



Productive Security: A Scalable Methodology for Analysing Employee Security Behaviours

Adam Beautement, Ingolf Becker, Simon Parkin, Kat Krol, and Angela Sasse,
University College London

<https://www.usenix.org/conference/soups2016/technical-sessions/presentation/beautement>

**This paper is included in the Proceedings of the
Twelfth Symposium on Usable Privacy and Security (SOUPS 2016).**

June 22–24, 2016 • Denver, CO, USA

ISBN 978-1-931971-31-7

**Open access to the Proceedings of the
Twelfth Symposium on Usable Privacy
and Security (SOUPS 2016)
is sponsored by USENIX.**

Productive Security: A scalable methodology for analysing employee security behaviours

Adam Beautement, Ingolf Becker, Simon Parkin, Kat Krol and M. Angela Sasse
University College London
{a.beautement, i.becker, s.parkin, k.krol, a.sasse}@cs.ucl.ac.uk

ABSTRACT

Organisational security policies are often written without sufficiently taking in to account the goals and capabilities of the employees that must follow them. Effective security management requires that security managers are able to assess the effectiveness of their policies, including their impact on employee behaviour. We present a methodology for gathering large scale data sets on employee behaviour and attitudes via scenario-based surveys. The survey questions are grounded in rich data drawn from interviews, and probe perceptions of security measures and their impact. Here we study employees of a large multinational company, demonstrating that our approach is capable of determining important differences between various population groups. We also report that our work has been used to set policy within the partner organisation, illustrating the real-world impact of our research.

1. INTRODUCTION

In order to express their preferences and requirements, security managers in organisations typically declare a centrally-managed security policy. This is then applied to all IT systems and individuals operating within the domain of the organisation. These policies are informed by the expertise, recommendations and regulatory requirements of the practitioner community, but ultimately must also fit to the working practices of the business itself. Effective security management must therefore tailor policies to both the operational and organisational contexts. For example, a commercial organisation aiming to maximise business opportunities will have a different security profile to a military organisation with a strong preference for confidentiality over availability. Likewise, policies must take in to account not only that the daily working life of an employee is not just about security [15], but also that employee populations are not homogeneous. A policy that is effective in theory may not translate into secure behaviour in practice, if it is not aligned with the productive processes of the organisation and the goals and capabilities of the employees to whom the policy

applies. This creates a need for security managers to empirically assess and compare the policies under their control, in order to determine how well they meet these goals [23]. It is toward this end that the Productive Security project has worked for the last four years.

While technically-focused sources of data – such as system logs – are commonly used to support analysis of policies, they do not provide an insight into employees' thought processes. Security systems are not just the sum of their technical components – user co-operation plays a critical role in providing organisational security, which highlights the need to consider the relationships between people, process, and technology [14]. In addition, an over-reliance on technical solutions can hinder an organisation's capacity to support employees in their productive tasks [30]. Behavioural data is therefore an important factor for effective security management, and a goal of this work has been to create a set of repeatable metrics capable of assessing employee attitudes and behaviour around security. In particular, we develop a methodology capable of identifying areas in which the security policy itself creates incentives for negative behaviour. Rigid systems can force compliance with policy but promote disgruntlement [6]. Where conflict exists between security systems and productive tasks, friction results. Workarounds and 'circumvention strategies' [1] are then likely to develop as users take advantage of system flexibility to modify how technology and procedures work. This reduces security effort but often introduces security vulnerabilities as a side-effect (e.g., using the same password for a number of accounts across both work and personal life). Managers may even be complicit in supporting workarounds if secondary tasks (such as security) stand in the way of business continuity [26]. Different threat models exist within different areas of life, so the vulnerabilities in one space can weaken security in others (e.g., carrying unencrypted USB data devices in transit between work and home) [5].

Balancing the demands of primary, productive tasks and secondary tasks – such as security – introduces cost-benefit dilemmas in which individuals are forced to choose between security and productivity. In particular, security that overburdens the user and is not aligned with their working practices can become less effective [6]. Security is presented to employees as being for their own good, but can introduce externalities, burdening the individual with indirect costs (e.g., changing an increasing number of passwords at regular intervals) [16]. Individuals may, rationally, perceive the personal cost of compliance as greater than the security be-

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.

Symposium on Usable Privacy and Security (SOUPS) 2016, June 22–24, 2016, Denver, Colorado.

benefits gained. As a result, detecting instances where security and business processes are in conflict is critical for both an organisation's security and productivity.

Our methodology is a multistage process designed to elicit realistic responses from employee populations at scale. As it is necessary for our data to support the decision-making process of different organisations of all sizes, both of these points are of great importance. Data that does not closely represent the operational reality of the organisation cannot be used to drive decision-making, as it is not a reliable predictor of future states and outcomes. Likewise a data collection method that is overly time-consuming or does not scale (potentially up to tens of thousands of participants) quickly becomes impractical for larger organisations. Both of these concerns are addressed by the Productive Security (ProdSec) methodology.

A programme of research as heavily based in the operational context of organisations as this could not be undertaken without the partnership of organisations representative of its target audience. Our research project has been fortunate to have such partners and this work could not have taken place without their support and cooperation.

There follows a review of the related literature in section 2. In section 3, we introduce the ProdSec methodology, which leads us to the results of our study in section 4. Discussion and Conclusion follow in sections 5 and 6.

2. RELATED LITERATURE

A number of existing works use surveys and/or interviews to explore the relationship between an organisation's information security policy and employee behaviour. These works examine the impact of attitudes and perceptions on behaviour and consider both intrinsic and extrinsic influencing factors. For example, Pahlila et al. [18] found that attitudes towards security and the habits of individuals can have a significant effect upon the intention to comply with security policies. They also assert that the social environment around an individual will have an effect upon their propensity to comply with policy.

Expanding on intrinsic motivators, Rhee et al. [25] use social cognitive theory to model the influence of experience with security incidents upon self-efficacy, and the role of self-determination upon the outcome of security-related scenarios. A large-scale survey completed by ~400 students found that individuals with high self-efficacy used more security tools and were more vigilant to security, and experience of security compromises negatively impacts self-efficacy.

The notion of competence was also investigated by Workman et al. [33], who explored the "knowing-doing" gap in individuals who have appropriate security skills and knowledge, but who do not apply these skills consistently. Based on the results of a survey in which 588 members of a technology services company participated, the paper concludes that security technology should be user-centred to avoid a tension between assessing threats and use of coping responses.

Siponen et al. [27] utilise Protection Motivation Theory (PMT) to reason about employee compliance with information security policies. The work considers component parts of PMT, namely threat appraisal and coping appraisal (where this includes response costs). A survey was conducted with

917 employees of Finnish companies. Amongst the findings, threat appraisal was found to have a significant impact on intention to comply with information security policies. Employee beliefs about their ability to adhere to policy influence their intention to comply. This finding stresses the importance of perception; the authors assert that if policies are not perceived as relevant by an employee, adherence to policy will be diminished.

Perception was also the focus of work by Bulgurcu et al. [11], who infer that the perceived costs and benefits of compliance (or non-compliance) are formed by the perceived consequences. The authors find that intention to comply is heavily influenced by attitude, beliefs and ability to comply. The relationships between these factors are explored using a survey of 464 employees across a number of organisations. The study identified three belief classes relating to consequences of compliance decisions – benefit of compliance, cost of compliance, and cost of non-compliance.

The prevalence of attitude and perception as themes throughout these works strongly influenced our survey design. However, these surveys all rely on some sort of rating (e.g., Likert) scale, or a sliding scale (e.g., keeping information safe is beyond, or within, a person's control). We build on these themes but opt to take a more immersive scenario-based approach.

A related approach is taken by Albrechtsen and Hovden [4], utilising the differences in skills, perceptions, and interpersonal relationships to characterise the 'digital divide' between information security managers and end-users. The researchers analysed interviews with 11 managers and 18 employees alongside complementary web-based surveys exploring how 87 managers and 151 users assess security threats and vulnerabilities. The study acknowledges that users prioritise other work tasks, that policy is potentially impenetrable and hard to find for the non-expert, and that security provisioning is often one-way. We extend this approach by grounding survey questions in interview outcomes, towards greater resonance with real-world user experiences.

Other methods of constructing scenario content have been attempted. Both D'Arcy et al. [13] and Parsons et al. [19] generate survey questions by drawing on existing literature and interviews with experts. While this makes good use of general information, it does not allow for surveys to be tailored to the specific context of deployment. Darcy et al. use their survey to explore links between stressful information security demands and intentional violation of security policies, to identify workplace factors which contribute to policy violation, including overload, complexity, and uncertainty. Stressful conditions contribute to security coping strategies, as behaviours are adapted in response to stress factors, which then have a knock-on effect on productivity. Where security requirements are perceived as overloading, complex and uncertain, users then become disengaged, implying that high-effort policies can themselves promote insecure behaviour. The inclusion of productivity as a consideration is of particular interest here, mirroring our goal of 'Productive Security'.

Counter to D'Arcy et al. [13], Guo et al. [15] propose a model of what is referred to as 'Non-Malicious Security Violation (NMSV)', validated by a survey, delivered in both paper




Figure 1: Overview of processes in our methodology

and web formats, of employees and their working conditions. The authors look beyond visible behaviours and instead examine the role of attitudes toward security policy violations, such as the productivity advantage of non-compliance, perceived risks, and workplace norms. As in D’Arcy et al.’s work, scenarios are developed based on related literature and interviews with security practitioners and experts, where end-users of policy are not directly engaged. Results imply that job performance advantages, perceived security risk and workgroup norms are key predictors of intention to engage in NMSVs, with users favouring business tasks. The study also found that attitudes toward security policy itself were not significant in driving non-compliant behaviour, in contrast to Bulgurcu et al. [11]. The authors recommend a user-centred security management strategy, where employees can satisfy productivity goals while also maintaining security.

A scenario-based approach is also taken by Blythe et al. [9], who study how individual and organisational factors in the workplace impact secure behaviours, using interviews based on 16 ‘vignettes’. These cover security behaviours identified from information security policies. Vignettes were effectively used as a device for building a rapport with participants and eliciting attitudes and beliefs relating to a specific subject, an approach reflected here in our interview technique. Results suggest that research should focus on individual security behaviours rather than beginning and ending with policy compliance, and that participants accepted responsibility for some elements of security, while leaving others to their organisation.

Building on the existing literature, the work presented here uses an immersive scenario-based survey, moving away from severity-based questionnaires to situate surveys in the environment in which they are deployed. Scenarios are derived from the results of a semi-structured interview process, based on common areas of friction. We discuss this methodology in more detail in the following section.

3. METHODOLOGY

The goal of the ProdSec methodology is to provide researchers studying organisations with a repeatable, scalable data gathering process that allows them to better understand the security-related issues facing their employees, and the behaviours and attitudes they adopt in response.

The full ProdSec methodology consists of two independent, iterative processes. Figure 1 illustrates the steps involved. The two cycles represent a *passive* data collection and monitoring phase on the left, and an *active* intervention phase on the right. This paper will focus on the passive cycle, although an overview of the active cycle is presented here for contextual clarity.

The passive cycle identifies the predominant security behaviours and attitudes within an organisation, along with specific points of friction between the business and security processes. In order to support real-world decision-making, data collection must both accurately represent the real-world environment where it is applied, and be sufficiently scalable so as to be of use to potentially very large, multi-national organisations. These two goals are to some degree in conflict. Rich, in-depth data capable of accurately representing a real-world context can require a greater investment of time and effort to collect, making it problematic at scale. ProdSec tackles this by utilising a two-stage method.

Firstly, semi-structured *Interviews* are conducted with a vertical cross-section of the organisation to capture attitudes and behaviours across as many roles, physical locations and demographic groups as possible. We discuss this aspect in section 3.1. Based on interview findings, we carefully craft a scenario-based survey that reflects dominant security-related issues. By tailoring our survey to each operational context, we ensure that survey questions are relevant and recognisable to participants, with the aim of eliciting more realistic and genuine responses.

Once this cycle is complete, security practitioners then have the choice of monitoring the situation over time by repeating the measurement cycle at some future interval (e.g., \approx 6 months later), or actively engaging with any uncovered problems. The right half of figure 1 describes the active *Intervention* phase. Based on the conclusions drawn from the passive cycle we work with the organisation to prioritise the issues identified, and design and deploy optimal intervention(s), taking into account business as well as socio-technical factors section 3.8. Direct collaboration is important as interventions also need to be suitably centred around the human.

Visualising the methodology as cyclical is essential for understanding its intent, reflecting the notion that security is a process and not a fixed state. It is our experience that organisations often see the implementation of an intervention as the final step, whereas we consciously position this step as part of an ongoing process. The passive cycle can therefore be used to track changes in attitude and behaviour over time, as a consequence of the organisation’s evolution or in response to specific interventions. Ongoing monitoring can inform decision-makers as to whether interventions are having the desired effect, and indeed whether interventions have themselves influenced security processes. Likewise, no one set of interventions will provide a ‘silver bullet’ solution, requiring repetition of the active cycle.

The analysis in this paper focuses on the data collection and analysis stages of the passive cycle. As this research spans three years, with two main phases of data collection, lessons have been learnt along the way. As such, we refined our methodology between the two rounds of data collections, both at multinational companies. In the interests of clarity, the methods given below are those used in the second, or most up to date, version of the methodology. Where lessons learnt between the two phases are relevant and of interest, they have been included in the discussion.

Sections 3.2 and 3.3 describe the methodology from semi-structured interviews to scenarios and the subsequent design of the scenarios. As a last step in the passive cycle we study the responses of employees to the scenarios, by way of the survey. We discuss the analytical approach in section 3.7 and boxout 4. The results from the survey and related analysis are described in section 4.

3.1 Semi-structured interviews

Interviews are resource-intensive to both conduct and analyse, which limits their use on a large scale. Interviews do however have several advantages, most notably that they provide the rich contextual information needed to ensure that more scalable forms of data gathering, such as surveys, can capture realistic and relevant data. The interactive nature of interviews also allows the researcher to probe more deeply on topics of interest during the data gathering process.

Although we have provided a set of questions in Appendix C, it should be noted that semi-structured interviews are not restricted to a fixed set of questions. In our work, the researcher attempts to build a rapport with the interviewee and guide the discussion through topics of interest, from the perspective of the individual. Each interview would last an hour and cover a range of security-related topics including: security awareness, data sharing, password management, laptops and removable media, remote working, clear desk policy, physical security, reporting and training.

We aim for ≈ 100 interviews to be conducted at each organisation, making it impossible for a single researcher to conduct them all. We recognise that this introduces problems of consistency – different interview techniques yield different levels of insight, engagement, and expression from participants. Interviewer training involving all interviewers focused on where and how to ask more follow-up questions through careful “probing”. We also attempted to standardise topic areas and question phrasing. This was adopted ex-

plicitly in our second round of data collection, having learnt from the first round that while we were able to identify where problems existed within the organisation, we did not have a good grasp on how frequently those problems occurred, or how widespread they were. The interview topics and probing questions can be found in section C.

Participant recruitment was managed in conjunction with the partner organisation. Our goal was to speak with a vertical cross-section of the organisation. Security problems are not confined to any one employee group; to develop a full understanding of the current organisational security culture, interviews should ideally be conducted within a variety of departments and with employees across a range of roles. This was not always possible due to internal pressures within the partner organisations. For example, in the company for which results are presented here, we were only able to interview within their Operations division as other divisions had been involved in a different survey process close to that time and there were concerns about data collection fatigue within the organisation.

In each case, we ask volunteers to take part in an interview, incentivising participation with – in the case of the survey presented here – a raffle prize. Making the process voluntary rather than mandatory carries both advantages and disadvantages. The success of semi-structured interviews is heavily dependent on the level of rapport the interviewer can develop with the participant – a participant that opens up to the interviewer is likely to give more honest and detailed responses. This is particularly true in the case of security interviews, given that discussion can potentially touch upon self-reporting of transgressions, rule-breaking and circumventions. A good rapport is then necessary to build the trust that is necessary between interviewer and participant.

The interview study was promoted in an item in the company newsletter, asking for employees to talk about security issues. Volunteers responding to the item may well have an agenda of their own. If an individual is keen to talk to the interviewer or otherwise be open about their views on security, this in itself may not represent a balanced view of the security culture in the organisation. It is then difficult to ascertain if their view is typical of the wider population or not. We favoured using volunteers as our methodology includes a survey stage, in part to ascertain the prevalence of any problems identified during the interview stage. We regarded it as preferable to ensure responsive and in-depth interviews, and aim for more representative participation at the survey stage.

3.2 Interview analysis

For interview findings to be of use throughout the data collection methodology, key insights must be extracted from the interviews and made available in a more usable form distinct from the original transcripts. This was achieved by a process of thematic analysis [10] conducted by three researchers. As with the interviews, we recognised the importance of consistency. To this end, we collaboratively developed a codebook that allowed us to systematically apply codes across the interviews. Initial coding was conducted by each interviewer on separate transcripts, with codes being added as necessary in order to capture new concepts as they arose. Regular coding meetings were then scheduled to merge and

Boxout 1: Expert- and Employee-driven Scenarios

The strict demand for the methodology to base survey questions on what emerges from the interviews derives from our experience with the first round of data collection with another company. Here, we included scenarios in the survey which were based on what in-house security experts asserted was happening. However, the results from these scenarios were less coherent and tended to be more polarised than those developed strictly from interview data, implying that they were less recognisable to participants. As the aim of the work is to improve both security and productivity, our survey with the second organisation was built solely on how security is perceived by employees, relative to their primary tasks.

prune these code sets with a view to avoiding instances of duplicate or very similar codes. Codes were also grouped into code families at this time, covering the major topic areas observed within the interview transcripts. After several iterations of this process, we arrived at a largely stable code set of approximately 120 codes that was flexible enough to accommodate the range of topics found in the interviews.

Once this coding process was complete we were then able to tally up the codes to identify the most common security-related issues. These then formed the foundation for a more scalable data collection method that was nevertheless grounded in real-world situations. It is important to note that this process is necessarily bespoke for each context. Interviews conducted in each organisation will reveal different problems, cultures, and technologies. Although we did find problems common to both organisations we assessed (see boxout 2) there were still many differences between how these issues were expressed and the contributing factors that surrounded them.

3.3 Online scenario-based survey

Our approach to scalable data collection was driven by an online survey. In order to efficiently reach large numbers of people, it was necessary to allow them to take part in the data collection exercise from any location, in particular from their usual work environment. Not only does this increase the response rate by minimising demands on participation, but it also furthers our aim of making data collection as naturalistic as possible, as the collection environment matches the operational environment being assessed. As our primary method of recruitment and communication with participants was through company newsletter emails, we embedded a link to the survey in an issue of the newsletter.

As described in section 2, many surveys present short questions with either a multiple choice or Likert scale-style answers. It is our view that this approach is unlikely to engage participants, in particular due to a widespread fatigue with questions of this style. Participants are likely to skim through the survey and apply little thought to their answers. Also, short questions would not allow us to utilise the full value of the interview information. As such, we elected to build scenario-based survey questions, in which participants

Boxout 2: Scenario Commonality

Although the companies that we worked with operate in different sectors there was some overlap in the issues that arose as the result of our interview analysis. Clear desk policies, tailgating through physical security and file sharing were present as issues in both environments. This suggests that as more companies are assessed, a database of scenarios could be developed that over time reduces or minimises the need for the interview stage. However, despite the similarities in topic area it was still necessary to alter key details in the text accompanying each scenario, in order to present to each participant set a scenario that approximated the reality of their environment. For example, one company observed tailgating through security doors, the other through turnstiles in their foyer. Accurately representing these details increases the realism of the scenario, with an aim to encourage honest responses.

were presented with one of the common situations identified via our interview analysis.

Once a topic was selected the scenario was written, using organisation-specific details and terminology from the interviews (see boxout 2). For each scenario, we also created four possible answers or outcomes, again drawing on the interview data to craft these so they appeared familiar and plausible to the participants. How these options were then used will be covered in more detail in section 3.6.1. Our goals here were to:

- Present scenarios to the participants that seemed both realistic and familiar,
- Offer answer options that were likewise realistic and familiar,
- Gather as much implicit data as possible to maximise the benefit of the survey while minimising the time taken per participant.

A potential problem with survey – especially one covering a sensitive topic such as security – is that participants attempt to give the answers they feel are expected of them, or are correct, rather than an honest reflection of their own thoughts and views. These deviations fall into two main categories, response bias and demand characteristics. We addressed these in different ways.

Response bias refers to biases introduced by the participant being influenced by sensory inputs and cognitive processes when answering the question, and thus unintentionally altering their response. For example, how a question is phrased can alter the response (in particular if it is a leading question). Aside from eliminating instances of leading or priming language we also took care to phrase our attitude scenarios from the point of view of a fictitious colleague, as participants are often more comfortable reporting on the actions of others. So rather than asking, ‘what would you do in this situation?’, we asked, ‘what would Jessica do?’, intending this to counter some aspects of the response bias. This

helps us obtain an accurate representation of the maturity levels of the employees.

Demand characteristics refer to the fact that research participants might speculate about the purpose of the research and give responses that they think align with what the researchers are trying to find out. This is something we were particularly concerned about, as security is a sensitive topic with potentially significant outcomes. As we were essentially asking people to report on their own rule breaking, the potential for participants trying to give ‘right’ answers was high. To tackle this we made sure that in each scenario there were no obviously correct options — each possible answer involved some difficulty or transgression.

We split the scenarios into two types, based on our interview analysis. When participants reported incidents it was either in the form of something they did themselves, or something they observed their colleagues doing. Our scenarios followed the same approach, and were divided into Behaviour and Attitude scenarios.

3.4 Behaviour-type scenarios

These scenarios present the actor with a situation that puts the requirements of the primary business process in conflict with some aspect of the security policy. Typically this involves the actor in the scenario needing to complete a specific task, the completion of which is being slowed down or prevented by a security process or mechanism. Four options were then given that presented courses of action that would resolve this conflict. Each of these options contained an element of non-compliance so as to avoid participants seeking to give the ‘right’ answer.

In pursuit of our goal to capture rich behavioural data, we linked each answer option as closely as possible to one of four behavioural risk types. This meant answer choices also allowed us to monitor the prevalent behaviour types. Crossler et al. [12] posit that cultural theory can be used as a predictor of the impact the norms of an organisation can have upon perceptions of security-related risks. The behavioural risk categories used in our surveys stem from Adams [2], and their characteristics are given below:

Individualists rely on themselves for solutions to problems,

Egalitarians rely on social or group solutions to problems,

Hierarchists rely on existing systems or technologies for solutions to problems,

Fatalists take a ‘naive’ approach to solving problems, feeling that their actions are not significant in creating outcomes.

Individualists may for instance feel less loyalty to others in the organisation and to policy, but may be more likely to report what they see as inappropriate behaviour to others [12]. The topics covered in the Behaviour-type scenarios were:

- Password manager,
- Use of the company VPN for remote working,
- File storage, and
- Conducting a credit check in a retail location.

Full texts for these scenarios can be found in appendix A.

3.5 Attitude-type scenarios

To further explore the security culture in an organisation, we complement the behaviour-type scenarios with attitude-type scenarios. Rather than presenting participants with a task, the actors described in the attitude-type scenarios observe an instance of non-compliance in their environment — such as finding a screen unlocked — and respondents are asked to indicate how they would react. The four options in this case represent distinct responses, such as to confront the transgressor, or dismiss the incident as commonplace. As with the behaviour-type scenarios, each response contained an element of non-compliance or an implicit cost. While it may seem like confronting a transgressor is an obvious right answer, there is in fact a high social cost associated with doing so (we find for instance that typically security-conscious individuals are regarded as paranoid by their peers).

The answers here were linked not to behavioural risk types but to a model of cultural maturity that has been developed in support of this work. The model considers security competence relative to an individual’s business tasks. Other works describe the need for competence in repeatable tasks which can form good security habits [31]. Here, we also consider the capacity to embody policy (where it is clear) and adapt it to new or complex situations that require a conscious response, a distinction that has been explored by Reason in the realm of safety [22].

Our model contains a series of levels (see section B) which attempt to articulate the maturing relationship between the individual and security policy. Those at the lower levels engage with security only as absolutely necessary, while those at the higher levels champion security in their local environment. The levels linked to the answers in the survey are as follows:

Level 1: Is not engaged with security in any capacity.

Level 2: Follows security policy only when forced to do so by external controls.

Level 3: Understands that a policy exists and follows it by rote.

Level 4: Has internalised the intent of the policy and adopts good security practises even when not specifically required to.

Level 5: Champions security to others and challenges breaches in their environment.

Although the model includes a level 1, practically speaking individuals at this level will not be found in an organisational environment, as there is typically infrastructure in place that at the least requires employees to have a registered username and password to facilitate access to IT resources. As such our survey utilises level 2 and upwards. Level 2 assumes that compliant behaviour must be imposed upon individuals to ensure that routine tasks remain secure, and so in turn the IT infrastructure constrains behaviour. This is analogous to ‘basic hygiene’ as described by Stanton et al. [29], acting to manage the ‘dangerous tinkering’ and ‘naive mistakes’ which might otherwise happen. However, organisational security is complex and technology-based solutions alone cannot anticipate and manage all situations. Level 3 assumes that employees have enough security knowledge to make some in-situ decisions, whereas toward level 5 employees know enough to apply their knowledge and skills to unforeseen situations,

Boxout 3: Assessing Non-Compliance

The first round of data collection through interviews gave us insights into why participants decided to break policy in order to complete a business talk. However we were unable to establish the prevalence of this behaviour, leading us to augment the survey. We added in an additional question to the business type scenarios that asked participants if they found it acceptable to prioritise the business process in this way.

Boxout 4: Hotspots

We refer to instances of ranking scores being positively correlated with severity rating scores, or negatively correlated with acceptability rating scores, as *'hotspots'*. Where these correlations are detected, it indicates that participants are favouring the use of options that they know carry high risk and which represent unacceptable forms of behaviour. Hotspots represent significant areas of concern for the organisation, as they show areas in which employees report that they have to choose (knowingly or unknowingly) insecure practices.

as well as to articulate workable solutions to those around them.

The topics covered by the attitude-type scenarios were:

- ID badges,
- Clear desk policy,
- Tailgating through physical barriers, and
- Secure disposal of confidential hardcopy.

Full texts for these scenarios can be found in appendix A.

3.6 Scenario tasks

For each scenario, a survey question would have two phases – participants would be asked to select their preferred option, and also complete a rating task.

In the case of the behaviour-type scenarios, the ranking task would involve asking participants to rank all four options in order of how likely they would be to take a particular course of action. For the attitude-type scenarios, this ranking would ask for participants to indicate how strongly they agreed with the response of the actor within the scenario.

With the ranking exercise complete, participants were asked to complete a rating task. Here a participant would be asked to rate on a Likert scale of 1 to 5 how severe a breach of security the behavioural-type options were to them, with 1 being *not severe at all* and 5 being *very severe*. In the case of the attitude-type scenarios, a participant would be asked to rate on a scale of 1 to 5 how acceptable the options were, with 1 being *not acceptable at all* and 5 being *very acceptable*. It should be noted that the survey software was set up to not allow participants to backtrack and change previous answers. This was an intentional design choice so participants could not adjust previous answers to align with subsequent ratings, allowing us to detect discrepancies between the two tasks as a way of highlighting areas significant friction between employees and security. We identify these areas by performing a statistical correlation between the rating and ranking scores for each scenario. We describe this analysis in boxout 4.

The behaviour-type scenarios had an additional question that preceded the ranking and rating tasks. The participant is asked to evaluate the severity of the scenario presented to him, allowing us to assess the participants willingness to trade of the completion of the business task and the preservation security.

3.6.1 Scenario selection and distribution

For each organisation, 8-10 scenarios were created. However, for several reasons we did not wish for each participant to complete the entire set of scenario questions. First, it would be too time-consuming; our partner organisations were generally concerned about the productivity impact of large-scale data collection involving any number of employees over a short period, and so we wished for our scenario survey to be completed in 10-15 minutes. Second, as the scenarios were tailored to specific topics not all of them would be relevant to all parts of the business. For example, giving a question about a retail environment to an engineering division will yield data based on guesswork rather than experience. During deployment, we would request a range of demographic information – including business role (the options for which were drawn from the company's structure) – at the beginning of the survey, then deploying a subset of 3-4 scenarios to each participant based on their responses. This is an example of where it is important – and necessary – to engage with a partner organisation at the right level, to ensure that survey tools etc. can be managed in a way that fits naturally with activities within the business.

3.7 Survey tasks – scoring method

As discussed in sections 3.4 and 3.5, each of the options given with each scenario was linked to either a behaviour type or a maturity level. In order to determine the prevalence of each categorisation, a scoring method linked to the ranking task (see section 3.6) was used. The position of the option in the ranking task determined the score of the associated type. This score was cumulative over the scenario questions. For example, if the first option was linked to Behaviour Type then the ranking of this option determined the score given to Type 1. The scoring was as follows:

Rank 1: 4 points

Rank 2: 3 points

Rank 3: 2 points

Rank 4: 1 point

As each participant answered a maximum of two behaviour and two attitude questions, these scores were normalised for each participant, enabling statistical analysis. This scoring system was also used to determine the popularity of the scenario options themselves.

3.8 Selecting interventions

Organisations may enact ‘interventions’ in order to influence a change in the regular security behaviours of employees. The rich picture of employee behaviours and attitudes provided by the interview and survey process means that the methodology described in this paper can support a more systematic and informed approach to the identification of interventions. While this paper does not cover the outcomes of this step in detail (the right-hand side of Figure 1), the intended use of the data and survey results are included here in the interests of completeness.

The interview and survey results provide security managers with information on the most pressing problems – the ‘hot-spots’ (see boxout 4) – encountered by employees in their own organisation’s IT environment, as well as an idea of the factors that underpin an issue. Interventions can then be targeted at the motivating factors of an issue, rather than the symptoms or the elements of it which most relate to specific regulatory expectations. It is intended that researchers would engage with the organisation to determine which category of intervention is optimal, drawing on system (re)design, awareness and training, or technical controls where appropriate. Means of addressing tensions between security and productivity are further discussed in [6]. Having some sense of the scale of a policy hotspot is also useful (how many employees regularly enact an insecure behaviour), as organisations can then invest resources proportionally, and crucially consider the intended scale of an intervention to ensure that it is properly implemented and does not introduce problems of its own (for instance by updating only a subset of the awareness materials which employees are directed to use, which can in turn introduce inconsistencies).

3.9 Research ethics and data handling

The study successfully went through an ethics approval process at our institution (approval number: 3615/002) and was registered with the Data Protection Act (registration number: Z6364106/2012/11/08). We had a written agreement with management – which was distributed with the recruitment email – that employees would not face negative consequences for policy violations they reported. The audio-recordings were transcribed by an external company under NDA. Transcripts were redacted to remove any identifying information such as names of people and locations. The original audio recordings were deleted.

4. STUDY RESULTS

In this section, we present the analysis of the survey data collected at the second large company we studied. We focus on the analysis of maturity scores (section 4.2) and behaviour types (section 4.3) between the different groups of the organisation: business division (sections 4.2.1 and 4.3.1), age group (sections 4.2.2 and 4.3.2) and location (sections 4.2.3 and 4.3.3).

4.1 Response rate

In total, 641 complete survey responses were recorded. The briefing document informed participants that any surveys completed in less than 5 minutes (minimum reading time in our pilot study) would not be included, which left us with 608 responses for analysis.

For the purposes of our study, the organisation is split up across 7 business divisions as well as a number of locations. The majority of responses originated from the *Sales & Services* (292), followed by *Operations* (152). The remaining divisions were all significantly smaller, ranging from 11 to 47 responses. Participation was more equally divided between the business locations surveyed, with locations 1: *HQ* and 5 (a large regional office) being the largest ones with 118 and 130 responses respectively. Further we analyse trends across 8 age groups. Survey respondents were approximately normally distributed across the age groups, with the age group 30 – 34 representing the largest share with 124 participants. The edge cases of < 25 and ≥ 55 were nonetheless sufficiently large with 53 and 22 responses respectively, to allow for potentially statistically significant results across all age groups.

The number of responses were sufficient to allow a factor analysis by business division (sections 4.2.1 and 4.3.1), age group (sections 4.2.2 and 4.3.2) and location (sections 4.2.3 and 4.3.3), with 8 factors each. A full factor analysis with 512 factors is outside the scope of this methodology at present as it would a sample size several orders of magnitude larger.

4.2 Maturity levels analysis

Each participant responded to at least one maturity type scenario (section A), by ranking the four options presented in order of preference as well as assigning an acceptability score on a Likert scale to each option (see section 3.3). A comprehensive statistical analysis was then carried out on these responses, with the results in table 1. In total, three such tables have been produced, but only the first one is shown here for brevity. The remaining diagrams are included in the supplementary material (see section 6.2).

The last line of table 1a shows the full organisation’s maturity level properties (please refer to the caption of table 1 for the details of the statistical analysis carried out). The ranking and acceptability score of each of the maturity levels are all statistically significantly separated and increasing with the maturity score. Level 5 has an average rank of 3.30 and acceptability score of 4.51. These ranks are high – a perfect score would represent an average maturity rank of 1, 2, 3 and 4 for levels 2 to 5 respectively. Further, there is a strong positive correlation between rank and acceptability score: the more acceptable the option, the more likely the participant is to choose it.

4.2.1 Maturity by division

Table 1a illustrates the relationship between maturity levels and business divisions. The data is shown in terms of variations from the organisation’s mean in order to facilitate comparisons across the business divisions. There are a number of interesting deviations from the organisational mean. Only the *Sales & Services* division ranks maturity level 5 statistically significantly above level 4, where the *Finance & Prof. Services* division ranks maturity level 4 highest and statistically significantly higher than level 5. The participants from the other divisions did not discriminate between level 4 and 5 options. Participants from *Sales & Services* opted for responses corresponding to level 5 statistically significantly more often than any other division in the organisation with a mean level 5 rank of 3.64.

Business Division	Level 2		Level 3		Level 4		Level 5		τ
	Rank**	Accept**	Rank**	Accept**	Rank**	Accept**	Rank**	Accept**	
Business	0.33**	0.29*	-0.23*	-0.32*	0.10	0.33**	-0.21**	-0.14*	0.61**
Finance & Prof. Services	0.37**	0.24**	-0.12*	-0.23*	0.13*	0.37**	-0.38**	-0.34**	0.59**
Human Resources	0.23**	0.18	-0.21	-0.47	0.14	0.20**	-0.16**	-0.15	0.66**
Marketing & Consumer	0.43**	0.37**	0.12*	-0.15	-0.10	0.29	-0.45**	-0.42**	0.50**
New Business	0.33*	0.32	-0.32*	-0.23	0.19	0.47**	-0.21*	0.04	0.55**
Operations	0.23*	0.14	0.25**	-0.34**	-0.13**	-0.23**	-0.34	0.06**	0.45**
Other	0.30**	0.30**	-0.33**	-0.42**	0.25**	0.27**	-0.22**	-0.38**	0.65**
Sales & Service	-0.31**	-0.23**	-0.03**	0.34**	-0.00**	-0.06**	0.34**	0.12**	0.76**
mean	1.48	1.50	2.13**	2.19**	3.08**	3.98**	3.30**	4.51**	0.62**

(a) Maturity level rankings and acceptability score split by business division.

Business Division	Scenario Sev**	Individualist		Egalitarian		Hierarchist		Fatalist		τ
		Rank**	Sev**	Rank**	Sev**	Rank**	Sev**	Rank**	Sev**	
Business	0.52**	0.17	-0.10**	0.59**	-1.18**	-0.53*	0.29	-0.23	0.10	-0.22**
Finance & Prof. Services	0.38**	0.34**	-0.24	0.50**	-0.76**	-0.67**	0.29	-0.16	0.15	-0.19**
Human Resources	0.62*	0.53*	-0.13	-0.09	-0.90**	-0.16	-0.41	-0.29	-0.73	0.08
Marketing & Consumer	0.84**	0.32	-0.74**	0.86**	-0.84**	-0.47*	0.34*	-0.71**	0.18	-0.22**
New Business	0.58	0.59*	-0.58	0.62	-0.39	-0.80*	-0.11	-0.41	0.01	-0.39**
Operations	0.03**	-0.02**	-0.33**	-0.34**	0.38**	-0.40**	-0.39**	0.76**	-0.96**	-0.48**
Other	0.24	0.35**	-0.46*	0.04	-0.53**	-0.26	0.11	-0.13	-0.16	-0.28**
Sales & Service	-0.28**	-0.18**	0.37**	-0.06	0.25**	0.48**	0.10*	-0.24**	0.50**	-0.17**
mean	2.24	2.68**	3.49	2.02	3.76**	2.80*	3.20	2.50**	3.44**	-0.20**

(b) Behaviour types rankings and behaviour severity score split by business division.

Table 1: The values in each cell of the tables above describe the variation from the mean in their column, with the mean being shown at the bottom (the mean is the value for the organisation as a whole). Based on the scoring in section 3.7, higher ranks imply more popular choices. Similarly, the higher the Accept/Sev score, the more acceptable/severe the participants take the option to be. In the second row, the **/* after Rank/Accept/Sev show statistical significant variations from the median rank or acceptability or severity score respectively based on the Kruskal-Wallis H-test for independent samples at $p < 0.01/p < 0.05$ confidence respectively. If this Kruskal-Wallis test shows statistical significance, for each subgroup a two-sided Mann-Whitney rank test between this subgroup and the union of all other subgroups is carried out; the results of these tests are shown by further **/* at each number, showing statistical significance at $p < 0.01/p < 0.05$ confidence respectively.

Further, the colours show the order of mean Rank/Accept/Sev for each of the groups (i.e., ranking them horizontally). The largest mean is given the darkest colour, and the colour changes to a lighter shade if there is a statistically significant difference between the distribution of ranks/scores of the current mean and the next largest mean, based on a one-sided paired Wilcoxon rank test. This statistical test is further shown by **/* at the value of the higher cell, showing $p < 0.01/p < 0.05$ confidence respectively. If more than one cell contains the same colour, there is no statistical significant variation between the ranks/scores for these options.

Lastly, the rightmost column τ lists Kendall's τ correlation coefficients between the rank and the acceptability/severity score respectively for each of the groups. Kendall's τ ranges from -1 (perfect anti-correlation) to 1 (perfect correlation). **/* signifies rejecting the null hypothesis of independence (i.e. $\tau = 0$) with statistical significance at $p < 0.01/p < 0.05$ confidence respectively.

The acceptability scores demonstrate a similar trend. Only *Operations* and *Sales & Service* discriminated between level 4 and 5. Yet none of the divisions inverted the ranking. *Operations* are noteworthy since they clearly distinguished between level 2 and 3 maturity as well as acceptability scores.

4.2.2 Maturity by age

There are three age groups that did not discriminate between level 4 and 5 maturity levels: 35 – 39, 50 – 54 and 55+. The 35 – 39 group also shows statistically significant lower average level 5 rank than the other age groups, but ranks level 4

statistically significantly higher than the other age groups. All age groups ranked the acceptability of the options according to the maturity levels.

4.2.3 Maturity by location

Responses from location 1: HQ rank maturity level 4 higher than level 5 as well rank level 2 significantly higher than all other locations. This is also evident in the acceptability score: level 5 is perceived as statistically significantly less acceptable and level 2 as more acceptable than at all other locations. Employees at locations 4, 5 and at minor offices

were unable to distinguish between levels 4 and 5. Location 3 achieved the highest average level 5 rank of 3.6, statistically significantly higher than the average.

Acceptability scores only varied significantly for staff at the minor offices, which collectively scored level 5 with an extremely high score of 4.91. Further, the level 3 score was significantly lower, with a mean of 1.48, making it indistinguishable from level 2's score.

4.3 Behaviour types analysis

The answer options of the four behaviour scenarios map to the four behaviour types. The participants were asked to rank the options in the order they would consider enacting them themselves as well as assign a severity score on a Likert scale to each answer option and to the scenario in general. The statistical analysis that follows is similar to the analysis of maturity levels described above. Again, we show only one analysis table here for brevity (table 1b).

The last line of table 1b shows the analysis of behaviour types for the organisation as a whole. The ranking of the behaviour types is Hierarchist (2.80 mean ranking), Individualist (2.68), Fatalist (2.50) and Egalitarian (2.02). All pairwise differences are statistically significant (see table 1b). The ranking of the severity of the options for each to the behaviour types is less clear as they can only be divided into 3 statistically distinguishable categories (as indicated by the use of three shades of colour only), although the Egalitarian option is seen as statistically significant most severe at 3.76. Further, there is a statistically significant negative correlation between severity score and behaviour type, implying that the employees rank less severe options higher, as may be expected.

It should be noted that there is no inherent ordering between the behaviour types (as it was the case between maturity levels), hence when analysing the data and table 1b, care has to be taken not to infer a ranking of the types themselves relative to each other, but rather work with the ranking of the types by the participants.

While at the level of the whole organisation there is a statistically significant ordering of the preferences of the behaviour types, this changes considerably when analysing across different subgroups as discussed in sections 4.3.1 to 4.3.3, where there are in many cases only 2 statistically different groups.

4.3.1 Behaviour types by division

In the *Business* division the Egalitarian and Hierarchist are ranked statistically significantly higher and lower, respectively. This is also the case in *Finance & Prof. Services*, but foremost the Individualist type is ranked highest here. *Marketing & Consumer* also agrees on the Egalitarian and Hierarchist differences, but here the Fatalist option is statistically significantly lower ranked than in the organisation as a whole. *Operations* are by far the most Fatalist: they rank this option statistically significantly highest and Egalitarian lowest, and are also much less Egalitarian and Hierarchist than the organisation generally. The *Sales & Services* team agree with the organisational ranking of the types, but they gave a significantly higher score to the Hierarchist option than any other division by at least 0.64.

The *Human resources* division represents the first Hotspot (see boxout 4), as the division shows a non-negative correlation between the option's severity score and rank. This implies that employees choose which option to prefer independent of the severity they assign to that option.

Analysing the severity scores, the *Operations* division is alone in perceiving a full ordering of the options, ranking the fatalist score third most severe and statistically significantly much less severe than the rest of the organisation. This is in stark disagreement with *Sales & Services*, who perceive the Fatalist option much more severely with a ranking difference of 1.44.

4.3.2 Behaviour types by age

There are no statistically significant variations between the different age groups for the Individualist and Egalitarian behaviour types. All the differences occur when considering the Hierarchist and Fatalist types: both the age groups 25 – 29 and 30 – 34 are statistically significantly more Hierarchist than all other age groups. The age group 50 – 54 shows the opposite, they are significantly less Hierarchist. When examining the Fatalist type, the picture changes: The 30 – 34 group is significantly less Fatalist, the 50 – 54 and the 55+ are significantly more so. In fact the 50 – 54 group rank Fatalist highest, followed by a statistically significant difference by the Hierarchist – an opposite order to the organisation at whole and unique to this group. It is interesting to note that the middle three age groups from 35 to 49 (as well as the under 25 group) have little or no preference between the behaviour types and also rank them nearly equally on the severity scales.

Between the different age groups there are no statistically significant variations of the severity scores for any of the behaviour types.

4.3.3 Behaviour types by location

The predominant behaviour types vary widely by business location. Both locations 1: *HQ* and *Homeworker* rank the Individualist options highest, in the case of 1: *HQ* because it ranks the Individualist and Hierarchist types statistically significantly higher and lower than the other locations, respectively.

Locations 5 and *Minor Offices* rank Fatalist first; this is followed by a statistically significant lower score for the Hierarchist type at these locations. Locations 2, 3, 4 and *Other* show an opposing trend, ranking the Hierarchist type higher than other locations and the Fatalist type lower. It is worth noting that the *Other* category represents mostly retail workers spread across the company's various sites.

Interestingly, there are also a large number of statistically significant variations in the severity scores, with all four types rejecting the null-hypothesis of equal distribution of the Kruskal-Wallis test. This is also reflected in strong variations in the severity score of the behaviour scenarios across the locations. Employees at location 3 saw all four options as significantly more severe, increasing the severity scores of each option by over 20%, but the scenario's severity score remains unchanged. The opposite effect is portrayed by *Homeworkers*, who rate the scenarios 0.51 more severe than the average, but show no variations for any of the behaviour type severity scores.

There is also a second hotspot present in this comparison: location 2 shows no statistically significantly negative correlation between severity scores and rank, implying that employees at this location choose which option to take independent of the severity they assign to the option.

5. DISCUSSION

Our research applied a scenario-based survey to assess both security maturity levels and self-reported security behaviours, and employee understanding of how risky certain behaviours are. A statistical analysis of the results of the survey conducted at the company allows us to draw several key conclusions regarding the security culture within the organisation. In line with the existing literature, we found that assessing attitudes provides a solid approach to understanding how employees interact with security policy. However, our scenario-based survey approach allows us to go further and detect intra-population differences within the organisation, showing that there are significant differences between different employee groups in how they respond to security-related challenges in the workplace. The salient outcomes are discussed below.

Our analysis of the survey has shown that the organisation in general has a very positive security posture: the majority of employees are at maturity level 5 and there is a downwards gradient of the ranking of the lower maturity levels. This combines well with a founded understanding of the acceptability and severity of the options presented to the employees of the organisation; employees in general choose what are in their opinion the more acceptable and less severe options. This strength is based on the willingness of the majority of the organisation's employees to engage actively with security. The predominant attitude within the company is to adopt good security practices, even when not specifically required to by technology or policy prescriptions. Many members of the organisation reported that they would challenge any breaches of policy they observe in their environment, with older employees being less likely to do so. Where friction exists between the business and security processes, employees take a predominantly Individualist approach to conflict resolution, meaning they rely on their own skills and knowledge. This echoes the results of Rhee et al. [25] and Siponen et al. [27] who both recognise the role of self-efficacy in decision making. Individually-derived approaches to security, driven by personal perception of what constitutes secure practice, can also manifest when policy and support is not known or visible to the individual [17].

The ranking of the behaviour types is also positive, but the differences in their respective rankings are weaker. Hierarchists are unlikely to challenge the existing structures, and while they may follow security policies to the letter, the Individualist that innovates may identify and solve new challenges before they become problematic [17]. Some CISOs might think that it is desirable if all employees were Hierarchists, but it could be argued that it would be counter-productive for an organisation to be exclusively one behaviour type, as there are many benefits in diversity. From a productivity point of view, the organisation requires diversity and even from a security point of view, variation has benefits. A diverse mix of behaviour types may even be essential, as security issues are embedded in, and deeply influenced by, social context such as corporate and national

culture [21]. In this sense, these issues have to be understood and addressed before any successful intervention program can be introduced.

Based on the results presented in the previous section 4, we will now discuss a number of areas of the organisation that are of particular interest. These areas hint at where interventions could be focused, or otherwise lessons learned and further studies conducted.

5.1 Targeting interventions

The *Sales and Service* division stands out by having significantly stronger maturity levels. They are also able to accurately assess the severity of the acceptability of the options presented. This is further accentuated by the extremely high rank of the Hierarchist type at the Sales and Services division. This alludes to a highly security competent division that is comfortable in its organisation structures. It should be an exemplary part of the organisation that should be able to provide a benchmark for the rest of the organisation.

Conversely, the *Finance and Professional* division rank maturity level 4 highest and further, this division ranks the Individualist type first. It is an interesting case of a combination of less mature security combined with an Individualist approach to security: interventions could focus here first, as this combination has the potential to create problems in the future. The *Operations* division is most Fatalist, and assesses the options as much less severe than all other employees, while maintaining a high maturity level. This suggests that many Operations employees may have given up trying to achieve their tasks using the organisational structures and policies, and instead attempt to fulfil their business goals as easily as possible. Their classification of the Fatalist options as much less severe implies that there are no negative effects of sidestepping the organisational structures. This division represents the 'disillusioned' section of the organisation. Their security maturity is in line with the organisation as a whole, but they feel that the organisation has ignored their needs. They are a primary target for engagement and it is paramount to find ways to make security fit better into their work.

The *human resources* division turned out to be a hotspot (section 4.3.1) that represents an interesting example of this organisation's security structure. While the employees choose highly mature options that were also most acceptable (with a very strong positive correlation), their choice of behaviour type is independent of the severity of each option. This may seem contradictory at first, but could stem from a diverse set of willful employees who are equally present in all four behaviour types and stand to their decisions. It may be argued that this is in fact a desirable property in a human resources division.

There are interesting variations between the different age groups of the employees. The young (25-34) are more Hierarchist, whereas older employees (50+) are more Fatalist. The middle age groups are split between all behaviour types. This could be interpreted as indicating that younger employees see the benefits of the organisation's structures and support, and might rely on them due to their lack of experience. Most younger people in the company are in the *Sales & Services* division where they experience fraud more directly. At the same time, older employees have diversi-

fied their beliefs and the oldest “have seen it all” and might have become disillusioned with their lack of influence and progress in the organisation. This is irrespective of their maturity ranking, as the differences in maturity ranking are only minor between age groups. In order to ensure that this trend is not repeated in the future, the organisation needs to take particular care to ensure that the voices of their young employees are heard and respected, especially as the older and therefore usually more respected employees portray a more challenging behaviour type that may often choose to ignore regulations. If breaking the rules is seen as a sign of seniority, it is toxic since older employees should be role models. Compliance should not be something that is only for those lower in the company’s structure.

There are clear cultural differences between the business locations. This manifests in strong variations in maturity levels as well as maturity types. The HQ uniquely ranks maturity level 4 higher than level 5. Separately, it also ranks the Individualist option first amongst the behaviour types. There is a strong absence of Hierarchists at this location. Interestingly, the organisation implemented a hot-desking policy here, which may explain the strong individualism that is present. While the low security maturity present at this location is suboptimal, the unique distribution of behaviour types may be positive and act as a foundation for involving all employees in security in diverse ways. For example, it is in the HQ where the organisation needs to constantly reinvent itself through use of new products and services, and a large number of Individualists may support this.

More worrisome for the organisation’s well-being are the employees at locations 5 and *minor offices*, who rank the Fatalist option highest. Further investigations may be required to find a solution to improve the distribution of behaviour types at these locations.

Lastly, location 2 represents a hotspot (section 4.3.3) that is similar to the *human resources* division mentioned above. Here, there is no correlation between an option’s rank and its severity score; that is, employees choose what option to take independent of how severe they perceive this option to be. This hints at an organisational site where the employees are well aware of the security impact of their options, but are inclined to choose the options that they have learned will work at their locations. While further investigations at this site may be a prudent course of action, interventions may be fruitful particularly as the Fatalist type is uncommon and maturity levels are above the organisation’s mean.

We were able to present our findings to the company at board level; as a result the security managers restructured security spending for the following year to target the locations, divisions and age groups we had identified as giving the most concern. Specifically, managers set targets to improve communication with these groups, and in particular to promote the need for leadership and to enable non-confrontational challenging of policies, amongst employees aged 25 to 45. Location 1 (HQ) was also targeted for special attention in this regard. This outcome showcases the real-world impact our methodology is capable of creating.

5.2 Limitations

We divide our discussion of limitations into those relating to methodology and findings. Each engagement with an or-

ganisation is time-consuming, involving interviews which are used to generate scenarios specific to the organisation. We envisage that the cost will decrease with further iterations of the methodology, but may present a high barrier of entry.

We would like to emphasise that from an organisational point of view however, employing our methodology is worthwhile because it creates a benchmarking tool that the organisation can use to re-evaluate and monitor over time to compare to previous iterations. As researchers working with many organisations, we envisage that the organisations where we conduct interviews yield a library of questions that we may be able to reuse for other organisations that are broadly similar. We also acknowledge that being able to benchmark a company’s security posture would feed the particular desire by some organisations to compare themselves with other organisations in the same sector. We would be hesitant to use our methodology to compare multiple organisations because it is difficult to obtain meaningful results as companies are complex and unique.

Acknowledging these demands, the authors are variously involved in efforts to simplify the extraction of meaningful results from interview data through additional tools (e.g., [7] and [8]).

These mappings from scenario option to maturity level and behaviour type need more extensive validation. As yet the mapping from scenario option to maturity level and behaviour type are not thoroughly validated; but we have attempted to make them match to [2] and section B as closely as possible. As the scenarios and options were created specifically for the organisation and the survey had to be conducted reasonably quickly after the interviews, a thorough validation was not possible. Validating these mappings is part of our ongoing work.

A large proportion of the respondents were from the *Sales & Services* division. While the statistical tests for table 1 have accounted for this, a large proportion of the *Sales & Services* staff worked at location *Other*, and fit in the younger age groups. These employees have more contact with customers and are more exposed to fraud and since they are younger, they tend to be more receptive to training. We were careful to make sure that this did not bias the results, but a perfect study may have sampled the organisation’s populations more carefully.

Our survey did not capture many contributing factors to the participants responses that may have helped to explain their answers. The respondents background (e.g., computer literacy, previous jobs, other relevant experience) would have provided hints at a number of other relationships worthwhile studying, and potentially allow us to tailor interventions even more specifically. We collected free-text responses at the end of the survey that we will analyse as part of future work, they might help us shed more light on employees’ reasoning and justification for their choices. Data collection does not stop once the intervention phase is reached – the methodology presented here supports decision-makers to identify broad employee categories and hotspots to target for improvements. A follow-up intervention may in itself involve data collection to identify contributing factors to particular behaviours.

6. CONCLUSION

The methodology presented here allows organisations to take steps towards empirically assessing the security culture, as well as gaining an understanding into the predominant behaviours and attitudes found within the organisation. We address the issue of scalability by deploying a scenario-based survey that employees can complete in 10-15 minutes but can therefore be deployed to a large enough fraction of the organisation to be representative. We ground all the scenario details, and answer options, in information gathered from a series of semi-structured interviews with employees of the organisation. We demonstrate that this approach allows us to detect statistically significant differences between employee groups that can inform targeted interventions. Business area, age, and geographical location all provide axis of differentiation. Giving an organisation an understanding of these details can potentially allow them to plan their future training, communication, awareness and policy making strategies more effectively. Enabling targeted interventions that focus on particular employee groups can save employees from both being involved in non-targeted interventions and needing to determine if they apply to them. Targeted interventions are then a good step towards reducing the draw on employees' compliance budget [6].

6.1 Future work

Tailoring our diagnostic tools to the operating context and working practices of the organisation provides meaningful results. Security awareness material can similarly be crafted to resonate with the experiences of employees in weaving security into their productive tasks. Tsohou et al. [32] discuss ways of interpreting cognitive and cultural biases – such as those described in the behaviour-type scenarios – to produce effective security awareness material. Awareness should be a two-way street: security specialists should use the understanding of what drives individuals' behaviour to engage with those individuals and be receptive and find collaborative solutions to conflicts between security and business processes.

Siponen and Vance [28] define conditions for field studies of policy violations – another avenue for future work could compare employee behaviour to the declared information security policy of an organisation. This will expose gaps in policy, and help to identify policies which are routinely ignored or misinterpreted, or communicated badly. For instance, Renaud and Gaucher [24] note a distinction between an intention to behave in a secure manner, and enacting a secure behaviour in practice – if an intention to comply is not supported by the infrastructure of the organisation the solution will not lie in the production of awareness material [3].

6.2 Acknowledgement

We would like to thank our partner company for making this research possible. Many thanks to Iacovos Kirlappos and many others for their help in data collection. Adam Beutement, Simon Parkin and Angela Sasse are supported by EPSRC and GCHQ, grant number: EP/K006517/1 (“Productive Security”). Ingolf Becker and Kat Krol are funded by EPSRC's grant to the Security Science Doctoral Training Centre, grant number: EP/G037264/1.

6.3 Supplementary material

The additional tables for section 4 (i.e., two more sets of tables similar to table 1 and the ipython notebook containing related statistical tests) can be accessed at <http://dx.doi.org/10.14324/000.ds.1496888>.

References

- [1] A. Adams and M. A. Sasse. Users are not the enemy. *Communications of the ACM*, 42(12):40–46, 1999.
- [2] J. Adams. Risk and morality: Three framing devices. *Risk and morality*:87–106, 2003.
- [3] E. Albrechtsen. A qualitative study of users' view on information security. *Computers & security*, 26(4):276–289, 2007.
- [4] E. Albrechtsen and J. Hovden. The information security digital divide between information security managers and users. *Computers & security*, 28(6):476–490, 2009.
- [5] A. Beutement, R. Coles, J. Griffin, C. Andronis, B. Monahan, D. Pym, M. A. Sasse and M. Wonham. Modelling the human and technological costs and benefits of USB memory stick security. *Managing information risk and the economics of security*:141–163, 2009.
- [6] A. Beutement, M. A. Sasse and M. Wonham. The compliance budget: managing security behaviour in organisations. In *New Security Paradigms Workshop (NSPW)*, 2008, pages 47–58.
- [7] I. Becker, S. Parkin and M. A. Sasse. Combining qualitative coding and sentiment analysis: deconstructing perceptions of usable security in organisations. In *Learning from Authoritative Security Experiment Results (LASER)*. IEEE, San Jose, California, US, 2016.
- [8] O. Beris, A. Beutement and M. A. Sasse. Employee rule breakers, excuse makers and security champions: mapping the risk perceptions and emotions that drive security behaviors. In *NSPW*. ACM, Twente, Netherlands, 2015.
- [9] J. M. Blythe, L. Coventry and L. Little. Unpacking security policy compliance: the motivators and barriers of employees' security behaviors. In *Symposium On Usable Privacy and Security (SOUPS)*. USENIX Association, Ottawa, 2015, pages 103–122.
- [10] V. Braun and V. Clarke. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2):77–101, 2006.
- [11] B. Bulgurcu, H. Cavusoglu and I. Benbasat. Information security policy compliance: an empirical study of rationality-based beliefs and information security awareness. *MIS quarterly*, 34(3):523–548, 2010.
- [12] R. E. Crossler, A. C. Johnston, P. B. Lowry, Q. Hu, M. Warkentin and R. Baskerville. Future directions for behavioral information security research. *Computers & security*, 32:90–101, 2013.
- [13] J. D'Arcy, T. Herath and M. K. Shoss. Understanding employee responses to stressful information security requirements: a coping perspective. *Journal of management information systems*, 31(2):285–318, 2014.

- [14] G. Dhillon and J. Backhouse. Current directions in is security research: towards socio-organizational perspectives. *Information systems journal*, 11(2):127–153, 2001.
- [15] K. H. Guo, Y. Yuan, N. P. Archer and C. E. Connelly. Understanding nonmalicious security violations in the workplace: a composite behavior model. *Journal of management information systems*, 28(2):203–236, 2011.
- [16] C. Herley. So long, and no thanks for the externalities: the rational rejection of security advice by users. In *New Security Paradigms Workshop (NSPW)*, 2009, pages 133–144.
- [17] I. Kirlappos, S. Parkin and M. A. Sasse. Shadow security as a tool for the learning organization. *ACM SIGCAS Computers and Society*, 45(1):29–37, 2015.
- [18] S. Pahlila, M. Siponen and A. Mahmood. Employees’ behavior towards IS security policy compliance. In *Annual Hawaii International Conference on System Sciences HICSS*. IEEE, 2007, 156b–156b.
- [19] K. Parsons, A. McCormac, M. Butavicius, M. Pattinson and C. Jerram. Determining employee awareness using the human aspects of information security questionnaire (HAIS-q). *Computers & security*, 42:165–176, 2014.
- [20] M. Paulk, B. Curtis, M. B. Chrissis and C. V. Weber. Capability Maturity ModelSM for Software. Version 1.1. CMU/SEI-93-TR-024. Technical report. 1993.
- [21] S. L. Pfleeger, M. A. Sasse and A. Furnham. From weakest link to security hero: transforming staff security behavior. *Jhsem*, 11(4):489–510, 2014.
- [22] J. T. Reason. *The human contribution: unsafe acts, accidents and heroic recoveries*. Ashgate Publishing, Ltd., 2008.
- [23] K. Renaud. Blaming noncompliance is too convenient: what really causes information breaches? *Security & privacy, IEEE*, 10(3):57–63, 2012.
- [24] K. Renaud and W. Goucher. 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*, pages 361–372. Springer, 2014.
- [25] H.-S. Rhee, C. Kim and Y. U. Ryu. Self-efficacy in information security: its influence on end users’ information security practice behavior. *Computers & security*, 28(8):816–826, 2009.
- [26] N. Röder, M. Wiesche, M. Schermann and H. Krcmar. Why managers tolerate workarounds—the role of information systems, 2014.
- [27] M. Siponen, S. Pahlila and A. Mahmood. Employees’ adherence to information security policies: an empirical study. In *New Approaches for Security, Privacy and Trust in Complex Environments*, pages 133–144. Springer, 2007.
- [28] M. Siponen and A. Vance. Neutralization: new insights into the problem of employee information systems security policy violations. *MIS quarterly*, 34(3):487, 2010.
- [29] J. M. Stanton, K. R. Stam, P. Mastrangelo and J. Jolton. Analysis of end user security behaviors. *Computers & security*, 24(2):124–133, 2005.
- [30] M. A. Tariq, J. Brynielsson and H. Artman. The security awareness paradox: A case study. In *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 2014, pages 704–711.
- [31] K.-L. Thomson and R. von Solms. Towards an information security competence maturity model. *Computer fraud & security*, 2006(5):11–15, 2006.
- [32] A. Tsohou, M. Karyda and S. Kokolakis. Analyzing the role of cognitive and cultural biases in the internalization of information security policies: recommendations for information security awareness programs. *Computers & security*, 52:128–141, 2015.
- [33] M. Workman, W. H. Bommer and D. Straub. Security lapses and the omission of information security measures: a threat control model and empirical test. *Computers in human behavior*, 24(6):2799–2816, 2008.

A. BEHAVIOUR AND ATTITUDE SCENARIOS

Below, we provide the texts of the behaviour and attitude scenarios used in the study. Labels are included to indicate which option related to which behaviour and attitude type, but these were not displayed to participants in the study and the order of the options was randomised as well.

A.1 Scenario C (Behaviour) – Password Manager

Hina, a member of the Operations division, has recently been required as part of her job to use a new piece of software about once a week. This requires her to log in to the service using a new username and password combination. Unfortunately the password manager does not work correctly with this new software and fails to store or enter her password. Because of the lack of support Hina is worried about being able to use the service as she struggles to remember infrequently used passwords.

Assuming that Hina decides to continue using the service without the support of the password manager, if you were Hina, what would you do in these circumstances?

- Individualist:** Store the password using a method of your own devising – you can be trusted to keep it safe.
- Egalitarian:** Share your password with a trusted member of your working group so that if you forget it they can remind you.
- Hierarchist:** Stop trying to remember the password and just use the password reset feature to generate a new password each time you need to use the service.
- Fatalist:** Re-use a password from another service that you have committed to memory.

A.2 Scenario D (Behaviour) – VPN

Robert, an analyst in the Operations team, has a set of logs from secure company hardware that he needs to upload to the manufacturer's website for analysis. He is working from home and unfortunately while connected to the VPN, he is unable to utilise the upload function on the manufacturer's site. It is necessary that the logs are analysed each day so he cannot wait until he is next in the office if he is to successfully complete this task.

Assuming that Robert decides to upload the logs via a different method, if you were Robert, what would you do under these circumstances?

- Individualist:** Make a local copy of the logs, disconnect from the VPN and upload the logs over your home connection.
- Egalitarian:** Give the password to the server to a trusted colleague not working from home and ask them to download the logs from the server before uploading to them to the manufacturer.
- Hierarchist:** Email the logs directly to the manufacturer's customer support email, and ask them to conduct the analysis and send the file back.
- Fatalist:** Email the logs to a colleague not working from home and see if they can upload the logs via a direct LAN connection.

A.3 Scenario F (Behaviour) – File Storage

Concerned about the safety of his current work, Shamal decides to back up his data, some of which is confidential. As he uses his own laptop under the 'bring your own device' scheme, he usually stores all his work on his drive on the central server but he wants to have a second copy just in case something happens or he loses connectivity to the company network. He thought about using one of the common drives but none of the ones he regularly uses have sufficient space.

- Individualist:** Create a local copy on the hard drive of your BYOD laptop, it is the only machine you work on so you know it will be safe and this ensures you will always have access to it if needed.
- Egalitarian:** Use a common drive that you used for an old project and still have access to, as your credentials were never revoked. It has enough space although you do not know who manages it now.
- Hierarchist:** Use an online service, such as Dropbox, to store the data as it is more under your control.
- Fatalist:** Back your work up onto a USB stick – you have ordered an encrypted one but while you wait for it to arrive you use a personal stick you have to hand.

A.4 Scenario H (Behaviour) – Credit Check

Karina works as a Sales Assistant in a company store. Her manager has asked her to increase her sales, in order to meet the store's monthly target. In her experience, customers can be put off by the need for credit and ID checks, and sometimes fail them altogether. She knows of a few unofficial ways of making the checks seem less of a problem, or to increase the chance of customers passing them.

- Individualist:** Attempt multiple credit checks in quick succession in order to try to figure out which details are causing the problem and amend them.
- Egalitarian:** Give information about the credit check to a few of your personal contacts so that they can prime potential customers on what they need to do to beat the system before referring them to the store.

- Hierarchist:** Use your employee discount to offer the customer a more attractive deal.
- Fatalist:** Give the benefit of the doubt when encountering IDs with indicators of possible fraud, such as dates of birth that do not seem to align with the apparent age of the customer, or addresses in different cities.

A.5 Scenario A (Attitude) – ID Badges

Jemima is a member of the Operations team working in Location 1. While sat working at her desk, she notices someone she doesn't recognise walk past without a visible ID badge. This prompts her to do one of the following:

- Level 2:** Nothing, the security badges are only used for accessing the building and once you are in serve no other real purpose.
- Level 3:** Nothing, although security badges are meant to be visible at all times it is a formality and it is the job of the security guards to check not hers.
- Level 4:** Make sure that her own ID badge is visible, seeing someone without theirs reminds her that she should have hers on display.
- Level 5:** Go and talk to the person and ask if they have a badge. If they have, remind them to have it on display, if not then politely escort them to security.

A.6 Scenario B (Attitude) – Clear Desk Policy

When leaving his desk to go for lunch with some colleagues Darren, a member of the HR team, notices that one of them has left his screen unlocked. The rest of the people he is with don't seem to have noticed, or seem to be OK with leaving it as it is. Darren got into the habit of locking his screen some years ago while working in a different company. As his colleagues start to walk away he decides to:

- Level 2:** Do nothing, there is no risk here as no-one could get into the office without passing through security. The screen locks are there just as a formality.
- Level 3:** Do nothing, the screen will automatically lock after a few minutes and this will keep things secure.
- Level 4:** Lock the screen himself.
- Level 5:** Quickly find out whose desk it is from the group and ask them to lock it before they leave for lunch.

A.7 Scenario E (Attitude) – Tailgating

Jessica is heading toward an access controlled entry door and notices a man she does not recognise gain entry by following close behind someone else who had tagged in at the door. The two men are walking close together although they do not appear to obviously be in conversation. The second man is holding a cup of coffee in one hand and his laptop in the other. His ID badge is not immediately visible. Jessica decides to:

- Level 2:** Return to her desk, she sees this sort of thing quite regularly and it is probably because his hands were full that he did not swipe through himself.
- Level 3:** Do nothing, if he is up to some mischief the security guards will catch him later on.
- Level 4:** Find a security guard at one of the manned turnstiles and tell them what happened.
- Level 5:** Follow the man and ask to see his ID badge.

A.8 Scenario G (Attitude) – Secure Disposal

John works as a Sales Advisor in a company store in London. During a busy period of the day he notices that a customer, served by one of his colleagues, has left their paperwork behind. John's colleague grabs the paperwork and throws it into a wastepaper bin under the desk. Seeing this John decides to:

- Level 2:** Carry on serving customers in the store, all the rubbish will be thrown out at the end of the day anyway so it is no big deal, and using the shredder in the back area, locked by a keypad, is inconvenient when the bin is right there.
- Level 3:** Make a note to check with his manager what the appropriate action would be, as it has been some time since he took the Data Protection training module and he cannot clearly remember the details.
- Level 4:** Go and grab the paperwork out of the bin when he has a spare moment and take it to the shredder in the back of the store.
- Level 5:** Go over immediately and ask his colleague to take the paperwork out of the bin and put it in the shredder, having documents lying around exposes both the store and the customer to the risk of identity theft.

B. MATURITY MODEL

This model expresses the maturity of the security culture within an organisation in terms of how aligned with the policy employee behaviour is, and also how integrated the policy is with the primary business process of the organisation. Most critically the model does not represent a checklist of required behaviours for employees, but aims to reinforce the synergy and co-operation required between employer and employee to deliver effective security. As such it is not possible to reach the highest levels of the model in an environment with an inefficient or poorly implemented policy that is in conflict with the primary process of the organisation. Thus the model is capable of guiding change both for the organisation and the individuals that work for it.

This model is based on the Carnegie Mellon Capability Maturity Model [20]. This model expresses the degree of formality associated with various processes. What we need from our security behaviour model is a characterisation of what represents effective employee security behaviour, as observed by the organisation. This will then act as a scale against which progress can be measured, as well as a tool for identifying the current state of security behaviour. The CMM consists of five levels, moving from unplanned/unmanaged through a managed state to one of optimisation through incremental innovation. These levels are listed below with definitions for reference.

Level 1 – Initial (Chaotic) It is characteristic of processes at this level that they are (typically) undocumented and in a state of dynamic change, tending to be driven in an ad hoc, uncontrolled and reactive manner by users or events. This provides a chaotic or unstable environment for the processes.

Level 2 – Repeatable It is characteristic of processes at this level that some processes are repeatable, possibly with consistent results. Process discipline is unlikely to be rigorous, but where it exists it may help to ensure that existing processes are maintained during times of stress.

Level 3 – Defined It is characteristic of processes at this level that there are sets of defined and documented standard processes established and subject to some degree of improvement over time. These standard processes are in place (i.e., they are the AS-IS processes) and used to establish consistency of process performance across the organization.

Level 4 – Managed It is characteristic of processes at this level that, using process metrics, management can effectively control the AS-IS process (e.g., for software development). In particular, management can identify ways to adjust and adapt the process to particular projects without measurable losses of quality or deviations from specifications. Process Capability is established from this level.

Level 5 – Optimizing It is a characteristic of processes at this level that the focus is on continually improving process performance through both incremental and innovative technological changes/improvements.

When considering a Security Behaviour version of this model we must consider how to convert these organisational indicators to indicators of personal behaviour. One approach is to consider how the individual is managing or motivating their own behaviour – what factors they are considering when planning their security actions. At the highest level, they will be actively working toward an improved and improving security culture. At the lower levels employees will be following the policy by rote (possibly reluctantly, ineffectively or incompletely) or simply taking actions as they see fit, based on their own internal security model with no input from the organisation. The following levels represent this range of behaviours.

Level 1 – Uninfluenced At this level, user behaviour is mediated only by their own knowledge, instincts, goals and tasks. Their actions will reflect only the needs of their primary task and will only deviate from that where their internal security schema conflicts with those actions. While some members of the organisation may be sufficiently knowledgeable to act securely it is expected that employees at this level will introduce a range of vulnerabilities in to the system. In practice this level can only exist where employees are working on non-organisational systems, as even the act of logging in to a managed network means that organisational security is exerting an influence.

Level 2 – Technically Controlled Employees at this level act as in level 1 except where technical controls exist that enforce policy on a case-by-case basis. Technically controlled employees will follow their own security rules except where they must use organisational systems in the execution of their primary task, and those systems enforce policy at the software or hardware level. Realistically, this is the lowest practical level that employees working in an office environment could function at.

Level 3 – Ad-hoc Knowledge and Application Employees at level 3 follow policy without necessarily a deep knowledge of what it contains. Their security knowledge comes from the ‘best practise’ or habits associated with their work environment, rather than from being aware of, and understanding, organisational policy.

Level 4 – Policy Compliant Level 4 behaviour demonstrates knowledge and understanding of the policy, and compliance with it, even in situations where the local work environment may include the use of workarounds and frequently made excuses. At Level 4, employees can be considered to be useful role models and guides for security culture within the organisation.

Level 5 – Active Approach to Security At Level 5, employees take an active role in the promotion and advancement of security culture within the organisation. They serve not just the letter of the policy but the intent as well and will challenge breaches at their level appropriately. They see security as a valuable part of the function of the organisation and have internalised this motivation. Level 5 employees are not security zealots, but rather understand the need to balance the security and business processes and champion the cause of security intelligently and effectively.

C. BASIC INTERVIEW QUESTIONS

C.1 Introductory questions

1. What do you do at the company?
2. How long have you been working at the company?
3. What does your usual day involve?

C.2 Security Awareness

1. How does security fit into your day?
2. Do you think your work has any security implications?
3. Do you encounter information that is in any sense confidential or sensitive?

C.3 Clear Desk Policy

1. Is there a policy that says what you should do with your desk when leaving in the evening?
2. Do you have a secure draw or storage area you can use?
3. Do you ever work on paper at all?

C.4 Laptops, Remote working and Removable Media

1. Do you ever use a laptop in the course of your work?
2. How do you share information with colleagues?
3. Do you ever use removable storage devices such as USB sticks?
4. When working from home what systems or technologies do you use?

C.5 Leadership and Management Roles

1. Do you supervise any other people?
2. Does your supervisor ever mention security issues to you?

C.6 Policies, Reporting and Training

1. How much would you say you know about the security policies at your company?
2. Have you ever received any security training?
3. Do you think people generally follow the policy rules?
4. Who would you report a security concern to?

C.7 Optional Topics

1. Compliance and security culture
2. Personal/mobile devices
3. Locking screens
4. Password behaviour
5. Password resets
6. Physical security
7. Customer data
8. Data classification
9. Trust