access with respect to the parties who have access rights is an effective method for understanding and evaluating the threat environment. Modelling and simulation can address how manipulation of these rights, or the access mechanism utilised in order to gain access, reduce the risk posed to the organisation. Such investigation benefits from an improved understanding of the user base and ultimately aids the decision making process through a more accurate assessment of the problem space.

Scope data

As stated, this extension definition is the work of Daniel Nesbitt [69]. It is included to aid the overall validity of our approach and is taken verbatim from Nesbitt & Turland [69]. The sub-heading 'Information Security Applicability' is my own work.

Scope data is defined as 'Data elements [that] can be defined which are accessible by a subset of the tasks in a case' [197]. Scope data is useful in a decision-making context as this notion can be used to account for private artefacts that a user or group possesses. Private artefacts could be defined as anything that holds information, such as an address book, report or anything else that can be considered private to a user or group. BPMN does not support scope data [238]. In order to visualise scope data, data objects must be placed in swim lanes or visually grouped with objects in order to imply where the object belongs.

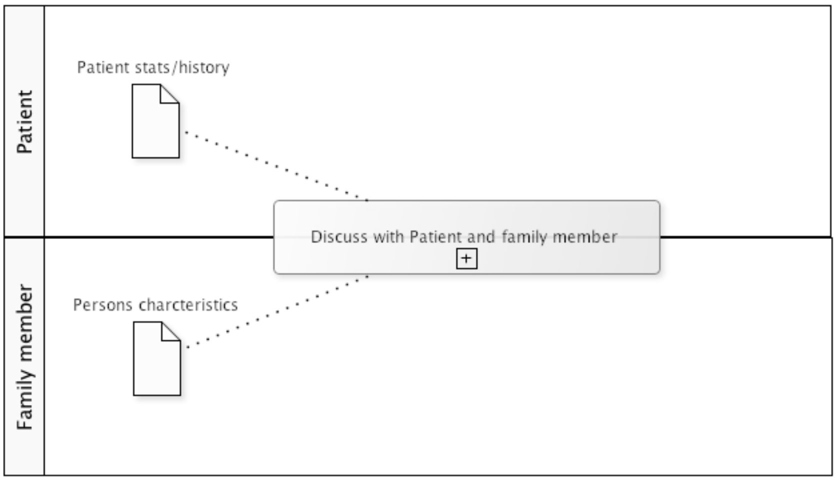

Fig. 7.5 Scope data object [69]

Figure 7.5 illustrates the use of scope data and how it can be passed from one subset to another. In the example, “Patient Stats/History” and “Persons characteristics” are data objects that belong to two actors, a Patient and a Family Member respectively. Scope data is only visible to the actor or group it belongs to unless the object is passed to another process as part of a flow. An example of this is the consultation process between a clinician, patient and the family member, which involves discussing the patient’s characteristics as well as adjusting the course of the consultation to minimise distress to the patient and family member by the clinician judging the emotional states of both of the other actors. Unless the object is explicitly duplicated by the clinician (i.e. writing down the necessary information), the clinician will retain the object for the length of the task it is assigned to, only after which it is assumed it is passed back to the sender process or destroyed.

The metamodel for Scoped and Workflow data is provided in 7.6.

Information Security Applicability

The use of scope data can be applied to our example in a number of ways. As scope data determines who the data belongs to, and moreover, who has access to it, there are many similarities with BYOD. This is increasingly relevant when dealing with data that is stored on a BYOD. As described in chapters 3 and 4 the role of device and data ownership greatly impacts the decision making process, similarly to location familiarity with respect to network selection. This impacts future security decisions and a modification of these factors can greatly enhance the security of the environment. Providing a visual tool for the assessment and representation of this is highly beneficial and allows future policies to more accurately

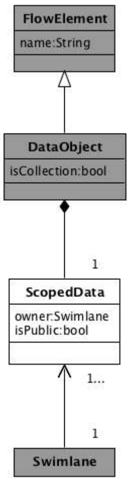

Fig. 7.6 Metamodel for Scoped and Workflow data [69]

target specific vulnerabilities. Such an approach could allow policies to nudge users towards more secure alternatives. The likelihood that these nudges are effective is increased through better expression of the target users and their environment.

From a modeller's perspective, the policy decision making process is enhanced as the practitioner is able to understand how the ownership of the data and the device impact user decisions (and subsequent node traversal). Omissions can be rectified through repeat investigation, obtaining a better understanding of how this effects the outcomes of security decisions. Being able to model potential conflicts and offer conflict resolution boosts the productivity and security of the environment through enhanced optimisation.

The movement of data and its subsequent access can also be modelled with areas of potential risk highlighted. The parties who have access to the data can be visually represented, as well as the fashion in which they do so. Once identified, it is possible to produce bespoke solutions and test them within the modelling environment until the desired level of access is met.

Decision points

As stated, this extension definition is the work of Daniel Nesbitt [69]. It is included to aid the overall validity of our approach and is taken verbatim from Nesbitt & Turland [69]. The sub-heading 'Information Security Applicability' is my own work.

Decision points (denoted by the use of hatched squares, typically used for grouping tasks in contemporary BPMN) are used to explicitly indicate the portions of a BPMNdm model

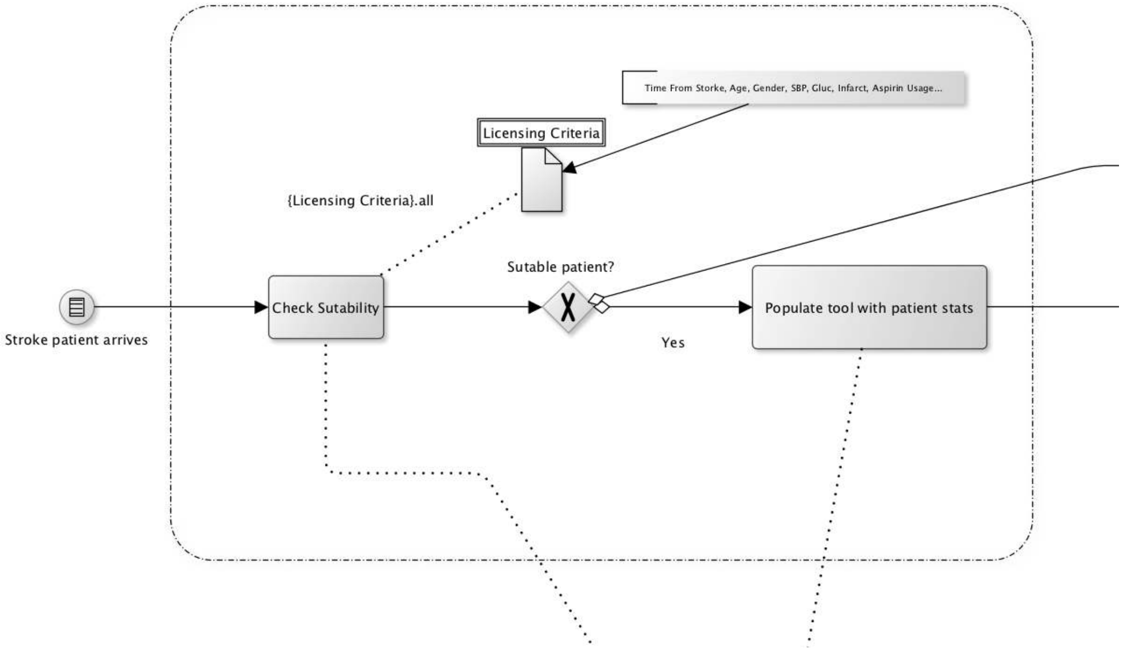
Time From Storke, Age, Gender, SBP, Gluc, Infarct, Aspirin Usage...
Licensing Criteria
{Licensing Criteria}.all
Sutable patient?
Stroke patient arrives
Check Sutability
Populate tool with patient stats
Yes

Fig. 7.7 Decision point [69]

at the point where a decision is being made as well as the critical tasks and documents that leads to this decision. With respect to tool support, this is useful as it provides a visual cue to the user about the most important tasks, users and elements that are associated with a specific decision point. In the above example, a decision point is formed where a user has to decide if a patient is suitable for treatment. This decision involves the user looking up the licensing criteria for the treatment and comparing it with the patient's stats. These steps are enclosed in the task 'Check Suitability'. If the patient is suitable, donated by the exclusive gateway, the next stage is to populate a decision making tool with the patient's stats, otherwise the treatment is terminated due to the patient being unsuitable. This could be achieved by implementing a algorithm which identifies exclusions gateways and iterates back through the incoming connections to identify tasks and/or data objects that influence the decision to be made (i.e. an assessment of tasks that use a licensing criteria to evaluate the patient's characteristics).

Decision points can also be used to identify when consent is required for a task to be undertaken, such as a medical procedure. Deciding which participants are involved in providing consent can be determined by the incoming edges and/or if the gateway is positioned between two swimlanes.

It is envisioned that the Decision Points would be generated automatically by BPMNdm based on the symbols chosen by the user. In the example above, BPMNdm would have identified the gateway ‘Suitable Patient?’ as a decision point and worked back to identify the preceding tasks/attributes/data sources that are closely attributed to this gateway. BPMNdm would then work forwards to identify the immediate tasks that are executed following the decision. This is necessary, as BPMNdm would then use the information contained in each activity/symbol to determine what tools/services would be most appropriate for this Decision Point.

The metamodel for Decision Points are provided in 7.8.

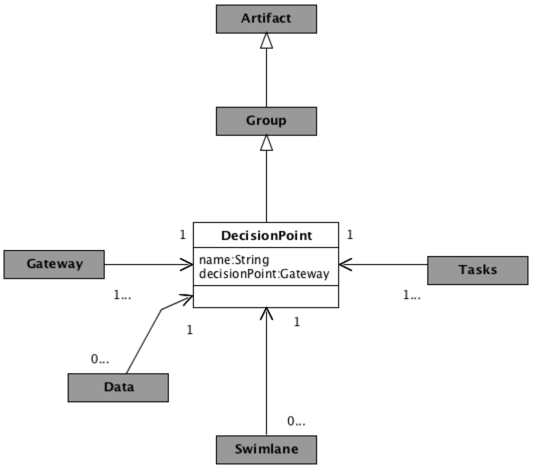

Fig. 7.8 Metamodel for Decision Points [69]

Information Security Applicability

The application to IS shares numerous similarities to the previously mentioned example. A decision point aids the modeller in the policy decision making process as it highlights the many inputs that form the decision. If we refer to chapters 4 and 5 we begin to understand the factors that influence user decisions. This may be past experiences or familiarity in the case of the Wi-Fi example, or the willingness to report an error or accept a cookie.

The added expressiveness is beneficial as we are able to map the individual elements that comprise the decision making process. Through the above noted investigations, we improve our understanding of these variables and ultimately iteratively refine our model. From a modellers perspective, this aid the policy decision making process as it highlights key areas of importance with respect to the users decision. This allows the policy maker to cater for specific requirements that have been identified, and if they are insufficiency understood, allows for subsequent investigation into these areas.

7.4.2 Weighted Paths

The current implementation of BPMN 2.0 does not provide any method for defining the likelihood of node traversal after a gate. Whilst this is satisfactory for defining a process model, it can be greatly improved to provide more detailed knowledge of a given process with tool support. Below is an example of a weighted path with the most likely route highlighted with a bold arrow and their respective probabilities.

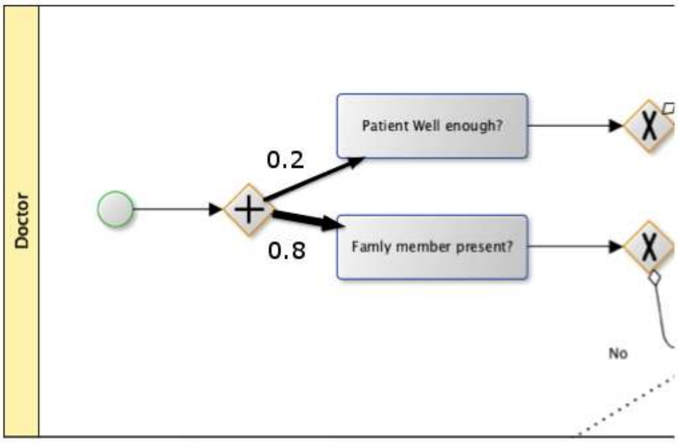

Fig. 7.9 Weighted Paths [69]

Weighted paths address the issue of unknown output from a gate by relying on assumptions or prior knowledge. In its most primitive form, a weighted path can be seen as the likelihood with which the next node is chosen. There are several benefits that such an extension provides.

Firstly, if data relating to the target environment is available (obtainable through experimentation similar to chapters 4 and 5) it will be possible to highlight particular trends or patterns in either the user's behaviour or how the system reacts to given stimuli. This likelihood can be statistically represented with a value between 0-1 within the model visualisation. With tool support, this helps to determine the most probable path and provides valuable insight that will help with optimisation and throughput.

With reference to our thrombolysis scenario, there are many examples of where an increased expressiveness benefits the decision maker. For example, weighted paths can be adopted to visualise the outcome of treatment based on previous patients, or visualise a preferred care pathway within a given demographic. With reference to figure 7.9 we witness the point at which a doctor will seek consent for administrating thrombolytic treatment after an event. Typically, the patient is unlikely to be sufficiently conscious or lacks the cognitive ability to formulate a reasoned response and thus the decision is often discussed with relatives. Therefore, consent from a relative tends to be the most likely path (exemplified by 0.8).

Specific treatments and medications also benefit significantly from the introduction of such a notation. As drug trials are extensively tested and their results statistically quantified it is a simple process to incorporate this probability into the XML. Patients will also have a record that can be visualized as a data object. In essence, this is a predefined document stating all medical knowledge pertaining to the individual. This data is incredibly useful in determining how treatment should progress and benefits similarly to the above. As we are more aware of the patient's ailments and previous treatments, documentation and medical trials will indicate the probability of how the patient responds and subsequently which steps should then proceed (represented as node traversal within the model).

In time critical scenarios such as thrombolysis, it allows for an optimization of node traversal relating to a reduced waiting time for the patient. This is both a time and cost saving feature as it is possible to plan future activities and processes based on a given probability (drug trials and patient records). For example, it enables the clinician or medical facility to be able to inform staff of potential surgery and aids with the general hospital logistics. Having such data available on a tool would empower medical staff to make these decisions in a mobile environment where decisions are based on prior knowledge and quantified statistical likelihood.

Financial allocation and budgeting is also an area where this notation change is beneficial. Through more effective resource management and finance allocation it would be possible to improve those services that are most commonly used in an effort to streamline and manage patient throughput.

The metamodel for Weighted paths is provided in 7.10.

Information Security Applicability

The introduction of weighted paths within the BPMN 2.0 formalism provides significant benefits to our problem space highlighted in chapter 3, and provides further rationale for the case studies documented in chapters 4 and 5.

Within IS, let us consider the wireless network selection problem (chapter 4). Through our experimentation we know that we are able to manipulate the wireless networks chosen by our participants. By nudging users we are impacting the network that they choose and ultimately changing the environment, moreover, the threat environment. Applying this to

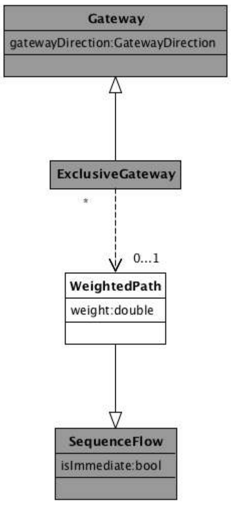

Fig. 7.10 Metamodel for Weighted Paths [69]

our current BPMN implementation, we can view the policy as scope data where there is a given probability (obtained from experimentation such as chapter 4) that users will select a given network. This network will have properties which can be expressed through the XML schema and subsequently visualised. As such, a weighted path would highlight the most likely node traversal after a security decision based on statistical testing of the policy. Being able to quantify user behaviour in such a fashion is a powerful tool and supports the validity of our methodology.

A similar adoption is possible when we consider the findings in chapter 5. As described in the literature review (2), risk and frequency are often known by the organisation either by past events or an assumption based on identified risks, and user behaviours can be examined more closely with studies such as 4 and 5. If we focus on error-reporting for example, the likelihood that a user experiences an error and subsequently reports it can be of high significance to the organisation and the user (identifying and reporting problems should lead to them being avoided in the future). By including a likelihood within a process that it may fail, the modeller is able to more accurately depict subsequent events based on behavioural observations and statistical analysis.

In both of these scenarios we witness an enhancement to the visualisation of the decision making process. This aids the decision maker to make a more informed choice based on the ability to utilise additional data sets. We validate the applicability of the findings within our user studies and ultimately aid the decision making process.

7.4.3 Cost

With respect to our application, cost is interpreted as a scalar quantity that apportions a given financial value to an item, process or activity. Cost is a vital component of any business process as it is a finite variable that is directly correlated with budgeting and financial health. Operating profits, specifically in the corporate world (and also in areas of resource management) affect success and overall functionality of the business. Controlled, precise allocation of these resources optimizes the business work flow, maintains peak efficiency, and promotes successful business.

Integrating cost into BPMN is a natural choice. Being able to attribute values to a given decision compliments tool support by providing clear, contextual information. The benefits of such an approach become increasingly clear with respect to budgeting. Node traversal can be impacted by a cost value that is declared within the XML. This will impact the model as pathing will be determined by checking the cost of a decision point or node and comparing this to the remaining budget. Below we see the possible application and impact of such a strategy.

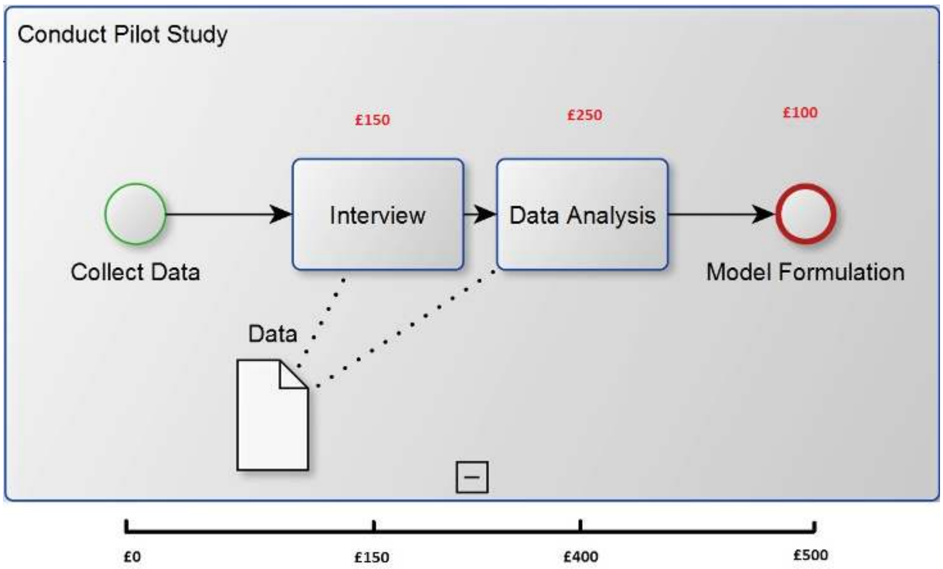

Fig. 7.11 Cost Notation [69]

Relating to thrombolysis, we visualize the cost function as an assessment of resource management and optimization. If a budget is allocated on a per person per operation/consultation basis and the cost of the activity is documented, BPMN with tool support is an effective method of aiding the DMP. Tool integration would allow for a visual representation of the problem space with the implications of specific decision outcomes clearly mapped. Choos-

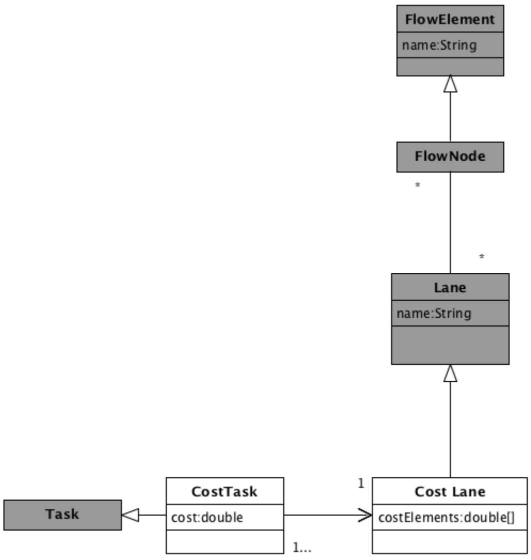

Fig. 7.12 Metamodel for Cost element [69]

ing a specific treatment for one patient may require altering subsequent treatments, or the treatments available for other patients. This granularity enables a more informed decision that allows for an assessment of the impact of the decision outcomes on a larger scale.

Within the United Kingdom, every resident citizen has access to free health care via the National Health Service (NHS). Within the NHS, there is a specific protocol that depicts how treatment is funded and paid for. This process is as follows:

- Patients are provided a HRG (Health Resource Group) code that is a record of procedures and treatments conducted.
- This document is then passed to your local General Practitioner (GP).
- The money for the procedure is then debited from the Primary Care Trust.
- These funds are controlled by Practice Managers whose duties involve careful monitoring of what has been spent. In documenting these expenditures, statistical auditing systems are used.
- In an effort to optimize expenditure, Practice Managers analyse the most costly procedures in order to allocate resources in an attempt to save money. Procedures with high capacity will often have more money invested in order to optimize the process. This process also involves Practice Managers looking at alternatives to treatments if a high

percentage of the budget is being spent on procedures. Examples include procedures being conducted at local clinics and home instead of the hospital.

The points noted above provide a clear example of why this approach is beneficial. We can visualise these steps as a business process that naturally benefits from the introduction of cost into BPMN. Being able to document cost at the atomic level for a given patient in a given scenario allows accurate modelling and understanding of where costs are accumulated and importantly where budget impacts care. This allows hospital managers to better allocate funds in an effort to save money, maximize patient throughput and streamline the care of patients.

Cost integration would enable a detailed, accurate representation of financial expenditure. The scale in Fig. 7.11 denotes a dynamic scale that updates in real-time to reflect the decisions made. This allows for a structured view of expenditure once you have exited the decision space. Maximum expenditure can be added to the XML schema and the policy that defines it may be amended as required.

Information Security Applicability

The benefits to IS are closely reflected within the thrombolysis example. Whilst decisions are unlikely to affect a patient's physical well-being, the corporate nature of the environment often requires a closer budgeting strategy. Any optimisations therefore, will have an impact on the profitability of a particular venture.

Consider the adoption of encrypted USB devices for staff who wish to transfer data. If this were to become mandatory, there are several implications with respect to the modelling process. Firstly, we can attribute the cost of acquiring the devices within the XML as a constant. Secondly, through modelling behaviour we are able to make assumptions as to the probability that a user conforms with the policy (do they use the device?). Thirdly, we can quantify the risks associated with non-compliance through an assessment of vulnerability (how much will a security breach cost? What will be leaked etc.). If we combine these elements and provide a visual feedback mechanism via tool support (BPMNdm), we aid the decision maker. We are able to highlight the cost of an action and can relate this to the fixed budget defined within our model. Node traversal will be impacted upon user decisions and subsequent actions.

7.4.4 Time

Time is a complex multidimensional value or quantity that often requires a bespoke definition through worked example. In a typical scenario we can summarise this as the indefinite progress of existence and events that occur in an irreversible succession from the past. In numerous scenarios, however, where such simple definitions are not appropriate or easily defined we must augment this. For example, in Computing Science, time is a variable that can be manipulated to facilitate specific needs. Time does not necessarily reflect the same time that is associated within the 'real world' (i.e. 24 hours in a day). Latency is a perfect example of time dilation where 1s in real-time does not reflect 1s in the system state. Instead, time can change at a given interval providing little to no continuity. Here we see a difference between the unit of time and the continued progression of time. As such, we need to understand the state of the system in order to define it.

This work proposes a new definition of time to be used within the BPMN modelling formalism. We define time as a meta value within the XML that represents the duration of a given activity. It is therefore a cumulative variable that defines the time taken to traverse a given network path. It is independent of any formally defined scale in reference to future paths that are not predefined. Therefore, altering the route through the network presents a new time-scale. Furthermore, time increments are non-uniform. Nodes are not uniformly distributed as activities have different associated completion times. This provides a highly dynamic environment where time is not a constant progression and reflects the decisions made by the individual.

Having defined the method in which time is understood with respect to BPMN, it is necessary to discuss how this will be utilized. The application of time will allow for an assessment of productivity and efficiency in terms of start - end completion time. From a modelling perspective, and further abstraction (system analysis), this will allow performance monitoring that has a plethora of real-world advantages. It is possible to identify particularly lengthy tasks and determine whether or not this will exceed time constraints. To discuss this, we must refer to our example scenario.

In a hospital environment, time (typical definition as denoted by real-time, real-world progression), budget and resources are intrinsically linked. Optimal and cost effective patient care is imperative to both the success and purpose of the service; there is a balance between caring exclusively for the individual and ensuring that other patients also receive equal care. In both these cases we assume that budget and resources are a subset of time by deducing that budgets are allocated at service level (with time scale and investiture) and resources are dynamically distributed with respect to demand (a value that changes with real-time, real-world time). With this assumption we can refer to a implementation of time into BPMN.

Fig. 7.14 demonstrates the implementation of time but it is necessary to discuss the specific adaptation. For tool support we see immediate benefits in terms of the real-world real-time representation of the model. Specific routes are defined by their timing and subsequently

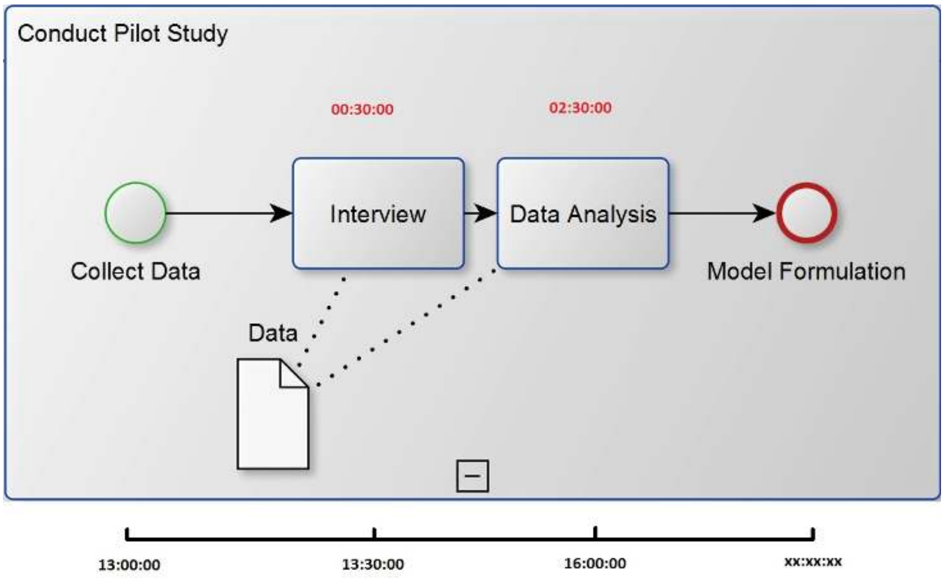

Fig. 7.13 Time Notation [69]

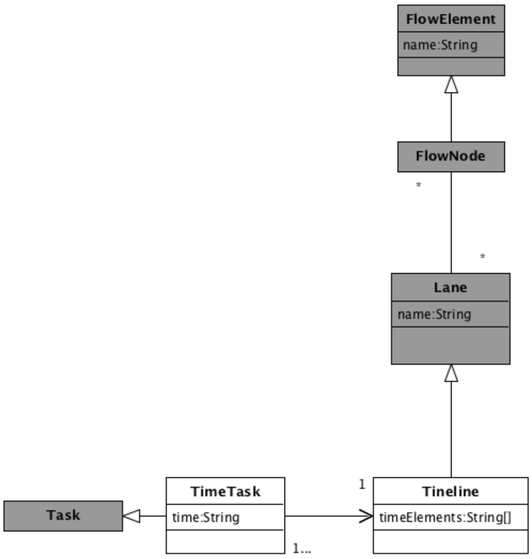

Fig. 7.14 Metamodel for Time element [69]

plotted on a time-line where scale indicates the distribution of events. This dynamically updates relating to any chosen path. In essence, the time-line provides an evaluative data set to be analysed for optimization purposes.

Within this example, node transitions are complimented with tool support. Subsequent route traversals are defined by given constraints (for instance a typical restraint may be that a patient cannot wait longer than 30 minutes when suffering fro ma particular incident).

In relation to thrombolysis, Decision Points and timing are critical. Stroke victims have a small window where action is necessary in order to prevent further medical complication. Translating this to a physical environment we can assume a party of 1 doctor, 1 patient, 1 other member (possibly family/friend). Given this, our model would contain a 'run-time time variable' that incrementally updates with node transition (this would naturally be bounded and relative to the patient condition). For example, a specific set of nodes (or task) cannot take longer than x minutes. Referring to the above example we can visualize a scenario where certain nodes are circumvented depending on the time. If the patient has too long to transition from the waiting room to consultation, it may be necessary for immediate medical intervention based on a doctor/consultant's approval. This is specifically where tool will benefit the decision making process.

As time progresses through the system, the optimal route for node transition also changes. This allows the tool to not only be a process model, but instead to dictate optimal transition based on predefined constraints. Mobile tool support would allow these decisions to be made in a mobile environment increasing productivity, potentially availability, and ultimately enhancing patient care. The addition of time is therefore highly beneficial to the decision maker and the case study provides clear evidence of its potential.

Information Security Applicability

The introduction of time within the BPMN formalism with respect to IS has numerous applications. Time is critical in relation to security incidents as response times can significantly reduce the impact of a breach or network infection. The need for a rapid response is paramount, but in doing so, there is a need to understand how events transpire temporally within the environment. Enabling a finer granularity allows for an assessment of each specific activity or decision. Human behavioural studies can identify the common response times of individuals (as seen in chapter 4 where the average response time was 10s) and the analysis of hardware and network infrastructure (related to computing power) can often dictate the speed at which events transpire.

With respect to aiding the policy decision making process, the introduction of timing to the modelling notation presents many benefits. An important part of any working environment is the optimisation and the streamlining of business processes (for example, logging onto a computer to send an e-mail). This often involves employee interaction where the task completion time is a combination of many variables (availability, competence and procedure etc.). Each of these variables can be assessed through user experimentation and such areas have been studied extensively as highlighted in chapter 2. Including time within the modelling process of policy testing allows the modeller to understand the implications of security interventions on the population. For example, the introduction of encrypted USB sticks may seem like a logical step towards resolving data loss. However, there are issues related to this such as the increased time it takes to access and use the device, along with the productivity losses that are associated with forgetting the password and thus rendering the device useless until an administrator is available. Such an implementation allows for an assessment of time related issues regarding policy implementation and allows for a more informed trade-off decision between security and productivity. When combined with the other notation changes, as well as the previous user behaviour studies it allows for an in-depth analysis of the costs and benefits associated with a given policy.

7.5 Evaluation

The notation changes detailed within this chapter highlight the importance of bespoke models that fully visualise the intricacies of the problem space. Many of the introductions here, however, have a multi-faceted application and benefit numerous scenarios from a breadth of disciplines (including the studies from chapters 4, 5). Our process of extending the formalism is well documented throughout external literature and many of our notations use the recommendations that are highlighted within these studies to aid the expressiveness and validity with which we can define our problem space.

Each notation change is documented with respect to its applicability to a given problem domain. From our examples, it is evident that the introduction of these notations provides a more expressive formalism for the modeller. This allows for a more bespoke representation of the problem space, and ultimately empowers the policy decision maker to design, test and implement more appropriate and effective policies. The use of tools provides a real-time decision support system that benefits both the modeller and the users within the environment through a more in-depth assessment of the decision space. The tool also benefits collaborative decision making through the inclusion of several new notations, namely ‘scope data’ and ‘workflow data’ that aim to encapsulate the collaborative nature of the working environment and the problems this can pose.

With reference to our problem space defined in chapter 3, we witness a genuine improvement with respect to the policy design and decision making processes compared to non-tool based solutions. Investing resources into understanding user behaviours yields valuable insight towards the variables which impact our environment. Utilising small user studies improves our understanding of key practices and highlights areas of risk and vulnerability. The ability to adequately model these as well as provide alternative routes to a certain goal is highly valuable.

The notations themselves benefit from being integrated with the use of the XML schema and the popular jBPMN environment. The ability to adapt the variables within the nodes allows a modeller to create bespoke, dynamic models that relate to specific actions. Enabling the modeller to create statements and conditions within the nodes that react to a given trigger or bound produces a dynamic tool that can react to changes in the current decision space. Furthermore, this enables alternative routes to be determined that may not have otherwise been obvious to the modeller.

7.5.1 Limitations

Limitations towards this approach are mainly reflected by the untested, abstract nature of the scenarios. Many of the scenarios may appear idealised, and without adequate testing to prove our assumptions, they lack the scientific rigour necessary to fully validate our approach. Furthermore, it would be necessary to investigate each specific behaviour in relation to the subsequent model it was aiding and the specific node it may benefit as the accuracy of the model is correlated with how well we can define our problem space.

To remedy this, it would be necessary to tailor specific investigations to fully ascertain human behaviours in order to most effectively represent them. One would also need to conduct a comprehensive evaluation of the current system and robustly determine critical variables such as asset value. The nodes, however, are fully customisable through the XML so their expressiveness is sufficient to encapsulate these findings.

7.5.2 Future Work

Having demonstrated the validity and benefits that our notations bring, and highlighted the outstanding limitations, it is necessary to begin the implementation phase within the jBPMN modelling environment. Through the implementation phase we will be able to conduct an empirical study of the process and be able to perform an in-depth analysis as to the accuracy and validity of our approach using a working model. We could further this by modelling a ‘real-world’ scenario and assessing the impact that the interventions and modelling extensions have on the decision making process. Data gained from this approach could then aid subsequent models by a process of iteration where we refine our model and ultimately are better able to articulate the problem space.

7.6 Chapter Contributions

This chapter documents how extending the BPMN 2.0 formalism provides a greater set of expressions with which to model a specific environment. We use real-world time-critical scenarios (thrombolysis) as well as promote the applicability towards designing our model based on previous experimentation detailed in prior chapters. We define how repeat small-scale user testing can aid this process and ultimately benefit the decision making process through greater understanding of our problem space.

Utilising commonly adopted environments and languages, we have created a simple yet intuitive method of creating bespoke models to aid the decision making process. The notations are customisable through the XML schema and are able to represent results obtained through experimentation into user behaviours and the environment.

Chapter 8

Conclusion

This thesis has investigated the application and effectiveness of nudges and behavioural change mechanisms with respect to Information Security practices. In doing, we have developed an iterative methodology based on user and CISO interaction in order to identify behaviours that are critical to understanding the security practices within our problem space. This has required the design of a pilot study aimed at identifying key areas of risk related to specific behaviours in an effort to better understand them. These behaviours have subsequently been investigated through bespoke experimental research in an effort to improve our articulation of the problem space. We utilise these findings to produce a model-based representation of our environment with the goal of aiding the policy decision making process.

Within this process we have contributed to the field of knowledge pertaining to user behaviours with respect to specific applications (see chapters 4, 5, 6). Through our experimentation, we have validated the effectiveness of our behavioural interventions and shown the benefits of adopting a pilot study based on CISO interaction. We have proven the validity of our work through the cross-reference and analysis of existing studies conducted within the respective fields, and designed our bespoke research approaches upon the recommendations found in such literature.

The remainder of the chapter is organised as follows: Section 8.1 highlights the contributions of the thesis with respect to each chapter. Section 8.2 details whether we met our aims and objectives. Section 8.3 highlights the limitations of our investigations and their applicability. Section 8.4 discusses future directions the work could follow in order to address these limitations and improve our understanding of the problem space.

8.1 Summary of Contributions

This section summarises the key contributions of the thesis with reference to their specific chapter.

- **A background and literature review of security risks, practices and user behaviours.**
  We provided a multi-disciplinary literature review of the problem space with respect to the broader domain highlighted in chapter 1 and more specifically to our scenario in chapter 3. Chapter 2 discusses and highlights the key topics and issues related to our scenario and highlighted relevant theories and practices adopted within the problem domain. Subsequent chapters (4, 5, 6 & 7) each contain specific reference to literatures and theories that are included due to their relevance pertaining to the investigated behaviour. These represent methodologies adopted and improved upon within our investigation, as well as recommendations from the authors. Our approach benefits from the multi-disciplinary nature of the literature review through the inclusion of work conducted in the social sciences in a typically computing science domain.

- **A problem formulation for expressing the current difficulties of the policy decision making process with respect to unidentified user behaviours.**
  In chapter 3 we document the key problems within the domain with respect to our aims and objectives and the concerns raised by our CISO interaction and IRIDIUM data. We formally document the problem and de-construct it, identifying smaller individual study areas related to specific human behaviours and practices.

- **Formulation of an iterative methodology incorporating a pilot study for problem identification and current security behaviours.**
  We document the design and adoption of an iterative methodology based on empirical data for the investigation of the problem space. This approach benefits from continued CISO interaction and the inclusion of data from the IRIDIUM study and our pilot study.

- **The design, implementation and analysis of three bespoke investigations targeting specific behaviours highlighted through conducting our initial pilot study and CISO interaction.**
  We design and conduct three separate user studies focusing on specific behaviours and practices highlighted within our problem definition process (see chapters 4, 5 & 6). These chapters are novel in that they examine the applicability of nudging and

behavioural interventions within a new Information Security domain, contributing specifically to the field of research which governs them. These chapters reference and build upon strategies from other studies, but apply and analyse them with respect to a new problem domain.

- **The design, implementation and analysis of tools used in conducting these investigations, and the formulation of a modelling strategy to map these behaviours.** The experiments detailed within chapters 4, 5, 6 and 7 required the design and development of specific tools and testing materials. These tools were built with recommendations and theories from our literature review and the results obtained from use mirrored those found in these studies. This highlights the validity of our tools and provides a further contribution to the research domain. The wide-ranging applicability of our tools also enables their adoption to other environments and problem spaces.

  Our modelling approach contributes to the BPMN formalism through the extension of the notation enabling greater expression of the problem domain. The benefit of these additions are demonstrated with respect to a real-world, time critical scenario and the additions are implemented using an open source environment. Our proposals are based upon existing literature and research recommendations and provide an effective method of modelling our environment using the data provided from our investigations.

- **A demonstration of the effect of our behavioural interventions on real data.** The interventions and recommendations we propose with respect to our study outcomes, benefit from their utilisation of real-world, non-simulated data. Using empirical data in this fashion further supports the validity of our approach as our results reflect real-time decision-making in genuine tasks. This benefits our methodology with respect to modelling as we are able to more accurately reflect the problem space.

8.2 Were Aims and Objectives Met?

8.2.1 Aim:

- ‘To design, develop and validate with scientific rigour, a robust methodology for aiding the chief information security officer in the policy decision making process’.

We formally met our aim by designing and conducting the bespoke experimentations detailed in previous chapters. To promote rigour, our methods adopted best practice and in several cases yielded statistically significant results that supported our hypotheses. Ultimately, our investigations provide a clearer understanding of our problem space enabling a more informed decision by the CISO with respect to Information Security Decisions.

8.2.2 Objectives:

- To understand the current (‘state of the art’) practices through a review of literatures and best practices.

  We provide an extensive literature review that targets many critical studies within the respective fields. This review highlights the applicability to our study along with formal limitations that may affect our results.

- To accurately interpret and document the issues and concerns from the CISO and narrow the problem space into a more focused investigative avenue.

  As part of the iterative methodology we outline a process for capturing our stakeholder requirements and use this to formulate an exploratory pilot study. This pilot study highlights many of the behaviours we subsequently investigate and further supports both the applicability and validity of our methodology.

- To identify and formulate solutions to specific problems with respect to CISO and user requirements.

  As detailed above, we defined specific behaviours that were of importance to both our stakeholder and the users within our test environment. We subsequently tested behavioural interventions designed to influence the user’s decision making process and validated the effectiveness of such interventions. The results of this process enables the CISO to develop more suitable policy that is holistic with user behaviours and practices ultimately leading to an increase in security policy compliance.

- To develop tool support and experimentation to aid in the assessment of these problems.

  All of the bespoke experiments required the design and implementation of specific tools. These tools were either web, mobile or model based.

- To evaluate the effectiveness of such tools and experimentation.

  The tools yielded several statistically significant results throughout our experimentations. This provides strong evidence for their validity but subsequent testing would be necessary to confirm this belief. It is possible that the tools did impact our results.

- To reflect on the overall impact of the approach, address its shortcomings and outline the direction of possible future work.

This objective is detailed within section 8.4.

8.3 Limitations

The limitations of this thesis in relation to the individual investigations are detailed respectively, in reference to the chapter in which they feature (see 'Limitations' in chapters 4, 5, 6 & 7). As such, the limitations in this section will feature on the overall methodology in relation to its applicability and validity.

The main limitation of the proposed methodology is the lack of subsequent testing. Whilst we firmly believe that this is an effective approach, a belief that stems from the knowledge that each investigation is designed and grounded with reference to important literature and theories (some of which follow a similar approach), and that our results through the design of bespoke methods emulate those of other surveys; it is impossible to validate this claim without testing. Similarly, interaction with our stakeholder and user-base, along with empirical study, has ensured that our approach remains focused on their needs whilst also enabling the identification of additional problem areas through the process.

To remedy this shortcoming, we require additional time in order to utilise our results, formulate a model to test policy adoption, and subsequently assess its impacts based on the claims of our behavioural intervention studies. In response to this, however, we see this work as an introduction to the overall methodology with the knowledge that it is a scalable solution that benefits from repeat investigation of highlighted behaviours (hence our iterative strategy). Given greater resources, it would be possible to investigate numerous areas simultaneously in an effort to optimise the modelling of the environment and ultimately increase the speed at which new policy can be designed, implemented and subsequently evaluated.

8.4 Future Work

The future direction of this thesis in relation to the individual investigations are detailed respectively, in reference to the chapter in which they feature (see 'Future Work' in chapters 4, 5, 6 & 7). As such, this section will discuss the direction of subsequent investigation with respect to the overall methodology.

The future work of this thesis typically relates to the limitations detailed in section 8.3, whereby additional studies related to specific user behaviours would benefit our understanding of the problem space, and ultimately provide a richer expression with which to formulate a solution. This process would also benefit from an increase in the number of participants within the pilot study as we believe more behaviours would be captured by increasing the number of users we investigate. In addition, future work would also document the design process of our environment using the BPMN extensions we propose from the stakeholder's perspective. Documenting this process would allow for an analysis of specifically how and where our methodology aids the policy decision-making process with respect to past practices.

In relation to specific behaviours and environmental factors, the following investigations would benefit the validity of the study:

- **Additional Devices:** It would be necessary to focus more specifically on behaviours with respect to device type. By understanding how a device choice may influences behaviours, we could examine intervention methods and nudges that aim to 'nudge' users towards a specific device that may be more appropriate for a given task with respect to security. For instance, the use of a laptop or personal computer may be preferable to a mobile device when checking the validity of a link in an email. It is often difficult to highlight the intended destination on a mobile device as their is no function pertaining to user input.

- **Collaborative Decision-Making:** As we discussed within our problem identification (chapter 3), the university environment encourages collaborative working. One future avenue of investigation would be to investigate how our results relate to the collaborative decision-making process, and what additional studies would need to be conducted if they do not apply. This would further benefit our modelling strategy highlighted in chapter 7 where we reference our development with respect to a group decision in thrombolytic care.

References

[1] Charles Abraham and Susan Michie. A taxonomy of behavior change techniques used in interventions. *Health Psychology*, 27(3):379, 2008.

[2] Alessandro Acquisti. Privacy and security of personal information. In *Economics of Information Security*, pages 179–186. Springer, 2004.

[3] Alessandro Acquisti. Nudging privacy: The behavioral economics of personal information. *Digital Enlightenment Yearbook 2012*, pages 193–197, 2012.

[4] Alessandro Acquisti and Ralph Gross. Imagined communities: Awareness, information sharing, and privacy on the facebook. In *Privacy enhancing technologies*, pages 36–58. Springer, 2006.

[5] Alessandro Acquisti and Jens Grossklags. Privacy and rationality in individual decision making. *IEEE Security & Privacy*, (1):26–33, 2005.

[6] Alessandro Acquisti and Hal R Varian. Conditioning prices on purchase history. *Marketing Science*, 24(3):367–381, 2005.

[7] Alessandro Acquisti, Idris Adjerid, and Laura Brandimarte. Gone in 15 seconds: The limits of privacy transparency and control. *IEEE Security & Privacy*, (4):72–74, 2013.

[8] Alessandro Acquisti, Leslie K John, and George Loewenstein. What is privacy worth? *The Journal of Legal Studies*, 42(2):249–274, 2013.

[9] Anne Adams and Martina Angela Sasse. Users are not the enemy. *Commun. ACM*, 42(12):40–46, December 1999. ISSN 0001-0782. URL http://doi.acm.org/10.1145/322796.322806.

[10] National Security Agency. Suite b cryptography. URL https://www.nsa.gov/ia/programs/suiteb_cryptography/. Visited on 02/01/13.

[11] AIRC. Attack intelligence research center annual threat report: 2008 overview and 2009 predictions. URL http://www.aladdin.com/pdf/airc/AIRC-Annual-Threat-Report2008.pdf.

[12] I. Ajzen. The theory of planned behavior. *Organizational behavior and human decision processes*, 50(2), 1991.

[13] I. Ajzen. Perceived behavioral control, self efficacy, locus of control, and the theory of planned behavior. *Journal of Applied Social Psychology*, 32(4), 2002.

[14] I. Ajzen and M. Fishbein. *Understanding attitudes and predicting social behaviour*. Englewood Cliffs, NJ: Prentice Hall, 1980.

[15] Icek Ajzen. The theory of planned behavior. *Organizational behavior and human decision processes*, 50(2):179–211, 1991.

[16] Anton Aluja, Michael Kuhlman, and Marvin Zuckerman. Development of the zuckerman–kuhlman–aluja personality questionnaire (zka–pq): A factor/facet version of the zuckerman–kuhlman personality questionnaire (zkpq). *Journal of personality assessment*, 92(5):416–431, 2010.

[17] Ian Anderson, Julie Maitland, Scott Sherwood, Louise Barkhuus, Matthew Chalmers, Malcolm Hall, Barry Brown, and Henk Muller. Shakra: tracking and sharing daily activity levels with unaugmented mobile phones. *Mobile Networks and Applications*, 12(2-3):185–199, 2007.

[18] Weston Anson, Donna P Suchy, and Chaitali Ahya. *Intellectual Property Valuation: A Primer for identifying and determining value*. American Bar Association, first edition, 2005.

[19] Apple. iOS 8, 2015. URL https://www.apple.com/uk/ios/. Visited on 02/01/15.

[20] Warwick Ashford. IT security awareness needs to be company-wide, says (isc)^2, 2012. URL http://www.computerweekly.com/news/2240163342/IT-security-needs-to-be-company-wide-says-ISC. Visited on 12/01/15.

[21] Terrence August and Marius Florin Niculescu. The influence of software process maturity and customer error reporting on software release and pricing. *Management Science*, 59(12):2702–2726, 2013.

[22] S. Aurigemma and R. Panko. A composite framework for behavioral compliance with information security policies. In *System Science (HICSS), 2012 45th Hawaii International Conference on*, pages 3248–3257, Jan 2012.

[23] Algirdas Avizienis, J-C Laprie, Brian Randell, and Carl Landwehr. Basic concepts and taxonomy of dependable and secure computing. *Dependable and Secure Computing, IEEE Transactions on*, 1(1):11–33, 2004.

[24] Ramakrishna Ayyagari, Varun Grover, and Russell Purvis. Technostress: Technological antecedents and implications. *MIS Q.*, 35(4):831–858, December 2011. URL http://dl.acm.org/citation.cfm?id=2208940.2208943.

[25] Sepideh Bahrani, Renaud Bougueng Tchemeube, Alain Mouttham, and Daniel Amyot. Real-time simulations to support operational decision making in healthcare. In *Proceedings of the 2013 Summer Computer Simulation Conference*, page 53. Society for Modeling & Simulation International, 2013.

[26] Rebecca Balebako, Pedro G Leon, Hazim Almuhimedi, Patrick Gage Kelley, Jonathan Mugan, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. Nudging users towards privacy on mobile devices. In *Proc. CHI 2011 Workshop on Persuasion, Nudge, Influence and Coercion*, 2011.

[27] A. Bandura. *Social Foundations of Thought and Action: A Social Cognitive Theory*. Prentice Hall, 1986.

[28] Albert Bandura. Social cognitive theory of moral thought and action. In *In*, pages 45–103. Erlbaum, 1991.

[29] Scott Bartis, Kate Szymanski, and Stephen G Harkins. Evaluation and performance a two-edged knife. *Personality and Social Psychology Bulletin*, 14(2):242–251, 1988.

[30] John W Beasley, Kamisha Hamilton Escoto, and Ben-Tzion Karsh. Design elements for a primary care medical error reporting system. *WMJ-MADISON-*, 103(1):56–59, 2004.

[31] Anne Beaudry and Alain Pinsonneault. Understanding user responses to information technology: A coping model of user adaption. *MIS Q.*, 29(3):493–524, September 2005. URL http://dl.acm.org/citation.cfm?id=2017264.2017271.

[32] Adam Beautement, M. Angela Sasse, and Mike Wonham. The compliance budget: Managing security behaviour in organisations. In *Proceedings of the 2008 Workshop on New Security Paradigms*, NSPW '08, pages 47–58, New York, NY, USA, 2008. ACM.

[33] France Belanger. Study reveals resistance to strong password security, 2011. URL http://cacm.acm.org/news/141643-study-reveals-resistance-to-strong-password-security/fulltext. Written on 15/11/11.

[34] France Belanger, Janine S Hiller, and Wanda J Smith. Trustworthiness in electronic commerce: the role of privacy, security, and site attributes. *The Journal of Strategic Information Systems*, 11(3):245–270, 2002.

[35] Richard E Bellman and Lotfi Asker Zadeh. Decision-making in a fuzzy environment. *Management science*, 17(4):B–141, 1970.

[36] Hasida Ben-Zur and Moshe Zeidner. Threat to life and risk-taking behaviors: A review of empirical findings and explanatory models. *Personality and Social Psychology Review*, 13(2):109–128, 2009. URL http://psr.sagepub.com/content/13/2/109.abstract.

[37] Rainer Bernnat, Olaf Acker, Nicolai Bieber, and Mark Johnson. Friendly takeover: The consumerization of corporate IT, 2010. URL http://www.booz.com/global/home/what-we-think/reports-white-papers/articledisplay/friendly-takeover-consumerization-corporate. Visited on 10/02/2013.

[38] J. S. Blumenthal-Barby. Seeking better health care outcomes: The ethics of using the "nudge". *American Journal of Bioethics*, 12(2):1–10, 2012.

[39] Edward S Boyle, Michael Wolfe, and Charles E Kimble. Overcoming groupthink bias with groupware. Technical report, DTIC Document, 1997.

[40] Tony Bradley. Android dominates market share, but apple makes all the money. *Forbes. November*, 2013.

[41] Marco Brambilla, Piero Fraternali, and Carmen Vaca. Bpmn and design patterns for engineering social bpm solutions. In *Business Process Management Workshops*, pages 219–230. Springer, 2012.

[42] Cristian Bravo-Lillo, Saranga Komanduri, Lorrie Faith Cranor, Robert W Reeder, Manya Sleeper, Julie Downs, and Stuart Schechter. Your attention please: designing security-decision uis to make genuine risks harder to ignore. In *Proceedings of the Ninth Symposium on Usable Privacy and Security*, page 6. ACM, 2013.

[43] Ian Brown. Britain's smart meter programme: A case study in privacy by design. *International Review of Law, Computers & Technology*, 28(2):172–184, 2014.

[44] Izak Benbasat Burcu Bulgurcu, Hasan Cavusoglu. Information security policy compliance: An empirical study of rationality-based beliefs and information security awareness. *MIS Quarterly*, 34(3):523–548, 2010. URL http://misq.org/information-security-policy-compliance-an-empirical-study-of-rationality-based-beliefs-and-informatio html.

[45] Franco Callegati, Walter Cerroni, and Marco Ramilli. Man-in-the-middle attack to the https protocol. *IEEE Security and Privacy*, 7(1):78–81, 2009.

[46] Vincent J Calluzzo and Charles J Cante. Ethics in information technology and software use. *Journal of Business Ethics*, 51(3):301–312, 2004.

[47] Giacomo Calzolari and Alessandro Pavan. Optimal design of privacy policies1. 2001.

[48] Cameron Camp. The byod security challenge: How scary is the ipad, tablet, smartphone surge? *WeLiveSecurity Blog post*, 2012.

[49] Michele Campagna, Konstatin Ivanov, and Pierangelo Massa. Orchestrating the spatial planning process: from business process management to 2nd generation planning support systems. In *accepted) Proceedings of the 17th AGILE Conference on Geographic Information Science Connecting a Digital Europe through Location and Place*, 2014.

[50] John M Carroll. Human computer interaction-brief intro. *The Encyclopedia of Human-Computer Interaction, 2nd Ed.*, 2013.

[51] John M Carroll, Wendy A Kellogg, and Mary Beth Rosson. *The task-artifact cycle*. Cambridge University Press, 1991.

[52] Aldo Cassola, William K Robertson, Engin Kirda, and Guevara Noubir. A practical, targeted, and stealthy attack against wpa enterprise authentication. In *NDSS*, 2013.

[53] H Cavusoglu, H Cavusoglu, JY Son, and I Benbasat. Information security control resources in organizations: A multidimensional view and their key drivers. Technical report, working paper, Sauder School of Business, University of British Columbia, 2009.

[54] Shelly Chaiken. Heuristic versus systematic information processing and the use of source versus message cues in persuasion. *Journal of personality and social psychology*, 39(5):752, 1980.

[55] Eun Kyoung Choe, Jaeyeon Jung, Bongshin Lee, and Kristie Fisher. Nudging people away from privacy-invasive mobile apps through visual framing. In *Human-Computer Interaction–INTERACT 2013*, pages 74–91. Springer, 2013.

[56] Robert B Cialdini and Noah J Goldstein. Social influence: Compliance and conformity. *Annu. Rev. Psychol.*, 55:591–621, 2004.

[57] Sunny Consolvo, Katherine Everitt, Ian Smith, and James A Landay. Design requirements for technologies that encourage physical activity. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 457–466. ACM, 2006.

[58] Sunny Consolvo, Predrag Klasnja, David W McDonald, Daniel Avrahami, Jon Froehlich, Louis LeGrand, Ryan Libby, Keith Mosher, and James A Landay. Flowers or a robot army?: encouraging awareness & activity with personal, mobile displays. In *Proceedings of the 10th international conference on Ubiquitous computing*, pages 54–63. ACM, 2008.

[59] David Constant, Lee Sproull, and Sara Kiesler. The kindness of strangers: The usefulness of electronic weak ties for technical advice. *Organization science*, 7(2): 119–135, 1996.

[60] Jorge L Contreras. Developing a framework to improve critical infrastructure cybersecurity (response to nist request for information docket no. 130208119-3119-01). *Available at SSRN 2248658*, 2013.

[61] Price Waterhouse Cooper. Information security breaches survey - survey conducted by PWC for UK government business and innovation department (2013), 2013. URL http://www.pwc.co.uk/assets/pdf/cyber-security-2013-technical-report.pdf. Visited on 04/01/15.

[62] Paul Cotofrei and Kilian Stoffel. Fuzzy extended bpmn for modelling crime analysis processes. *Data-Driven Process Discovery and Analysis SIMPDA 2011*, page 13, 2011.

[63] Lynne Coventry, Pam Briggs, Debora Jeske, and Aad van Moorsel. Scene: A structured means for creating and evaluating behavioral nudges in a cyber security environment. In *Design, User Experience, and Usability. Theories, Methods, and Tools for Designing the User Experience*, pages 229–239. Springer, 2014.

[64] SK Crook, J Jaffe, R Boggs, and SD Drake. Worldwide mobile worker population 2011-2015 forecast. *Abstract. International Data Corporation (IDC). http://www. idc. com/getdoc. jsp*, 2011.

[65] Robert E Crossler, Allen C Johnston, Paul Benjamin Lowry, Qing Hu, Merrill Warkentin, and Richard Baskerville. Future directions for behavioral information security research. *computers & security*, 32:90–101, 2013.

[66] Parker. D, B. *Fighting Computer Crime: A New Framework for Protecting Information*. John Wiley & Sons, Inc., New York, NY, USA, 1998.

[67] Norman L Chervany D Harrison McKnight. What trust means in e-commerce customer relationships: an interdisciplinary conceptual typology. *International journal of electronic commerce*, 6(2):35–59, 2001.

[68] Olof Dahlbäck. Personality and risk-taking. *Personality and individual differences*, 11 (12):1235–1242, 1990.

[69] Nesbitt Daniel and James Turland. Bpmndm – extending the bpmn formalism to aid the decision making process. *European Journal of Information Systems*, 2015. In submission.

[70] John D'Arcy, Anat Hovav, and Dennis Galletta. User awareness of security countermeasures and its impact on information systems misuse: A deterrence approach. *Information Systems Research*, 20(1):79–98, 2009. URL http://pubsonline.informs.org/doi/abs/10.1287/isre.1070.0160.

[71] John D'Arcy, Tejaswini Herath, and Mindy K. Shoss. Understanding employee responses to stressful information security requirements: A coping perspective. *Journal of Management Information Systems*, 31(2):285–318, 2014. URL http://www.tandfonline.com/doi/abs/10.2753/MIS0742-1222310210.

[72] Sauvik Das, Adam DI Kramer, Laura A Dabbish, and Jason I Hong. Increasing security sensitivity with social proof: A large-scale experimental confirmation. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 739–749. ACM, 2014.

[73] Robson de Oliveira Albuquerque, Luis Javier García Villalba, and Tai-Hoon Kim. Gtrust: Group extension for trust models in distributed systems. *International Journal of Distributed Sensor Networks*, 2014, 2014.

[74] Robson de Oliveira Albuquerque, Luis Javier García Villalba, Ana Lucila Sandoval Orozco, Fábio Buiati, and Tai-Hoon Kim. A layered trust information security architecture. *Sensors*, 14(12):22754–22772, 2014.

[75] Tamara Denning, Adrienne Andrew, Rohit Chaudhri, Carl Hartung, Jonathan Lester, Gaetano Borriello, and Glen Duncan. Balance: towards a usable pervasive wellness application with accurate activity inference. In *Proceedings of the 10th workshop on Mobile Computing Systems and Applications*, page 5. ACM, 2009.

[76] Gerardine DeSanctis and Brent Gallupe. Group decision support systems: a new frontier. *ACM SIGMIS Database*, 16(2):3–10, 1984.

[77] Morton Deutsch and Harold B Gerard. A study of normative and informational social influences upon individual judgment. *The journal of abnormal and social psychology*, 51(3):629, 1955.

[78] Scott J Dickman. Functional and dysfunctional impulsivity: personality and cognitive correlates. *Journal of personality and social psychology*, 58(1):95, 1990.

[79] Oxford English Dictionary. Definition of privacy, 2015. URL http://www.oxforddictionaries.com/definition/english/privacy. Visited on 21/04/15.

[80] Tamara Dinev and Paul Hart. Internet privacy concerns and their antecedents-measurement validity and a regression model. *Behaviour & Information Technology*, 23(6):413–422, 2004.

[81] A. Dix. *Human-computer interaction*. Prentice hall, 2004.

[82] Paul Dolan, Antony Elliott, Robert Metcalfe, and Ivo Vlaev. Influencing financial behavior: from changing minds to changing contexts. *Journal of Behavioral Finance*, 13(2):126–142, 2012.

[83] Paul Dolan, Michael Hallsworth, David Halpern, Dominic King, Robert Metcalfe, and Ivo Vlaev. Influencing behaviour: The mindspace way. *Journal of Economic Psychology*, 33(1):264–277, 2012.

[84] Paul Dolan, Michael Hallsworth, David Halpern, Dominic King, and Ivo Vlaev. Mindplace: influencing behaviour through public policy. 2014.

[85] Frank A Drews and Adrian Musters. Individual differences in interrupted task performance: One size does not fit all. *International Journal of Human-Computer Studies*, 79:97–105, 2015.

[86] James N Druckman. Evaluating framing effects. *Journal of Economic Psychology*, 22 (1):91–101, 2001.

[87] Geoffrey B Duggan, Hilary Johnson, and Petter Sørli. Interleaving tasks to improve performance: Users maximise the marginal rate of return. *International Journal of Human-Computer Studies*, 71(5):533–550, 2013.

[88] Vikki A Entwistle and Ian S Watt. Patient involvement in treatment decision-making: the case for a broader conceptual framework. *Patient education and counseling*, 63(3): 268–278, 2006.

[89] D. Eskins and W.H. Sanders. The multiple-asymmetric-utility system model: A framework for modeling cyber-human systems. In *Quantitative Evaluation of Systems (QEST), 2011 Eighth International Conference on*, pages 233–242. IEEE, 2011.

[90] Sybil BG Eysenck, Paul R Pearson, G Easting, and John F Allsopp. Age norms for impulsiveness, venturesomeness and empathy in adults. *Personality and individual differences*, 6(5):613–619, 1985.

[91] Ana Ferreira, Jean-Louis Huynen, Vincent Koenig, Gabriele Lenzini, and Salvador Rivas. Socio-technical study on the effect of trust and context when choosing wifi names. In *Security and Trust Management*, pages 131–143. Springer, 2013.

[92] Peter Fischer, Andreas Kastenmüller, and Kathrin Asal. Ego depletion increases risk-taking. *The Journal of Social Psychology*, 152(5):623–638, 2012.

[93] Brian J Fogg. Persuasive technology: using computers to change what we think and do. *Ubiquity*, 2002(December):5, 2002.

[94] Hadasch Frank, Benjamin Mueller, and Alexander Maedche. Changing employees' security behavior with technology enforcement of information systems security policies. *Working Paper Series in Information Systems*, (2), 2012.

[95] Steven Furnell, Rehan Shams, and Andy Phippen. Who guides the little guy? exploring security advice and guidance from retailers and isps. *Computer Fraud & Security*, 2008(12):6–10, 2008.

[96] Feng Gao, Maciej Zaremba, Sami Bhiri, and Wassim Derguerch. Extending bpmn 2.0 with sensor and smart device business functions. In *Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE), 2011 20th IEEE International Workshops on*, pages 297–302. IEEE, 2011.

[97] Roland Gasser, Dominique Brodbeck, Markus Degen, Jürg Luthiger, Remo Wyss, and Serge Reichlin. Persuasiveness of a mobile lifestyle coaching application using social facilitation. In *Persuasive Technology*, pages 27–38. Springer, 2006.

[98] GCHQ. 10 steps to cyber security, 2012. URL https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/73128/12-1120-10-steps-to-cyber-security-executive.pdf.

[99] Simona Gilboa, Arie Shirom, Yitzhak Fried, and Cary Cooper. A meta-analysis of work demand stressors and job performance: examining main and moderating effects. *Personnel Psychology*, 61(2):227–271, 2008.

[100] Carroll J Glynn, Michael E Huge, and Carole A Lunney. The influence of perceived social norms on college students' intention to vote. *Political Communication*, 26(1): 48–64, 2009.

[101] Go-Gulf. Cyber crime statistics and trends. URL http://www.go-gulf.com/blog/cyber-crime/. Visited on 04/04/15.

[102] Google. Android developers. URL http://developer.android.com/index.html. visited on 16th February 2015.

[103] Google. Google play store, 2015. URL https://play.google.com/store?hl=en_GB. Visited on 04/02/2015.

[104] Tyrone Grandison and Morris Sloman. A survey of trust in internet applications. *Communications Surveys & Tutorials, IEEE*, 3(4):2–16, 2000.

[105] Miriam Greenspan. *Healing through the Dark Emotions The Wisdom of Grief, Fear, and Despair*. Shambhala Publications, 2003.

[106] The Guardian. Universities need to plug into threat of cyber-attacks. URL http://www.theguardian.com/education/2015/mar/31/universities-cyber-attacks-research-criminals. Visited on 04/06/15.

[107] Tipton. H, F and Krause. M. *Information Security Handbook*. CRC Press, sixth edition, 2012.

[108] Diane F Halpern. *Thought and knowledge: An introduction to critical thinking*. Routledge, 2002.

[109] Catherine Hanssens. Legal and ethical implications of opt-out hiv testing. *Clinical Infectious Diseases*, 45(Supplement 4):S232–S239, 2007.

[110] Marian Harbach, Markus Hettig, Susanne Weber, and Matthew Smith. Using personal examples to improve risk communication for security & privacy decisions. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, pages 2647–2656. ACM, 2014.

[111] Stephen G Harkins. Social loafing and social facilitation. *Journal of Experimental Social Psychology*, 23(1):1–18, 1987.

[112] Kerm Henriksen, James B Battles, Margaret A Keyes, Mary L Grady, Ranjit Singh, Wilson Pace, Ashok Singh, Chester Fox, Gurdev Singh, et al. A visual computer interface concept for making error reporting useful at the point of care. 2008.

[113] Tejaswini Herath and H Raghav Rao. Protection motivation and deterrence: a framework for security policy compliance in organisations. *European Journal of Information Systems*, 18(2):106–125, 2009.

[114] L Herbert and Robin Sharp. Quantitative analysis of probabilistic (bpmn) workflows. In *Proc. for ASME 2012 International Design Engineering Technical Conferences and computers and information in Engineering Conference (IDETC/CIE2012)*, 2012.

[115] Cormac Herley. So long, and no thanks for the externalities: the rational rejection of security advice by users. In *Proceedings of the 2009 workshop on New security paradigms workshop*, pages 133–144. ACM, 2009.

[116] Randy Y Hirokawa and Marshall Scott Poole. *Communication and group decision making*. Sage Publications, 1996.

[117] Richard J Holden and Ben-Tzion Karsh. A review of medical error reporting system design considerations and a proposed cross-level systems research framework. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 49(2):257–276, 2007.

[118] Adele E Howe, Indrajit Ray, Mike Roberts, Malgorzata Urbanska, and Zinta Byrne. The psychology of security for the home computer user. In *Security and Privacy (SP), 2012 IEEE Symposium on*, pages 209–223. IEEE, 2012.

[119] Robert Hurling, Michael Catt, Marco De Boni, Bruce William Fairley, Tina Hurst, Peter Murray, Alannah Richardson, and Jaspreet Singh Sodhi. Using internet and mobile phone technology to deliver an automated physical activity program: randomized controlled trial. *Journal of medical Internet research*, 9(2), 2007.

[120] Princely Ifinedo. Understanding information systems security policy compliance: An integration of the theory of planned behavior and the protection motivation theory. *Computers & Security*, 31(1):83–95, 2012.

[121] Infonetics. Enterprises rate mobile device security vendors, reveal byod concerns, 2012. URL http://www.infonetics.com/pr/2012/Enterprise-Mobile-Security-Strategies-Survey-Highlights.asp. Visited on 07/02/15.

[122] GSMA Intelligence. Definitive data and analysis for the mobile industry, 2015. URL https://gsmaintelligence.com/. Visited on 07/02/15.

[123] ISO. ISO/IEC 27001:2013 - information technology – security techniques – information security management systems – requirements, 2013. URL http://www.iso.org/iso/catalogue_detail?csnumber=54534. Visited on 04/01/15.

[124] Jeffrey M Jackson and Kipling D Williams. Social loafing on difficult tasks: Working collectively can improve performance. *Journal of Personality and Social Psychology*, 49(4):937, 1985.

[125] Irving Lester Janis. *Groupthink: Psychological studies of policy decisions and fiascoes*. Houghton Mifflin Boston, 1982.

[126] Luke Jefferson and Richard Harvey. An interface to support color blind computer users. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 1535–1538. ACM, 2007.

[127] Debora Jeske, Lynne Coventry, Pam Briggs, and James Turland. Less redundancy and increased relevance: Encouraging technical and security error reporting. *International Journal of Human-Computer Studies*, 2015.

[128] Debora Jeske, James Turland, Pam Briggs, and Lynne Coventry. Personality and framing factors in privacy decision-making: A study on cookie acceptance. *ACM TOCHI*, 2015.

[129] Eric J Johnson and Daniel G Goldstein. Do defaults save lives. In *SCIENCE*. Citeseer, 2003.

[130] M. Eric Johnson and Eric Goetz. Embedding information security into the organization. *IEEE Security and Privacy*, 5(3):16–24, May 2007. ISSN 1540-7993. URL http://dx.doi.org/10.1109/MSP.2007.59.

[131] Allen C. Johnston and Merrill Warkentin. Fear appeals and information security behaviors: An empirical study. *MIS Q.*, 34(3):549–566, September 2010. URL http://dl.acm.org/citation.cfm?id=2017470.2017478.

[132] Adam N Joinson, Carina Paine, Ulf-Dietrich Reips, and Tom Buchanan. Privacy and trust: The role of situational and dispositional variables in online disclosure. In *Workshop on Privacy, Trust and Identity Issues for Ambient Intelligence*, 2006.

[133] Stølen. K, Winsborough. W, H, Martinelli. F, and Massacci. F, editors. *Trust Management, 4th International Conference, iTrust 2006, Pisa, Italy, May 16-19, 2006, Proceedings*, volume 3986 of *Lecture Notes in Computer Science*, 2006. Springer.

[134] Yee. K. Guidelines and strategies for secure interaction design. *Security and Usability: Designing Secure Systems That People Can Use*, pages 247–273, 2005.

[135] Daniel Kahneman. *Thinking, fast and slow*. Farrar, Straus and Giroux, 2011.

[136] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*, pages 263–291, 1979.

[137] John A Kalapurakal, Aleksandar Zafirovski, Jeffery Smith, Paul Fisher, Vythialingam Sathiaseelan, Cynthia Barnard, Alfred W Rademaker, Nick Rave, and Bharat B Mittal. A comprehensive quality assurance program for personnel and procedures in radiation oncology: Value of voluntary error reporting and checklists. *International Journal of Radiation Oncology* Biology* Physics*, 86(2):241–248, 2013.

[138] Steven J Karau and Kipling D Williams. Social loafing: A meta-analytic review and theoretical integration. *Journal of personality and social psychology*, 65(4):681, 1993.

[139] Steven J Karau and Kipling D Williams. Social loafing: Research findings, implications, and future directions. *Current Directions in Psychological Science*, pages 134–140, 1995.

[140] Frank H Katz. The effect of a university information security survey on instruction methods in information security. In *Proceedings of the 2nd annual conference on Information security curriculum development*, pages 43–48. ACM, 2005.

[141] Marilyn Kingston, Susan Evans, Brian Smith, and Jesia Berry. Attitudes of doctors and nurses towards incident reporting: a qualitative analysis. *Medical Journal of Australia*, 181(1):36–39, 2004.

[142] Predrag Klasnja, Sunny Consolvo, and Wanda Pratt. How to evaluate technologies for health behavior change in hci research. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3063–3072. ACM, 2011.

[143] Judy Kopp. Self-monitoring: A literature review of research and practice. In *Social Work Research and Abstracts*, volume 24, pages 8–20. Oxford University Press, 1988.

[144] Coventry. L, Briggs. P, Jeske. D, and van Moorsel. A. Scene: A structured means for creating and evaluating behavioral nudges in a cyber security environment. In Aaron Marcus, editor, *Design, User Experience, and Usability. Theories, Methods, and Tools for Designing the User Experience*, volume 8517 of *Lecture Notes in Computer Science*, pages 229–239. Springer International Publishing, 2014.

[145] Pradip Lamsal. Understanding trust and security. *Department of Computer Science, University of Helsinki, Finland*, 2001.

[146] Doo-Hee Lee, Seunghee Im, and Charles R Taylor. Voluntary self-disclosure of information on the internet: A multimethod study of the motivations and consequences of disclosing information on blogs. *Psychology & Marketing*, 25(7):692–710, 2008.

[147] Jintae Lee and Younghwa Lee. A holistic model of computer abuse within organizations. *Information Management and Computer Security*, 10(2):57–63, 2002. URL http://dx.doi.org/10.1108/09685220210424104.

[148] Jeffery A. Lepine, Nathan P. Podsakoff, and Marcie A. Lepine. A meta-analytic test of the challenge stressor–hindrance stressor framework: An explanation for inconsistent relationships among stressors and performance. *Academy of Management Journal*, 48(5):764–775, 2005. doi: 10.5465/AMJ.2005.18803921. URL http://amj.aom.org/content/48/5/764.abstract.

[149] Irwin P Levin, Sandra L Schneider, and Gary J Gaeth. All frames are not created equal: A typology and critical analysis of framing effects. *Organizational behavior and human decision processes*, 76(2):149–188, 1998.

[150] Irwin P Levin, Gary J Gaeth, Judy Schreiber, and Marco Lauriola. A new look at framing effects: Distribution of effect sizes, individual differences, and independence of types of effects. *Organizational behavior and human decision processes*, 88(1): 411–429, 2002.

[151] Huigang Liang and Yajiong Xue. Avoidance of information technology threats: A theoretical perspective. *MIS Q.*, 33(1):71–90, March 2009. URL http://dl.acm.org/citation.cfm?id=2017410.2017417.

[152] Huigang Liang and Yajiong Xue. Understanding security behaviors in personal computer usage: A threat avoidance perspective. *Journal of the Association for Information Systems*, 11(7):394–413, 2010.

[153] Dan Lockton, David Harrison, and Neville Stanton. Design with intent: Persuasive technology in a wider context. In *Persuasive technology*, pages 274–278. Springer, 2008.

[154] K. Louise Barriball and A. While. Collecting data using a semi-structured interview: a discussion paper. *Journal of Advanced Nursing*, 19(2):328–335, 1994.

[155] Patrick D Lynch, Robert J Kent, and Srini S Srinivasan. The global internet shopper: evidence from shopping tasks in twelve countries. *Journal of advertising research*, 41 (3):15–24, 2001.

[156] Zuckerman M. Dimensions of sensation seeking. *Journal of Consulting and Clinical Psychology*, 36:45–52, 1971.

[157] Lena Mamykina, Elizabeth Mynatt, Patricia Davidson, and Daniel Greenblatt. Mahi: investigation of social scaffolding for reflective thinking in diabetes management. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 477–486. ACM, 2008.

[158] Jon K Maner and Mary A Gerend. Motivationally selective risk judgments: Do fear and curiosity boost the boons or the banes? *Organizational Behavior and Human Decision Processes*, 103(2):256–267, 2007.

[159] Max-Emanuel Maurer, Alexander De Luca, and Sylvia Kempe. Using data type based security alert dialogs to raise online security awareness. In *Proceedings of the Seventh Symposium on Usable Privacy and Security*, page 2. ACM, 2011.

[160] Roy A Maxion and Robert W Reeder. Improving user-interface dependability through mitigation of human error. *International Journal of human-computer studies*, 63(1): 25–50, 2005.

[161] Susan Michie, Marie Johnston, Jill Francis, Wendy Hardeman, and Martin Eccles. From theory to intervention: mapping theoretically derived behavioural determinants to behaviour change techniques. *Applied psychology*, 57(4):660–680, 2008.

[162] Microsoft. Microsoft privacy and security, 2015. URL https://msdn.microsoft.com/en-us/library/ms976532.aspx. visited on 21/04/15.

[163] K. D. Mitnick and W. L. Simon. *The art of deception: Controlling the human element of security*. John Wiley & Sons, Inc, 2002.

[164] Frank Mols, S Alexander Haslam, Jolanda Jetten, and Niklas K Steffens. Why a nudge is not enough: A social identity critique of governance by stealth. *European Journal of Political Research*, 54(1):81–98, 2015.

[165] Margaret E Morris, Qusai Kathawala, Todd K Leen, Ethan E Gorenstein, Farzin Guilak, Michael Labhard, and William Deleeuw. Mobile therapy: case study evaluations of a cell phone application for emotional self-awareness. *Journal of Medical Internet Research*, 12(2), 2010.

[166] Bill Morrow. Byod security challenges: control and protect your most sensitive data. *Network Security*, 2012(12):5–8, 2012.

[167] Siponen M. T. Pahnila S. Vartiainen T. Myyry, L. and A. Vance. What levels of moral reasoning and values explain adherence to information security rules? an empirical study. *European Journal of Information Systems*, 18(2):126–139, 2009.

[168] NCSA. Symantec national small business study. *National Cyber Security Alliance, Symantec, JZ Analytics*, 20, 2012.

[169] Rosemery O Nelson. Assessment and therapeutic functions of self-monitoring. *Progress in behavior modification*, 5:263–308, 1977.

[170] Boon-Yuen Ng, Atreyi Kankanhalli, and Yunjie (Calvin) Xu. Studying users' computer security behavior: A health belief perspective. *Decision Support Systems*, 46(4):815 – 825, 2009. ISSN 0167-9236. doi: http://dx.doi.org/10.1016/j.dss.2008.11.010. URL http://www.sciencedirect.com/science/article/pii/S0167923608002157. {IT} Decisions in Organizations.

[171] Donald A Norman. *The design of everyday things*. Basic books, 2002.

[172] Jum C Nunnally, Ira H Bernstein, and Jos MF ten Berge. *Psychometric theory*, volume 226. McGraw-Hill New York, 1967.

[173] Jason RC Nurse, Sadie Creese, Michael Goldsmith, and Koen Lamberts. Guidelines for usable cybersecurity: Past and present. In *Cyberspace Safety and Security (CSS), 2011 Third International Workshop on*, pages 21–26. IEEE, 2011.

[174] Jason RC Nurse, Sadie Creese, Michael Goldsmith, and Koen Lamberts. Trustworthy and effective communication of cybersecurity risks: A review. In *Socio-Technical Aspects in Security and Trust (STAST), 2011 1st Workshop on*, pages 60–68. IEEE, 2011.

[175] House of Commons Science and Technology Committee. Responsible use of data: Fourth report of session 2014–1. URL http://www.publications.parliament.uk/pa/cm201415/cmselect/cmsctech/245/24502.htm.

[176] Cambridge Dictionaries Online. Cambridge dictionaries online: English definition of "trust", 2015. URL http://dictionary.cambridge.org/dictionary/british/trust. Visited on 15/01/15.

[177] OpenTracker. Third-party cookies vs first-party cookies. URL http://www.opentracker.net/article/third-party-cookies-vs-first-party-cookies.

[178] Pfleeger P. *Security in Computing*. Prentice Hall, Inc, first edition, 1997.

[179] Seppo Pahnila, Mikko Siponen, and Adam Mahmood. Employees' behavior towards IS security policy compliance. In *System Sciences, 2007. HICSS 2007. 40th Annual Hawaii International Conference on*, pages 156b–156b. IEEE, 2007.

[180] Marco Perugini, Marcello Gallucci, Fabio Presaghi, and Anna Paola Ercolani. The personal norm of reciprocity. *European Journal of Personality*, 17(4):251–283, 2003.

[181] Richard E Petty, John T Cacioppo, and Rachel Goldman. Personal involvement as a determinant of argument-based persuasion. *Journal of personality and social psychology*, 41(5):847, 1981.

[182] Shari Lawrence Pfleeger, Joel B Predd, Jeffrey Hunker, and Carla Bulford. Insiders behaving badly: addressing bad actors and their actions. *Information Forensics and Security, IEEE Transactions on*, 5(1):169–179, 2010.

[183] Nathan P. Podsakoff, Jeffery A. LePine, and Marcie A. LePine. Differential challenge stressor-hindrance stressor relationships with job attitudes, turnover intentions, turnover, and withdrawal behavior: A meta-analysis. *Journal of Applied Psychology*, 92(2):438–454, 2007. ISSN 1939-1854, 0021-9010. doi: 10.1037/0021-9010.92.2.438. URL http://doi.apa.org/getdoi.cfm?doi=10.1037/0021-9010.92.2.438.

[184] Ponemon. Confidential documents at risk study, 2012. URL http://www.ponemon.org/local/upload/file/WatchDoxWhite_Paper_FINAL.pdf. Visited on 12/02/15.

[185] James O Prochaska and Wayne F Velicer. The transtheoretical model of health behavior change. *American journal of health promotion*, 12(1):38–48, 1997.

[186] The IRIDIUM Project. About our project, 2011-2013. URL https://research.ncl.ac.uk/iridium/aboutourproject/. Visited on 26/04/13.

[187] Lazarus. R, S. *Psychological Stress and the Coping Process*. New York: McGraw-Hill, 1966.

[188] Lazarus R, S and S. Folkman. *Stress, Appraisal, and Coping*. New York: Springer, 1984.

[189] T. S. Ragu-Nathan, Monideepa Tarafdar, Bhanu S. Ragu-Nathan, and Qiang Tu. The consequences of technostress for end users in organizations: Conceptual development and empirical validation. *Information Systems Research*, 19(4):417–433, 2008. URL http://pubsonline.informs.org/doi/abs/10.1287/isre.1070.0165.

[190] Fahimeh Raja, Kirstie Hawkey, Steven Hsu, Kai-Le Clement Wang, and Konstantin Beznosov. A brick wall, a locked door, and a bandit: a physical security metaphor for firewall warnings. In *Proceedings of the Seventh Symposium on Usable Privacy and Security*, page 1. ACM, 2011.

[191] Jan C Recker. Bpmn modeling–who, where, how and why. *BPTrends*, 5(3):1–8, 2008.

[192] Paul Resnick et al. The social cost of cheap pseudonyms. *Journal of Economics & Management Strategy*, 10(2):173–199, 2001.

[193] Jens Riegelsberger, M. Angela Sasse, and John D. McCarthy. The mechanics of trust: A framework for research and design. *International Journal of Human-Computer Studies*, 62(3):381 – 422, 2005. URL http://www.sciencedirect.com/science/article/pii/S1071581905000121.

[194] Tina L Robbins. Social loafing on cognitive tasks: An examination of the "sucker effect". *Journal of Business and Psychology*, 9(3):337–342, 1995.

[195] Jessica B Rodell and Timothy A Judge. Can "good" stressors spark "bad" behaviors? the mediating role of emotions in links of challenge and hindrance stressors with citizenship and counterproductive behaviors. *Journal of Applied Psychology*, 94(6): 1438, 2009.

[196] Yvonne Rogers, Helen Sharp, and Jenny Preece. *Interaction design: beyond human-computer interaction*. John Wiley & Sons, 2011.

[197] N. Russell, A.H.M. Ter Hofstede, D. Edmond, and W.M.P. van der Aalst. Workflow data patterns. Technical report, QUT Technical report, FIT-TR-2004-01, Queensland University of Technology, Brisbane, 2004.

[198] Pahnila. S, Siponen. M, and Mahmood. A. Employees' behavior towards is security policy compliance. In *System Sciences, 2007. HICSS 2007. 40th Annual Hawaii International Conference on*, pages 156b–156b, Jan 2007.

[199] Pfleeger. S, L and Caputo. D, D. Leveraging behavioral science to mitigate cyber security risk. *Computers & Security*, 31(4):597 – 611, 2012. ISSN 0167-4048.

[200] Pamela Samuelson. Privacy as intellectual property? *Stanford Law Review*, pages 1125–1173, 2000.

[201] SANS. The business justification of data security. URL http://www.sans.org/reading-room/whitepapers/dlp/business-justification-data-security-33033. Visited on 12/06/2012.

[202] Harald Schmidt, Kristin Voigt, and Daniel Wikler. Carrots, sticks, and health care reform — problems with wellness incentives. *New England Journal of Medicine*, 362 (2):e3, 2010. ISSN 0028-4793. doi: 10.1056/NEJMp0911552.

[203] Stefano Scoglio. *Transforming privacy: A transpersonal philosophy of rights*. Greenwood Publishing Group, 1998.

[204] IEEE Cyber Security. Avoiding the top 10 security flaws, 2014. URL http://cybersecurity.ieee.org/center-for-secure-design/avoiding-the-top-10-security-flaws.html. Visited on 22/01/15.

[205] James A Shepperd. Social loafing and expectancy-value theory. In *Multiple perspectives on the effects of evaluation on performance*, pages 1–24. Springer, 2001.

[206] Mikko Siponen and Anthony Vance. Neutralization: New insights into the problem of employee systems security policy violations. *MIS Q.*, 34(3):487–502, September 2010. URL http://dl.acm.org/citation.cfm?id=2017470.2017475.

[207] Skift. Top 25 online booking sites in travel for october 2013. URL http://skift.com/2013/11/11/top-25-online-booking-sites-in-travel/.

[208] TomTom Skype, Norton. Survey finds nearly half of consumers fail to upgrade software regularly and one quarter of consumers don't know why to update software. URL http://about.skype.com/press/2012/07/survey_finds_nearly_half_fail_to_upgrade.html. Visited on 04/02/14.

[209] Brian K Smith, Jeana Frost, Meltem Albayrak, and Rajneesh Sudhakar. Integrating glucometers and digital photography as experience capture tools to enhance patient understanding and communication of diabetes self-management practices. *Personal and Ubiquitous Computing*, 11(4):273–286, 2007.

[210] H Jeff Smith, Tamara Dinev, and Heng Xu. Information privacy research: an interdisciplinary review. *MIS quarterly*, 35(4):989–1016, 2011.

[211] Ryan Spellecy. Reviving ulysses contracts. *Kennedy Institute of Ethics Journal*, 13(4): 373–392, 2003.

[212] Matthew S Stanford, Charles W Mathias, Donald M Dougherty, Sarah L Lake, Nathaniel E Anderson, and Jim H Patton. Fifty years of the barratt impulsiveness scale: An update and review. *Personality and Individual Differences*, 47(5):385–395, 2009.

[213] K E Stanovich and R F West. Individual differences in reasoning: Implications for the rationality debate. *Behavioral and Brain Sciences*, 23:645–726, 2000.

[214] Jeffrey M. Stanton, Kathryn R. Stam, Paul Mastrangelo, and Jeffrey Jolton. Analysis of end user security behaviors. *Comput. Secur.*, 24(2):124–133, March 2005. URL http://dx.doi.org/10.1016/j.cose.2004.07.001.

[215] StaySafeOnline. Assess your risk, 2016. URL https://staysafeonline.org/business-safe-online/assess-your-risk. visited on 22/02/16.

[216] F Marijn Stok, Denise TD De Ridder, Emely De Vet, and John BF De Wit. Minority talks: the influence of descriptive social norms on fruit intake. *Psychology & health*, 27(8):956–970, 2012.

[217] Luis Jesús Ramón Stroppi, Omar Chiotti, and Pablo David Villarreal. Extending bpmn 2.0: method and tool support. In *Business Process Model and Notation*, pages 59–73. Springer, 2011.

[218] Cass R Sunstein. Nudging: a very short guide. *Journal of Consumer Policy*, 37(4): 583–588, 2014.

[219] Symantec. Symantec internet security threat report: Trends for 2008. URL http://eval.symantec.com/mktginfo/enterprise/white_papers/b-whitepaper_internet_security_threat_report_xiv_04-2009.en-us.pdf.

[220] Symantec. Internet security threat report 2011, 2011. URL https://www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_2011_21239364.en-us.pdf. Visited on 04/02/2015.

[221] Symantec. Exploiting the business potential of BYOD (bring your own device), 2014. URL https://scm.symantec.com/en/d/eu/mobility/resources/wp_apj_exploiting_the_business_potential_of_byod.pdf/. Visited on 21/10/14.

[222] Monideepa Tarafdar, Qiang Tu, and T. Ragu-Nathan. Impact of technostress on end-user satisfaction and performance. *J. Manage. Inf. Syst.*, 27(3):303–334, January 2010. URL http://dx.doi.org/10.2753/MIS0742-1222270311.

[223] Curtis R Taylor. Private demands and demands for privacy: Dynamic pricing and the market for customer information. 2002.

[224] TechNet. Windows error reporting and the problem reports and solutions feature in windows vista. URL https://technet.microsoft.com/en-us/library/cc709644%28v=ws.10%29.aspx. Visited on 21/05/15.

[225] Herath. Tejaswini and Rao. H, Raghav. Protection motivation and deterrence: a framework for security policy compliance in organisations, 2009. URL http://www.palgrave-journals.com/ejis/journal/v18/n2/full/ejis20096a.html.

[226] Tenable. Mobile device vulnerability management flagged as top concern for security professionals in 2012, 2012. URL http://www.tenable.com/press-releases/mobile-device-vulnerability-management-flagged-as-top-concern-for-security. Visited on 04/02/2015.

[227] Deborah J Terry and Michael A Hogg. Group norms and the attitude-behavior relationship: A role for group identification. *Personality and Social Psychology Bulletin*, 22(8):776–793, 1996.

[228] Richard H Thaler and Cass R Sunstein. *Nudge*. Yale University Press, 2008.

[229] Richard H. Thaler, Cass R. Sunstein, and John P. Balz. *Choice Architecture*. The Behavioral Foundations of Public Policy, 2014.

[230] Dianne M. Tice, Roy F. Baumeister, Dikla Shmueli, and Mark Muraven. Restoring the self: Positive affect helps improve self-regulation following ego depletion. *Journal of Experimental Social Psychology*, 43(3):379 – 384, 2007. URL http://www.sciencedirect.com/science/article/pii/S0022103106000862.

[231] James Turland, Lynne Coventry, Debora Jeske, Pam Briggs, and Aad van Moorsel. Nudging towards security: Developing an application for wireless network selection for android phones. *British HCI 2015*, 2015.

[232] Amos Tversky. Elimination by aspects: A theory of choice. *Psychological review*, 79 (4):281, 1972.

[233] Amos Tversky and Daniel Kahneman. The framing of decisions and the psychology of choice. *Science*, 211(4481):453–458, 1981.

[234] Gov UK. Information security: Making sure user data stays secure, 2014. URL https://www.gov.uk/service-manual/making-software/information-security.html. Visited on 04/01/2015.

[235] Children's Services U.K. Office for Standards in Education and Skills. The safe use of new technologies. *Report 090231*, 2010. UK: Ofsted; February 2010.

[236] Blase Ur, Patrick Gage Kelley, Saranga Komanduri, Joel Lee, Michael Maass, Michelle L Mazurek, Timothy Passaro, Richard Shay, Timothy Vidas, Lujo Bauer, et al. How does your password measure up? the effect of strength meters on password creation. In *USENIX Security Symposium*, pages 65–80, 2012.

[237] Marleen H Van den Berg, Johannes W Schoones, and Theodora PM Vliet Vlieland. Internet-based physical activity interventions: a systematic review of the literature. *Journal of medical Internet research*, 9(3), 2007.

[238] Wil MP van der Aalst, Arthur HM ter Hofstede, Nick Russell, Bartek Kiepuszewski, Alistair Barros, Marlon Dumas, and Petia Wohed. Workflow Patterns | Patterns | Data - Workflow Data. http://www.workflowpatterns.com/patterns/data/visibility/wdp7.php, 2010. URL http://www.workflowpatterns.com/patterns/data/visibility/wdp7.php.

[239] Anthony Vance, Mikko Siponen, and Seppo Pahnila. Motivating is security compliance: Insights from habit and protection motivation theory. *Information & Management*, 49(3):190–198, 2012.

[240] Peter J Veazie. An individual-based framework for the study of medical error. *International Journal for Quality in Health Care*, 18(4):314–319, 2006.

[241] VHM Visschers, PM Wiedemann, H Gutscher, S Kurzenhäuser, R Seidl, CG Jardine, and DRM Timmermans. Affect-inducing risk communication: current knowledge and future directions. *Journal of Risk Research*, 15(3):257–271, 2012.

[242] Kathleen D. Vohs and Todd F. Heatherton. Self-regulatory failure: A resource-depletion approach. *Psychological Science*, 11(3):249–254, 2000. URL http://pss.sagepub.com/content/11/3/249.abstract.

[243] Kevin G Volpp, George Loewenstein, Andrea B Troxel, Jalpa Doshi, Maureen Price, Mitchell Laskin, and Stephen E Kimmel. A test of financial incentives to improve warfarin adherence. *BMC Health Services Research*, 8(1):272, 2008.

[244] Basie Von Solms. Information security—a multidimensional discipline. *Computers & Security*, 20(6):504–508, 2001.

[245] Yang Wang, Pedro Giovanni Leon, Kevin Scott, Xiaoxuan Chen, Alessandro Acquisti, and Lorrie Faith Cranor. Privacy nudges for social media: an exploratory facebook study. In *Proceedings of the 22nd international conference on World Wide Web companion*, pages 763–770. International World Wide Web Conferences Steering Committee, 2013.

[246] Merrill Warkintin and Robert Willison. Behavioral and policy issues in information systems security: The insider threat. *European Journal of Information Systems*, 18(2): 101–105, 2009. ISSN 0960-085X. doi: 10.1057/ejis.2009.12.

[247] Molly McLure Wasko and Samer Faraj. Why should i share? examining social capital and knowledge contribution in electronic networks of practice. *MIS quarterly*, pages 35–57, 2005.

[248] Dirk Weirich and Martina Angela Sasse. Pretty good persuasion: A first step towards effective password security in the real world. In *Proceedings of the 2001 Workshop on New Security Paradigms*, NSPW '01, pages 137–143, New York, NY, USA, 2001. ACM. URL http://doi.acm.org/10.1145/508171.508195.

[249] Dirk Weirich and Martina Angela Sasse. Persuasive password security. In *CHI'01 Extended Abstracts on Human Factors in Computing Systems*, pages 139–140. ACM, 2001.

[250] T. Wengraf. *Qualitative research interviewing: Biographic narrative and semi-structured methods*. Sage Publications Ltd, 2001.

[251] Lawrence R Wheeless. A follow-up study of the relationships among trust, disclosure, and interpersonal solidarity. *Human Communication Research*, 4(2):143–157, 1978.

[252] Stephen White. Introduction to bpmn, 2004. URL http://yoann.nogues.free.fr/IMG/pdf/07-04_WP_Intro_to_BPMN_-_White-2.pdf. BPTrends.com.

[253] Alma Whitten and J Doug Tygar. Why johnny can't encrypt: A usability evaluation of pgp 5.0. In *Usenix Security*, volume 1999, 1999.

[254] C Arthur Williams. Attitudes toward speculative risks as an indicator of attitudes toward pure risks. *The Journal of Risk and Insurance*, 33(4):577–586, 1966.

[255] Michael Workman. Gaining access with social engineering: An empirical study of the threat. *Information Systems Security*, 16(6):315–331, 2007. URL http://dx.doi.org/10.1080/10658980701788165.

[256] Michael Workman, William H. Bommer, and Detmar Straub. Security lapses and the omission of information security measures: A threat control model and empirical test. *Comput. Hum. Behav.*, 24(6):2799–2816, September 2008. URL http://dx.doi.org/10.1016/j.chb.2008.04.005.

[257] Michael Workman, William H Bommer, and Detmar Straub. Security lapses and the omission of information security measures: A threat control model and empirical test. *Computers in Human Behavior*, 24(6):2799–2816, 2008.

[258] Iryna Yevseyeva, Charles Morisset, James Turland, Lynne Coventry, Thomas Groß, Christopher Laing, and Aad van Moorsel. Consumerisation of it: Mitigating risky user actions and improving productivity with nudging. *Procedia Technology*, 16:508–517, 2014.

[259] Ernst & Young. Bring your own device: Mobile security and risk, 09/2013. URL http://www.ey.com/Publication/vwLUAssets/EY_-_Bring_your_own_device:_mobile_security_and_risk/$FILE/Bring_your_own_device.pdf. Visited on 01/02/15.

[260] Ernst & Young. Moving beyond compliance: Ernst and Young's 2008 global information security survey. 2008. URL http://www.ey.com/security.

[261] Stephen J Zaccaro. Social loafing the role of task attractiveness. *Personality and Social Psychology Bulletin*, 10(1):99–106, 1984.

[262] Tomasz Zaleskiewicz. Beyond risk seeking and risk aversion: Personality and the dual nature of economic risk taking. *European Journal of Personality*, 15:S105–S122, 2001.

[263] Marcel Zeelenberg, Jane Beattie, Joop Van der Pligt, and Nanne K de Vries. Consequences of regret aversion: Effects of expected feedback on risky decision making. *Organizational behavior and human decision processes*, 65(2):148–158, 1996.

[264] Leah Zhang-Kennedy, Sonia Chiasson, and Robert Biddle. Stop clicking on "update later": Persuading users they need up-to-date antivirus protection. In *Persuasive Technology*, pages 302–322. Springer, 2014.

[265] Bin Zhao and Fernando Olivera. Error reporting in organizations. *Academy of Management Review*, 31(4):1012–1030, 2006.