

Employee Rule Breakers, Excuse Makers and Security Champions: Mapping the risk perceptions and emotions that drive security behaviors

Odette Beris
University College London
Department of Computer Science
Malet Place, London, WC1E 6BT
Tel: +44 (0)20 7679 2788
odette.beris.12@ucl.ac.uk

Adam Beautement
University College London
Department of Computer Science
Malet Place, London, WC1E 6BT
Tel: +44 (0)20 7679 0353
a.beautement@cs.ucl.ac.uk

M. Angela Sasse
University College London
Department of Computer Science
Malet Place, London, WC1E 6BT
Tel: +44 (0)20 7679 7212
a.sasse@cs.ucl.ac.uk

ABSTRACT

We introduce a new methodology for identifying the factors that drive employee security behaviors in organizations, based on a well-known paradigm from psychology, the *Johari Window*. An analysis of 93 interviews with staff from 2 multinational organizations revealed that security behavior is driven by a combination of *risk understanding* and *emotional stance* towards security policy. Furthermore, we found that a quantitative analysis of these dimensions is capable of differentiating between the staff populations of the two organizations. Organization B showed a healthier set of security behaviors, as a result of its employees having better risk understanding and a more positive emotional stance. The framework distinguishes between 16 theoretical behavioral types, (3 of which are *rule breakers*, *excuse makers* and *security champions*). It can be used to identify groups of employees that potentially pose a risk to the organization, as well as those with beneficial skills and expertise. This allows highly specific messages to be targeted to change the risk perception and emotional stance of such groups. Assuming the organization has ensured *security hygiene* (i.e. its policies can be complied with in the context of productive activity), this can shift behavior towards compliance. Our framework thus offers diagnostic and intervention-shaping tools for the next step in improving security culture.

Categories and Subject Descriptors

General Terms

Risk Perception, Emotion, Information Security,

Keywords

Risk Perception, Risk Understanding, Affect, Emotion, Information Security, Security Policy.

1. INTRODUCTION

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

NSPW'15, September 8-11, 2015 Twente, The Netherlands
ISBN 978-1-4503-3743-0
Copyright 2015 ACM.

Employees who do not comply with security policies are a key risk for organizations trying to protect systems and data. Schneier is often quoted as saying people are “*the weakest link in the security chain*” [1]. Whilst this statement has some essence of truth, it is too often used as a justification for blaming people for being uneducated or lazy. However, this assumption does nothing to improve the situation; employees are an essential component of an organization’s security posture, and their behavior is a key resource that must be effectively managed effectively, just like any other resource [2]. As Pallas [3] states “*security management is all about human behavior*”. Even when organizations recognize that successful security management involves managing undesirable security behavior, what constitutes an effective response is less clear. They typically treat their staff as a homogenous group, rather than a collection of individuals with differing security attitudes, levels of knowledge, goals and tasks. However, pushing the same messages to all staff creates information overload and additional security tasks. Where security-related work distracts users from their main productive activity friction between business and security is created. People have a limited tolerance for friction and disruption. When this tolerance, their *Compliance Budget* [4] is exceeded, they will be tempted not to comply. Most non-compliance is not lazy or malicious, but occurs because security is seen as an onerous overhead and a distraction from an employee’s primary task or day-to-day role [4]. Employees may also circumvent security rules because the policy itself, and the associated technical mechanisms, are not fit for purpose in the business environment [5]. Recent research has identified a pattern of behavior referred to as ‘shadow security’ where employees create workarounds when ‘official’ security is too burdensome, yet are still security-conscious and take other measures to protect against the risks they understand. Management is often unaware that employees are operating in this fashion [6].

While compliance with policy is no guarantee of security, in many cases a failure to comply, and the associated workarounds that replace sanctioned processes, creates new vulnerabilities. We agree that compliance is desirable, but trying to enforce policies and mechanisms that are unworkable in the context to which they are deployed is futile. Organizations must perform essential *security hygiene*, a process of identifying and re-designing high-friction security [7]. Security hygiene is a necessary, but not sufficient condition for compliance: staff may still be tempted to cut corners where they perceive risks as negligible, or think the organization does not ‘deserve’ their contribution to security. It is this ‘next layer’ of influencing security behavior our paper targets.

Security managers typically only consider lack of knowledge - specifically not appreciating the severity of a risk as a driver of

security behavior. Thus, their current efforts in ‘security education’ consist of repeating all policies and rules to everyone. This is the equivalent of shouting louder at someone who does not understand your language; we need a smarter, targeted approach if we want to meaningfully change behavior towards manageable compliance.

For that, we must consider the factors that shape employee perceptions of risk and security, and recognize that these may suggest behavioral types. Being able to systematically identify and categorize meaningful heterogeneous characteristics within an employee population represents a critical first step. Measures of behavioral types in the social sciences focus on aspects of personality, and do not consider knowledge or expertise. None of these factors in isolation are sufficient, but combining them can facilitate effective targeting. In particular the perceived *cost* of security compliance does lead to emotional, or *affective* responses to security requirements amongst employees. These emotional responses consciously or unconsciously shape employees’ general attitude towards security, and their risk perception. Risk perception is also based on an individual’s skill at assessing risk, backed by the relevant information or knowledge they may have. Thus, security behavior results from 1) an individual’s affective responses to security, and 2) their competence in assessing risk. Organizations with a healthy security culture are likely to have high levels of risk understanding, combined with positive emotion towards security.

The specific aim of this research is primarily practical in nature; our goal is to provide organizations with a means of capturing the dynamic between staff affect and risk understanding, and categorizing the resulting behavioral types. This will allow organizations to monitor the ‘health’ of their security culture and identify areas of both strength and weaknesses. The information gathered can be then used to target risk security communication messages, security training and/or information security policies more adaptively and to greater effect. In this paper we suggest a framework for assessing types of employee security behavior based on their emotional response to security, the background for which we review in Section 2, and their degree of risk understanding. Our framework is based on a revised version of the Johari Window [8], which we review in Section 3. An empirical basis for this work is provided by an analysis of 93 interview transcripts taken from two separate organizations. In Section 4 we test our approach on this corpus of interviews, seeking to identify differences between them which a security manager should take in to account. Our findings, and a discussion thereon, are presented in Sections 5, 6 and 7.

2. BACKGROUND

Existing psychological research explicitly links emotion with risk perception, suggesting that emotion plays a fundamental role in influencing perceptions of risk [9]. Slovic argues that individuals’ positive and negative feelings about an event, referred to as *affect*, are often experienced unconsciously by the individual yet play an important role in driving their risk perception and assessment and subsequent decision-making [10].

The affect heuristic is also applied to both conscious and subconscious modes of thinking. Kahneman suggests that we are likely to default to automatic and intuitive processing in risk assessments particularly under pressure, referred to as System 1, rather than a more analytical approach, referred to as System 2 [11]. Security author Bruce Schneier, has also highlighted the importance of these two modes of thinking in relation to how individuals may

assess security risk [12]. These systems of thought are as applicable to the information security context as anywhere else; individuals who prefer working in a non-compliant fashion may discount any potential security risks because a given security task requires too much effort [13]. Consequently, the “risk as feelings” hypothesis which emphasizes the impact of *emotional* rather than cognitive responses to judging perceived risks underpins our research approach [14]. It is also useful to mention the “affect as information” hypothesis [15] which posits that feelings *directly* impact decision-making under risk. The affect-as-information hypothesis predicts that feelings during the decision-making process influences individuals’ choice of behavior [15], highlighting the importance of emotion in shaping risk assessments.

In relation to risk perception studies within the information security literature, Farahmand et al., developed a model based on the consequences and understanding of security risks [16] which attempts to integrate the affective components associated with risk assessment, reflecting the link in literature between risk decision-making and emotion. We go a step further and consider emotion and risk understanding as separate dimensions.

In particular, this work is inspired by an existing framework, the Johari Window [8]. In its original form [8] the Johari Window is a psychological framework used to facilitate a better understanding of an individual’s relationship with themselves and others. It takes the form of a 2 x 2 grid which expresses four states of awareness, combining what is known and not known by the self and what is known or not known by others. The Johari Window framework has been widely used in conceptualising risk in other domains such as space exploration for instance. Massie and Morris’ risk model [17] builds on the Johari Window to explore how known and unknown information influences decision-making under conditions of risk.

Of specific relevance to the development of our framework, the BSG, are the four states of awareness incorporated into the Johari Window which are referred to as: Open, Blind, Hidden and Unknown. Briefly, the Open area refers to what is known by both the self and others, the Blind area refers to what others know about the person but they are not aware of themselves, the Hidden area refers to what the person knows about themselves but others are not aware of and finally the Unknown area refers to what is not known by self and others. We considered that the quadrants of the Johari Window, Open, Blind, Hidden and Unknown offered us a basic heuristic to express the employee’s style or mode of security behaviour.

Given that this work aims to better understand the relationship between individuals and organizational security policy, the quadrants of the Johari Window provides a useful framework to represent differences in security behavior. However, though the Johari quadrants enabled us to express different modes of employee security behaviour, we had to discard the Johari Window axes relating to the self and others, since this did not fit our model. Figure 1 presents a revised version of the Johari Window [8], our prototype psychological framework, referred to as the Behavioral Security Grid (BSG) for the purposes of this research. This is discussed further in Section 3.

3. METHODOLOGY

To provide an empirical grounding for our framework we analyzed 93 semi-structured interviews. These were randomly sampled from a

larger corpus of interview data collected by researchers at UCL as part of a study looking at security behavior in organizations. While these interviews have been previously used as the basis for published work [18][6] they were re-analyzed completely for this work and no prior coding was used. The interviews were undertaken at 2 major multinational corporations; Company A, a utility company, and Company B, a telecommunications organization. The participants represent a vertical cross-section of employees of both companies, consisting of varying seniority and departments, including IT and Operations. Confidentiality and anonymity were assured, in line with the principles of ethical research and in order to promote an open and frank discussion. As the interviews were not conducted expressly for this study the questions focused more generally on security attitudes and behaviors. As a result any statements made during the interview that express either risk understanding or an affect toward security formed a natural part of the conversation, rather than being directly elicited. This assists in avoiding biases which may otherwise exist, ensuring the data set more closely represents the real-world views of the participants.

Of the 93 interviews 48 were conducted at Company A and 45 and Company B. Each interview was independently coded by two researchers and the analysis was split into two stages. The first stage of qualitative analysis used Applied Thematic Analysis [19] where the authors' coded inductively, working independently to identify themes related to security behavior. This process generated a large list of codes, many of which were duplicated between authors, or represented very similar concepts. This led to the second stage of the analysis where the authors worked collaboratively to reduce the codebook to a more usable size by deleting duplications and grouping similar codes together. This collaborative approach allowed the authors to resolve any disagreements in code types of application. The result was a list of less than 20 core code families that could be used to delineate different aspects and levels of the relevant concepts. Most significantly, based on this first stage of coding, we identified two major themes relating to employee risk understanding and their emotional responses to security. These themes around affect and risk perception added additional insight in to understanding how employees' perceived their own security behavior. As such, emotion and risk themes formed the basis of our methodology.

It was at this point that we also recognized the potential value of the Johari Window [8] as a framework to better understand security behavior. From the initial coding process, we observed that security behavior seemed to be broadly consistent with the Johari Window quadrants (Open, Hidden, Blind and Unknown) allowing organizations a simple heuristic to classify behavior. For instance, staff that fit in to the 'Open' quadrant reported security behavior that was openly aligned with security policy and security-related tasks. They demonstrated an understanding of security risks and held a generally positive view about security, based on a mutual understanding of security requirements. Others were behaving inconsistently with the policy because primarily they were not aware of the risks (the Blind quadrant), or else because they understood the risks but were negative about the security provision in their organization (Hidden). As previously mentioned, hidden security behavior has been referred to as 'shadow security' [6] where employees may engage in circumventing security policy in order to achieve their primary task. The unknown quadrant is where the employees believe organizational security is poor, while being unaware of many risks themselves. As such is likely to be sparsely populated, unless risks that were unknown to all parties are identified

retrospectively.

We were also interested in the Johari Window's use of two dimensions to inform resulting behaviour – in particular that categorisation in to one of the quadrants is based on a comparison between how the individual perceives themselves and knowledge contained within the environment. Having identified emotion (something personal to the individual) and risk (an aspect of the environment) as two key dimensions in the interviews the Johari Window offered a natural synergy that we set out to explore.

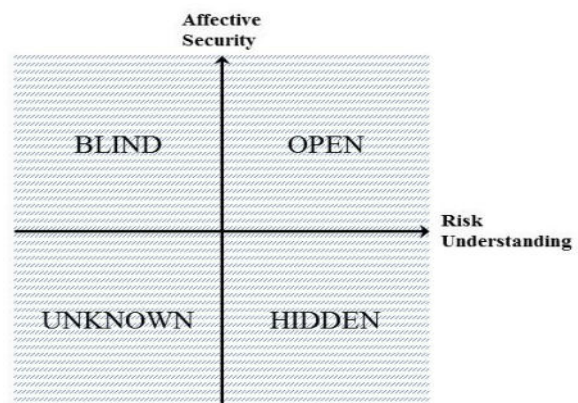
3.1 Revising the Johari Window

While the Johari Window [8] offers us a starting point for categorizing members of a population it has no capacity to offer a more granular analysis beyond its four quadrants. Additionally, the spatial relationship between the quadrants lacks significant meaning and there is no way of organizing individuals within the quadrants themselves. What does it mean to be more or less 'Open'? How can we determine this and what are the implications? Johari in its current form cannot tell us. In order to address these limitations we assigned a set of axes to the Johari framework that would allow us to create meaningful spatial relationships. First and foremost this meant reorienting the Johari Window such that the 'Open' quadrant now occupied the upper right, as would seem more intuitive under a typical Cartesian topology (see figure 1).

Having recognized that end user behavior is the resultant of emotion or affect (see section 2), and competence, as described by their ability to recognize security risks, we utilize these two concepts as the basis of our axes. The emotional dimension we label as '*Affective Security*' (AS). AS deals with the individual's emotional response to security, as represented by the organization's security policy. It is worth noting at this point that for the purposes of this paper we do not consider the quality of an organization's security policy, assuming that the policy is an effective one, thus making 'security' and 'security policy' interchangeable. AS is assigned to the y-axis.

The dimension of competence we label as '*Risk Understanding*' (RU). RU denotes the individual's ability to accurately perceive the existence and severity of the risks associated with the actions they take themselves, as well as those they observe in the surrounding environment. RU is assigned to the x-axis. The application of these axes, along with the re-orientation of the window, results in the Behavioral Security Grid (BSG), as seen in figure 1.

Figure 1: *The Behavioral Security Grid*



After developing the BSG framework we moved on to the second stage of our analysis, using the revised Johari Window to categorize members of two different organizations in order to identify differences between their populations. As we now wished to code for AS and RU specifically, rather than looking for general themes, we adjusted our coding strategy. Now we reanalyzed the transcripts, only coding statements that contained either affective content, or an indication of risk understanding. This required us to first create definitions for AS and RU.

3.2 Affective Security

In order to meet our goal of being able to spatially locate individuals within the BSG we need to be able to position individuals along the AS. As such we considered the characteristics of both strong and weak positive AS, and strong and weak negative AS, and the sort of phrases and terminology they might use. These are discussed below:

- **Strong Positive (AS++):** these individuals regard security as their personal business and responsibility. In addition they feel that the organization has effectively designed and implemented its security strategy. They may also act as leaders and have the capacity to positively influence those around them. We categorized statements as AS++ if they contain a clear indication that the individual personally takes action to comply with, or support, the security policy of the organization, such as adopting practices aligned with the policy, or challenging non-compliant practices they observe in their environment. Statements containing strong positive language, such as *“this was very important,”* or *“security in this case was essential,”* even when not linked to an instance of individual behavior, were also included.

- **Weak Positive (AS+):** Individuals in this category show a positive inclination toward security and therefore statements reflecting a reasonably, but not strongly, positive stance will be coded weak positive. In some instances, while they express a view that organizational policy is useful, they do not necessarily see it as their personal responsibility. They appreciate the need for security in a general sense but are less likely to take personal initiative to ensure security. We coded statements as AS+ where they expressed a favorable view of security but did not report action taken personally by the participant such as, *“We have shared drives for a reason. I saw a lot of that, so I think people are getting better”*. Another exemplar from the transcripts describes how the interviewee noticed a colleague challenging someone about their security pass: *“people are aware,...I saw it once that, um, the guy in front of me actually stopped another person and said, ‘Hey, you have to use your badge.’”* This demonstrates that the interviewee recognised good security awareness and practice but didn’t necessarily adopt those behaviours personally. Thus, coding for the weak positive category included statements made about peers, whether supporting their positive security practices or criticizing negative ones.

- **Weak Negative (AS-):** Individuals making these statements think security processes are useful to the organization in the abstract, but when it comes to applying personal effort to the task they frequently make excuses. These typically take the form of saying security tasks take up too much time, or effort, because organizational policy is not as effective as it could be. Statements were coded as AS- where security was referred to as a hindrance or burden, although not something that should be circumvented if possible, for example, *“I’ve had some problems where I need to get*

files from external companies”.

- **Strong Negative (AS--):** Individuals making strongly negative statements are typically highly frustrated by the current security policy and seek to implement ad hoc workarounds that minimize their involvement with it. By taking direct action on their own behalf they may also set unwanted precedents for others (particular those falling in the weak negative category). Statements were coded as AS-- where they contained examples of intentionally circumventing the policy, such as, *“everybody I know either has a cheat sheet or something written down, all their passwords, because there’s just too much”*. As in some cases a sufficiently comprehensive security implementation prevents the possibility of workarounds we also include statements that expressed a desire to circumvent, even if it was not actually feasible to do so. Statements that contained strong negative language regarding security were also included.

During the coding process each statement with emotional content relating to security made by the interview participants were assigned one of these codes.

3.3 Risk Understanding

As with AS, in order to position individuals on this axis we considered the characteristics of strong and weak, and positive and negative aspects of RU. These are discussed below:

- **Strong Positive (RU++):** Individuals in this category display a comprehensive understanding of risk factors, including the ability to understand the causal relationship between their actions, risk, and any associated outcomes. Statements were coded as RU++ where they showed that the participant understood not only that a risk exists, but what causes the risk and the impacts associated with it being realized, such as, *“I think the biggest concern around [a system] is really the malicious attacker who is looking to get full access all the way into the operation center. You’ve now potentially provided a direct path”*.

- **Weak Positive (RU+):** Here the existence of risks is recognized but individuals are less clear about what causes them, or do not demonstrate an understanding of the relationship between their actions and the risk (or its mitigation). Statements were coded as RU+ where risks are correctly identified, either explicitly or implicitly, but no further discussion is offered as to their causes or impacts, for example, *“Even though we’re using it less and less now, you have to print out stuff and it just sits there and, you know, it’s not safe”*.

- **Weak Negative (RU-):** Individuals in this category are characterized by omissions in their ability to recognize risk. While what knowledge they do claim to have may be accurate, it appears incomplete, leading them to make errors in judgment, or be uncertain as to how to proceed in a given situation. Statements were coded as RU- where the participant does not mention common risks when discussing courses of action, or demonstrates uncertainty. Unusually, this required an element of researcher discretion and understanding, as we used our own knowledge of risks to identify when participants were not demonstrating the expected level of risk awareness. We assume that if no risk is mentioned that the participant does not currently recognize that the risk exists, although they may do when prompted on the topic. Statements expressing a lack of knowledge on the part of the participant, such as *“I don’t know what the clear desk policy is,”* were also included here.

- **Strong Negative (RU--):** Individuals in this category actively

hold misconceptions about risk, they do not just fail to mention that they exist but make statements that are incorrect. RU -- statements are exemplified by individuals that believe they are right while making significant mistakes. Statements such as, “*I guess it is not too serious if it leaks out,*” where participants dismissed a known risk as unimportant, or not present, were assigned this code.

To ensure qualitative coding consistency, code review meetings were conducted with other members of the research team to discuss the various definitions used for RU and AS. Codes were then applied to the interview transcripts using the ATLAS.ti software package.

3.4 Code Tallies

To begin locating individuals within the BSG the codes assigned to each participant needed to be condensed into a single RU and AS score that reflected the security behavior of that participant. We adopted a relatively simple method of obtaining the scores for each participant; the codes relating to the AS and RU axes were summed, with strong statements being worth twice as much as weak statements. Therefore:

$$\text{AS or RU} = 2(\text{strong positive}) + (\text{weak positive}) - (\text{weak negative}) - 2(\text{strong negative})$$

We decided not to normalize the AS and RU scores for each individual in order to preserve the differences in the frequency of codes between interviews. While normalizing the score would provide a measure of how strong each comment was on average, we took the view that code frequency was a salient factor in the analysis. Our assumption here is that participants with stronger affective security and risk understanding are more likely to have risk and security at the front of their mind, and therefore bring them up in conversation, generating a higher number of coded statements during their interview. The higher scores resulting simply from raw frequency of statements are therefore a reflection of this more active security-related mind set.

A weakness with this approach is that due to the absence of a ‘neutral’ code, the use of which we considered but ultimately decided not to use, participants making many strong statements evenly split between positive and negative will end up with a similar score to a participant making a few weak statements. We recognize that the assumptions we have made here may need to be revisited in the future in order to improve the sensitivity of the measure to cases like that described above.

4. RESULTS

Using the above methodology 1874 codes were applied across the 93 interviews. The distribution of the codes is shown in tables 1 and 2.

Table 1: Affective Security code distribution

	Company A	Company B
AS++	113	342
AS+	147	181
AS	121	136
AS--	16	32
Total	397	691

Table 2: Risk Understanding code distributions

	Company A	Company B
RU++	119	249
RU+	100	162
RU-	64	72
RU--	10	10
Total	293	449

Using the code tallying method outlined above a pair of AS and RU scores were generated for each participant, allowing them to be plotted on our axes.

Figure 2: Scatter plot of AS and RU for Company A

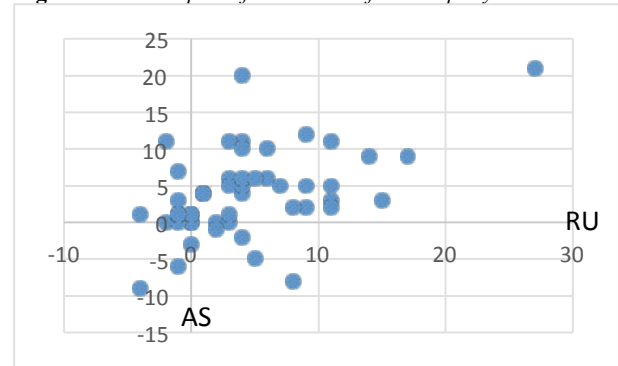


Figure 3: Scatter plot of AS and RU for Company B

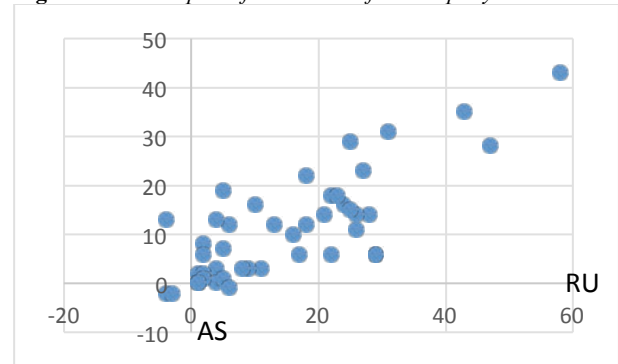
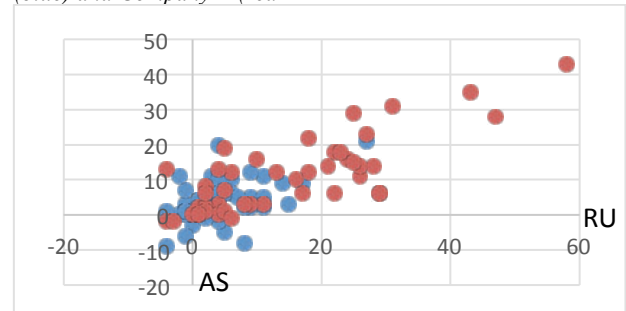


Figure 4: Combined scatter plot of AS and RU for Company A (blue) and Company B (red)



In order to verify that a significant difference exists between companies A and B an independent samples Mann-Whitney U test was conducted to compare the strong positive, weak positive, strong negative and weak negative means of both AS and RU between companies A (n = 48) and B (n = 45). 4 significant differences were identified between the two groups:

For AS++ a significant effect of company was found (Company A and B means were 35.53 and 59.77 respectively; U = 516, P < 0.05)

For RU++ a significant effect of company was found (Company A and B means were 37.63 and 57.43 respectively; U = 619, P < 0.05).

For RU+ a significant effect of company was found (Company A and B means were 41.24 and 53.41 respectively; U = 796, P < 0.05).

For RU-- a significant effect of company was found (Company A and B means were 42.06 and 52.50 respectively; U = 836, P < 0.05).

5. DISCUSSION

Our method has identified a clear difference between companies A and B. The frequency of coding is higher for Company B, indicating that discussion of risk and emotive responses to security were more prevalent during the semi-structured interviews conducted at this organization. Employees of both organizations are predominantly found in the 'Open' quadrant, indicating that their cultures are on the whole positive about security. However, by separating AS and RU we are able to meaningfully compare how 'Open' the organizations were. Company B employees scored significantly higher in the positive categories of AS++, RU++ and RU+. This indicates an overall higher level of risk understanding, as well as an attitude toward security more closely aligned with organizational policy. Company B contains several individuals that scored unusually high in both RU and AS, suggesting that they may represent security leaders that drive the culture within the organization. Company A lacks such standout individuals, and this represents the key difference between the two organizations. Company B did also score more highly in RU-- but it is likely this is a product of the small number of codes used (10 for each organization).

Our methodology has allowed us to meaningfully compare the two populations, but as yet does not afford us the power to make an objective assessment of a single organization as the range of the axes is determined by the organizations being assessed, rather than against a standardized scale. In part this is due to the size of the data set available, and the speed of the assessment. While interviews provide a rich source of anecdotal data they are not scalable to the size of a large organization. This yields a data set that is informative but lacks sufficient numbers to show clusters within the population, or to truly represent the organization as a whole. As such, our future work will be to develop a scalable metric that is capable of assessing a much larger population with the same time and effort investment. With the collection of multiple large data sets, a standardized axes will also be produced.

That said the interview data does provide us with examples of the sorts of behavioral types that a larger scale study is likely to identify. The strength of this approach derives from the use of the two dimensions of measurement, AS and RU, which form the axes of the grid. Having established that these are a valid means with which to differentiate between populations, we can extrapolate with some confidence from our earlier weak and positive definitions to consider the characteristics of each quadrant. We suggest the following 16 categories of individual, summarized in Figure 3 and discussed below, that are may potentially be identified within employee

populations. Supporting quotations from the interviews are used throughout.

While our axes afford us the power to reason about the areas they encompass we do not suggest that all areas will be equally represented in any given population. Indeed we expect that clusters will form in a few common areas with, in particular, the 'Unknown' quadrant (bottom left) being sparsely populated. This is borne out by the application of a crude clustering approach to our data set. In order to subdivide the quadrants we considered the positive and negative scores for AS and RU separately and took the mean of each group. This gave us 4 values, which were then assigned to the interim lines, as shown in Figure 2. Individuals were then grouped into the types according to their RU and AS scores. For example, an individual with a positive AS score above the positive AS score mean, and a positive RU score above the positive RU scores mean, would be sorted in to the 'Champion' category, whereas an individual with a positive AS score above the mean but a positive RU score *below* the mean would be sorted in to the 'Follower' category. The distribution of participants is also shown in Figure 2. This is again a relative, rather than absolute, method of categorization, meaning that individuals scoring as Abdicators are doing so when compared to others in their organizations, rather than the wider range of possible behaviors. It should be noted that the following security types are purely theoretical extensions of the grid at this early stage of the work. It is intended in future work that they will be subject to further testing and validation.

Figure 5: Categorization of types within the BSG

$\mu(RU-ve)$		AS	$\mu(RU+ve)$		
Gung-Ho [1]	Uncertain [2]		Willing [10]	Champion [24]	$(\sigma_A + \sigma_V)/n$
Naïve [0]	Passive [3]		Follower [35]	Expert [9]	
Reckless [0]	Apathetic [0]		Excuse Maker [2]	Circumventer [0]	$(\sigma_A - \sigma_V)/n$
Abdicators [2]	Rule Breakers [2]		Disaffected [3]	Shadow Agent [0]	

5.1 Blind

1) Strong Positive AS & Strong Negative RU: “Gung Ho”

Individuals of this type can pose a significant, if unintentional, threat to the organization. They see security as something they should be personally involved in, but are burdened by inaccurate risk perception. This may lead them propagate undesirable culture traits as they will seek to take a leadership role, but will not have a clear view of what constitutes effective action. While they will be keen to follow the existing policy their lack of understanding regarding the risks it addresses may lead them to perceive some or all of it as arbitrary, increasing their likelihood of non-compliance.

When they do decide to use workarounds their misplaced risk beliefs may result in them pursuing options they believe to be safe but in fact create significant vulnerabilities. In the quotation below we can see that the participant maintains a “rule of thumb” that they believe keeps their data secure, however this causes them to store data locally on their laptop, where it is more exposed, than on the company’s network drive.

Interviewer: “So how do you decide what to put on locally and what goes on the network drive?”

Participant: “If it was really sensitive I would keep it locally quite frankly...it’s my personal rule of thumb”

2) Strong Positive AS & Weak Negative RU: “Uncertain”

‘Uncertain’ members of an organization are strongly motivated by security. However, they are unaware of the risks they may encounter, leading them to be unsure as to why certain policies may be in place, or unclear as to the consequences of any potential workarounds. While they may wish to play a role in creating a positive security culture they lack the knowledge to consistently choose between good and bad, leaving them uncertain of where to place their effort. Their lack of certainty makes them less likely to take action, meaning they are less of a risk than ‘GungHo’ individuals. The following quote illustrates the essence of this type; they want to do the right thing but aren’t sure what that might be:

“I don’t know what the guidance is on restoring passwords, no. Actually, I should probably find out to make sure I’m not in breach of it. But I suspect I probably am, ‘cause they’d probably say don’t write it down.”

3) Weak Positive AS & Strong Negative RU: “Naïve”

AS+ individuals hold a generally positive outlook toward security, but are more likely to contravene security policy when it negatively impacts their primary task. Here this is combined with active misconceptions regarding what constitutes risky behavior. The risk for the organization created by this group is that, like others with a strong negative risk understanding score, their misconceptions regarding risk can lead them to adopt highly insecure behaviors, sometimes under the misguided assumption that they are acceptable. The following quote illustrates this attitude. The participant sees the password manager provided by the organization as a useful addition to their working day, but worries about its security. This incorrect assessment of the risks could lead them to use their own, less secure, approach of reminders.

“There’s this auto-sign in tool, I don’t know what it’s called. ‘Remember Me’, something like that. So it remembers your password for that thing, and I think that’s useful and handy, but I also think it’s a bit insecure.”

4) Weak Positive AS & Weak Negative RU: “Passive”

Individuals in this group feel that security is necessary for the organization, although not something they themselves should have to put time in to. While they are aware of the policy they are not always clear why it exists, leaving them following rules by rote. Lacking strong convictions they are the group most susceptible to outside influences – if placed with ‘Champions’ they will pursue better habits; conversely if surrounded by frequent rule breaking they will likely follow suit. An organization can work to surround this group with positive influences, which will likely move them, at worst, to being ‘Followers’. In the following quotation the participant recognizes that the laptop screen should be locked some of the time, indicating that they are aware of the need for security. However, they are uncertain as to when, and thinks that there are times it is acceptable to leave it unlocked, showing that they do not have a full grasp on why locking the screen is advised.

Interviewer: “You’re supposed to lock your laptop every time you leave your desk aren’t you?”

Participant: “Yeah I’d say most of the time, err but not all of the time I would say.”

Similarly, this following quotation also reflects a lack of risk understanding around locking screens, where the individual is inconsistent in their approach to security tasks because they lack understanding of what the risks actually are:

“I do a little work and get some coffee or something, but if I go to the bathroom I don’t put it on lock.”

Suggested ‘Blind’ Interventions: Those in this quadrant should be given targeted education and training. Most types, particularly the Gung-Ho and Uncertain, recognize the value of security but lack the knowledge to act appropriately. This knowledge can be drawn from training packages, or from others in their environment, such as ‘Champions’, where such individuals exist.

5.2 Open

The ‘Open’ quadrant represents the part of the population where risks are known to both the organization and the individual. That is to say that an effective policy exists to address the risks faced by the organization and that the individuals understand these risks sufficiently to understand and comply with the policy.

5) Strong Positive AS & Weak Positive RU: “Willing”

‘Willing’ individuals are those with a desire to take a full part in the security processes of the organization, but have only a limited understanding of the risks. In particular they do not fully grasp the causal relationship between their actions and the associated risks. This limits their ability to act securely in situations outside of those specifically covered by the policy. This also means they may lack the confidence to challenge non-compliance in their immediate environment, unless they feel they have the backing of clear rules, or

the support of more senior members. Their weak risk understanding holds them back from being true leaders of a good security culture. In the following quote a member of staff discusses apprentices within the organization. While they strongly state that security is important they only relate it to the rules laid out by the policy, indicating that they may not understand the risks that have driven the creation of the policy.

“They vary from about 16 right the way up to about 27, but they’re not allowed to have a different view on security. There is a policy to be followed. That’s the way it’s done [here] and that’s the way they need to do it.”

Another exemplar of the ‘Willing’ category is reflected in the following quote where the individual clearly demonstrates a strongly positive stance towards security whilst being less specific about associated risks:

“I don’t find security a problem, it’s probably as rigid as it needs to be, I can’t think of anywhere where there isn’t security that there should be, let’s put it that way. So from my point of view, I guess if I was a regular user of customer data, that’s more sensitive than the kind of stuff I do, but nah I don’t find it obtrusive, it’s what it is, it’s necessary, and I understand why.”

6) Strong Positive AS & Strong Positive RU: **“Champion”**

‘Champions’ are ideal members of staff. They combine a high level of motivation regarding security with a good understanding of both the risks they are likely to face, and the implications of those risks. This makes them able to not only comply with policy, but ensure that they remain secure even in situations where no explicit policy exists, as well as carrying over good habits in to their personal lives. In addition, they will promote a positive culture around them, acting as security leaders with the ability to influence others. In the following quotation the participant discusses the use of social media and clearly demonstrates both an understanding of the risks involved in its use, and their own role in reviewing its use within their team, reflecting the attributes of a Champion:

“If someone on my team was putting something negative on there then that would be something we’d have a full review on. You’ve got to be very conscious about what you’re writing; all these things are stored and can be re-used. Basically you put in an evidence trail on something whereas when something’s verbal, we could be having a conversation and nobody would be any the wiser.”

It is worth noting that an effective security culture does not require every member to be a security champion, only that a critical mass exists, sufficient to ensure that any insecure behaviors that crop up never take hold and become institutionalized. It is critical to understand that ‘Champions’ can only exist in organizations that have a policy worth championing! Policies that do not take in to account the business process of the organization, or those that overlook known risks, will not be promoted by ‘Champions’. Under such conditions these individuals will attempt to remain secure even if this means going against the existing (ineffective) policy – placing them in the ‘Shadow Security’ category. This is illustrated in the following quote, in which the individual is positive about security but raises concerns about the security and privacy policy with the organization.

“To be honest, I generally think we do a fairly good job around

[business activity]. The only thing that I would say is that, going forward, I think we’ve decided to monetise the data we hold about people, which is going to be a very, very difficult thing to do without a leak in secure ... leak in data, giving out too much data, or that side of things. The privacy of our customers, I think there’s a very fine line; I think there’s a bit of a balancing act coming up if we’re going to go down that route...”

7) Weak Positive AS & Weak Positive RU: **“Follower”**

Individuals in this group will follow the prevailing security culture within the organization, without taking much initiative of their own. They understand enough of the risks to see security as important, but are not sufficiently invested in the process to pursue secure options if they come at too high a cost to themselves. As a group they will likely form a large part of any well-developed security culture, providing a stable core of behavior as long as the policy is well communicated. However, ‘Followers’, like ‘Passive’ individuals, are easily influenced, and do not promote or maintain a positive culture. As such without the influence of ‘Champion’ or ‘Willing’ individuals they can over time adopt insecure habits. This is clearly illustrated in the quote below, in which the participant take a generally positive view of security, but states that the behavior of his colleagues is driven more by what others do than by their own motivations.

“I think people generally follow the norm. We’re sort of alright, but I wouldn’t say we’re security conscious. Round where I sit we all lock our desks, but ... I don’t think people say ‘I’m going to come in and be security conscious.’ They say, ‘I’m going to come in and do what everyone else does.’”

This category would also include individuals who, whilst reasonably positive about security, demonstrate weak awareness of both the risks and security policy, as evidenced in the following quote by their comment about security being ‘common sense’:

“For me it’s, this is a big organisation, it has a big security policy in place, but it’s just, you don’t know what it is, but you, you use common sense /okay/ that you’re following rather than the policy that’s on there, if you know what I mean...”

8) Weak Positive AS & Strong Positive RU: **“Expert”**

‘Expert’ users possess the same level of risk knowledge as ‘Champions’ but are not as motivated by either security, or the organization, or both. This means that whilst they are inclined to see security as a positive part of the organization, their time and effort is more likely to be spent pursuing their own goals rather than seeking to promote a wider culture of security. Their own security practices will be technically competent but tailored to their own use. ‘Expert’ individuals pose no risk to the organization, which also has no pressing need to attempt to shift them to another category. However, the organization must be careful not to push these individuals past their Compliance Threshold [4] as they have the knowledge and skills to effectively circumvent the policy. The silver lining is that their understanding of the implications of their actions will likely lead them use workarounds that are relatively secure. The following illustrates the ‘Expert’ type as they are dismissive of the current policy while being fully aware of its technical limitations.

“That’s another thing, all the security, you can put all the security

you want on the computers, but they're enforced by the network. I could just take my laptop home and plug it in to my wireless at home and do whatever I want on the computer and bring it back in."

Suggested 'Open' interventions: rather than intervening to change behavior here the organization should look to take advantage of the positive attributes of the 'Open' quadrant, in particular Champion and Willing individuals, to influence others (such as 'Passive' members of staff) and drive forward culture change initiatives. This category of behavior types illustrates the value of having the cooperation of members of staff when designing and implementing security.

5.3 Unknown

9) Weak Negative AS & Strong Negative RU: "Reckless"

Individuals in this category feel that security is more of a hindrance than a benefit to their primary task, while also actively misunderstanding what constitutes risky behavior. This makes them likely to seek workarounds, which may introduce significant vulnerabilities to the system as they do not have the necessary knowledge to understand the possible consequences of their actions. The following quotation from Company A exemplifies the 'Reckless' type. They are happy to reveal that they keep their passwords on a post-it note near their computer screen, and seem completely unaware of the risk this poses.

Interviewer: "When you say a note, is that a piece of paper note? Or..."

Participant: "Yeah post it."

I: "...Okay in your desk or on the desk or on the screen?"

P: "Yeah near my desk. It's not on the screen, it's near the screen."

10) Weak Negative AS & Weak Negative RU: "Apathetic"

'Apathetic' individuals are primarily motivated to just keep their heads down and get their jobs done. They do not see the benefit of security, and are unaware of some of the risks, making them prone to committing errors, but do not hold any serious misconceptions about what constitutes insecure behavior. Of all the groups they are the least involved with the security process. Over time, if immersed in a positive security culture, 'Apathetic' individuals may move into either 'Passive' or 'Follower' groups, sharing as they do their susceptibility to outside influence. In the following quotation we can see the participant respond to a discussion about security policy with a clear statement expressing his lack of interest.

"I mean, I read it when they, they send it to me, but that's not, I don't need to memorize it, so I don't memorize it."

Another exemplar that demonstrates employees' apathy towards security and any associated risks is as follows:

Interviewer: "If you work from home, where do you put your laptop at home?"

Participant: "Generally it's just in the living room, sort of thing; yes, just behind the sofa"

I: "Do you lock it away for the night?"

P: "...No."

I: "...You're not afraid of your kids doing something to it?..."

P: "No, I'm more afraid of them doing something to my own laptop,

which I'd have to pay to replace." (Laughs)

An organization is not necessarily at significant extra risk due to the presence of 'Apathetic' individuals, but should take their presence as a warning sign that their existing security policy and communication strategy may not be as effective as it should be.

11) Strong Negative AS & Strong Negative RU: "Abdicators"

'Abdicators' represent a serious concern for any organization. Not only do they have active misconceptions about the level of risk associated with a given course of action, but they also do not see any value in organizational policy. They feel that security hampers their own goals and seek to go their own way more often than not. For an organization the first step to dealing with such individuals is to remove them from the organization. If they are considered essential for other reasons then it will take a considerable level of effort to move them in to a non-threatening category. This quotation demonstrates both a lack of understanding of the risks associated with data sharing via unencrypted USB sticks, and also a dismissive attitude toward the existing policy.

"We couldn't send databases, we couldn't email a database. Which is stupid, because you could just zip the database and send it as a zip file and get around it anyway. But it was quicker sometimes to just throw it on a flash drive and chuck it over the cube wall. We didn't see any, honestly, just didn't really see the point in needing to buy an \$80 [encrypted] flash drive to do that..."

12) Strong Negative AS & Weak Negative RU: "Rule Breakers"

This group is highly dissatisfied with the current security policy, seeing it is strongly negatively impacting their primary task. This may lead them to break the rules whenever they feel it would benefit their productivity. Alternatively, they may negatively influence security culture by being overly critical of certain policies. While they are not completely unaware of the risks they face they do lack certain key pieces of information, meaning that their rule breaking is likely to introduce vulnerabilities in to the system.

An organization's primary focus when attempting to shift this group into another category would be to direct their efforts into both reducing the impact of security on their primary task, and increase their perception of the value of security to both themselves and the organization. The following is an example of a member of staff who feels justified in going against the security policy as they feel the organization does not provide an adequate solution that allows them to remain secure and achieve their work goals:

"I know that security frown quite a lot on that, they don't like that and their principle is if you need the tools for the job, the company should be providing those tools, which is a fair statement, but sometimes I know that I can do things with my personal computer ... different operating system, I can do something with my own computer and do it very quickly, that I just can't do at all on my work one."

Suggested 'Unknown' interventions: While this group is unlikely to be heavily populated, its members pose the highest risk to the organization. As they are not motivated by security, education and training will have little effectiveness. A sustained effort will be needed to both provide the necessary knowledge as well as to address the negative emotional response to security. It may be simpler in the

more extreme cases to remove such individuals from the organization entirely. Due to the antagonistic relationship with the organization, making use of 'Open' employees to shift the culture in a positive direction is likely to be more effective than a centralized initiative.

5.4 Hidden or 'Shadow'

A significant number of members of staff being identified as members of the 'Hidden' quadrant most likely indicates that the organization's current policy is not suited to the context in which it is being applied.

13) Weak Negative AS & Weak Positive RU: "Excuse Makers"

'Excuse Makers' feel that security is a hindrance to their primary process, despite understanding that risks are associated with non-compliance. They will circumvent the policy when it gets in the way of their goals, especially in environments when rule-breaking is the norm. They tend to excuse their rule-breaking by referring to the costs associated with compliance. Like the 'Passive', 'Follower', and 'Apathetic' groups, 'Excuse Makers' can be influenced by those around them, and can be seen as an indicator of wider problems, rather than as a serious problem in and of themselves. This quotation demonstrates the consequences of having security tasks that hinder normal activities. The staff in question was trying to solve a network adapter issue but lacked the rights to do so. There were no effective procedures in place to help them and so the delays became an excuse not to bother:

"I knocked around for a while trying to find an answer. I restarted my machine a couple of times and that probably would have taken an hour. They are slow to boot. Does it stop me? I guess there's two ways. It's either it does slow you down or it makes something so difficult that you think, "Do you know what? I won't bother."

14) Weak Negative AS & Strong Positive RU: "Circumventers"

This category of individuals share many characteristics with 'Experts'. However, unlike experts they see security as a barrier to achievement and use their skills and knowledge to circumvent policy when it exceeds their limited tolerance. As 'Circumventers' have a strong understanding of the consequences of their actions an organization that identifies these individuals in their staff population should assume that their current policy has not been effectively designed as it is forcing people to choose non-compliance even in the face of a full understanding of the risks. This is illustrated in the quotation below:

"I know the policy is that there is certain particular...secure flash drive that's got to be used. But it's terrible to use. It's got a...part of it is some virtual thing where you have to log into it...at crucial moments you're taking this flash drive to a presentation and you need to get it to work and at a crucial moment, the damn thing doesn't work. So people, what people do is that they...don't use them because of the...the problem. Because it, uh, we can, I think people can understand the rationale behind the policy, but if the thing doesn't work in practice every time, it gets abused."

15) Strong Negative AS & Weak Positive RU: "Disaffected"

A more extreme case of the 'Excuse Makers', this group feels

strongly that security is a hindrance and rather than making excuses for their circumvention will feel fully justified in their non-compliant behavior. They lack the deep knowledge to circumvent security effectively, but will instigate non-compliant habits and propagate those through the organization. This quote reflects a member of staff's frustration with the security provision in the organization whilst recognizing the risk:

Interviewer: "How do you devise passwords?"

Participant: "I would make them as simple as the system would let me get away with and I would record them in a word document. So it's not all that secure but if there's so many of them and 'remember me' won't work. I can't remember them all...."

16) Strong Negative AS & Strong Positive RU: "Shadow Agent"

'Shadow Security' practitioners seek to completely step outside the company policy as much as possible. They have a full understanding of what constitutes secure and risky behavior and likely possess the technical skills to implement sophisticated workarounds. Members of this group are arguably the most dangerous to the organization as it will include both malicious insiders and those individuals that, while currently still wishing to complete their work tasks, regard the organization as largely incapable of supporting those tasks through an appropriate security policy. Organizations identifying 'Shadow Agents' within their organization should consider themselves at significant risk of either insider attack, or losing highly skilled staff through high levels of dissatisfaction with the organization. In the following quotation we can see this view clearly expressed. The participant is clearly very aware of the risks of keeping password files, but is critical of the current sign on procedures leading them to adopt their own methods for solving the problem. This quote also illustrates the point that individuals in this category are not necessarily malicious, but are simply working outside of the organizational system.

"Single sign-on should mean that one password is used, but the way I'm talking about is actually like an assisted sign-on that repeats that method but, we're not maintaining that properly so what you're finding is you go back to the old method of logging in directly and keeping passwords in and to be honest, I think what you find a lot is that people either break those passwords down. They keep them in a file on their desktop or they use the same password for everything which defeats the purpose of it I think. The days are passed where this was ok and we need catch up."

Suggested 'Hidden' interventions: Education and training will not be useful in shifting this group, as the problem is not one of knowledge, but rather of decision-making. Instead, the organization should identify which security-related factors contribute toward non-compliance and look to remove or redesign them. 'Fix the human' approaches will be particularly ineffective in adjusting the behavior of individuals in this quadrant, and the organization should look to make changes themselves.

6. TOOLS

The long-term goals of the project is to create a set of tools to detect, measure and visualize actual staff attitudes to risk and security and behavior, using the Behavioral Security Grid, our revision of the Johari Window. The optimal area for employees to reside in is within the 'Open' quadrant, which is positioned on the upper-right hand of

the grid and which implies the staff has scored above average in terms of both risk understanding and affective security. Our intent is to develop a scalable metric, not reliant on interview analysis (a current bottleneck in the methodology), that will allow us to rapidly plot whole staff populations on the BSG, providing organizations with a complete 'snapshot' of their current security culture.

Subject to testing and validation, this snapshot would be used as an organizational diagnostic audit tool to identify information about staff security risk behavior which will i) facilitate understanding of the different risk profiles of staff members within the organization and their behavioral responses to real-world security scenarios, ii) identify and locate where the risks reside within the organization, iii) plan and execute appropriate, targeted interventions and iv) by re-measuring after a period of time, track changes in the risk attitude within the staff population following an intervention. Using this visualization method will enable organizations to adopt a scientific approach to tracking staff security risk perception enabling them to tailor the deployment of their resources more effectively.

7. CONCLUSION

Treating staff as a homogeneous group damages an organization's ability to provide effective security as policies that do not take the different attitudes, competencies and resulting behaviors of staff populations in to account promote non-compliance. This comes as a result of the burden such policies place on employees. By recognizing that security behavior is driven by both affect and risk understanding we provide a framework based on these dimensions that allows organizations to map the heterogeneity of their staff populations. We also suggest categories based on our understanding of the axes and provide practical suggestions to organizations as to how to identify and manage these individuals. As our categorization includes both positive and negative individuals, organizations can not only identify potential problems that require intervention, but also where they have beneficial expertise and cultural elements that can be harnessed as part of those interventions. Additionally, we recognize the difference in intervention methods needed by the 'Blind' quadrant, where targeted training is a viable strategy, and the 'Hidden' quadrant, which requires organizations to focus on 'security hygiene' by re-designing high-friction security. A successful management strategy will require the right combination of interventions, which must be determined by the specific characteristics of the target population.

Of course, the BSG does not make the assumption that the organizational security policy in place is always beneficial. Indeed, if employees are clustered in the Hidden quadrant and report narratives of not being able to complete their primary task as well as comply with policy, it is likely that the existing policy may be less than optimal. While understanding how appropriate an organisation's policy is can add context to the results of the BSG, this early iteration is designed primarily as a diagnostic tool and does not attempt to formally assess the efficacy of security policy. It instead seeks to reflect how employees feel about security and the extent to which they understand the risks associated with their roles. Nevertheless, it is a useful starting point for an organization to explore in further depth which issues or productivity blockages, if any, exist in relation to the security policy.

A limitation of this work, however, is that the current data set is small. This will be addressed in future work where the model will be analyzed using larger population samples. Additionally, the security

types are, at this stage, purely theoretical hypotheses based on logical extensions of the framework. Future work will include further validation and refinement of the suggested security types. The collection of data from other organizations is currently underway, which will allow us to begin validation for the framework itself as well as the security types.

The authors are aware that the coding methodology to date is not sensitive to certain patterns of response, particular those with polarised views, in that it has applied strong and weak positive and negative dimensions to the qualitative data but does not make use of neutral codes. This means that passive participants making few statements and active participants making many statements but split evenly between positive and negative will achieve similar resulting scores. In response, the inclusion of neutral codes may be incorporated into the framework as part of the validation stage of both the BSG framework itself and the behaviour types. Further research is required to explore the efficacy of this approach.

It should also be noted that although the BSG attempts to measure employees' knowledge around security risks it does not explicitly incorporate a measure for knowledge of the security policy. It is anticipated analysis of the qualitative data from the interviews should provide organizations with further information about this. However, the value of the BSG is in the clarity offered by the axes of Risk Understanding and Affective Security, on which it is not possible to incorporate every aspect of employee security.

In addition, we recognize that at this stage our approach does not provide a means of objectively assessing organizational culture against a standardized baseline. It does however effectively allow comparison between two data sets. These could be drawn from two separate populations in order to compare different organizations, or to compare different departments or geographical locations within a single organization. Alternatively the data sets could be taken from the same population but at different times, allowing changes to be tracked and acted on. This information directly contributes to an organization's ability to secure itself as it both reduces the amount of resources wasted on ineffective training, and decreases the burden placed on users, leading to an increase in compliant behaviors.

Acknowledgements: The authors would like to thank Anthony Morton, for his initial contribution to the project and Kat Krol and Simon Parkin for their invaluable assistance in preparing this paper.

Copyright: Figures 1-5 © Odette Beris, Adam Beaument and M. Angela Sasse

REFERENCES

- [1] B. Schneier. *Secrets and Lies: Digital Security in a Networked World*. 1st ed. Wiley, 2004.
- [2] I. Kirlappos and M. A. Sasse. What usable security really means: Trusting and Engaging Users. In *Human Aspects of Information Security, Privacy, and Trust*, T. Tryfonas and I. Askoxylakis, Eds. pages 69–78. Springer International Publishing, 2014.
- [3] F. Pallas. Information Security Inside Organizations - A Positive Model and Some Normative Arguments Based on New Institutional Economics. *Social Science Research Network*, Rochester, NY, SSRN Scholarly Paper ID 1471801, 2009.

- [4] A. Beutement, M. A. Sasse, and M. Wonham. The compliance budget: managing security behaviour in organisations. In *Proceedings of the 2008 workshop on New security paradigms*, pages 47–58. ACM 2009
- [5] A. Adams and M. A. Sasse. Users are not the enemy. *Communications of the ACM*. 42 (12) pages. 40–46, 1999.
- [6] I. Kirlappos, S. Parkin, and M. A. Sasse. Learning from “shadow security”: Why understanding non-compliance provides the basis for effective security. In *USEC*, 2014.
- [7] S. L. Pfleeger, M. A. Sasse, and A. Furnham. From Weakest Link to Security Hero: Transforming Staff Security Behavior. *Journal of Homeland Security and Emergency Management*. 11 (4), 489-510, 2014.
- [8] J. Luft and H. Ingham. The Johari Window: a graphic model of awareness in interpersonal relations. *NTL Reading Book for Human Relations Training*. pages 32–4, 1982.
- [9] P. Slovic, E. Peters, M. L. Finucane, and D. G. MacGregor. Affect, risk, and decision making. *Health Psychology*, 24 (4,Suppl) S35-S40, 2005.
- [10] P. Slovic, M. L. Finucane, E. Peters, and D. G. MacGregor. The affect heuristic. *European Journal of Operational Research*, 177 (3) 1333–1352, 2007.
- [11] D. Kahneman. *Thinking Fast and Slow*. Penguin Books, 2012.
- [12] B. Schneier. *The Psychology of Security January 21, 2008*.
- [13] C. Herley. So long, and no thanks for the externalities: the rational rejection of security advice by users. In *Proceedings of the 2009 workshop on New security paradigms workshop*, pages 133–144, NY, USA, 2009.
- [14] G. F. Loewenstein, E. U. Weber, C. K. Hsee, and N. Welch. Risk as feelings. *Psychological bulletin*, 127 (2): 267-286. 2001.
- [15] G. L. Clore, K. Gasper, and E. Garvin. Affect as information. In J.P. Forgas, (Ed). *Handbook of affect and social cognition*, pages 121-144, Mahwah, NJ, Lawrence Erlbaum Associates, 2001.
- [16] F. Farahmand, M. J. Atallah, and E. H. Spafford. Incentive Alignment and Risk Perception: An Information Security Application. *IEEE Transactions on Engineering Management*, IEEE 60(2), 238-246, 2012.
- [17] M. J. Massie and A.T.Morris. Risk Acceptance Personality Paradigm: How we view what we don't know we know. American Institute of Aeronautics and Astronautics, 1-18, 2011
- [18] I. Kirlappos, A. Beutement, and M. A. Sasse. “Comply or Die” Is Dead: Long Live Security-Aware Principal Agents. In *Financial Cryptography and Data Security*, pages 70–82, Springer Berlin, 2013.
- [19] G. Guest, K. M. MacQueen, and E. E. Namey. *Applied Thematic Analysis*. Sage Publications, 2012.