You've left me no choices: Security economics to inform behaviour intervention support in organizations

Albesë Demjaha[1†], Simon Parkin[2†], and David Pym[1]

[1] University College London and Alan Turing Institute, UK
[2] University College London, UK
{*albese.demjaha.16, s.parkin, 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 due to 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 neoclassical 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. We also point to applications of the framework in negotiating sustainable security behaviours, such as policy concordance and just security cultures.

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 *security function* of the organization. The security function is the part of the organization recognised as having the expertise to identify and manage the security technologies and processes necessary to protect the organization from threats that relate to its

† Authors contributed equally.

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 of ways to promote behaviour change [15], with challenges in guaranteeing that behaviours are changed successfully [53]. Declaring a behaviour in a security policy is then not an assurance that the behaviour will happen. This reality has drawn increasing attention to the need to manage behaviour effectively. Consideration of behaviour change theory and *behavioural economics* [13] is one such approach.

Both research and practice have shown that behaviours 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 [36] (where they may well be [31]). Employees may create their own alternative behaviours [12], sometimes in an effort to approximate secure working, rather than abandoning security [37]. Organizational support can be critical to whether secure practices persist [22], where individuals may assume that others with relevant knowledge and resources will manage the problem for them.

Rational security micro-economics has proved useful for explaining the interaction between organizational security policies and behaviours [8], where security ecosystems are otherwise too complicated to study directly in this way. Herley posits that the rejection of advocated security behaviours by citizens exhibits traits of rational economic behaviour [28].

Security managers must have a strategy for how to provision for security, provide workable policy, and support user needs. In early workshops on the Economics of Information Security, Schneier advocated consideration of trade-offs [14, p. 289]; more than 15 years later this is not happening sufficiently in organizations. Here we revisit principles of information economics and behavioural economics in tandem, identifying contradictions which point to gaps in support. After reviewing the capacity for economics to explain a range of security-related behaviours (Section 2), we demonstrate how current approaches to infrastructure and provisioning of security mirror rational-agent economics, even when behavioural economics is applied to promote individual behaviours (Section 3). We show through examples how these align with regularly cited causes of security non-compliance from the literature.

We present a framework (Section 4), based on consolidated economics principles, with the following goal:

> *Better support for 'good enough' security-related decisions, by individuals 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, which 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.

We then apply the framework to one of the most widely promoted security behaviours (Section 5), the maintenance of up-to-date device software, demonstrat-

ing through comparison with independent user studies where the consolidated economics approach can anticipate organizational support requirements. We consider how the framework can be situated to support practitioners (Section 6), before concluding with a summary and future work (Section 7). A supporting glossary of foundational economics terminology is detailed in the Appendix.

2 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 [26], where here we apply a similar approach to a combination of economic models, to reposition investment challenges related to security behaviour management. Beautement et al. [8] 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, to identify ways to support individuals as they engage in privacy-related decision-making; similarly, Baddeley [6] 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, such as the *endowment effect* [58] and *framing* within the domain of information security and privacy [27, 3]. 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 [25]. Verendel [61] 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* [1, 59]. Through empirical modelling of behavioural economics, Redmiles et al. [49] effectively advocate for identifying and presenting options which 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 these capabilities, which must be closed in order for organizations to support secure behaviours.

In terms of capturing the dynamic between a decision-maker (here, an employee), and the security function – an 'influencer' – Morisset et al. [42] 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.

In summary, there is a need to reconcile the advancements in the application of economics to security with how management of behaviour change strategies in organizations are conceptualised. Here we fill in the gaps, where currently there

are contradictions and shortcomings which act against both the organization and the individual decision-maker.

3 Applying economics to organizational security

Pallas [43] applies institutional economics to revisit information security in organizations, developing a structured explanation of how the centralised security function and decentralized groups of employees interact in an environment of increasingly localised personal computing. Pallas delineates three forms of security apparatus for achieving policy compliance in organizations (as in Table 1): *architectural means* (which prevent bad outcomes by strictly controlling what is possible); *formal rules* (such as policies, defining what is allowed or prohibited for those in the organization); and *informal rules* (primarily security awareness and culture, as well as security behaviours). We demonstrate how a strategic approach is lacking in how to manage the 'medium/high' *marginal costs* of realizing the informal rules which are intended to support formal rules.

Table 1. Costs of hierarchical motivation (reproduced from Pallas [43]).

Meta-measure	**Fixed costs**	**Marginal costs**	**Enforcement costs (single case)**	**Residual costs**
Architectural means	high	medium	none/negligible	none/negligible
Formal rules	low	medium	high	medium
Informal rules	medium	medium/high	low	high

3.1 Rational vs. bounded decision-making

In traditional economics, a decision-making structure assumes a rational agent [56, 57]. The rational agent is equipped with the capabilities and resources to make the decision which will be most beneficial for them. The agent knows all possible choices, and is is assumed to have complete information when evaluating those choices, as well as a detailed analysis of probability, costs, gains, and losses [57]. A rational agent 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 agents make fully rational decisions. Instead, the field refers to the concept of *bounded rationality* which explains that an agent's rationality is bounded due to cognitive limitations and time restrictions. These considerations also challenge the plausibility of complete information, which is practically unrealistic for a bounded agent. According to these restrictions, the bounded agent turns instead to 'rules of thumb' and makes ad hoc decisions based on a quick evaluation of *perceived* probability, costs, gains, and losses [33, 56].

Table 2. Rationality vs. bounded rationality in decision-making.

Traditional economics	*Behavioural economics*
RATIONAL AGENT	**BOUNDED AGENT**
- detailed evaluation of costs, gains, and losses - complete information - careful calculation of potential investment ↓ chosen outcome ↓ optimal decision	- brief consideration of perceived costs, gains, and losses - incomplete information - insufficient skills, knowledge, or time - quick evaluation of risks driven by loss aversion ↓ decision fatigue ↓ satisfactory decision

Table 2 outlines the differences between the decision-making process of a rational agent and that of a bounded agent. The classical notion of rationality (or, rather, *the neoclassical assumption of rationality* [57]) is quite unachievable outside of its theoretical nature. From the standpoint of neoclassical rationality, the decision-making agent 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 agent, it is expected that the taken decision will be the one which provides maximal utility for the agent.

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 [57]. 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 more constructive decision-support 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 [57] argues that people tend to make decisions by *satisficing* [33] 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, people's resources become strained and *decision fatigue* [62] often contributes to poor choices. This leads to our goal to: *better support 'good enough' decisions which best approximate secure behaviours under constraints such as limited time or knowledge.*

3.2 Why we are here, with too few choices

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 [41]: *Capability*, *Opportunity*, and *Motivation*, which are all required to support a change to a particular *Behaviour*. We discuss how each pillar is represented in the two economic approaches.

Traditional economics. The move from centralized to decentralized computing [43] has resulted in an imposed information asymmetry of having a recognized security function distinct from everyone else in the organization. The security function may declare formal rules and informal rules (training, behaviours), 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 then cannot be assumed. Motivation comes from formal policies, and architectural means which force certain behaviours; 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 [11] (also impacting Opportunity). As the security function is distinct from the rest of the decentralized 'PC-computing' organization, it is often assumed that information about advocated behaviours has been sufficiently communicated to the decision-maker (where the Opportunity also cannot be assumed, because the 'trigger' does not match the employee's current Ability and Motivation [19, 46]).

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 closes the information asymmetry). An Opportunity for a new behaviour may be created, through training or shaping of the environment, and assumed to be a nudge toward a behaviour beneficial to the decision-maker [53]. If a behaviour is framed like a 'nudge', but accounts only for what is desirable for the influencer without checking also that it is desirable to the decision-maker, it is a 'prod' which cannot rely on the decision-maker's own resources and willingness to ensure that it works, such that Motivation will fail. If the provisioned choices (the Capability) are no more beneficial than what the decision-maker already has available to them, they may resort to 'shadow security' behaviours [37].

4 A framework for security choices

4.1 Toward a consistent strategy

Current approaches to security provisioning in organization appear as if to support the rational decision-maker, as per traditional economics. We outline the

'contradictions' that currently exist in how the two economic models are being brought together as follows, where examples of 'contradictory' and 'better' approaches to supporting secure behaviours in organizations are illustrated through real-world examples in Table 3.

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' [8] of complying with security, we must consider the *endowment effect* – as also applies to security [35] – and acknowledge that for the busy decision-maker, doing security requires 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' [8] suggests to reduce the demands of security expectations, but does not define an upper bound on expectations.

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., not knowing where to find personal firewall settings [47]). Unchecked, this leads to satisficing. Current approaches appeal to the skillful user, or assume 'non-divisable' target behaviours [4] (with only one way to do what is being asked).

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 unknown to the security function). We should note also that a choice is often *perceived*, and so elements of a choice can impact the 'gulf of evaluation' [52]. Example: users forming incomplete/incorrect understanding of two-factor authentication technology options [23]).

Edit out the old, edit in the new. More security advice is often presumed to be better for security, but is not [29], and can create confusion. Stale advice can persist unless it is curated – 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. Example: hosting both obsolete and new security policies (without time-stamps).

4.2 Bounded security decision-making

Security research increasingly focuses on organizational security and the interaction between managers, policies, and employees. Principles from economics have

Table 3. 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 [37]	Assumes knowledge of policy and time to find it
Secure passwords	Not communicating the rules for a secure password [44]	Assumes expert knowledge about passwords
Authentication choice	Integrating a suite of options into log-on without explaining the options [23]	Lacking support for making reasoned decision
Do secure work	Advocating generic security practices [36]	Staff must relate practices to work
Security training	Provide training but no time to do it [8]	Staff must negotiate the time themselves
Behaviour	**Better Approach**	**Successes**
Policy compliance	Ensure that the environment naturally supports policy-compliant behaviour [37]	Does not assume any extra effort from staff
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 [21], sector-specific tailoring (e.g., differentiated NCSC Guidance for Small Biz. and Small Charities)	Supports interpretation of a perceived choice
Security training	Agree a fixed window of paid time to complete training	Cost to (pri. and sec.) tasks negotiated for staff

been deemed useful in security [14], and concepts from behavioural economics further support understanding of security behaviours in an organizational context [13]. For security policies to be effective, they must align with employees' limited capacity and resources for policy compliance [16].

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 from the tendency to apply behavioural intervention concepts to security while assuming the intervention targets to be rational agents. This is in itself a contradiction because a rational agent would by default make the optimal choice and would not require any behavioural aid or intervention (as explored in Section 4.1). Similarly, employees cannot possibly dedicate sufficient

time or resources for every single task or policy [16]. 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 [16], which is constructed within the modelling framework described in [18, 17]. This model explicitly considers the decision-point for an agent (the decision-maker), and incorporates elements of the decision-making process (where we reconcile elements of behavioural economics), and available choices provided by the organization (the influencer). We adapt this framework to consider factors which should be considered when provisioning security choices, toward supporting the decision-maker to choose 'good enough' behaviours under constraints on knowledge and resources.

Fig. 1. A decision point in a decision-maker's process *bounded security decision-making* (adapting elements from Caulfield and Pym [16]).

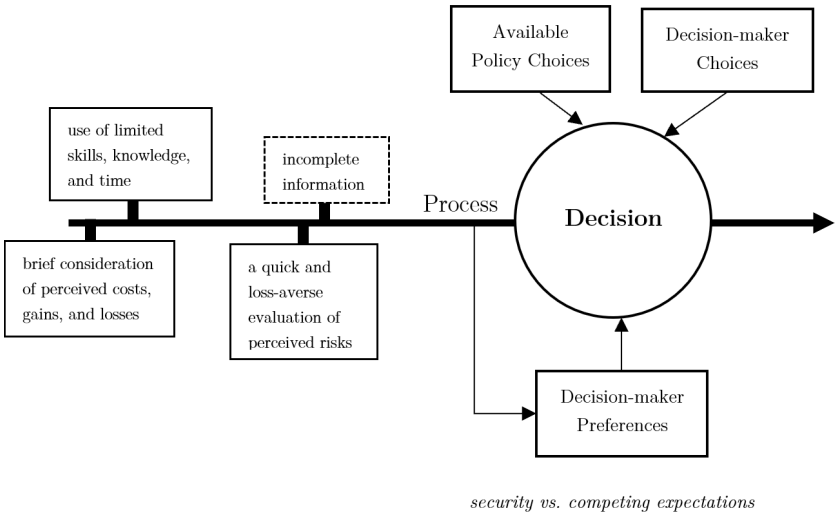

Figure 1 illustrates the components and processes which must be considered at policy design. *Influencer* refers to the security policy-maker in the organization, and *decision-maker* (DM) the bounded agent (the employee).

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) which 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.

When evaluating the risks that come with a choice 'losses loom larger than gains' [34, 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 (which potentially only the influencer – the overseer and expert of security – can see).

Information asymmetry. Information asymmetry regular occurs between the influencer and the decision-maker. In the context of security policies and policy compliance, the following are examples of information asymmetry:

- The recognised differentiation of the influencer being more knowledgeable and capable in security than the decision-maker (as security is the influencer's primary task);
- The influencer's lack of knowledge about the decision-maker's context, and pressures which factor into their choice-making process (resulting in the influencer seeming to perceive the decision-maker as a rational agent 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).

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

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 (such as 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 e.g., an HR policy requires less technical engagement (and time investment), this will factor into the preferences.

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 provided by the policy itself. However, as literature shows, a choice may be to circumvent the policy [12, 38], or to attempt to work in a way that best approximates compliance with secure working policies, in the best way the decision-maker knows to [37]. 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 which introduces risks for the organization [37]). 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.

Moral hazard. When a number of information asymmetries exists 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.

Choice architecture. The circle in Figure 1 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 provisioned security behaviours are unusable, the set of choices comprises 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 which is not being followed.

4.3 Framework implementation

Fig. 2. Implementation steps of the *bounded security decision-making* framework.

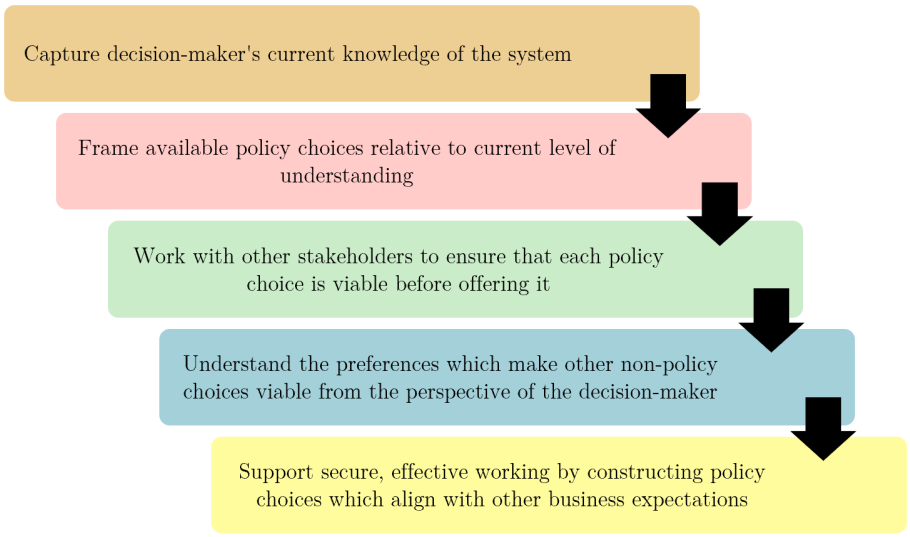

Framework components map to practical implementation steps (Figure 2):

1. **Capture the *process***: Influencers must understand the decision-maker's process (as defined in Figure 1) and consider their current knowledge of the system — either as individuals or discernible groups of users. This may also be influenced by any cognitive limitations [9];
2. **Adapt *available policy choices***: Policy choices must be adapted to the decision-makers current level of understanding and supported with concrete information — working from the decision-maker's current state (of knowledge and resources) rather than the desired security end-state;
3. **Validate *policy choices* with stakeholders**: Collaboration with stakeholders must be established before policy choices are offered so that the decision-maker is not left responsible of ensuring that it is a possible choice amongst other imperatives;
4. **Acknowledge *decision-maker preferences* and choices**: Decision-maker preferences (including their motivations) must be utilized rather than ignored — knowledge of these can aid in aligning policy choices with decision-maker preferences;
5. **Align choices with *competing expectations***: Influencers must ensure that security policy choices do not — at the very least — interfere with other business expectations.

5 Worked example – software security updates

Here we apply the framework to a pertinent case study – keeping software up-to-date. This is selected from the top online security controls advocated by security experts (as prompted by Reeder et al. [50]). This is also the top piece of advice advocated by e.g., the UK government[3].

5.1 Process

Skills, knowledge, and time. Applying updates as soon as possible is seen as achieving the best results [32]. 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).

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.

Perceived costs, gains, and losses. In organizations, system patches are first deployed to a test-bed [32], to ensure that they do not create problems (losses); advice to 'keep systems up-to-date' ignores this, and also does not declare the cost, in terms of time, for a user to achieve this. This would then be concise, high-level advice which inadvertently assumes that a user knows already how to do this, and how often to do it. An employee may not feel that updates are a concern for them [60], so may not be motivated to do it at all.

A *bounded security decision-making* approach would need to 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 this for themselves (and to avoid loss of cognitive automation and a need to rebuild cognitive maps [10]). It would also be necessary to convey that an up-to-date system protects specific assets that the decision-maker wants to minimize losses for (where business management / asset-specific communications could help).

Incomplete information. The minimal advice does not declare how to check or how often, assuming a rational approach. If the 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 [63]). There is also an assumption that the user may know the changes that updates will create in advance, when it may instead impact them in a range of ways [10].

A *bounded security decision-making* approach could involve informing the user of how long each update takes to install [40] (especially if a restart is required), based on testing on a comparable setup (including machine performance,

[3] As at the National Cyber Security Centre (NCSC) website.

available disk space [40] and provisioned software). It may be that updates can be scheduled centrally [40], for instance to occur when employees are most likely to have their computer on, but not be using it (if the organization has scheduled workplace lunch breaks, for instance). Ultimately, finding a time to install updates and avoid disruption is increasingly difficult to find in a PC-computing work environment.

Loss-averse evaluation of risks. A rational approach does not accommodate the chance that the user has had prior bad experiences with updates [60]. 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 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 [60], 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.

5.2 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, so advice may need to consider specific operating system software (for instance). Unless an OS 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.

A *bounded security decision-making* approach would acknowledge how updates work on the system the decision-maker us 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).

5.3 Decision-maker choices

Because choices framed for a rational decision-maker are not made explicit and compared meaningfully, the *bounded security* decision-maker may construct the set of choices in an ad hoc fashion, with little to no information about the consequences of taking action or not doing so (the expertise that the security function has which they personally do not have). In an environment of incomplete information, the security function may not know this either (as may be the case with many policy mandates [29]).

6 Future directions

Informed by recent user-centred security research, we outline directions for how a security manager in an organization (and by extension, the recognized security

function) can consider the proposals we have made (Section 4). Security managers cannot be assumed to have in-depth knowledge of the human aspects of security, but may nonetheless value it in security policy decision-making [45]. They then require methods and tools to do so [51].

6.1 A security diet

A 'security diet' would document perceived occurrence and costs of advocated behaviours (for instance through a typical working day). Questions can then be asked to reconcile these costs with expected behaviour elsewhere in the organization [35], 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 [48]. An individual arguably *should not* be expected to commit more than 100% of their working day to all tasks including security. Security is then self-defeating if it leaves the decision-maker to figure out how to make this possible. Consideration of how to manage security with other pressures can act to reduce the 'gulf of execution' [52].

6.2 Just culture and the genuine choice architecture

If we are to involve the decision-maker in shaping viable options, we would want to find a way to acknowledge the choices employees make which are outside of policy, to include them alongside advocated choices for clear comparison. This does however 'declare' unsecure options, though this aligns with the practice of a 'blame-free', *just culture* [20], toward learning from shortcomings. By defining associated properties of these two sets of choices, support can be negotiated to shape solutions which allow productive and secure working.

6.3 Policy concordance

The 'Security Dialogues' research [5] promotes a move toward policy concordance — 'mutual understanding and agreement' on how the decision-maker will behave. In medicine [30], concordance occurs at the point of consultation, to incorporate the respective views of the decision-maker and influencer.

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 which are regarded as viable.

6.4 Security investment forecasting

Security modelling can begin to forecast the impact of investments in complex environments, before making infrastructure and provisioning changes (e.g., [16]). 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.

7 Conclusion

We have shown how current approaches to security provisioning and infrastructure reflect neoclassical economics, even when concepts from behavioural economics are applied to 'nudge' individual security behaviours. We have constructed a framework that accommodates a set of security behaviours, as a continuous programme of choices which must be provisioned for to adequately support 'good enough' behaviour decisions. We then apply our framework to one of the most advocated security behaviours — software patching — and demonstrate that the rational-agent view is incompatible with the embrace of behaviour change practices in isolation.

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, and; in organizations, security expertise can exist in places recognized by the organization and others not - constructed information asymmetries ought to be accounted for when assessing user behaviours. Future work can involve situated studies in organizations, including participatory design with security managers to develop viable and sustainable security behaviour interventions.

Acknowledgements

Demjaha is supported through a Doctoral Studentship granted 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. 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)
5. Ashenden, D., Lawrence, D.: Security dialogues: Building better relationships between security and business. IEEE Security & Privacy **14**(3), 82–87 (2016)
6. Baddeley, M.: Information security: Lessons from behavioural economics. In: Workshop on the Economics of Information Security (2011)

7. Bateman, H., McAdam, K.: Dictionary of Economics. A & C Black Publishers Ltd (2003)
8. Beautement, 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)
9. Benenson, Z., Lenzini, G., Oliveira, D., Parkin, S., Uebelacker, S.: Maybe poor johnny really cannot encrypt: The case for a complexity theory for usable security. In: Proceedings of the 2015 New Security Paradigms Workshop. pp. 85–99. ACM (2015)
10. Bergman, O., Whittaker, S.: The cognitive costs of upgrades. Interacting with Computers **30**(1), 46–52 (2017)
11. Beris, O., Beautement, 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)
12. Blythe, J., Koppel, R., Smith, S.W.: Circumvention of security: Good users do bad things. IEEE Security & Privacy **11**(5), 80–83 (2013)
13. Briggs, P., Jeske, D., Coventry, L.: Behavior change interventions for cybersecurity. Behavior Change Interventions for Cybersecurity pp. 115–136 (2017)
14. Camp, L.J., Lewis, S.: Economics of information security, vol. 12. Springer Science & Business Media (2006)
15. Caraban, A., Karapanos, E., Gonçalves, D., Campos, P.: 23 ways to nudge: A review of technology-mediated nudging in human-computer interaction (2019)
16. Caulfield, T., Pym, D.: Improving security policy decisions with models. IEEE Security & Privacy **13**(5), 34–41 (2015)
17. 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)
18. Collinson, M., Monahan, B., Pym, D.: A Discipline of Mathematical Systems Modelling. College Publications (2012)
19. Das, S., Dabbish, L.A., Hong, J.I.: A typology of perceived triggers for end-user security and privacy behaviors
20. Dekker, S.: Just culture: Balancing safety and accountability. CRC Press (2016)
21. Demjaha, A., Caulfield, T., Sasse, M.A., Pym, D.: 2 fast 2 secure: A case study of post-breach security changes (2019)
22. 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)
23. Dutson, J., Allen, D., Eggett, D., Seamons, K.: "don't punish all of us": Measuring user attitudes about two-factor authentication. EuroUSEC 2019 (2019)
24. Friedman, J.P.: Dictionary of business and economic terms. Simon and Schuster (2012)
25. 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)
26. 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)
27. 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)

28. Herley, C.: So long, and no thanks for the externalities: the rational rejection of security advice by users. In: Proceedings of the 2009 workshop on New security paradigms workshop. pp. 133–144. ACM (2009)
29. Herley, C.: More is not the answer. IEEE Security & Privacy **12**(1), 14–19 (2013)
30. 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)
31. Information Security Forum: From promoting awareness to embedding behaviours: Secure by choice, not by chance (2014)
32. Ioannidis, C., Pym, D., Williams, J.: Information security trade-offs and optimal patching policies. European Journal of Operational Research **216**(2), 434–444 (2012)
33. 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)
34. 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)
35. Karlsson, F., Karlsson, M., Åström, J.: Measuring employees compliance–the importance of value pluralism. Information & Computer Security **25**(3), 279–299 (2017)
36. Kirlappos, I., Beautement, 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)
37. 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)
38. 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)
39. Mankiw, N., Taylor, M.: Microeconomics: Thomson learning (2006)
40. 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)
41. 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)
42. 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)
43. Pallas, F.: Information security inside organizations-a positive model and some normative arguments based on new institutional economics. TU Berlin - Information Systems Engineering (2009)
44. 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)
45. 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)

46. Parkin, S., Redmiles, E.M., Coventry, L., Sasse, M.A.: Security when it is welcome: Exploring device purchase as an opportune moment for security behavior change. In: Proceedings of the Workshop on Usable Security and Privacy (USEC'19). Internet Society (2019)
47. 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)
48. Reason, J.: Human error. Cambridge university press (1990)
49. 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)
50. Reeder, R., Ion, I., Consolvo, S.: 152 simple steps to stay safe online: Security advice for non-tech-savvy users. IEEE Security & Privacy (2017)
51. Reinfelder, 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)
52. 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)
53. Renaud, K., Zimmermann, V.: Ethical guidelines for nudging in information security & privacy. International Journal of Human-Computer Studies **120**, 22–35 (2018)
54. Richard H. Thaler, C.R.S.: Nudge: Improving decisions about health, wealth, and happiness (2008)
55. Shafir, E.: The behavioral foundations of public policy. Princeton University Press (2013)
56. Simon, H.A.: Rational choice and the structure of the environment. Psychological review **63**(2), 129 (1956)
57. Simon, H.A.: Models of bounded rationality: Empirically grounded economic reason, vol. 3. MIT press (1997)
58. Thaler, R.: Toward a positive theory of consumer choice. Journal of Economic Behavior & Organization **1**(1), 39–60 (1980)
59. 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)
60. 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)
61. Verendel, V.: A prospect theory approach to security. Chalmers University of Technology (2008)
62. 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. (2014)
63. 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)

Appendix. Glossary of terminology, derived from [24, 39, 7, 54, 55, 62, 33, 34].

Economics Terminology	
Term	Definition
gain	*A gain is an increase in the value of an asset*
loss	*A loss is a decrease in the value of an asset*
cost	*A cost signifies the using up of assets*
investment	*The allocation or use of goods with the expectation of some benefit in the future*
rationality	*The idea that an individual takes into account all information, probability, potential costs, gains or losses in order to take the most beneficial decision*
decision	*The choice that results in the optimal level of benefit for the decision-maker*
rational decision-making	*The process of making a choice that results in the optimal level of benefit for the decision-maker*
information asymmetry	*When one party has more or better information about something than the other party*
moral hazard	*When an individual takes more risks because someone else is responsible for bearing those risks*
principal-agent problem	*When one individual has the ability to make decisions on behalf of another*
Behavioural Economics Terminology	
Term	Definition
perceived gain	*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)*
perceived loss	*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)*
perceived cost	*A perceived cost signifies a subjective value of an asset as according to limitations of bounded rationality*
prospect	*The likelihood or possibility of some event occurring in the future*
risk	*The possibility or likelihood of losing something valuable*
co-dependent risks	*When the likelihood of two or more risks are dependent on each other*
loss aversion	*The concept that people are far more psychologically affected by a loss rather than a gain*
bounded rationality	*The idea that an individual's rationality is limited when making a decision because of cognitive limitations and time restriction*
choice architecture	*The practice of influencing an individual's choice by organising the context in which they make decisions*
satisficing	*The act of making a decision which is satisfying and sufficient (given the constraints) rather than optimal*
decision fatigue	*Fatigue caused by the difficulty and effort required to make a choice*