



Delft University of Technology

The boundedly rational employee

Security economics for behaviour intervention support in organizations

Demjaha, Albesë; Parkin, Simon; Pym, David

DOI

[10.3233/JCS-210046](https://doi.org/10.3233/JCS-210046)

Publication date

2022

Document Version

Accepted author manuscript

Published in

Journal of Computer Security

Citation (APA)

Demjaha, A., Parkin, S., & Pym, D. (2022). The boundedly rational employee: Security economics for behaviour intervention support in organizations. *Journal of Computer Security*, 30(3), 435-464. <https://doi.org/10.3233/JCS-210046>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

The boundedly rational employee: Security economics for behaviour intervention support in organizations

Albesë Demjaha ^a,

^a *University College London and The Alan Turing Institute,
London, United Kingdom
E-mail: albesa.demjaha.16@ucl.ac.uk*

Simon Parkin ^b,

^b *Delft University of Technology, Delft, The Netherlands
E-mail: s.e.parkin@tudelft.nl*

David Pym ^c,

^c *University College London and The Institute of Philosophy,
University of London, London, United Kingdom
E-mail: d.pym@ucl.ac.uk*

Abstract.

Security policy-makers (influencers) in an organization set security policies that embody intended behaviours for employees (as decision-makers) to follow. Decision-makers then face choices, where this is not simply a binary decision of whether to comply or not, but also *how* to approach compliance and secure working alongside other workplace pressures, and limited resources for identifying optimal security-related choices. Conflict arises because of information asymmetries present in the relationship, where influencers and decision-makers both consider costs, gains, and losses in ways which are not *necessarily* aligned. With the need to promote 'good enough' decisions about security-related behaviours under such constraints, we hypothesize that actions to resolve this misalignment can benefit from constructs from both traditional economics *and* behavioural economics. Here we demonstrate how current approaches to security behaviour provisioning in organizations mirror rational-agent economics, even where behavioural economics is embodied in the promotion of individual security behaviours. We develop and present a framework to accommodate *bounded security decision-making*, within an ongoing programme of behaviours which must be provisioned for and supported. Our four stage plan to Capture, Adapt, Realign, and Enable behaviour choices provides guidance for security managers, focusing on a more effective response to the uncertainty associated with security behaviour in organizations.

Keywords: security decision-making, security economics, security policy, security behaviour modelling

1. Introduction

Information security in larger organizations is often managed by an information security manager and/or a security team — the (information) *security function* of the organization. The security function is recognized as having the expertise to identify and manage the security technologies and processes

necessary to protect the organization from threats to its assets. Outwardly, this is embodied in controls and procedures, often detailed in the organization's security policy (or policies).

Policy may dictate specific *security-related behaviours*, which employees are expected to adopt. There are myriad ways to promote adoption of a behaviour [22], with challenges in guaranteeing successful change [80]. Declaring a behaviour in a security policy is not an assurance that the behaviour will happen. This reality has drawn increasing attention to the need to manage behaviour change effectively. Consideration of behaviour change theory and *behavioural economics* [20] is one such approach.

Both research and practice have shown that behaviours defined in policy may not be adopted in organizations. Employees may not see how policy applies to them, find it difficult to follow, or regard policy expectations as unrealistic [52] (where they may well be [46]). Rather than abandoning security, employees may create their own alternative behaviours [18], in an effort to approximate secure working [53]. Organizational support can be critical to whether secure practices persist [35] or whether security-related issues begin to be delegated to others.

To compound these challenges, security policies in organizations may not fit the security needs and ways of working of different employee groups; this extends to how they perceive and address security risks [13, 19]. General advice may be relevant but not actionable, and specific advice only applicable in very particular circumstances [77]. It is then in the security function's interest to address the distinct security needs of (official or informal) groups of users. For these reasons, the sense of how to work securely can also come from colleagues [56] or managers [52]. If solutions are more specific to a discernible group, the more immediately usable they are for that group, albeit with a need to tailor those solutions upfront. There are, however, benefits to understanding group-specific needs, as with phishing guidance [91], where distinct business functions have their own challenges (e.g., a finance team discerning a genuine invoice email from a fake one, where other functions would simply not expect to receive invoices).

The security function must have a strategy for how to provision for security, provide workable policy, and support user needs. Rational security micro-economics has proved useful for explaining the interaction between organizational security policies and behaviours [4, 14], where security ecosystems are otherwise too complicated to study directly in this way. We revisit principles of information economics and behavioural economics in tandem, identifying contradictions which point to gaps in support.

The paper is structured in the following way:

- In Section 2, we provide the foundational terminology and definitions that underpin bounded rationality for security in organizations.
- We review the capacity for various economics principles to explain a range of security-related behaviours (Section 3).
- We then demonstrate how current approaches to infrastructure and provisioning of security mirror unbounded, rational-agent economics, even when techniques more familiar to bounded-agent behavioural economics are applied to promote individual behaviours (Section 4). We show through examples how these contradictions align with regularly cited causes of security non-compliance from the literature, and point to more appropriate solutions.
- We present a framework (Section 5), based on consolidated economics principles, with the following goal:

Better support for 'good enough' security-related decisions, by individuals and groups within an organization, that best approximate secure behaviours under constraints, such as limited time or knowledge.

This requires us to identify the factors affecting security behaviours that should be considered by the organization in order to inform policy design, support the identification of provisioning requirements, and describe expectations of users. The framework is intended to underpin provisioning to reach this goal. In considering the factors affecting security behaviours, a clearer consideration can be conducted for establishing assumptions about target groups within an organization. At present, security controls are generally applied to all users, and interventions are targeted at all users — such assumptions are less likely to be reliable as organizations become larger and more complex.

- We then apply the framework to one of the most widely promoted security behaviours (Section 6), the maintenance of up-to-date device software, demonstrating through comparison with independent user studies where the consolidated economics approach — *bounded security decision-making* — can anticipate organizational support requirements.
- We consider how the framework can be situated to support practitioners (Section 7), before a review of related work (Section 8) and a concluding summary (Section 9).
- A supporting glossary of economics terminology is detailed in the Appendix.

2. Terminology

Before we present our conceptual framework, we must make a choice of terminology and definitions. It is important to define explicitly the central concept of this paper, that of *bounded rationality*. We turn to the definition of Herbert Simon, who introduced the concept of bounded rationality with the following motivation:

‘Broadly stated, the task is to replace the global rationality of economic man with the kind of rational behavior that is compatible with the access to information and the computational capacities that are actually possessed by organisms, including man, in the kinds of environments in which such organisms exist.’ [88, p.99]




Fig. 1. An illustration of Bounded Rationality.

We define bounded rationality as *the idea that humans have certain constraints — cognitive, time, and information/knowledge — that limit their decision-making processes; as such, bounded rational*

agents turn to ‘satisficing’ rather than ‘optimizing’ [10, 88–90]. The components of our definition are illustrated in Figure 1.

Cognitive constraints may include limits on memory or numerical processing, as well as limited skills or knowledge. Having *time constraints* means that a bounded rational agent does not have unlimited time to make a decision — and may therefore be required to sometimes rush a decision. *Incomplete information* suggests that in the real world, complete information is unattainable and the bounded rational agent must take decisions without necessarily knowing everything relevant about that decision. *Satisficing* — in contrast to optimizing — is a decision-making strategy which accounts for the above limitations and accepts a ‘good enough’ choice rather than aiming for one, unrealistic, ‘optimal choice’ [89]. When mentioning an ‘optimal choice’ in the paper, we are referring to it within the context of traditional economics rationality, and not within the context of security policy. When referring to the latter, we will explicitly mention that it is a security choice, or behaviour, mandated by policy, or the organization.

3. Towards a consistent strategy

Existing research has identified that organizational security approaches may ask too much of staff [4, 14], or leave staff to adapt advice to their local context [52]. Organizations often have a recognized security function, as an individual or team. We assume for simplicity that this function is also responsible for influencing the behaviour of employees toward it being more secure. This involves defining rules which employees must follow. Employees are typically assumed to be able to apply that advice as it is, and that it both addresses their security needs and does not impact any other, non-security needs that the decision-maker has [14]. That is, that the decision is natural, easy to make, and carries no unexpected consequences. We show that this is not a robust set of assumptions, through the following dilemmas. Where such dilemmas emerge, it creates a situation where an employee has a decision to make which draws further on their resources rather than being natural and free of effort.

Respect me and my time, or we are off to a bad start. Security behaviour provisions tend to imply that the decision-maker has resources available to complete training and policies, but in an organization the decision-maker is busy with their paid job.

To avoid ‘decision fatigue’ and the ‘hassle factor’ [4, 14] of complying with security, we must acknowledge that for the busy decision-maker, moving resources to security can require a loss to *something else*. This requires an institutional view to helping the decision-maker to negotiate where that cost will be borne from. The notion of a ‘Compliance Budget’ [4, 14] suggests to reduce the demands of security expectations, where here we note the need for an upper bound on expectations. This then leads to a design principle, that **any additional time needed to identify and apply a security-related behaviour must be negotiated for employees, not by employees.**

If this is guidance, be the guide. The security function must assume that employees are (security) novices. They then will need to be told the cost of security and exactly what the steps are. Otherwise, the novice must guess the duration of an unfamiliar behaviour, and exactly what constitutes the behaviour in its entirety (e.g., knowing where to find personal firewall settings [74]). An employee is also likely to base their security behaviours on their interaction with colleagues or managers [68] and follow their behaviours instead of the policy advocated ones. Unchecked, this leads to satisficing (see Section 2). Current approaches appeal to the skillful user, or assume ‘non-divisible’ target behaviours [5] with only

one, clear way to do what is being asked (that establishing what must be done is not itself a draw on resources). Another design principle is then to **move toward providing a clear way to apply a security-related behaviour , that will match known work-related situations or group dynamics and social influences on policy compliance** — or otherwise that help is available to find an alternative.

Frame a decision to make, not a decision made. Advice is given assuming that what is advised is the best choice, and there is no other choice to be articulated. The advocated choice is rarely, if ever, presented alongside other choices (such as previous sanctioned behaviours, or ad hoc, ‘shadow security’ behaviours that were not sanctioned but have been used in the past). A choice is often *perceived*, meaning there can be a ‘gulf of evaluation’ [79] between perception and reality, which must be navigated. An example is when users form incomplete/incorrect understanding of how provisioned two-factor authentication technology options work in reality [36]). A principle here is to **reduce the need for the decision-maker to invest resources in establishing which perceived choice is appropriate**, so that one advisable, workable behaviour becomes apparent, preferred, and approachable, and hence more likely to be adopted. Such framing is one means to de-emphasise harmful behaviours in the environment [28]. This also requires understanding what appeals in non-advocated behaviours and accounting for it in the advocated behaviour, so that there is a workable choice available [54].

Edit out the old, edit in the new. Providing more security advice is often presumed to be better for security, but is not [42], and can instead create confusion. Stale advice can persist unless it is curated [43] – an employee may do the wrong thing which is insecure, or the wrong thing which *was* secure but now is not. When policies and technologies change, the decision-maker is often left to do the choice-editing. An example is when old and new security policies are accessible but without time-stamps. The principle here is twofold — if the security function is committed to defining security behaviours, they should also **commit to managing the communication of secure behaviours**, rather than leaving it to employees to identify the right behaviour to follow. This includes explicitly communicating behaviours which are no longer advocated choices. It is unrealistic to expect employees to go direct to the source every time, where they may instead ask managers or peers how to address dilemmas [53].

3.1. Why we are here, with too few choices

To explore how employee security behaviours relate to security provisioning in organizations, we consider traditional economics and behavioural economics in the context of supporting effective behaviour change. We derived the ‘pillars’ of behaviour change from the COM-B model [64]: *Capability*, *Opportunity*, and *Motivation*, which are all required to support a change to a particular *Behaviour*. We discuss how each pillar is represented in organizations according to the two economic approaches.

Traditional economics.

The move from centralized to decentralized computing [69] has resulted in an imposed information asymmetry of having a recognized security function distinct from everyone else in the organization. An asymmetry may occur between the organization and the employee if the security function declares formal rules and informal rules (training, behaviours) [69], assuming that the decision-maker (individual employee) has the same knowledge that they do.

Conversely, the security function does not know about expectations placed on the decision-maker by other functions, assuming they have the capacity to approximate the same knowledge; Capability to enact a behaviour then cannot be assumed. Motivation comes from formal policies, and architectural means

Table 1

Examples of ‘contradictory’ and ‘better’ approaches to supporting secure behaviours in organizations (derived from experiences reported in real-world settings, and relevant studies).

Behaviour	Contradictory Approach	Failures
Policy compliance	Publishing policy without communicating location to staff [53]	Assumes knowledge of policy and time to find it
Secure passwords	Not communicating the rules for a secure password [70]	Assumes expert knowledge about passwords
Authentication choice	Integrating a suite of options into log-on without explaining the options [36]	Lacking support for making reasoned decision
Do secure work	Advocating generic security practices [52]	Staff must relate practices to work
Trained-for security	Provide training but no time to do it [4, 14]	Staff must negotiate the time themselves
Top-down training	Provide training which suspects colleagues (e.g., screen-locking)	Staff must reconcile training with local/team culture
Behaviour	Better Approach	Successes
Policy compliance	Ensure that the environment naturally supports policy-compliant behaviour [53] / relay policy directives to managers	Does not assume any extra effort from staff / managers can relay policy to their team
Secure passwords	Examples of ‘strong’ passwords (CyberAware UK)	Assumes little-to-no prior knowledge
Authentication choice	Communicating the different options in a suite of options at the point of configuration	Puts choices side-by-side
Do secure work	Visible board-level support [33], sector-specific tailoring (e.g., differentiated NCSC Guidance for Small Biz. and Small Charities)	Supports interpretation of a perceived choice to business conditions
Trained-for security	Agree a fixed window of paid time to complete training	Cost to (pri. and sec.) tasks negotiated for staff
Top-down training	Provide training which respects trust and culture between colleagues [56]	Staff do not have to feel awkward applying training

which force certain behaviours [69]. However, if Motivation to follow security rules is not sufficiently related to the assets which the decision-maker cares about, it will not support the recognition of risks which require the behaviour [17] (also impacting Opportunity).

Behavioural economics. In organizations, capabilities must be supported, but this is often approached in a ‘one-size-fits-all’ way, such that the decision-maker is forced, through the Motivation of enforced

formal rules, to seek out the knowledge to develop the Capabilities they need. however, they may not know if they have the complete and correct knowledge unless someone with that knowledge checks (and acts to reduce the information asymmetry). An Opportunity for a new behaviour may be created, through training or shaping of the environment, and *assumed* to be the provision of behaviour beneficial to the decision-maker [80]. If a choice of a particular behaviour is beneficial to a decision-maker, they may be assumed to gravitate toward it, even when that choice is offered alongside existing behaviours. If a behaviour is only assumed to be beneficial by the entity creating it, but whether it is that is not to say that it is also guaranteed to be beneficial according to the decision-maker's view and needs – Motivation is then not guaranteed. If the provisioned choices (the assumed Capability) are not perceived as beneficial they may instead adopt 'shadow security' behaviours [53] which better match their available resources.

Examples of 'contradictory' and 'better' approaches to supporting secure behaviours in organizations are illustrated through real-world examples in Table 1.

4. Applying economics principles to organizational security

We demonstrate how a strategic approach is lacking in how to manage the relatively high *marginal costs* of realizing the informal rules (awareness and culture) [69] which are intended to support formal rules in an organization. Although Enforcement costs are low for informal rules [69], there is a reliance on 'tacit consent' to govern appropriate behaviour. For example, an employee may have freedom to access the Internet as they see fit while working, but it would be assumed that they will not visit dubious websites. This then raises the issue that the centrally defined security policy may not be the only set of rules that informs employee behaviour, and that cues as to how to address security may come from colleagues as well, or be discussed or agreed implicitly within groups or indeed by non-work cultural aspects. Critically, culturally-derived behaviours cannot be managed hierarchically, but instead through 'informal control', with every employee influencing to some extent how others around them behave. These aspects of organization behaviour indicate that it is not always obvious how to follow appropriate, doable security behaviours.

4.1. Rational vs. bounded decision-making

In traditional economics, a decision-making structure assumes a rational agent [89, 90]. The rational decision-maker is equipped with the capabilities and resources to make the decision which will be most beneficial for them. The decision-maker knows all possible choices, and is assumed to have complete information when evaluating those choices, as well as a detailed analysis of probability, costs, gains, and losses [90]. The decision-maker is then capable of making an informed decision that is simultaneously the optimal decision for them.

Behavioural economics, on the other hand, challenges the assumption that decision-makers make fully rational decisions. Instead, the field refers to the concept of *bounded rationality*, which explains that a decision-maker's rationality is bounded because of cognitive limitations, resource constraints and time restrictions. These considerations also challenge the plausibility of complete information, which is practically unrealistic for a bounded decision-maker. According to these restrictions, the bounded decision-maker turns instead to 'rules of thumb' and makes ad hoc decisions based on a quick evaluation of *perceived* probability, costs, gains, and losses [48, 89]. It is this quick evaluation that must be supported.

Table 2 outlines the differences between the decision-making process of a rational decision-maker and that of a bounded one. The neoclassical assumption of rationality [90] is quite unachievable outside of

its theoretical nature. From the standpoint of traditional rationality, the decision-maker is assumed to have an objective and completely true view of the world and everything in it. Because of this objective view, and the unlimited computational capabilities of the decision-maker, it is expected that the decision they take will be the one that provides maximal utility. However, in reality — decisions can be bounded, imperfect, and prone to unknown implications. Such implications include *externalities*, which are situations imposed on others and can be both positive or negative [9].

Table 2

Rationality vs. bounded rationality in decision-making.

<i>Traditional economics</i>	<i>Behavioural economics</i>
RATIONAL DECISION-MAKER	BOUNDED DECISION-MAKER
- detailed evaluation of costs, gains, and losses	- brief consideration of perceived costs, gains, and losses
- complete information	- incomplete information
- careful calculation of potential investment	- cognitive and time constraints
	- quick evaluation of risks driven by loss aversion
↓	↓
chosen outcome	decision fatigue
↓	↓
optimal decision	satisficing

It is a common misconception that behavioural economics postulates irrationality in people. The difference in viewpoint arises from how rationality was originally defined, rather than from the assumption that people are rational beings. It is agreed upon that people have reasons, motivations, and goals when deciding to do something — whether they do it well or badly, they do engage in thinking and reasoning when making a decision [90]. However, it is important to denote in a more realistic manner how this decision-making process looks for a bounded agent. It is by considering these principles that we explore a more constructive approach to decision-support by employees in organizations.

While an objective view of the world always leads to the optimal decision (Table 2), a bounded agent often settles for a satisfactory decision. Simon [90] argues that people tend to make decisions by *satisficing* [48] rather than optimizing. They use basic decision criteria that lead to a combination of a satisfying and sufficient decision, which, from their perspective, is ‘good enough’ considering the different constraints. Furthermore, when faced with too many competing decisions, a person’s resources become strained and *decision fatigue* [99] often contributes to poor choices. In organization security behaviours, it has been shown that this fatigue can lead to decisions being circumvented or delayed [14].

4.2. Group decisions, herding, and culture

The recent consideration of socio-psychological factors in economics helps deconstruct the expectation of rationality being a dichotomous concept. While traditionally there may be an expectation to categorize behaviours such as herding in a binary manner and label it as either rational or irrational [7], the concept of bounded rationality [89] provides perspective and context into the complexity of individual and group decision-making. Generally there has been greater focus on modelling individual decision-making [57], although many important daily decisions are often taken by *groups*. Decision-making in

groups is also likely to occur when faced with uncertainty and people start equating their beliefs with those of others. [9].

Engaging in herding and replicating others' decisions is a way of social learning [7], and a part of the heuristics and 'rules of thumb' applied by bounded decision-makers [9]. *Herding* occurs when people do not act completely independently but rather follow others and copy their behaviours [7]. Groups or individuals may think that herd behaviour is a rational choice because they believe that others are better informed than they are. However, herd behaviour may sometimes be perceived as irrational if only engaged in for conformity, rather than an actual belief that the right behaviour is being followed [87].

Herding is often used interchangeably with the term *information cascade*, but they are in fact significantly different — an information cascade happens when people follow the behaviours of others with zero regard to their private information [26]. Another difference between the two is that while engaging in herding, individuals can still receive and consider private information or a signal that may lead to a different behaviour than that of the herd [11]. Thus, herding still allows a process of social learning, which, by comparison, ceases during an information cascade because of the uninformative nature of the process [26]. Information cascading may then lead to negative externalities if the copied behaviour is wrong and no individual or group questions it — regardless of their private signals [11].

The role of culture has been noted as important generally for understanding complex socio-technical systems and identifying causes of problems involving people and processes [30]. Organizational culture is commonly described as *the way we do things around here* [31] to reflect employees' common behaviours in an organizational setting. An individual's or a group's propensity for engaging in herding or information cascades may be impacted by the culture of the organization. The culture consolidates people's responses to chaos and uncertainty through shared belief systems and actions [97] – with herd behaviour being one such response to uncertainty [7].

Decision-making in a security context is also influenced by factors such as uncertainty and limitations, often leading to sub-optimal decisions [73]. Parsons et al. found a positive relationship between security culture and security decision-making, suggesting that an improvement in culture would lead to an improvement in security behaviours [73]. The value of security culture is increasingly highlighted when there is a tendency for herd behaviour in the organization. Shao, Siponen, and Liu found that when security managers face uncertainty about their reputation, and making security investments, they are more likely to ignore their own information and follow others' decisions instead [86]. Furthermore, a study by Yazdanmehr, Wang, and Yang found that social influence between co-workers can significantly impact compliance with security policy [103].

Such findings align with our goal to *better support 'good enough' decisions that best approximate secure behaviours under various constraints*. Ideas from both traditional and behavioural economics can be used as 'rules of thumb' when characterizing various constraints and social phenomenon that impact security behaviours and compliance with security policy in organizations [23]. Findings about group decision-making and the influence of culture on behaviours shed light on our understanding of constraints imposed upon bounded decision-making. The culture and group dynamics within an organization may affect whether herding or information cascades might supplant some of the criteria in bounded decision-making.

5. A framework for security choices

Current approaches to security provisioning in organizations appear as if to support the rational decision-maker, as per traditional economics. Such an approach does not support non-expert employ-

ees to find and follow appropriate security behaviours. We outline how to address the contradictions that currently exist in how traditional and behavioural economics have up to now been selectively represented, as follows below.

5.1. Bounded security decision-making

Security research increasingly focuses on organizational security and the interaction between managers, policies, and employees. Principles from economics have been deemed useful in security [21], and concepts from behavioural economics further support understanding of security behaviours in an organizational context [19]. For security policies to be effective, they must align with employees' limited capacity and resources for policy compliance [24]. Furthermore, policy-makers should be aware of other recurring security behaviours in the organization that are not aligned with policy to understand why such behaviours persist [53].

We use the term *bounded security decision-making* to move away from any ambiguity that arises when merging concepts from traditional and behavioural economics. This distances us from the tendency to apply behavioural intervention concepts to security while assuming the intervention targets to be rational decision-makers. This is in itself a contradiction because a rational decision-maker would by default know of and make the optimal choice, and would not require any behavioural aid or intervention (as explored in Section 3). Similarly, employees cannot possibly dedicate sufficient time or resources for every single task or policy to account for this [24]. This is a consideration that must be acknowledged at the point of security policy design.


To represent these concepts within an information security strategy model, we adapt the security investment model developed by Caulfield and Pym [24], which is constructed within the modelling framework described in [25, 29]. This model explicitly considers the decision-point for an employee (the decision-maker), and incorporates elements of the decision-making process (where we reconcile elements of behavioural economics), and available choices provided by the policy-maker/organization (the influencer). We adapt this framework to outline factors that should be considered when provisioning security choices, toward supporting the decision-maker(s) to choose 'good enough' behaviours under constraints on knowledge and resources. An influencer can then explore assumptions they are making for these factors, toward identifying where they can reduce the draw on decision-maker resources to find an appropriate and secure behaviour to adopt.

Figure 2 illustrates the components and processes that must be considered in security policy design. *Influencer* refers to the security policy-maker in the organization, and *decision-maker* (DM) refers to an employee engaging in bounded decision-making. A decision-maker can represent both an individual or a group.

5.1.1. Process

On the left-hand side we consolidate factors in decision-making from behavioural economics into the decision-making process that informs a decision (the arrow on the left-hand side). We outline the restrictive factors (limited skills, knowledge, time, and incomplete information) that characterize a bounded decision-maker. We acknowledge that the decision-maker is bounded in several ways, from individual skills and knowledge to temporal restrictions set by the organization. Our bounded decision-maker has incomplete information about the world and others, and must make do with information available within their abilities; they can only consider the perceived costs, gains and losses and prioritize subjective interests when faced with a choice.

Fig. 2. A decision point in a decision-maker's process of *bounded security decision-making* (elements adapted from Caulfield and Pym [24]).



When evaluating the risks that come with a choice — ‘losses loom larger than gains’ [49, p. 279], and the decision-maker tries harder to avoid losses rather than to encounter gains. This then puts the expectations of the influencer at a loss, as the decision-maker may be more concerned with the loss of productivity than with a potential security gain, where the latter may be all that the influencer — the overseer and expert of security — can see.

Decision-maker preferences. The restrictive factors on the left hand side of Figure 1 influence the decision-maker’s preferences. Using these factors as a reference point, the DM may have preferences over complying with one behaviour over another. Advocated security behaviours compete with other behaviours (e.g., compliance with HR policies or work deadlines) for the DM’s choice of preference, where that preference impacts their final decision. If compliance with, for example, an HR policy requires less technical engagement (and time investment), this will factor into the preferences. To consider the social element, costs may include any loss of trust from colleagues if following policy [56]; for example, a screen-locking policy in an office shared with trusted colleagues. There are also social costs associated with asking for help — such as admitting incompetence and dependence to others, which may provoke a feeling of powerlessness [61]. Another is how difficulty in completing a task may be perceived, where seeking help on an easy task may incur higher social costs; for example, embarrassment of failure in front of peers and business partners/customers [4, 14], while asking for help on more difficult tasks would be less embarrassing [34].

Choices and decision. The two boxes above the Decision circle represent the type of choices available to the decision-maker. Available policy choices consist of the rules listed in the security policy by the influencer, but also any included advice on what to do and solutions provided. In organizations with security policies, the influencer usually assumes that the only choices available to the decision-maker are the ones noted by the policy itself. However, as literature shows, a choice may be to circumvent the policy [18, 58], or to attempt to work in a way that best approximates compliance with secure working policies, in the best way the decision-maker knows how to [53]. Though workarounds and circumventions of policy predominantly go unnoticed in organizations, this does not eliminate them from the set of choices available to the decision-maker. Behaviours regarded as choices by the decision-maker — but which are hidden to the influencer — are another information asymmetry (one that introduces risks for the organization [53]). By assuming that the only available choices come from the security policy, the influencer indirectly undermines policy by having less predictable control over policy compliance decisions in the organization.

Information asymmetry. Information asymmetry regularly occurs between the influencer and the decision-maker. Such asymmetry can negatively impact the influencer, decision-maker, as well as the organization. Although many asymmetries in security occur because of the influencer's lack of awareness about the decision-maker's restrictive factors (limited skills, knowledge, time, and incomplete information), there are additional reasons for their existence. The security function is often unaware of the security decisions outside of policy that are taken by the decision-maker — as well as the sources of information and influence upon such decisions. An expected source of influence on security behaviours and decisions is social interaction with others in the organization [68]. In the context of security policies and policy compliance, the following are examples of information asymmetry:

- The recognized differentiation of the influencer being more knowledgeable and capable in security than the decision-maker (as security is arguably the influencer's primary task);
- The influencer's lack of knowledge about the decision-maker's context, and pressures that factor into their choice-making process (resulting in the influencer seeming to perceive the decision-maker as a rational decision-maker with motivation and resources dedicated to security);
- The influencer's lack of awareness about competing company policies with which the decision-maker must *also* comply;
- The decision-maker's lack of information about why security restrictions matter to the organization (overly demanding policies may cause decision-makers to lose sight of why the policies exist in the first place);
- The influencer's lack of awareness about the impact of social learning and group decision-making on security policy compliance;
- The influencer's insufficient understanding of *other* non-policy behaviours (shadow security behaviours), which may or may not be secure.

Such discrepancies in knowledge and information between the influencer and the decision-maker cause friction and create a power imbalance. Asymmetries can be identified and understood, toward reducing the gap between influencer and decision-maker perceptions (which is engineered by having a distinct, designated security function).

Moral hazard. When a number of information asymmetries exist in the organization, a moral hazard is likely occurring. A common example of a moral hazard is that of the principal-agent problem, when one

person has the ability to make decisions on behalf of another. Here, the person making the decisions (the agent) is the decision-maker, and decisions are being made on behalf of the influencer (the principal) who represents the organization's security function. However, problems between the agent and the principal arise when there are conflicting goals *and* information asymmetry.

If we go back to the decision-maker's perceived risks, we argue that these are not synonymous with the risks that the influencer knows of or is concerned with. Hence, when the decision-maker enacts behaviours, they do so by prioritising their interests and aiming to reduce their perceived risks. Because of the information asymmetry that persists between the decision-maker and the influencer, as well as the decision-maker's hidden choices driven by personal benefit, the influencer cannot always ensure that decisions are being made in their best interest. The moral hazard here is that the decision-maker can take more (security) risks because the cost of those risks will fall on the organization rather than on the decision-maker themselves.

Information asymmetries also impact any decisions to allow for moral hazard in the workplace; for example, if forming an infrastructure where employees should keep their own work machine secure (for instance to keep software up-to-date), and at the same time be held personally responsible if it is not secure. The influencer decides that the employee has to take action (i.e., accept and manage the risk), but the action requires more commitment and expertise than the employee is assumed to have in making this decision. It is part of the contradiction of an organisation having an engineered information asymmetry, where the expert is there to support the non-expert. In comparison, having a 'weak' email filter creates an externality of more effort to detect phishing emails that have arrived in an Inbox; having a stronger filter creates a positive externality *unless* it removes legitimate emails. A moral hazard here would be to move the risk to the individual as if spotting a phishing email is easier than it actually is (so the risk is high but the 'right' behaviour is very difficult to sustain). This becomes a career-related moral hazard if security performance is linked to the terms by which a decision-maker is judged if things go wrong, and what 'doing it right' looks like.

Moral hazards in organizations are likely to occur unconsciously — the influencer does not consciously make the decision to delegate the risk to the decision-maker, because they assume and expect the decision-maker to be (traditionally) rational. On the basis of that assumption, a rational decision-maker would know how to enact the given security behaviour without error, leaving no risk in the end. Thus, the way in which an influencer is to become aware of such a potential moral hazard, is by accepting employees as bounded decision-makers. Realising the possibility of a moral hazard would be the first step — the second step would be to safeguard against it so that the risk does not fall on the non-expert. Given the current way security non-compliance is commonly handled in organizations, the risk would very likely fall on the non-expert decision-maker.

Choice architecture. The circle in Figure 2 signifies the decision made by the decision-maker. In our framework, we refer to the circle by using the term 'decision' rather than 'choice architecture' for the following reasons: (1) while unusable advocated security behaviours persist, the set of choices is a composite of choices created by both the influencer and the decision-maker, which does not correspond to the accepted nature of a curated choice architecture, and; (2) referring to a choice architecture implies an intention to nudge decision-makers towards a particular choice, which also implies that there exists one optimal choice. As we have mentioned previously, a single optimal choice cannot exist for bounded decision-makers because they have perceived costs, gains, and losses individually; a more helpful approach would be to accommodate a range of choices rather than strictly advocate for one choice that is

not being followed.

Policy and decision-maker choices. Figure 3 provides a zoomed in view of the available policy and decision-maker choices previously shown at a glance in Figure 2. The purpose of outlining these choices in a more detailed way is to really emphasize the variety of behaviours that may go unnoticed in an organization — as well as the sources and influences of those behaviours. The discussion of the outlined behaviours reflects the concepts introduced in Section 4.2.

The circle on the left outlines the available policy choices that the influencer provides for the decision-maker, often being under the impression that the formal security policy is the only available choice in the organization. The circle on the right outlines potential security behaviours that the decision-maker engages in — not necessarily aligned with the security policy. Lastly, in the middle of the figure there is an intersection between the available policy choices and the decision-maker choices, which are labelled as *perceived policy choices* because although they are not confirmed as being viable security policy choices, the decision-maker could perceive them as such, because of the source of these choices.




Fig. 3. A zoomed-in view of the *available policy choices* and *decision-maker choices*.

The decision-maker choices outlined in Figure 3 differ in the following way:

- **Shadow security behaviour.** This type of behaviour refers to security practices outside the authority of the organization, developed by employees who do not intend to neglect security [53]. Shadow security behaviours emerge because of security-incurred costs such as cognitive load, disruption, and time [54], which are in line with the restricted factors of a bounded decision-maker outlined in Figure 2. The decision-maker choosing to enact a shadow security behaviour may be conscious of making an *insecure* or a *less* secure choice [53] — but is forced to approximate a ‘good enough’ behaviour because of the difficulty in complying with policy.

- 1 • **Information cascade behaviour.** When during a decision-making process the decision-maker no
2 longer considers their private information or signal, and simply follows the identical behaviour of
3 others, an information cascade emerges [26]. During an information cascade, the chosen behaviour
4 may contradict the decision-maker's private information but because this information is ignored,
5 decision-makers are likely to make sub-optimal choices [11]. Negative externalities may appear
6 where in security that could put the organization at risk if the decision-maker is following a choice
7 that actively ignores policy or security breaches [23].
- 8 • **Herding behaviour.** In comparison to information cascade behaviours, herd behaviour occurs when
9 the decision-maker can make an identical decision to others, but not necessarily by ignoring their
10 own information [26]. Hence, herd behaviour supports the opportunity for social learning and access
11 to social information about others' decisions [23]. However, herding behaviours become difficult
12 when decision-makers follow the herd even when they have unambiguous information about the
13 behaviour being incorrect [7]. Reasons for following an incorrect herd may include social pressure
14 or a decision-maker's preference for conformity [11].

15 Figure 3 also points out behaviours that may be perceived as viable policy choices by the decision-
16 maker. Informal resources in a social environment may influence the decision-maker's policy compliance
17 when such resources originate from peers or managers within the organization [101]. A decision-maker's
18 perception of security may be changed by the viewpoints of co-workers or supervisors, which may then
19 be perceived as directives [45]. We distinguish between two types of perceived policy choices:

- 21 • **Communicated behaviour.** This behaviour originates from various communications the decision-
22 maker may have received about how to do something. It primarily concerns new decision-makers
23 in an organization, dependent on the instructions of others at the time of their arrival. The decision-
24 maker may observe the behaviour from the person leading their induction day, or from the colleague
25 who is doing their handover for a new role. For a new employee, the absence of security advice — or
26 a dismissive comment about policy compliance, may be perceived as a cultural cue of 'how things
27 are done around here'.
- 28 • **Inherited hierarchical behaviour.** This behaviour may be perceived — although not verified as
29 such — as a viable policy choice primarily because it is inherited from a person of authority; a
30 supervisor, a manager, or even a significantly more experienced colleague. The decision-maker may
31 assume that more senior employees 'know what they're doing' and are a sufficiently trustworthy
32 source of security behaviours. Also, these inherited behaviours might have potentially been made
33 more specific to the local working conditions in order to make directives actionable [77] and are
34 perhaps easier for the decision-maker to follow.

35 Preferences also factor into the final choice of a decision-maker. A decision-maker's preferences can
36 be influenced by their lack of knowledge or inexperience relative to a particular decision [23]. This
37 may then prompt them to use information from other, more senior decision-makers. When faced with
38 uncertain decisions, people tend to believe that others are better equipped than they are to make the right
39 choice [7]. The bounded decision-maker might follow others because they trust that choice more than
40 their own interpretation of the security policy.

41 Decision-maker preferences can also be influenced by cultural factors such as shared values, beliefs,
42 or even tasks [83]. Behaviours are often created based on social interactions with others [68]. In orga-
43 nizations where social bonding between coworkers is stronger, individuals or groups are more likely to
44 adopt identical behaviours because of social pressure [45, 81].

In cultures where teams or groups engage in communication and social interactions, some security choices may be a result of discussions between several non-experts. While such security choices can be risky and potentially cost the influencer, the decision-maker may prefer to avoid social costs at the expense of a security gain. Bearing the social cost of ‘not fitting into the culture’ may be worse for the decision-maker than the potential costs of non-compliance with security.

5.2. Framework implementation

We describe a framework that can be applied as part of a strategy to better anticipate employee security-related behaviours under bounded security decision-making conditions. A goal here is not to dictate how decision-makers (the employees) should behave, but to provide solutions in a consistent manner, with the aim of reducing the likelihood that the decision-maker be in a position of having to use their limited resources to make up for gaps and inconsistencies in security provisioning. If an influencer can understand pertinent decision points in work processes, and identify ways to support both productivity and security needs, then they can create a system of mutual advantage, rather than a potentially hierarchical approach of nudging — or ‘prodding’ — employees toward an outcome that benefits only security [80]. In this way, the intervention can address conscious decisions about their security behaviours, with the aim of achieving a natural fit with working practices. In the spirit of mutual advantage [92], the security influencer can only succeed if their security provisions benefit decision-makers.

Here we describe steps for applying the framework (as in Figure 2). We note that smaller organizations may not have the resources to maintain an overview of systems and system usage (more so if elements are outsourced [71]).

(1) Capture decision-maker process.

Justification: Influencers must understand the decision-maker’s process and consider the decision-maker’s current knowledge of the system (Figure 2), whether it is as individuals or discernible groups of users. Security managers in organizations may be unaware of the range of factors in bounded security decision-making if they assume that the only possible choices are those provided in the formal policy. Thus, influencers should acknowledge that (1) there is in fact a decision-making process for employees that impacts the outcome of security behaviours, and (2) there are co-existent factors that impact the bounded rational decision-maker (other imperatives alongside security). In short, assumptions about the technical expertise of the decision-maker must be checked against what a representative sample of employees/groups is capable of doing.

Implementation steps:

- 1 Depending on the size of the organization, set up an initiative to engage with individual or groups of employees;
- 2 Acknowledge the existing information asymmetries and moral hazard in the organization to gain a better overview of who is currently carrying security risks (e.g., employees carrying risk on behalf of the organization) — this will help clarify, determine, and communicate risk ownership within the organization;
- 3 Discuss with employees the challenges of behaving securely in order to understand and capture their decision-making process (including here perceived costs, gains, losses, as well as employee preferences);




Fig. 4. Implementation steps of the *bounded security decision-making* framework.

- 4 Evaluate the skills, knowledge, available resources, and time that employees have for security;
- 5 Discuss and identify available, perceived, decision-maker choices with employees (including here examples of communicated or inherited behaviours, as well as shadow security, information cascade, and herding behaviours).

(2) Adapt policy choices to decision-maker.

Justification: Policy choices must be adapted to the decision-maker's current level of understanding and supported with concrete information. This requires a focus on the decision-maker's current state of knowledge and resources, rather than the desired security behaviour end-state. It may be that the influencer will not acknowledge any viable behaviours outside of the formal security policy [6, 78], in which case a way to make the policy more flexible should be explored. For example,

more than one single choice for a security behaviour could be offered (as happens with fallback authentication methods, for example). This rests on the influencer having ensured that the policy choices are genuinely feasible for employees in light of the information gathered in the first step.

Implementation steps:

- 1 Consider the bounded rationality of employees by simplifying policy choices so that they correspond with the knowledge and skills of a representative sample of employees;
- 2 Utilize previously collected information about individual and groups of employees such as their cognitive, time, resource, and information constraints to better align policy choices to their decision-making processes and constraints;
- 3 Evaluate the perceived and decision-maker choices that have been identified, and consider whether they can qualify as viable policy choices and be supported by the organization.

(3) Realign policy choices with stakeholders.

Justification: Collaboration must be established with stakeholders to determine which other imperatives exist for decision-makers, where these are other pressures on decision-making resources alongside security. If security and other competing imperatives are not aligned, there is a risk that the decision-maker is left responsible for ensuring that an advocated security behaviour is a possible choice in the presence of other imperatives. In this step it would be crucial to consult with authoritative employees such as managers or supervisors [53], or any employees who act as a local security champion representing a team of employees [15]. These stakeholders are likely to have more visibility of barriers and ‘dilemmas’ [19] experienced in adapting security policy choices to work activities, as well as more insight on how to tailor security behaviours to fit local working conditions [77]. In practice, this may be a cycle of learning and adapting [53, 104].

Implementation steps:

- 1 Collaborate with relevant stakeholders from other parts of the organization — such as managers — to understand competing policy expectations;
- 2 In collaboration with other stakeholders, identify, based on different employee roles and responsibilities, potential compromises for policy requirements, such that time and resources are allocated adequately;
- 3 Revise available policy choices and requirements with stakeholders and employees.

(4) Enable business objectives.

Justification: Security influencers must be supported by the larger organization to ensure that security policy choices do not interfere with other business expectations. When business and security goals are misaligned, decision-makers experience friction [33]. Employees should not be forced into a position where they would have to break security rules to maintain productivity [55] or complete other primary tasks. Having competing policies and expectations is essentially a failure in organizational strategy, which is bound to occur if security requirements are not aligned with

business expectations at an executive level [55, 82]. There is a need to support security management teams to understand and resource security policies alongside non-security workplace policies.

Implementation steps:

- 1 Assure employees that aligning business objectives with security requirements is a priority for the organization;
- 2 Address time and resource constraints of bounded decision-making by budgeting time and money for required security activities such as training, or even mandatory software updates (e.g., by creating a designated time-allocation code that employees can use to log these hours);
- 3 Address employees' cognitive constraints by empowering IT and security staff to proactively support them in their security efforts;
- 4 Revisit security expectations and the decision-making processes of boundedly rational employees continuously to align policy choices accordingly.

6. Worked example: software security updates

We apply the framework to a pertinent case study — keeping software up-to-date — and demonstrate examples of both individual and group decision-making. This case study is derived from the top online security controls advocated by security experts (as prompted by Reeder et al. [77]). At the time of writing, this is also the top piece of advice advocated by, for example, the UK government¹.

We have chosen the behaviour of enacting software security updates as there is evidence that knowing what impact an update will have and when to install it is a complex decision for employees [14]. Decisions around installing updates are also fraught with intrinsic uncertainty, especially concerning the impact an update can have on an employee's ability to perform their immediate work tasks if they follow the advice. For instance, employees may not be able to work while a software or security update is installing [4, 14]. Specifically to workplaces, if employees are busy, it is arguably difficult to find an appropriate time to install an update as far as work imperatives are concerned [4, 14], unless the update activity is surfaced, discussed, agreed, and facilitated within the organization.

6.1. Process

Skills, knowledge, and time. Applying updates as soon as possible is seen as achieving the best results [47]. However, advocating to 'keep software up-to-date' or to 'apply updates immediately' does not accommodate consideration of preferences for committing time to other tasks (such as primary work tasks). Some employee groups may not have sufficient time to retrieve and install software updates during work hours. It may also be that some employee groups are using a complex collection of applications, all impacted by an update to another application [67], or the underlying operating system, making the update process complicated to assess in advance.

A *bounded security decision-making* approach would provide step-by-step guidance to match skill levels, and potentially the version of software that is currently on a device. Automation could also be considered, if the update process is complex or requires technical skill. Where there are different groups

¹ As at the National Cyber Security Centre (NCSC) website.

of employees, it may be possible to ‘fingerprint’ a range of typical software configurations, and provide clear instructions for each.

Especially if staff use a range of applications, it would be helpful to calculate the time it would take to install an update and be able to use the machine again, alongside a checklist of steps for staff to follow, ideally separated into discrete steps as much as possible to allow tasks to be followed whenever time is available. An employee can then plan for following a more predictable process. This already differs from existing advice, which predominantly assumes in its framing that software updates have no associated cost, that security imperatives are ahead of all other objectives (such as completing work) or that there are no impacts of installing an update.

If it is not possible to make the update process more predictable for employees, a simpler option would be to provide an estimate of how much time each critical update would take, so that at least this estimated time could be negotiated for, or by, staff and their managers. Staff could then be asked to report if the estimate was accurate. In terms of informing ‘good enough’ decisions, the approaches above would eliminate guesswork in planning use of time and employees’ computing resources.

Perceived costs, gains, and losses. In organizations, system patches are first deployed to a test-bed [47], to ensure that they do not create problems (losses). Direct advice to employees to ‘keep systems up-to-date’ ignores this consideration to test patches first. Simplified advice about patching neglects the need for a user to already know how to install patches, and when. It may be that users also do not consider updates as a concern [96], so may not be motivated to do it at all unless the connection to their primary work is made clear.

A *bounded security decision-making* approach could provide an assurance that the latest updates have been tested on a system similar to the one the receiver of the advice is using for their work. This is so that they do not have to establish whether the patch works without problems themselves. Ideally a patch would also come with sufficient explanation of any features or applications which have changed due to the patch, reducing any sense of potential loss of cognitive automation, which would require a user to rebuild their cognitive maps of how software works [16]).

It would be necessary to convey that an up-to-date system protects specific assets that the decision-maker wants to minimize losses for (where top-management or asset-focused messaging could help). One approach here is to explicitly state that an update has been tested and will work and, if necessary, list any changes to core applications that employees can expect to see (e.g., a menu option has moved to a different menu or looks different). This does not completely eliminate employees feeling a need to check, but may at least encourage them that the cost of installing the update is more predictable. Additionally, if an update is being rolled out across many employees’ machines, providing dedicated technical support for that update – and communicating that this has been provided – would reduce uncertainty as to whether employees will be able to use computers for their work.

Incomplete information. The minimal advice does not declare how to check or how often, assuming a rational approach. If an update seems to be taking a long time, a decision-maker may not know if the problem is with the machine (requiring support) or personal expectations (and not being able to troubleshoot problems [102]). There is also an assumption that the user may know the changes that updates will create in advance, when it could impact them in a range of ways [16]. In a corporate environment, an employee cannot know if all networked services or systems they wish to interact with are also up-to-date, or will work with a newer version of the software they are using. Indeed, if other systems are not kept up-to-date — which an employee cannot know — this also could prevent different

systems from working together. Conversely, the provisions of the organization may lag behind what an individual is capable of doing — corporate systems may be behind on versions, limiting how ‘far ahead’ an employee can go in keeping their own software updated.

A *bounded security decision-making* approach could involve informing the user of how long each update takes to install [63] (especially if a restart is required), based on testing on a comparable setup (including machine performance, available disk space [63] and provisioned software). It may be that updates can be scheduled centrally [63], for instance to occur when employees are most likely to have their computer on but it not be in use (if the organization has scheduled workplace lunch breaks, for instance). Efforts in this direction may reduce the perceived costs to work activities when a patch is installed.

Ultimately, finding a time to install updates and avoid disruption is increasingly difficult to find in a PC-computing work environment. The above approaches are subtly different to the standard ‘progress bar’ and timing information provided by operating systems, by working to reduce individual costs and also reduce uncertainty about those costs. This moves in the direction of better matching updates to the work context, i.e., how operating system updates interact with corporate applications and systems, and how updates can fit into the work day of employees. This would be akin to knowing employees’ ‘Active Hours’ (as is already seen within Windows 10 home installations [66]), and assessing whether these hours make it feasible to install the updates within the working day.

Loss-averse evaluation of risks. A rational approach does not accommodate the chance that the user has had prior bad experiences with updates [96]. It also does not provide assurances that the update will not cause software to cease working properly, and does not declare how much (paid/salaried) time the update will take (assuming this to be none/negligible).

A *bounded security decision-making* approach would provide support for creating backups before updates, and point to the existence of the backups (to assuage concerns about losses). A user may simply choose to delay or ignore the installation of an update [96], so there would be a need to convey or imply why this is *not* an appropriate option to consider — this is most readily achieved by presenting the options that the user perceives relative to each other. Across different groups of employees, it may be possible — if not necessary — to ensure that the organization can still function while updates are occurring, and that there is declared support in case of problems that are discovered after the update has been installed.

This does not remove the uncertainty of whether problems could occur, but changes the deliberation of the employee to one of whether they are willing to risk needing help after an update. This is opposed to a user risking installation of an update and then needing to fix it on their own. The delegation of security to others seen as better-placed to manage it has been seen elsewhere [35], and can inform ‘good enough’ decisions around installing updates by reducing the perceived costs of potential outcomes.

6.2. Available policy choices

Rational advice to keep a system up-to-date does not consider that modern systems may already be doing (some or all of) this without user involvement, so advice may, for example, need to consider specific operating system software. Unless an operating system or application provides separate feature updates and security updates, the value of updates for security may not allow a decision-maker to consider clear choices [66].

A *bounded security decision-making* approach would acknowledge how updates work on the system the decision-maker is using. It would also recognize the other options that are available to the decision-maker, from the perspective of their personal preferences and not solely the one ideal preference of the security function (influencer).

To better approach ‘good enough’ behaviours, the benefit of installing an update must consider the (declared) costs and risks of installing the update (such as time waiting for installation), but also the cost of delaying or ignoring the update. If the machine continues to function without the update, an employee may see that an update only brings costs.

With security controls such as forced software updates, an information cascade behaviour may be imposed on decision-makers by the technical policies of the organization, and in turn the collective workplace practices. This would leave no other available choices, especially if it is necessary to maintain software that is compatible with the applications used between employees or with collaborators elsewhere.

As a result, employees would be forced into an information cascade of simply updating their device, even if their own individual judgement indicates that they would want to do otherwise. This creates a negative externality that imposes additional costs upon the decision-maker’s primary task.

6.3. Decision-maker choices

If workplace security choices are framed for a rational decision-maker, and are not made explicit and compared meaningfully, the *bounded security* decision-maker may construct the set of choices in an ad-hoc fashion from their own perspective, with little to no information about the consequences of taking action or not (based on the non-expert knowledge they personally have [53], which is not as informed as that of that the security function).

Competing work and social pressures may urge the decision-maker to make improvised choices. When working in a group or a team, the decision-maker attempting to comply with the security policy may be instructed by colleagues or more senior members of their team to delay updates until work activities are completed. This can then create a moral hazard because the choice benefits the decision-maker (who manages to both complete their work and seem to comply with policy), but imposes the risk of not updating the software in a timely fashion which may in turn pose a risk to the organization. Another option would be for the IT team to ensure that updates can be installed in such a way that prior operating system and application state would be reinstated once the update has been installed. This opens up options to identify tasks where a computer would not be used, for instance prompting staff to install updates together before an in-person team meeting or similar event where computers would not be used.

7. Future directions

Informed by user-centred security research, we outline directions for how a security function in an organization can consider the proposals we have made (Section 5). A security function cannot be assumed to have in-depth knowledge of the human aspects of security, but may nonetheless value it in security policy decision-making [72], and benefit from methods and tools to do so [78].

7.1. A security diet

A ‘security diet’ would document perceived occurrence and costs of advocated behaviours (for example, through a typical working day). Questions can then be asked to reconcile these costs with expected

behaviours elsewhere in the organization [50], to determine if time for security tasks is being taken from elsewhere.

If security behaviours add to an already busy schedule, then time constraints, pressure, and stress increase the likelihood of errors [75]. This is akin to the Compliance Budget [4, 14], which supposes that there is a limit to the effort an employee will commit to security before they choose not to comply (and crucially, beyond which the organization needs to commit additional resources to encourage a return to compliance). The ideal would be that security is progressively revisited and redesigned to be achieved as part of core business tasks, targeting the burden of identifying appropriate behaviours — this would begin to achieve something closer to ‘Productive Security’ [4, 13].

Security diets may be populated by talking to a representative sample of employees across different teams, for instance, or team managers (rather than all employees for large organizations).

An individual arguably *should not* be expected to commit more than their full working day to all tasks including security. As mentioned, critically, security mandates demand the very same qualities from employees that diminish when available resources are exceeded. Security is then self-defeating if it leaves the decision-maker to figure out how to make this possible, as the effort involved in identifying and following behaviours can induce errors from already-burdened employees. Consideration of how to manage security with other pressures can reduce this ‘gulf of execution’ [79].

7.2. *Just culture and the genuine choice architecture*

The intentions of a choice architecture continue to be divisive, in terms of whether they are a tool for paternalism, soft paternalism, addressing specific roles in a system akin to a market and mutual advantage, etc. [92]. Here we argue for the last of these in organizational security, primarily in the sense that policy and company directives may all inform an employee’s work, but as an employee decision-maker, they are the one needing to make sense of all of this within a given situation.

If we are to involve the decision-maker in shaping viable options, we would want to find a way to acknowledge the choices seen or made by individuals or teams within the organization, from their bottom-up perspective and not only the top-down view of policy and policymakers, to include them alongside advocated choices for clear comparison. An implication of this is that it will also ‘declare’ non-secure options for them to be compared to advocated choices, requiring the practice of a ‘blame-free’, *just culture* [32], toward learning from shortcomings. Workarounds created by staff, or behaviours outside of policy generally, are in the first instance treated as efforts to subvert good behaviour, but can in fact be attempts to balance security and productivity [53] (as basic ‘preferences’ of an employee who wants to do the right thing by their colleagues and the company).

Ultimately, influencer and decision-maker do not see the same choice architecture, as highlighted by the need for care in choice editing. By defining associated properties of these two sets of choices, support can be negotiated to shape available solutions that allow productive and secure working. Seeing security from the employee’s perspective would also have the advantage of identifying forgotten policy artefacts [43]; for example, older versions of policy documents, which remain on a staff intranet and which have not yet been removed. A genuine choice architecture could leverage ‘all in the room’ consultation methods with a cross-section of employees [83] or security dialogues [6].

7.3. *Policy concordance*

‘Security Dialogues’ research [6] promotes a move toward policy concordance — ‘mutual understanding and agreement’ on how the decision-maker will behave. In medicine [44], concordance occurs at the

point of consultation between a medical practitioner and a patient, to incorporate the respective views of the decision-maker and influencer. This also incidentally reveals the aims of both sides, and encourages them to reach a shared goal.

The definitions of distinct behaviour choices can be considered by both sides when negotiating a solution for security concordance. This then further leverages the co-developed choice architecture. This could ‘zoom in’ further on decision options, to examine properties of individual choices according to the decision-maker’s preferences, comparing to other options that are regarded as viable. An example might be where an employee respects the need to browse safely, but is often visiting unfamiliar websites; for example, if conducting market research; that is, they cannot always know the provenance of the websites they visit for their work. Employees in such a role may need distinct browsing solutions, separate from those with more internal-facing (and in this dimension, more predictable) roles.

An approach of policy concordance would not simply defer to employee preferences, but input these preferences into a negotiation process alongside the expectations of the influencer. Our framework then becomes useful for discovering the preferences and resource requirements of each option, collating all options, and potentially exploring alternative options. The goal would be to find options that match employee needs to the extent that they are *sustainable* and can be repeatedly enacted within workplace conditions.

7.4. Security modelling and investment forecasting

Given the complexity of security policy choices and all the factors impacting them, it is difficult to predict the consequences different policy choices may have on the security of a system. Models are helpful in predicting the implications of design decisions. System models are particularly suitable here as they can represent the behaviour, structure and environment of a system, as well as the behaviour of employees interacting with it [23]. Lessons from behavioural economics are useful for improving modelling of people’s decision making.

Security modelling can begin to forecast the impact of investments in complex environments, before making infrastructure and provisioning changes (e.g., [24]). Security deployed is not security as designed; contact with the complex organizational environment will alter how successful a control is in practice, and how well it fits with other practices in the organization. Incorporating employee perspectives into structured economic models will inform the viability of new controls before their deployment.

8. Related work

There is a growing body of research advocating the application of economics concepts to security generally, as a means to understand complex challenges. Foundational work by Gordon and Loeb asserted that traditional economics can inform optimal investment in security [40], whereas here we apply a similar approach to a combination of economic models in order to reposition investment challenges related to security behaviour management. Beauteament et al. [14] articulate how employees have a restricted ‘compliance budget’ for security, and will stop complying once they have reached a certain threshold.

Acquisti and Grossklags [2] apply behavioural economics to consumer privacy in order to identify ways in which to support individuals as they engage in privacy-related decision-making. Similarly, Baddeley [8] applies behavioural economics in a management and policy setting, finding, for example, that

loss-aversion can be leveraged in the design of security prompts. Other concepts from behavioural economics have been explored within the domain of information security and privacy, such as the *endowment effect* [93] — a bias that causes people to value something more if they already own it, and *framing* — another bias that influences people's choices depending on whether the options are presented in a negative or positive manner [3, 41]. Anderson and Agarwal [3] identify potential in the use of goal-framing to influence security behaviour, where commitment devices have since been explored as a way to influence behaviour change [38]. Verendel [98] applies behavioural economics principles to formalize risk-related decisions toward predicting decision-making problems, positing that aspects of usable security must also be explored.

In addition to understanding security and privacy behaviour through behavioural economics, some have advocated the *influencing* of such behaviour through the application of *nudge theory*, which suggests the use of positive reinforcement and a choice architecture to influence people towards better decisions [1, 95]. Through empirical modelling of behavioural economics, Redmiles et al. [76] effectively advocate for identifying and presenting options that are optimal for the decision-maker, and making the risk, costs, and benefits of each choice transparent. Here we explore where there are 'gaps' in realising these capabilities, which must be closed in order for organizations to support secure behaviours.

With regard to capturing the dynamic between a decision-maker (here, an employee), and the security function — an 'influencer' — Morisset et al. [65] present a model of 'soft enforcement', where the influencer edits the choices available to a decision-maker toward removing bad choices. Here we acknowledge that workarounds and changes in working conditions occur regularly, proposing that the range of behaviour choices is in effect a negotiation between the two parties.

Research has often observed that people have a tendency to cluster their behaviours [11]. Engaging in collective behaviours — often referred to as herding, frequently arises from communication between individuals [11]. Despite these findings, in economics, the decision-maker is most commonly modelled as an individual, rather than as a member of a group [57]. In real life, however, the decision-maker is normally part of a group, such as a board of directors, a family, or a work team [57, 60]. Group decisions are not necessarily smarter or more rational than individual ones, but groups can learn faster than individuals [57], and experiments have shown that individuals sometimes follow others' decisions despite having unambiguous private information that contradicts that decision [11].

In the context of organizational security, groups could make suboptimal decisions [11] in relation to security policy compliance, which can in turn have negative implications for security if herd behaviour ensues [23]. To motivate adequate security behaviours in an organization, it is important to consider the security culture [100]. Schein's model of organizational culture [83], which has been largely applied to security [39, 59, 84], refers to the shared values and behaviours of a group. According to Vroom and Von Solms [100], in order to change the security culture of the organization, it must be changed both at the individual and group level. If group behaviour begins to alter, it can then influence individual behaviour too [100]. Therefore, both levels of behaviour should be evaluated when designing targeted behavioural interventions.

There is a need to reconcile the advancements in the application of economics to security with how management of behaviour change strategies in organizations is conceptualized. Here we fill in the gaps, where currently there are contradictions and shortcomings that act against both the organization and the individual decision-maker whether that individual makes choices for themselves or with input from others in their official or informal group.

9. Conclusion

We have shown how current approaches to security provisioning and infrastructure reflect expectations from traditional economics, even when concepts from behavioural economics are applied to encourage individuals to adopt specific security behaviours. We have constructed a framework that can be used to improve the process of providing a set of advocated security behaviours. This is more doable for employees who work under conditions of uncertainty and workplace constraints. The framework encourages a continuous programme of security behaviour choices, which must be provisioned for in order to support adequately ‘good enough’ behaviour decisions. We have also considered the dimension of inter-personal decision-making and security culture and the implications of social interactions on available policy and decision-maker choices. We have applied our framework to a regularly advocated security behaviour — software patching — to demonstrate that the rational-agent view is incompatible with the embrace of isolated behaviour change activities.

Our work identifies considerations for researchers working in organizational security:

- the importance of capturing where a decision-maker is, alongside where an influencer wants them to be;
- that a security choice architecture is essentially decentralized and cannot be wholly dictated by any one stakeholder;
- that, in organizations, security expertise can exist in places recognized by the organization and also in places that are not; and
- that information asymmetries — for example, as constructed by information security teams — ought to be accounted for when assessing user behaviours.

Future work might involve situated studies in organizations, including co-design with the security function to develop viable and sustainable security behaviour interventions.

Acknowledgements

Demjaha is supported through a Doctoral Studentship awarded by the Alan Turing Institute.

References

- [1] Acquisti, A.: Nudging privacy: The behavioral economics of personal information. *IEEE security & privacy* **7**(6) (2009)
- [2] Acquisti, A., Grossklags, J.: What can behavioral economics teach us about privacy. *Digital Privacy: Theory, Technologies and Practices* **18**, 363–377 (2007)
- [3] Anderson, C.L., Agarwal, R.: Practicing safe computing: a multimedia empirical examination of home computer user security behavioral intentions. *MISQ* **34**(3), 613–643 (2010)
- [4] Anderson, G., McCusker, G., Pym, D.: A logic for the compliance budget. In: *International Conference on Decision and Game Theory for Security*. pp. 370–381. Springer (2016)
- [5] Ashenden, D., Lawrence, D.: Can we sell security like soap?: a new approach to behaviour change. In: *Proceedings of the 2013 New Security Paradigms Workshop*. pp. 87–94. ACM (2013)
- [6] Ashenden, D., Lawrence, D.: Security dialogues: Building better relationships between security and business. *IEEE Security & Privacy* **14**(3), 82–87 (2016)
- [7] Baddeley, M.: Herding, social influence and economic decision-making: socio-psychological and neuroscientific analyses. *Philosophical Transactions of the Royal Society B: Biological Sciences* **365**(1538), 281–290 (2010)
- [8] Baddeley, M.: Information security: Lessons from behavioural economics. In: *Workshop on the Economics of Information Security* (2011)

- [9] Baddeley, M.: Herding, social influence and expert opinion. *Journal of Economic Methodology* **20**(1), 35–44 (2013)
- [10] Baddeley, M.: Behavioural economics: a very short introduction, vol. 505. Oxford University Press (2017)
- [11] Baddeley, M., Parkinson, S.: Group decision-making: An economic analysis of social influence and individual difference in experimental juries. *The Journal of Socio-Economics* **41**(5), 558–573 (2012)
- [12] Bateman, H., McAdam, K.: *Dictionary of Economics*. A & C Black Publishers Ltd (2003)
- [13] Beautelement, A., Becker, I., Parkin, S., Krol, K., Sasse, A.: Productive security: A scalable methodology for analysing employee security behaviours. In: Twelfth Symposium on Usable Privacy and Security (SOUPS). pp. 253–270 (2016)
- [14] Beautelement, A., Sasse, M.A., Wonham, M.: The compliance budget: managing security behaviour in organisations. In: Proceedings of the 2008 workshop on New security paradigms. pp. 47–58. ACM (2009)
- [15] Becker, I., Parkin, S., Sasse, M.A.: Finding security champions in blends of organisational culture. EuroUSEC 2017 (2017)
- [16] Bergman, O., Whittaker, S.: The cognitive costs of upgrades. *Interacting with Computers* **30**(1), 46–52 (2017)
- [17] Beris, O., Beautelement, A., Sasse, M.A.: Employee rule breakers, excuse makers and security champions:: Mapping the risk perceptions and emotions that drive security behaviors. In: Proceedings of the 2015 New Security Paradigms Workshop. pp. 73–84. ACM (2015)
- [18] Blythe, J., Koppel, R., Smith, S.W.: Circumvention of security: Good users do bad things. *IEEE Security & Privacy* **11**(5), 80–83 (2013)
- [19] Blythe, J.M., Coventry, L., Little, L.: Unpacking security policy compliance: The motivators and barriers of employees' security behaviors. In: Eleventh Symposium On Usable Privacy and Security ({SOUPS} 2015). pp. 103–122 (2015)
- [20] Briggs, P., Jeske, D., Coventry, L.: Behavior change interventions for cybersecurity. *Behavior Change Interventions for Cybersecurity* pp. 115–136 (2017)
- [21] Camp, L.J., Lewis, S.: *Economics of information security*, vol. 12. Springer Science & Business Media (2006)
- [22] Caraban, A., Karapanos, E., Gonçalves, D., Campos, P.: 23 ways to nudge: A review of technology-mediated nudging in human-computer interaction. In: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. pp. 1–15 (2019)
- [23] Caulfield, T., Baddeley, M., Pym, D.: Social learning in systems security modelling. *constructions* **14**(15), 3 (2016)
- [24] Caulfield, T., Pym, D.: Improving security policy decisions with models. *IEEE Security & Privacy* **13**(5), 34–41 (2015)
- [25] Caulfield, T., Pym, D., Williams, J.: Compositional security modelling. In: International Conference on Human Aspects of Information Security, Privacy, and Trust. pp. 233–245. Springer (2014)
- [26] Çelen, B., Kariv, S.: Distinguishing informational cascades from herd behavior in the laboratory. *American Economic Review* **94**(3), 484–498 (2004)
- [27] Chamley, C.P.: *Rational herds: Economic models of social learning*. Cambridge University Press (2004)
- [28] Clear, J.: *Atomic habits: An easy & proven way to build good habits & break bad ones*. Penguin (2018)
- [29] Collinson, M., Monahan, B., Pym, D.: *A Discipline of Mathematical Systems Modelling*. College Publications (2012)
- [30] Davis, M.C., Challenger, R., Jayewardene, D.N., Clegg, C.W.: Advancing socio-technical systems thinking: A call for bravery. *Applied ergonomics* **45**(2), 171–180 (2014)
- [31] Deal, T.E., Kennedy, A.A., et al.: Corporate cultures: The rites and rituals of corporate life. *Business Horizons* **26**(2), 82–85 (1983)
- [32] Dekker, S.: *Just culture: Balancing safety and accountability*. CRC Press (2016)
- [33] Demjaha, A., Caulfield, T., Sasse, M.A., Pym, D.: 2 fast 2 secure: A case study of post-breach security changes. 4th European Workshop on Usable Security (EuroUSEC) (2019)
- [34] DePaulo, B.M., Fisher, J.D.: The costs of asking for help. *Basic and Applied Social Psychology* **1**(1), 23–35 (1980)
- [35] Dourish, P., Grinter, E., Delgado De La Flor, J., Joseph, M.: Security in the wild: user strategies for managing security as an everyday, practical problem. *Personal and Ubiquitous Computing* **8**(6), 391–401 (2004)
- [36] Dutson, J., Allen, D., Eggett, D., Seamons, K.: “Don’t punish all of us”: Measuring user attitudes about two-factor authentication. EuroUSEC 2019 (2019)
- [37] Friedman, J.P.: *Dictionary of business and economic terms*. Simon and Schuster (2012)
- [38] Frik, A., Malkin, N., Harbach, M., Peer, E., Egelman, S.: A promise is a promise: The effect of commitment devices on computer security intentions. In: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. p. 604. ACM (2019)
- [39] Furnell, S., Thomson, K.L.: From culture to disobedience: Recognising the varying user acceptance of it security. *Computer fraud & security* **2009**(2), 5–10 (2009)
- [40] Gordon, L.A., Loeb, M.P.: The economics of information security investment. *ACM Transactions on Information and System Security (TISSEC)* **5**(4), 438–457 (2002)
- [41] Grossklags, J., Acquisti, A.: When 25 cents is too much: An experiment on willingness-to-sell and willingness-to-protect personal information. In: WEIS (2007)
- [42] Herley, C.: More is not the answer. *IEEE Security & Privacy* **12**(1), 14–19 (2013)

- [43] Hielscher, J., Kluge, A., Menges, U., Sasse, M.A.: “taking out the trash”: Why security behavior change requires intentional forgetting. In: New Security Paradigms Workshop. pp. 108–122 (2021)
- [44] Horne, R., Weinman, J., Barber, N., Elliott, R., Morgan, M., Cribb, A., Kellar, I.: Concordance, adherence and compliance in medicine taking. London: NCCSDO **2005**, 40–6 (2005)
- [45] Ifinedo, P.: Information systems security policy compliance: An empirical study of the effects of socialisation, influence, and cognition. *Information & Management* **51**(1), 69–79 (2014)
- [46] Information Security Forum: From promoting awareness to embedding behaviours: Secure by choice, not by chance (2014)
- [47] Ioannidis, C., Pym, D., Williams, J.: Information security trade-offs and optimal patching policies. *European Journal of Operational Research* **216**(2), 434–444 (2012)
- [48] Johnson, E.J., Shu, S.B., Dellaert, B.G., Fox, C., Goldstein, D.G., Häubl, G., Larrick, R.P., Payne, J.W., Peters, E., Schkade, D., et al.: Beyond nudges: Tools of a choice architecture. *Marketing Letters* **23**(2), 487–504 (2012)
- [49] Kahneman, D., Tversky, A.: Prospect theory: An analysis of decision under risk. In: *Handbook of the Fundamentals of Financial Decision Making: Part I*, pp. 99–127. World Scientific (2013)
- [50] Karlsson, F., Karlsson, M., Åström, J.: Measuring employees’ compliance—the importance of value pluralism. *Information & Computer Security* **25**(3), 279–299 (2017)
- [51] Keynes, J.M., et al.: *Treatise on money* (1930)
- [52] Kirlappos, I., Beutement, A., Sasse, M.A.: “Comply or die” is dead: Long live security-aware principal agents. In: *International Conference on Financial Cryptography and Data Security*. pp. 70–82. Springer (2013)
- [53] Kirlappos, I., Parkin, S., Sasse, M.A.: Learning from “shadow security”: Why understanding non-compliance provides the basis for effective security. In: *Workshop on Usable Security (USEC) 2014* (2014)
- [54] Kirlappos, I., Parkin, S., Sasse, M.A.: Shadow security as a tool for the learning organization. *ACM SIGCAS Computers and Society* **45**(1), 29–37 (2015)
- [55] Kirlappos, I., Sasse, M.A.: What usable security really means: Trusting and engaging users. In: *International Conference on Human Aspects of Information Security, Privacy, and Trust*. pp. 69–78. Springer (2014)
- [56] Kirlappos, I., Sasse, M.A.: Fixing security together: Leveraging trust relationships to improve security in organizations. In: *Proceedings of the Workshop on Usable Security and Privacy (USEC’15)*. Internet Society (2015)
- [57] Kocher, M.G., Sutter, M.: The decision maker matters: Individual versus group behaviour in experimental beauty-contest games. *The Economic Journal* **115**(500), 200–223 (2005)
- [58] Koppel, R., Smith, S.W., Blythe, J., Kothari, V.H.: Workarounds to computer access in healthcare organizations: you want my password or a dead patient? In: *ITCH*. pp. 215–220 (2015)
- [59] Kraemer, S., Carayon, P.: Computer and information security culture: Findings from two studies. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* **49**(16), 1483–1488 (2005)
- [60] Kugler, T., Kausel, E.E., Kocher, M.G.: Are groups more rational than individuals? a review of interactive decision making in groups. *Wiley Interdisciplinary Reviews: Cognitive Science* **3**(4), 471–482 (2012)
- [61] Lee, F.: When the going gets tough, do the tough ask for help? help seeking and power motivation in organizations. *Organizational behavior and human decision processes* **72**(3), 336–363 (1997)
- [62] Mankiw, N., Taylor, M.: *Microeconomics: Thomson learning* (2006)
- [63] Mathur, A., Engel, J., Sobti, S., Chang, V., Chetty, M.: “They keep coming back like zombies”: Improving software updating interfaces. In: *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. pp. 43–58 (2016)
- [64] Michie, S., Van Stralen, M.M., West, R.: The behaviour change wheel: a new method for characterising and designing behaviour change interventions. *Implementation science* **6**(1), 42 (2011)
- [65] Morisset, C., Yevseyeva, I., Groß, T., van Moorsel, A.: A formal model for soft enforcement: influencing the decision-maker. In: *International Workshop on Security and Trust Management*. pp. 113–128. Springer (2014)
- [66] Morris, J., Becker, I., Parkin, S.: In control with no control: Perceptions and reality of windows 10 home edition update features. In: *Workshop on Usable Security and Privacy (USEC) (2019)*
- [67] Morris, J., Becker, I., Parkin, S.: An analysis of perceptions and support for windows 10 home edition update features. *Journal of Cybersecurity* **6**(1), tyaa017 (2020)
- [68] Pahlila, S., Siponen, M., Mahmood, A.: Employees’ behavior towards is security policy compliance. In: *2007 40th Annual Hawaii International Conference on System Sciences (HICSS’07)*. pp. 156b–156b. IEEE (2007)
- [69] Pallas, F.: Information security inside organizations—a positive model and some normative arguments based on new institutional economics. TU Berlin - Information Systems Engineering (2009)
- [70] Parkin, S., Driss, S., Krol, K., Sasse, M.A.: Assessing the user experience of password reset policies in a university. In: *International Conference on Passwords*. pp. 21–38. Springer (2015)
- [71] Parkin, S., Fielder, A., Ashby, A.: Pragmatic security: modelling it security management responsibilities for sme archetypes. In: *Proceedings of the 8th ACM CCS International Workshop on Managing Insider Security Threats*. pp. 69–80. ACM (2016)

- [72] Parkin, S., van Moorsel, A., Inglesant, P., Sasse, M.A.: A stealth approach to usable security: Helping it security managers to identify workable security solutions. In: Proceedings of the 2010 New Security Paradigms Workshop. pp. 33–50. NSPW '10, ACM (2010)
- [73] Parsons, K.M., Young, E., Butavicius, M.A., McCormac, A., Pattinson, M.R., Jerram, C.: The influence of organizational information security culture on information security decision making. *Journal of Cognitive Engineering and Decision Making* **9**(2), 117–129 (2015)
- [74] Raja, F., Hawkey, K., Jaferian, P., Beznosov, K., Booth, K.S.: It's too complicated, so I turned it off!: expectations, perceptions, and misconceptions of personal firewalls. In: Proceedings of the 3rd ACM workshop on Assurable and usable security configuration. pp. 53–62. ACM (2010)
- [75] Reason, J.: Human error. Cambridge university press (1990)
- [76] Redmiles, E.M., Mazurek, M.L., Dickerson, J.P.: Dancing pigs or externalities?: Measuring the rationality of security decisions. In: Proceedings of the 2018 ACM Conference on Economics and Computation. pp. 215–232. ACM (2018)
- [77] Reeder, R.W., Ion, I., Consolvo, S.: 152 simple steps to stay safe online: security advice for non-tech-savvy users. *IEEE Security & Privacy* **15**(5), 55–64 (2017)
- [78] Reinfeldt, L., Landwirth, R., Benenson, Z.: Security managers are not the enemy either. In: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. p. 433. ACM (2019)
- [79] Renaud, K., Goucher, W.: The curious incidence of security breaches by knowledgeable employees and the pivotal role a of security culture. In: *Human Aspects of Information Security, Privacy, and Trust*. pp. 361–372. Springer International Publishing, Cham (2014)
- [80] Renaud, K., Zimmermann, V.: Ethical guidelines for nudging in information security & privacy. *International Journal of Human-Computer Studies* **120**, 22–35 (2018)
- [81] Robinson, S.L., O'Leary-Kelly, A.M.: Monkey see, monkey do: The influence of work groups on the antisocial behavior of employees. *Academy of management journal* **41**(6), 658–672 (1998)
- [82] Rothrock, R.A., Kaplan, J., Van Der Oord, F.: The board's role in managing cybersecurity risks. *MIT Sloan Management Review* **59**(2), 12–15 (2018)
- [83] Schein, E.H.: *Organizational culture and leadership*, vol. 2. John Wiley & Sons (2010)
- [84] Schlienger, T., Teufel, S.: Information security culture. In: *Security in the Information Society*, pp. 191–201. Springer (2002)
- [85] Shafir, E.: *The behavioral foundations of public policy*. Princeton University Press (2013)
- [86] Shao, X., Siponen, M., Liu, F.: Shall we follow? impact of reputation concern on information security managers' investment decisions. *Computers & Security* **97**, 101961 (2020)
- [87] Shiller, R.J.: Conversation, information, and herd behavior. *The American economic review* **85**(2), 181–185 (1995)
- [88] Simon, H.A.: A behavioral model of rational choice. *The quarterly journal of economics* **69**(1), 99–118 (1955)
- [89] Simon, H.A.: Rational choice and the structure of the environment. *Psychological review* **63**(2), 129 (1956)
- [90] Simon, H.A.: *Models of bounded rationality: Empirically grounded economic reason*, vol. 3. MIT press (1997)
- [91] Steves, M.P., Greene, K.K., Theofanos, M.F.: A phish scale: rating human phishing message detection difficulty. In: *Workshop on usable security (USEC)* (2019)
- [92] Sugden, R.: Why incoherent preferences do not justify paternalism. *Constitutional Political Economy* **19**(3), 226–248 (2008)
- [93] Thaler, R.: Toward a positive theory of consumer choice. *Journal of Economic Behavior & Organization* **1**(1), 39–60 (1980)
- [94] Thaler, R.H., Sunstein, C.R.: *Nudge: Improving decisions about health, wealth, and happiness* (2008)
- [95] Turland, J., Coventry, L., Jeske, D., Briggs, P., van Moorsel, A.: Nudging towards security: Developing an application for wireless network selection for android phones. In: Proceedings of the 2015 British HCI conference. pp. 193–201. ACM (2015)
- [96] Vaniea, K.E., Rader, E., Wash, R.: Betrayed by updates: how negative experiences affect future security. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. pp. 2671–2674. ACM (2014)
- [97] Vasu, M.L., Stewart, D.W., Garson, G.D.: *Organizational Behavior and Public Management, Revised and Expanded*. Routledge (2017)
- [98] Verendel, V.: A prospect theory approach to security. Chalmers University of Technology (2008)
- [99] Vohs, K.D., Baumeister, R.F., Schmeichel, B.J., Twenge, J.M., Nelson, N.M., Tice, D.M.: Making choices impairs subsequent self-control: a limited-resource account of decision making, self-regulation, and active initiative. *Journal of Personality and Social Psychology* (2014)
- [100] Vroom, C., Von Solms, R.: Towards information security behavioural compliance. *Computers & security* **23**(3), 191–198 (2004)
- [101] Warkentin, M., Johnston, A.C., Shropshire, J.: The influence of the informal social learning environment on information privacy policy compliance efficacy and intention. *European Journal of Information Systems* **20**(3), 267–284 (2011)

- [102] Wash, R., Rader, E., Vaniea, K., Rizor, M.: Out of the loop: How automated software updates cause unintended security consequences. In: 10th Symposium On Usable Privacy and Security (SOUPS 2014). pp. 89–104 (2014)
- [103] Yazdanmehr, A., Wang, J., Yang, Z.: Peers matter: The moderating role of social influence on information security policy compliance. *Information Systems Journal* **30**(5), 791–844 (2020)
- [104] Zimmermann, V., Renaud, K.: Moving from a ‘human-as-problem’ to a ‘human-as-solution’ cybersecurity mindset. *International Journal of Human-Computer Studies* **131**, 169–187 (2019)

Appendix. Glossary of terminology, derived from [7, 10–12, 26, 27, 37, 48, 49, 51, 62, 85, 88–90, 94, 99].

Traditional and Behavioural Economics Terminology	
Term	Definition
gain	<i>A gain is an increase in the value of an asset</i>
loss	<i>A loss is a decrease in the value of an asset</i>
cost	<i>A cost signifies the using up of assets</i>
investment	<i>The allocation or use of goods with the expectation of some benefit in the future</i>
rationality	<i>The idea that an individual takes into account all information, probability, potential costs, gains or losses in order to take the most optimal decision</i>
decision	<i>The choice that results in the optimal level of benefit for the decision-maker</i>
rational decision-making	<i>The process of making a choice that results in the optimal level of benefit for the decision-maker</i>
information asymmetry	<i>When one party has more or better information about something than the other party</i>
moral hazard	<i>When an individual takes more risks because someone else is responsible for bearing those risks</i>
principal-agent problem	<i>When one individual has the ability to make decisions on behalf of another</i>
externality	<i>A cost or benefit that is imposed on a third party who did not agree to incur that cost or benefit</i>
information cascade	<i>When an infinite sequence of individuals ignore their private information when making a decision</i>
herding	<i>When an infinite sequence of individuals make an identical decision, not necessarily ignoring their private information</i>
perceived gain	<i>A perceived gain is an increase in the value of an asset that is important and subjective to the decision-maker (as according to limitations of bounded rationality)</i>
perceived loss	<i>A perceived loss is a decrease in the value of an asset that is important and subjective to the decision-maker (as according to limitations of bounded rationality)</i>
perceived cost	<i>A perceived cost signifies a subjective value of an asset as according to limitations of bounded rationality</i>
risk	<i>The possibility or likelihood of losing something valuable</i>
loss aversion	<i>The concept that people are far more psychologically affected by a loss rather than a gain</i>
bounded rationality	<i>The idea that humans have certain constraints — cognitive, time, and information/knowledge — that limit their decision-making processes; as such, boundedly rational agents turn to ‘satisficing’ rather than ‘optimizing’</i>
choice architecture	<i>The practice of influencing an individual’s choice by organising the context in which they make decisions</i>
satisficing	<i>The act of making a decision which is satisfying and sufficient (given the constraints) rather than optimal</i>
social learning	<i>When a person’s beliefs, decisions, or behaviour are affected or altered by some form of social interaction</i>
decision fatigue	<i>Fatigue caused by the difficulty and effort required to make a choice</i>