

Northumbria Research Link

Citation: Nicholson, James, Coventry, Lynne and Briggs, Pamela (2018) Introducing the Cybersurvival Task: Assessing and Addressing Staff Beliefs about Effective Cyber Protection. In: Proceedings of the Fourteenth Symposium on Usable Privacy and Security (SOUPS) 2018: Baltimore, MD, USA August 12-14, 2018. USENIX Association, Berkeley, pp. 443-457. ISBN 9781931971454

Published by: USENIX Association

URL: [https://www.usenix.org/conference/soups2018/technical...](https://www.usenix.org/conference/soups2018/technical-sessions)
<<https://www.usenix.org/conference/soups2018/technical-sessions>>

This version was downloaded from Northumbria Research Link:
<http://nrl.northumbria.ac.uk/id/eprint/34526/>

Northumbria University has developed Northumbria Research Link (NRL) to enable users to access the University's research output. Copyright © and moral rights for items on NRL are retained by the individual author(s) and/or other copyright owners. Single copies of full items can be reproduced, displayed or performed, and given to third parties in any format or medium for personal research or study, educational, or not-for-profit purposes without prior permission or charge, provided the authors, title and full bibliographic details are given, as well as a hyperlink and/or URL to the original metadata page. The content must not be changed in any way. Full items must not be sold commercially in any format or medium without formal permission of the copyright holder. The full policy is available online: <http://nrl.northumbria.ac.uk/policies.html>

This document may differ from the final, published version of the research and has been made available online in accordance with publisher policies. To read and/or cite from the published version of the research, please visit the publisher's website (a subscription may be required.)

Introducing the Cybersurvival Task: Assessing and Addressing Staff Beliefs about Effective Cyber Protection

James Nicholson

PaCT Lab
Northumbria University
Newcastle, UK

james.nicholson@northumbria.ac.uk

Lynne Coventry

PaCT Lab
Northumbria University
Newcastle, UK

lynne.coventry@northumbria.ac.uk

Pam Briggs

PaCT Lab
Northumbria University
Newcastle, UK

p.briggs@northumbria.ac.uk

ABSTRACT

Despite increased awareness of cybersecurity incidents and consequences, organisations still struggle to convince employees to comply with information security policies and engage in effective cyber prevention. Here we introduce and evaluate *The Cybersurvival Task*, a ranking task that highlights cybersecurity misconceptions amongst employees and that serves as a reflective exercise for security experts. We describe an initial deployment and refinement of the task in one organisation and a second deployment and evaluation in another. We show how the Cybersurvival Task could be used to detect ‘shadow security’ cultures within an organisation and illustrate how a group discussion about the importance of different cyber behaviours led to the weakening of staff’s cybersecurity positions (i.e. more disagreement with experts). We also discuss its use as a tool to inform organisational policy-making and the design of campaigns and training events, ensuring that they are better tailored to specific staff groups and designed to target problematic behaviours.

1. INTRODUCTION

The number and scale of cyber-attacks targeted at organisations over the past few years is unprecedented. These include hackers compromising 55 million voter records in the Philippines, hospitals worldwide hit by ransomware attacks, 33 million Twitter user names and passwords being compromised, and 11.5 million documents relating to offshore accounts of international politicians, business leaders and celebrities being leaked from a law firm [51]. Many major breaches still go unreported, with only a quarter of businesses in the UK reporting their major breaches last year [33]. Business email compromise, ransomware, and phishing are cited across industries as the top vector of compromise. In many of these cases, the attack vector involves the employee. Organisations and their employees understand that they have a responsibility to change employee behaviour as an important tool in their defence strategy, yet there is very little consensus about exactly what protective behaviours are to be advocated and prioritised. Security practitioners, policy-makers, managers, and employees tend to

advocate different approaches and the end result is that users receive conflicting advice, become sceptical about the information they are given, and are consequently less proactive in cyber defence than they might otherwise be [9, 35].

Organisations typically have one or more policies addressing appropriate cybersecurity behaviour, referred to as *security policies* from here on. There is now significant literature that describes those factors that influence employees’ intentions to comply with security policies [21, 31, 37, 46] and further literature documenting poor outcomes from cybersecurity awareness campaigns and organisational training initiatives [5, 41, 49, 55]. Sometimes the reason for these failures is straightforward. For example, security policies are often inaccessible or buried deep within an organisation’s website, tend to be over-complex, incomprehensible and/or poorly tailored to staff needs and workload [39]. They are generally poor calls to action, not least because of the aforementioned confusion about the protective actions they promote. This is a particular problem when we consider the psychology of threat, where we know that highlighting the threat to a user, without also offering them a simple, consistent response to that threat, produces ‘defensive’ reactions that can include simply ignoring the problem and continuing to engage in old behaviours [29].

One example is the conflicting advice surrounding the password, where standard advice was once to create strong, unique passwords for every user account involving combinations of letters, numbers and ‘special’ characters. Recently, this advice has been supplanted (e.g. by NIST and GCHQ) with a ‘three random words’ instruction for password creation [25]. This would seem to constitute an advance, but can lead to greater confusion on the part of the end user as many current accounts still enforce ‘strong’ passwords requiring multiple character types, effectively rendering GCHQ and NIST advice useless in that particular context.

In this paper, we focus on the consensus problem in cyber protection and describe a tool (*The Cybersurvival Task*) that highlights the many different behaviours encompassed by a cybersecurity policy and the mental models held by members of an organisation. The task requires users to rank protective behaviours in terms of their effectiveness as a cybersecurity defence. Unlike other self-report measurement tools (e.g. [19]), these rankings provide a means for staff to disclose their assumptions in a structured way, so that organisations can understand where employee confusion and associated defensive responding might be taking place. Most importantly, the process allows for organisational security experts to reflect upon their policy and training priorities, based on direct feedback from their own

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.

USENIX Symposium on Usable Privacy and Security (SOUPS) 2018, August 12 -- 14, 2018, Baltimore, MD, USA.

employees. The ultimate aim is that the Cybersurvival Task could inform the development of organisational policy-making and the design of campaigns and training events, ensuring that they are better tailored to specific staff groups and/or misconceptions. Here, we describe the development of the task and describe a process whereby we piloted the task in one institution, made some refinements, and then conducted an evaluation of the final task in a second institution. We show how the task highlighted misconceptions and revealed behavioural discrepancies between experts and employees, and between different employee groups, and discuss how organisations can benefit from the Cybersurvival Task.

2. BACKGROUND

2.1 The Human Factor in Cyber Protection

Organisations face a growing range of security threats, including denial of service (DoS) and ransomware attacks that aim to take down a business or service, as well as social engineering attacks that are designed to obtain and exploit private information. While it may be possible to stay safe from such attacks by improving and maintaining the organisation's technical defences – e.g. firewalls and anti-virus software – employees have often been labelled as the 'weak link' in the security ecosystem (e.g. [53]). In recent years, this weak link argument has been replaced by an understanding that humans, far from being 'the enemy' [2] are an integral part of the *whole system* and that a proper understanding of human behaviour and of employee motivation should inform the cybersecurity design process [45].

Much of the work in this space has focused upon the fact that cybersecurity does not comprise the primary task for most employees. Unsurprisingly, people attend to their primary work tasks and tend to overlook security actions. Beauteament, Sasse, & Wonham [8] have suggested that employees have a relatively small 'compliance budget' that they can allocate to security procedures and that this can shrink when job demands are particularly high or the protective behaviours demanded of users are too onerous. There are unrealistic expectations that users will create a strong password for every unique account they have [56], that they will be vigilant in checking for phishing emails they receive [54] or that they will simply not click on any links or open any email attachment in the workplace [27]. The reality is that the vast majority of people reuse simple passwords [1, 26] and that almost half of all users are likely to fall for phishing emails, with some 17% on average entering credentials on phishing websites [12].

2.2 The Non-Compliance Problem

There is often a disconnect between how organisations would *like* their employees to behave and how the employees *actually* behave and this is an important consideration for computer security (e.g. [8, 30]). Much of the existing organisational research tends to focus upon this as a 'policy compliance problem' rather than see it more holistically as an issue around the ways that employees come to understand both the cybersecurity threat and the kinds of protective security behaviours they can use to ameliorate that threat. This is important, because employees do not typically gain their understanding directly from security policies, but rather from their work peers and from the media, building up a set of *shadow security* beliefs and behaviours [34] that deviate from company policy. In other words, employees reach a compromise between security and productivity that allows them to achieve their work goals by utilising non-compliant but sufficient security behaviours.

Regardless, many organisational policies and procedures are simply not fit for purpose. There are issues with policies that are too dense and contain tracts of information that are irrelevant for many users. There are also issues with policies that are too vague and provide very little in the way of useful information [3]. Unsurprisingly there are also many organisations that have no security policies in place and many users who are simply unaware of their own organisation's stance on cybersecurity behaviour. In short, it is not easy for organisations to develop usable cybersecurity policies to keep employees safe.

2.3 Choosing the 'Right' Behaviour

We noted that users often struggle to protect themselves and their organisation online, and part of the problem is that they are given inconsistent advice about what actions to take. As cyber security experts differ in their opinion of the skills and behaviours that are important [13], so too do the security policies they create. This means security policies vary between organisations and include many different behaviours associated with accessing, categorising, storing, and transferring data – but may also cover general computer user policies including internet and email behaviours and use of external devices (USBs, personal devices). With this in mind, researchers at Google distributed a survey to both security experts (those having at least 5 years of experience working or studying in computer security) and security non-experts (Mechanical Turk workers) and found a discrepancy between online security behaviours reported as essential between the expert and non-expert group [32]. Most importantly, the researchers compiled a list of advice considered 'good' by experts consisting of 20 items. While this list constitutes a step in the right direction for identifying security behaviours that are important for staying safe online – for both policy creation and advice-generation – this advice is based on both academic and industry experts which may have contrasting views on a number of topics [32]. Additionally, this list is based on 'good' advice, defined as advice that is both effective and realistic, which potentially means that security behaviours that are very important for the organisation may have been pushed down the list. Finally, the list was compiled for the average internet user, meaning that some behaviours may not apply to everyone and this is already a problem faced by users who are overloaded with occasionally irrelevant advice [30]. In a corporate environment where job roles are clearly defined and responsibilities differ across individuals, such a generic list will likely offer excessive or irrelevant advice to individuals.

2.4 Measuring Security Behaviours and Beliefs

We have highlighted the problems that organisations face when writing security policies, so it is no surprise that enforcing the policy becomes even more challenging. But how can organisations understand what their employees are doing in the security spectrum?

Direct measurement of actual security behaviour in a live environment has proved elusive for cybersecurity researchers and many have adopted self-report scales as workable alternatives. These, of course, measure intentions to behave in a certain way and assume there are no barriers to converting these intentions into actual behaviour. A range of psychometric scales have been developed and these typically include different behavioural items where participants are asked to rate the likelihood of complying or agreement with the behavioural statements. For example, Egelman

& Peer [19] start off with 30 items which they reduce to 16 security behaviours covering 3 security topics, whereas Parsons et al. [48] list 63 different behaviours covering 7 topics. However, Wash et al. [61] found that people are poor at self-reporting security behaviours, as they may not understand what the behaviours are and may underreport less salient behaviours. This has important implications for the validity of such scales.

With this in mind, a different approach to measuring and observing behaviour may be necessary and in this paper we consider the advantages of ranking behaviours instead of rating them. The inspiration from this work comes from two seminal examples of ranking tasks used both to facilitate group discussions and to study group dynamics in occupational settings: The Desert Survival Situation [38] and the Moon Landing Task [16]. While these tasks do not measure organisationally-relevant behaviours and beliefs, they are worth considering here as they have been used for over four decades to understand the kinds of decision-processes individuals and groups make within the work context and to determine which factors are most likely to shape attitudes within the workplace.

2.5 Ranking Tasks as a Measure of Work-Related Behaviour

The Desert Survival Task [38] places participants in a simulated scenario where they are stranded in the desert after a plane crash and must rank 15 items in order of importance for survival. Participants' answers are then compared to the 'correct' answers – i.e. the rankings offered by experts – in order to indicate the accuracy of the individual and group rankings. The task has been a popular tool for understanding the behaviour of leaders in groups (e.g. [23, 42, 52]), evaluating group facilitation techniques (e.g. [58]) and exploring both individual and group decision making and problem-solving processes (e.g. [15, 24, 44]). The Desert Survival Task has also been used in disciplines other than management as a tool for understanding gender differences in schools [6], understanding what features of embodied conversation agents are most important for communicating feedback [40] and for understanding reactions to different computer personalities [20] amongst many others.

Similarly, the Moon Landing Task [16] requires participants to rank 15 items in order of importance for surviving a trip to a rescue vessel off the moon's surface. Individual and group rankings are then compared with an expert list compiled by the National Aeronautics and Space Administration (NASA). The Moon Landing Task has been used largely as a problem-solving task in studies, e.g. for understanding role of stereotypical context on the judgement of groups [4], cognitive busyness [17, 28], and teasing [10]. The task has also been used to understand group interactions amongst children [16] and as a tool for facilitating intelligence expectancy judgements on peers [43].

3. THE CYBERSURVIVAL TASK

The Cybersurvival Task asks *participants* (employees in an organisation) to *rank* the security behaviours that would best help protect their own organisation. This process is different from other security questionnaires that operate on a self-report basis, where users are asked to disclose whether or not they perform certain behaviours [19, 48]. By asking users to rank behaviours, we ensure that participants prioritise certain behaviours over others. By asking users to justify these rankings, we ensure that they articulate their beliefs about the benefits and drawbacks of these behaviours.

Table 1: Overview of the Cybersurvival Task stages.

Stage	Approx. Duration
Generate appropriate list of behaviours with the organisation's security experts tailored to workplace	30 minutes
Workshops with employees	60 minutes (each)
Reflection with experts	45 minutes

The task involves a process similar to the Moon Landing and Desert Survival tasks – in which participants engage in both individual and group ranking decisions and compare them against previously-obtained expert rankings. The major difference in this implementation is that the task items are highly salient to the cybersecurity context. In other words, the Moon Landing and Desert Survival tasks allowed exploration of a problem that was not directly relevant to the organisation in order to understand group dynamics in a 'neutral' problem space. In contrast, the Cybersurvival Task is highly relevant and allows not only the exploration of group dynamics, but the elicitation of specific mental models (at group and individual level) that are cybersecurity relevant. Critically, the Cybersurvival Task also incorporates a final reflection stage (see Table 1) not present in similar ranking tasks, where *experts* (those responsible for setting the security agenda in an organisation) can be presented with data capturing employee rankings, assumptions and beliefs.

Table 2: Overview of workshop activities.

Activity	Approx. Duration
Introduction by Facilitator	2 minutes
Individual ranking of Cybersurvival Sheet	10 minutes
Reveal of top 3 and bottom 3 behaviours (from individual rankings), plus suggestions for new behaviours	10 minutes
Group ranking of Cybersurvival Sheet – assisted by facilitator	10 minutes
Group ranking of Cybersurvival Sheet - independent	15 minutes
Reveal of expert rankings & scoring	10 minutes
Debrief	3 minutes

The task itself is simple: each participant is initially presented with a sheet (the *Cybersurvival Sheet*) consisting of n relevant security behaviours (agreed in advance with the organisation's security experts), listed in a random order, and is required to rank those behaviours in order of importance for staying safe online. The task is conducted individually, then conducted as a group, where participants are encouraged to discuss and agree on the importance

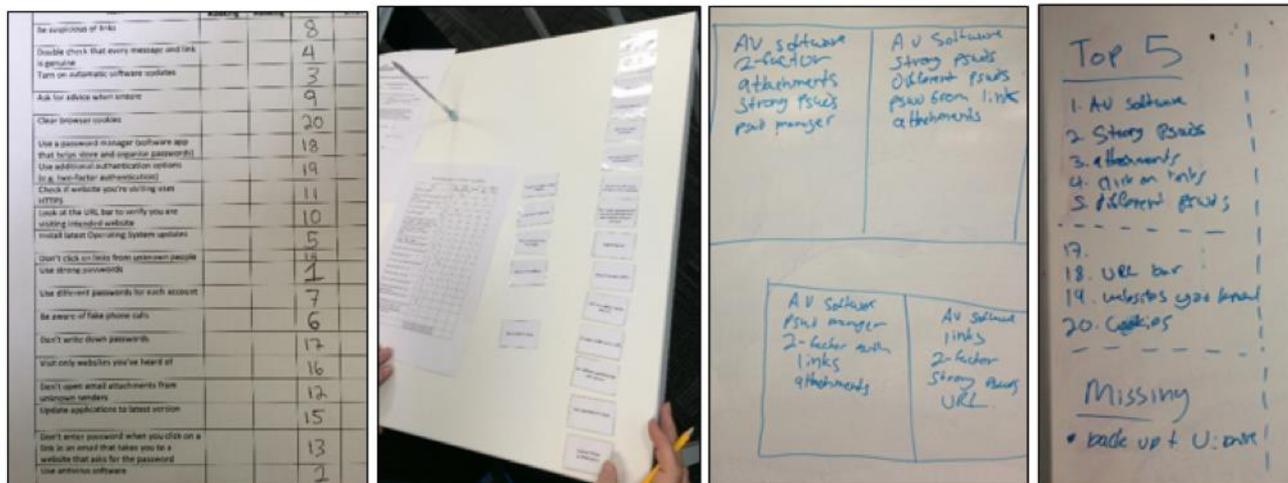


Figure 1: (a) Example of the ranking sheet used in Phase II; (b) laminated note cards used to support the individual ranking process; (c) the ‘individual reveal’ from each participant and (d) the group’s top five agreed behaviours.

of the different behaviours. The discussion through which participants come to a group consensus is as important as the rankings themselves, which are then compared with those derived from security experts in their organisation. Each participant is given a set of laminated note cards with the printed behaviours that they can use to facilitate both the individual and group ranking process (e.g. by arranging the notes before committing pen to Sheet (see Figure 1).

The security experts’ rankings are obtained via a similar process where each expert is asked to rank an initial list of behaviours (see 3.1 below) individually, followed by a group discussion where all experts have to agree on an order that suits their organisation. Experts are allowed to add, rename, and remove any behaviours at any time during the process. The initial expert ranking exercise lasts approximately 30 minutes, while the final reflection stage lasts approximately 45 minutes.

This final reflection stage highlights a striking difference between the Cybersurvival Task and the Dessert Survival or Moon Landing tasks. While the latter two tasks operate under an absolute and ‘best set’ of rankings, the Cybersurvival Task challenges the quality of the expert rankings in the final stage where they are encouraged to reflect on (and re-assess) their priorities and training programmes. The reflection stage consists of the researchers presenting the findings to the experts and allowing them to seek clarification on any of the findings (or specifics on behaviour choices). See Section 3.3 for more information on this stage.

We acknowledge that experts can be wrong (as we will show later), and by no means do we believe that the expert rankings from each institution necessarily represent ‘best practice’, but we do see the value in comparing employee rankings to their institutional experts as they have been tasked with setting and enforcing the security culture within their organisation.

Below we describe the multi-phase process undertaken to refine and evaluate the Cybersurvival Task, comprising a first deployment, task refinement and second deployment and evaluation in an institution of similar character and size.

3.1 Phase I Deployment

The first Cybersurvival Task deployment was in a large academic institution (approximately 3,000 members of staff). The goal was to understand the ‘face validity’ of the task from the point of view of experts and employees and to see whether any improvement should be made to its structure, activities, and delivery. We were also interested in whether the organisational experts and employees believed there was any value in engaging with the Task.

We first needed to develop a list of protective behaviours that were deemed relevant to the organisation, and so we conducted an initial workshop with two security experts from the organisation (the Head of IT Security and the Head of IT Services). We began with an initial list comprising the 20 behaviours from Ion et al.’s [32] study described above (see Appendix A). The two experts were asked to work individually and to rank the list of behaviours in order of their importance *for protecting their organisation*, and they were also given the chance to add and remove behaviours. Both experts were then asked to work together to rank the complete set of behaviours, including any new ones they had added. Their final ranked list, the ‘expert agreed list’, presented in randomised order, formed the Cybersurvival Sheet for employees (see Appendix B). We then used this sheet to run the Cybersurvival Task in four workshops (see Table 2 for activities) with staff in the same organisation, followed by one final workshop with the same experts who generated the initial list. Both this and the subsequent deployment received ethical approval from our university.

Twenty employees were recruited using strategically-located flyers and email distribution lists. There were 13 support staff with roles ranging from procurement to personal assistants and 7 academic staff responsible for either research or student learning. The 20 participants were split into four workshops of five participants each. One workshop consisted of solely support staff and one of solely academic staff, with the remaining two mixed. The activities and procedures in all four sessions were identical (see Table 2). Each workshop involved the participants ranking the behaviours on their own, discussing any additional behaviours with the group, and then ranking the behaviours again as a group, with a final ranking order agreed by all members of the group (see Figure 1).

Participants were then shown the agreed ‘expert rankings’ and were given the opportunity to discuss any differences between their rankings and those of the security experts. The sessions lasted approximately one hour. Thus, we collected the ranked list of behaviours for every participant (n=20) and the ranked list of each group (n=5) as well as the qualitative discussions during the group ranking activity (n=5).

Finally, the organisation’s security experts were briefed on the findings and allowed to reflect on these (see Section 3.3).

3.1.1 Lessons Learned

The Phase I deployment of the Cybersurvival Task provided us with very valuable feedback and led us to improve upon the procedures and materials for Phase II. Below we cover the most important lessons that we learned from Phase I.

Experts expressed major interest in the reflection stage and viewed this as the most valuable aspect of the Task. However, its importance was not evident at the beginning of the task, thus leading to lower engagement with the initial ranking task. Therefore, in Phase II we were clearer with experts upfront about the entire process and highlighted the benefits of tailoring the initial set of behaviours to their own organisation.

In Phase I we focussed on the ranking of the top 5 behaviours, which meant that subsequent discussion between the group members centred around those 5 behaviours with less discussion around the lowest-ranked items. In Phase II we decided to facilitate the ranking of the top and bottom 3 behaviours, thus resulting in a more balanced discussion of ‘good’ and ‘bad’ behaviours.

We also improved the presentation of the Cybersurvival Sheet based on feedback from participants. The most important change was numbering the behaviours on the sheet to facilitate discussion amongst participant (e.g. “cookies, number 7, should go below two factor authentication, number 17”). Feedback from participants also highlighted their appreciation for the laminated cards, so we continued using these facilitators during Phase II.

Finally, we observed from the initial set of 4 workshops that the most insightful data was generated by the group ranking activity, where participants were forced to directly compare the pros and cons of behaviours which led to uncovering flawed mental models and/or shadow security measures, as well as exposing issues with the security policy. Thus, we altered the timings during Phase II to allow staff more time in the group ranking activity and less time in the expert reveal, where participants predominantly dismissed the expert rankings.

3.2 Phase II Deployment

The second phase involved a deployment of the revised Cybersurvival Task in a larger but structurally similar university (approximately 5,300 members of staff). This meant that the lessons learned from Phase I were appropriate to the new context and that the kinds of attacks and protective behaviours described were appropriate, recognising that universities are prime targets for attackers due to publicly-available information [50].

The list of behaviours for Phase II was again developed through an initial workshop with security experts from that organisation – the Chief Information Security Officer and a Faculty deputy (see Table 3 for ranked list). The format of the workshop was similar to that used in Phase I, with a greater emphasis on the potential benefits of the task during the initial briefing. The initial list of behaviours

comprised the list from Phase I, plus additional behaviours that were recommended in the new organisation’s security policy. This resulted in the experts spending more time adding, removing, and rewording behaviours on the list in order to tailor it to their specific organisation.

Again, the participants were 20 non-expert employees who were split into 4 groups of 5 participants each. In Phase II, we kept support and academic staff separate – with two groups of each. Note that while we chose to separate academic and support staff due to differences observed during Phase I, it is possible for organisations to separate staff as they see appropriate (e.g. by job role or subjective experience). In fact, the Cybersurvival Task can serve as an exercise for identifying potential subgroups of employees who may share similar misconceptions.

Table 3: Final ranking of behaviours by the security experts for Phase II. Keys correspond to Figure 3.

Ranking	Behaviour	Key
1	Ask for advice	ASK
2	Save files to the network	SAV
3	Use different passwords for accounts outside the organisation	DIF
4	Keep passwords safe if written down	WRI
5	Report any data loss incidents	REP
6	Turn on automatic software updates	AUT
7	Do not disclose your personal password, even to the IT department	DIS
8	Use anti-malware software and keep it up to date	ANT
9	Use strong passwords	STR
10	Educate yourself on how to avoid fraud	EDU
11	Use additional authentication options (e.g. two-factor authentication)	ADD
12	Restrict physical access to computers and removable media	PHY
13	Check if website you’re visiting uses HTTPS	HTT
14	Look at the URL bar to verify you are visiting intended website	URL
15	Don’t open attachments from unknown senders	UNK
16	Don’t open unnecessary attachments	UNN
17	Don’t click on links from unknown senders	LIN
18	Don’t enter password when you click on a link in an email that takes you to a website that asks for the password	PAS
19	Clear browser cookies	COO

Participants were recruited via snowball emails across all Faculties of the university with the exception of Computing Science (who were excluded on the basis that they may have had particular cybersecurity expertise). Academic participants included PhD students, researchers and lecturers, while support participants included receptionists and staff in finance and human resources departments. All these ‘non-expert’ participants were compensated with a £10 voucher.

Behaviour	Individual Range	Individual Mean Rank	Group Mean Rank	Expert Rank	Key Reason
Ask for advice	7-19	18	17	1	-Not practical to do due to time required for response. Use own experience instead to make the correct judgement. -More likely to ask a colleague for advice than ask IT
Save files to network	1-16	9	4	2	Default "best practice" behaviour; pragmatic: do not want to regenerate data if lost.
Use different passwords for accounts outside organisation	2-18	10	12	3	Acknowledged as an important behaviour, but practically not feasible as cannot be expected to remember 20 unique passwords.
Keep passwords safe if written down	3-19	8	15	4	Tensions between "never" and "limited access" points of view, but generally seen as a poor security behaviour.
Report any data-loss incidents	5-15	15	13	5	Important as if data has been lost then that is an indication that there is something wrong (and it may happen again).
Turn on automatic software updates	4-19	17	16	6	-Seen as necessary for functional improvements but not security patches -Those aware of security implications also rated behaviour low due to unlikelihood of being hacked due to out-of-date software
Do not disclose personal password, even to IT	1-17	6	6	7	Very important, as disclosing password means any effort put into the creation of the password is rendered useless.
Use anti-malware software and keep it up to date	3-19	13	5	8	Seen as IT's responsibility (installation and maintenance). Although they do not need to actively manage it, AMS is seen as a very important safeguard.
Use strong passwords	1-10	1	7	9	Worthless if given away or stolen.
Educate yourself on how to avoid fraud	8-18	16	11	10	Seen as losing battle against fraudsters as they are always ahead of the curve, so time-intensive to stay up-to-date.

Figure 2: Screenshot of the report produced for experts.

The sessions consisted of a quick introduction by the facilitator, an individual ranking task followed by a ‘reveal’ of each participant’s top and bottom three behaviours (see Table 2 for activities and timings). Staff were given a chance to suggest new behaviours to add to the list. These were written on the board by the facilitator (see Figure 1). Participants were then asked to rank all the behaviours as a group with everyone having to agree on the final list at the end of the process. The group discussion was facilitated by a researcher for the top 3 and bottom 3 behaviours, and once those were agreed participants were allowed to continue with the group ranking activity unassisted. Once all participants were in agreement, the expert agreed list was shown to the group and they were allowed to discuss discrepancies both with the group and with the facilitator. Finally, participants were debriefed and allowed to go.

We collected the ranked list of behaviours for every participant (n=20) and the ranked list of each group (n=5) as well as the qualitative discussions during the group ranking activity (n=5).

Once all data was analysed, experts were briefed on the findings during the Reflection Stage (see below).

3.3 Reflection

The purpose of the reflection stage was to brief the organisational security experts on the findings from the workshops and collect their thoughts on the process and understand their reaction to the findings. In total, the session lasted 45 minutes.

Half of the session consisted of an oral presentation describing the methodology of the Task, a reminder of their rankings, and an overview of the main findings including the graphs in Figure 3.

A brief physical report was generated for the experts that summarised the purpose of the Task, the methodology used to collect the data, and the most salient findings (see Figure 2 for example). The main section contained a table that included each behaviour (ordered according to the expert ranking), the individual range for the employee scores, the individual mean rank for the scores, the group mean rank, and the expert rank (for easy comparison). The key reasons for the overall scores were also included. The behaviours were highlighted where the individual and group scores were markedly different – in green if the change resulted in a higher score, or red if it resulted in a lower score. In this specific case, different tables were created for academic and support staff to highlight the differences between the groups. Finally, a section with the main takeaways (summarising the most controversial opinions or differences) closed the report.

Following the presentation, experts were engaged in a brief semi-structured interview where they were asked to comment on the Cybersurvival Task and reflect on the findings. Experts were also encouraged to seek clarifications on conflicting behaviours and were asked about future actions based on the presented data.

4. RESULTS

Below we present both quantitative and qualitative results from Phase II, including insights from both employees and experts. The quantitative data was analysed by averaging the scores across groups for each behaviour (e.g. Ask for Advice). All tests carried out were two-tailed. The qualitative data was obtained from the employee discussions during the group ranking activities and was analysed using thematic analysis.

4.1 Comparison of Rankings Between Experts and Staff

The rankings of experts were plotted against those given by staff (academic and support). These are presented in Figure 3. The identity line (dotted) shows perfect calibration between experts and staff. However, the further away the behaviours are from the identity line, the bigger the discrepancy between staff and experts’ security priorities. Behaviours above the identity line represent those that are most important to staff, while those below the line represent behaviours that are most important to experts. Figure 3 also shows the difference between those rankings made as individuals and those made following group discussion with arrows indicating the shift between mean individual and group scores. One important thing to note here is that group discussion seldom moves staff towards better agreement with the experts. This is important given the way that social norms can intervene in determining staff security priorities (e.g. [35]). This will be explored in more detail below.

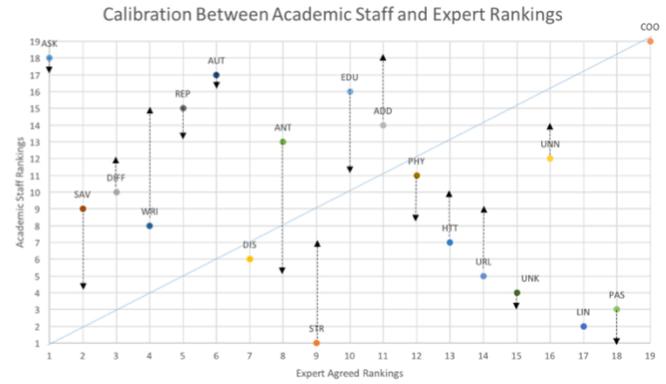
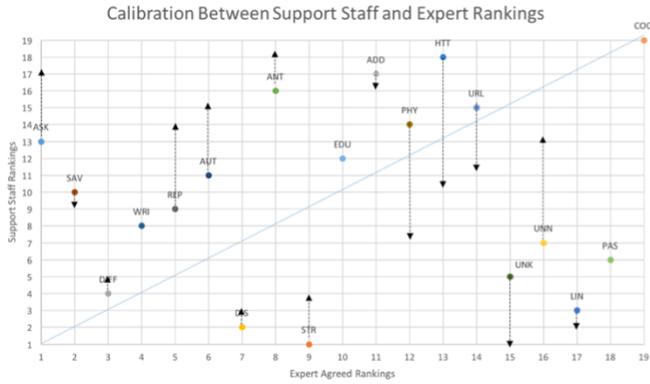


Figure 3: Scatter Plots comparing the rankings of experts (X-axis) against academic staff (right) and support staff (left). Arrows show the shift from mean individual rankings to final group rankings (dots). See Appendix C for high quality graphs.

4.2 Discrepancies in Rankings

At first glance, our graph shows poor calibration between experts and staff, with experts wanting staff to prioritise *asking for advice* and with staff (both academic and support) opting for the creation of *strong* passwords as their number one priority (again, an interesting issue in the light of recently shifting password advice).

Strong passwords have been the subject of many campaigns and are one of the few ‘engineered’ behaviours that staff are likely to encounter (as password systems often force the inclusion of upper and lower case, numerical and special characters as a means of creating stronger passwords). There is a certain irony here, given the new advice of *three random words* issued by GCHQ. Interestingly, support staff were more aware than academic staff of the need to use different (unique) passwords for each account, something that could possibly be tied to their use of systems that could hold sensitive data (e.g. finance, student performance, student identity, etc.).

A set of behaviours around not opening attachments or links from unknown senders were also seen as very important by academic and support staff but not so much by experts (who tended to place more trust in automated detection of malware). Similarly, checking URLs and checking HTTPS (to a lesser extent) were seen as important behaviours by the academic staff while they were rated low by the experts, although these behaviours were ranked as being less important after the group discussion and more in line with the expert scores.

Experts prioritised *asking for advice* as the single most important behaviour, but this was very poorly ranked by staff who generally believed that asking for advice was unnecessary and they could not envisage a scenario when that would happen (see below). *Reporting data loss* and *turning on automatic updates* were also seen as less important by academic staff when compared with the expert agreed rankings.

4.3 Individual vs. Group Rankings

In order to understand the differences in individual and group rankings, we calculated the absolute difference between the staff scores and the expert agreed scores (i.e. expert ranking minus staff ranking) for each of the subgroups: academic staff’s individual rankings, academic staff’s group rankings, support staff’s individual rankings, and support staff’s group rankings. The lower the added score, the closer the rankings were to the expert ones (0

= perfect, 361 = complete opposite). We then ran a Wilcoxon Signed Ranks Test between the individual scores and the group scores to measure any significant changes to ranking scores that emerged as a function of group discussion (in relation to the expert ones).

Table 4: Mean scores in the Cybersurvival Task (0 = perfect score; 361 = worst score).

	Individual Score	Group Score
Academic Staff	131.13	131
Support Staff	117.11	162

We found no significant difference in individual and group rankings for academics, $Z=-.140, p=.889$. We did, however, find a significant difference in individual and group rankings for support staff where individual scores were higher (i.e. more secure) than group scores, $Z=-2.668, p=.008$.

These results are worrying as they show that a group discussion about the importance of different cyber behaviours led to a weakening of the support staff’s cybersecurity position (i.e. more disagreement with experts). This finding is also reflected in a statistical comparison of academic and support staff, where a Mann-Whitney U test did not find a statistically significant difference in the performance of *individuals*, $U=21, p=.167$ but where there was a significant difference in *group performance*, with academic staff generating rankings that were much more closely aligned with experts, $U=10, p=.011$.

4.4 Expert Assumptions

Here we report some of the qualitative data from the discussions within the different groups. We start by detailing some of the assumptions made by security experts about staff behaviour. Firstly, experts were adamant that there was an onus on employees to learn about security threats and to educate themselves. This notion was thoroughly rejected by our employees who felt that such behaviour would be too time consuming:

Academic Group 1 (Male): “Yeah, I think it’s one of the things on my list that, I would really like to do, but you never get the time to actually get

around to it. Presumably the fraudsters are getting cleverer and cleverer, so you have to keep up to date with new ways of helping and keeping yourself stay safe”

This is possibly one of our most predictable findings, given the extensive research literature on ‘productive security’ that notes the unrealistic and unacceptable ‘cost’ of cybersecurity policy compliance [7].

Secondly, experts assumed that users would save all their work regularly to the network drive in order to allow immediate restoration in the case of infections or attacks. In reality, this was common practice, but many staff chose convenience over security and downloaded a local working copy, which would then be uploaded to the network once access was no longer required.

Support Group 2 (Female): Force of habit, it’s just a habit. I don’t not for any particular reason, just lazy I guess ‘cause it saves me the click for going into that, then going into that – and instead I’m like ‘it’s there on the desktop’.

Support Group 2 (Male): A lot of the time for me it’s something that I’ll only need access to for a limited time so once I’m done with it I’ll just delete it.

Finally, experts believed that users would report any data breaches immediately. Employees, however, questioned how they would know if a data breach had occurred:

Support Group 2 (Female): “But how do you know that you’ve lost something? I’m not sure I would recognise a data loss unless it said to me, ‘you’ve lost some data’.”

This assumption highlights an important problem for experts: employees do not possess a concrete understanding of the consequences associated with cybersecurity – e.g. what actually happens when you have suffered a data breach? A possible remedy would appear to be for experts to contextualise advice and policy in order to encourage compliance.

4.5 Employee Misconceptions and Disagreements

Next, we explore employees’ misconceptions about security behaviours and their failures to come to any agreement about appropriate actions. Firstly, staff believed that software updates – whether applications or an operating system – were primarily a means to access new features, arguing that updates could be delayed without any adverse impact, a finding previously reported by Vaniea et al. [59].

Additionally, they erroneously believed that if the update was important, it would get pushed through by the IT staff regardless.

This misconception is in line with work showing how updating software was rarely seen as a key security behaviour [60].

Academic Group 2 (Female 2): “I’ve got the turn on automatic software updates, because I thought software was quite general and there’s the other one that covers the anti-malware software – so any software updates could be anything. Uhm, that’s why I thought it was not specific to internet security”

There was extended discussion regarding the threats from email attachments and links. While staff were generally aware that clicking or downloading items from emails could harm their computer, the exact nature of the harm was disputed. Some employees believed that links were more dangerous than attachments as clicking them automatically compromised the computer, while others argued that attachments were harmless if you did not allow them to install. While most points argued were true to an extent, it was worrying how varied their perspectives of the threats were.

Academic Group 2 (Female 1): “But if you opened it, I wouldn’t anyway, but open an attachment from someone I didn’t know – I would just delete it – but if I did open it I would assume that unless I clicked on a link within that attachment then the attachment couldn’t, unless, you know like a Word attachment, if they sent me some kind of attachment that could be actually downloading a virus.”

Academic Group 2 (Male 1): “I think an attachment is more important because that’s a file that you download to your computer and could potentially run directly on your computer”

Academic Group 1 (Female 1): “to actually open an attachment itself may be important, because I know that you don’t need to put your password in and malware starts to come, and there are many of those everyday. So if we put that as a priority behaviour then we can prevent a lot of malware from coming in. And it’s very simple as well – that’s my opinion.”

Academic Group 2 (Female 2): “The more that I talk the more I realise I don’t know”

This last observation is important. Employees lacked a good mental model of the nature of the threat and the way that they could realistically guard against it. This led to disagreements about the most effective forms of protection. For example, there were heated discussions about writing passwords down, with the majority of participants agreeing that it was a ‘must not do’ behaviour and should be avoided at all costs. In the meantime, password reuse was seen as a negative, but necessary, behaviour – especially given that mapping personal accounts to work accounts would be difficult for attackers. This demonstrates a mental model where most staff prioritise the need to protect themselves against colleagues rather than against external threats.

“Academic Group 2 (Female 3): See it’s funny because I’ve put keep passwords safe if written down before I’ve put use strong passwords because obviously if you have it written down it doesn’t matter how strong it is – people can get it.

Academic Group 2 (Male 1): But if you’ve got it written down there is maybe only a handful of corrupt people who could get their hands on it...”

While previous literature (e.g. [56]) has reported this as a flawed mental model, other work [14] argues that this might not actually be a serious security threat, while Zhang-Kennedy et al. [62] suggest that this rule should be changed, promoting the keeping of written down passwords secure. Again, these academic disagreements demonstrate the difficulties with generating security advice. Ultimately, both GCHQ and NIST have taken the stance of promoting secure storage of written down passwords in their new guidelines.

4.6 The Sources of Guidance

We now look at some of the issues around where employees would turn to for education and guidance. Firstly, as we noted, participants were reluctant to ask experts for advice as they felt it was time consuming and unnecessary. They seldom knew who they could turn to for advice either within the Department, the Faculty or the University. Participants generally agreed that learning from each other or from their own personal experience was more realistic than asking for advice from an expert:

Academic Group 2 (Male 1): “I think people are more likely to ask their immediate colleagues for advice about things.”

Support Group 2 (Female 1): “Would you not just tend to ask for advice once you’ve done something wrong or something bad has happened?”

This is a problem when local knowledge is based upon poor mental models of both threat and effective deterrence. The tendency to rely upon peers and to trust social norms is a known problem in cybersecurity research, leading to the development of shadow security cultures within an organisation [36]. We know that teams do have an important role to play in the development of security behaviours, but we also know that these teams can appropriate security behaviours and practices, moulding them to better fit their own work context, but occasionally introducing vulnerabilities and misconceptions as a result [47].

Our employees felt that they could not be expected to stay on top of the latest advice and information. They were aware of certain ‘rules’ (such as not opening attachments and clicking on links) but they felt that they should not be held responsible for cyber defence as they could not be expected stay current with that knowledge and were unwilling to put extra time into learning.

Support Group 2 (Female 3): “Yeah, just come and ask us – spend an hour educating you. I mean, nobody has that time, so...”

Finally, employees recognised that certain issues were out of their control. They believed that the IT department was responsible for

cyber defence and that this defence was primarily undertaken with automated detection and control systems. This is an interesting issue as it reflects the kinds of culture that evolves around staff who have restricted access in relation to installing or updating software. Knowing that IT services have control over such matters brings with it the assumption that staff have no real responsibilities in this area.

Academic Group 1 (Male 2): “So the anti-malware thing, because it’s the university computer I just take it that’s it’s all sorted out anyway. It’s not like you’re meant to keep it up to date yourself personally.”

Again, this speaks to the way that employees are empowered in the cybersecurity space. We know from the psychology literature on social loafing that in the presence of others, an individual user may not react to a request, assuming that others will make the required response [11, 22].

4.7 Feedback to Experts

The final step in the Cybersurvival Task was to present the findings to the university experts. Below we cover the lessons learnt from that session as well as feedback regarding the findings and the methodology.

Firstly, the experts were surprised at some of the misconceptions shown by employees. They had made assumptions that certain behaviours or terms were common knowledge, and the results of the exercise made them realise that extra effort was required to better understand their audience.

CISO: “It forced us to re-evaluate our desired behaviours. Because, I have, based on years of experience, developed a prejudice towards certain desired behaviours that I now think, based on this, perhaps I’ve allowed that prejudice to drive my own personal baseline. And I think this tool helps break that and forces me to re-evaluate my concept of desired behaviours.”

Secondly, experts took the output from the Cybersurvival Task as evidence that their one-size-fits-all training approach was failing the university.

CISO: “One size fits all is a fallacy. It’s not going to work. You need to cater your risk management programmes specifically to the people within their respective work areas. I think that’s what I’m taking from this.”

While this school of thought is not necessarily new for the academic security community, it is important to note that it is still being employed in organisations (this was a common finding across both Phases I & II). By utilising this tool, the CISO was able to make this realisation for himself and thus can seek more effective ways of promoting secure behaviours.

Thirdly, they argued that the task would be an excellent tool for establishing a baseline prior to undertaking training development and then using this baseline data to deliver more targeted training:

Faculty Deputy: “It’s critical because this provides a mechanism for determining, not to find out whether our programmes are successful, but whether our programmes are correctly designed and catered for the intended audience. Because that’s the initial hurdle. Because if the programme isn’t adapted for the culture then it will fail”

Finally, experts expressed their support for the Cybersurvival Task, focusing on the fact that the issues raised by the tool were specific to their organisation:

CISO: “I don’t think I’ve come across a tool that’s quite so powerful. I’ve come across metrics. I’ve challenged metrics, but this tool is different because it’s using my metrics that I’ve provided and compared them against other people’s metrics to see how they match and there’s no way I can argue against that data because it’s data that I’ve provided, as an individual, and data that other people have provided. I can’t see any weakness in there. I’m struggling to find a weakness. I think it’s a very powerful tool”

We note here, that one useful aspect of the Cybersurvival Task is that the output for different staff groups can be easily quantified in terms of the kinds of visualisations shown in Figure 3. This was important as it is not easy to use purely qualitative data to illustrate discrepancies between the beliefs of different groups, but we found these illustrations, used in combination with the discussion data, were very effective as a means of organisation-specific highlighting issues.

4.8 Summary of Findings

The Phase II deployment of the Cybersurvival Task in a large institution involving two organisational security experts and 20 employees demonstrated the benefits of this tool by highlighting differences between the cybersecurity beliefs and attitudes of security experts and employees. We also found that a group discussion around desired security behaviours actually led to *less* agreement between employees and experts, which raises interesting questions regarding the social construction of cybersecurity within workgroups and related issues of how best to disseminate security information in organisations.

A follow up session with the organisational security experts found that they valued the information uncovered by the tool, and they had a clear understanding of how that information could be used to improve their organisation in the future – for example in understanding what content should be covered in mandatory training sessions.

5. DISCUSSION

In this paper, we described two deployments of the Cybersurvival Task in two large universities and showed how the task revealed security misconceptions of staff and some behavioural discrepancies between security experts and employees. We specifically highlight how the organisation’s security experts were able to reflect upon flawed assumptions regarding certain employee behaviours, as well as realising how their approach to training was not fit for purpose. While the reported employee misconceptions

are not all novel, it is important for the organisational experts to know which security issues exist within their realm so that they are able to address problematic behaviours or beliefs. Importantly, the fact that the task has highlighted some well-known behavioural issues serves as a sanity check that participants were being truthful and that the task is externally valid.

Here, we discuss the benefits of using the Cybersurvival Task over other existing security behaviour measurement tools and explore how organisations can use the tool to improve training programmes, tailor their security policies, and understand the development of non-compliant attitudes and shadow security behaviours. We should note that the Cybersurvival Task has two quite discrete functions. Firstly, in keeping with the Desert Survival task and the Moon Landing task, the Cybersurvival Task can highlight individual and group opinion differences between staff groups and see how they are resolved. Secondly, the task can produce useful cybersecurity data about staff behaviours, understanding and possible compliance with security policies. We will explore these functions in more detail below.

5.1 Measuring Individual and Group Decision-Making

In terms of the first function – to observe the processes of individual and group decision making – it was very interesting to note the differences between groups within the organisation, but perhaps more intriguing to note that group discussion *never* resulted in more secure rankings overall, when compared to individual rankings. Indeed, in the case of support staff, group discussion resulted in a set of beliefs that were less secure (i.e. less aligned with expert opinion). Earlier we talked about this in relation to the development of a shadow security culture within the organisation in which social norms can come to dominate [36]. However, we should also note that this resonates with other studies using ranking tasks to measure group behaviour, when the dynamics of the group can result in sub-optimal decisions. For example, in a ‘Desert Survival’ study involving mixed gender groups, expertise tended to be ignored in group settings if the experts were women, resulting in poor group performance, but not if they were men [57].

While it is certainly interesting to observe the differences between individual and group scores, some may argue that cybersecurity is predominantly an individual task. We disagree given the social nature of organisations and the data suggesting that users are more likely to turn to colleagues rather than experts for advice. However, it is possible to build the visual representations (e.g. Figure 3) using the individual scores, although we would recommend running the group ranking sessions regardless due to the insights they generate (see below).

5.2 Measuring Cybersecurity Attitudes and Behaviours

In terms of the second function of the task – to measure cybersecurity attitudes and behaviours – we should ask how the Cybersecurity Task compares with other available measures. The most obvious point of comparison – albeit serving a different purpose – is the Security Behaviour Intentions Scale (SeBIS) which was initially developed in 2015 with the aim of becoming the standard tool for assessing the security behaviours of end-users [19]. This has since been validated to show how some security behaviours can be reliably predicted using the scale [18]. One of the interesting differences between SeBIS (and self-reporting

questionnaires in general) and the Cybersurvival Task is the request in the latter to *rank* behaviours, rather than indicate compliance level. There is no obvious ‘correct’ ranking and so we can attenuate the problem of ‘social desirability’ (giving the ‘right’ answers to questions irrespective of behaviour). Additionally, by having to justify priorities, participants in the Cybersurvival Task reveal underlying assumptions and/or flawed mental models that can then be used by experts to deliver appropriate remediation.

However, there are two further issues that come to light when comparing tasks. The SeBIS is not resource heavy – it can be completed quickly and can therefore give organisations rapid, actionable data about the beliefs and reported behaviours of their staff. In contrast, the Cybersurvival Task when done properly (involving both individual and group stages) can be quite resource intensive but also allows for training opportunities while also providing a baseline measure of security knowledge within the organisation. This is not a negative thing if it results in greater understanding and ownership of the problem. In addition, the SeBIS has a static set of items to be used in any organisation, despite the fact that there are always disagreements over what items should be included and prioritised depending on the context (e.g. [32]). In contrast, the Cybersurvival Task, as we have described it, sees cybersecurity as an evolving process and adapts this list to those set up by a specific organisation. At the beginning of our study, we asked the organisation’s CISO to generate 19 important behaviours and compared these with the 16 items on SeBIS [19] and the 20 “good” behaviours identified by Ion et al. [32]. There was a substantial overlap, but our CISO added certain behaviours (*ask for advice* and *educate yourself on how to avoid fraud*) which he ranked very highly. Staff did not prioritise these items and so it would be easy to argue that they were unimportant, but this would be missing the point. The Cybersurvival Task is designed to show differences between the beliefs and opinions held by the CISO and those held by employee groups throughout the organisation. Where there is disagreement, then there is an opportunity to consider whether staff communication has been adequate or whether expectations are unrealistic.

5.3 How Can Organisations Benefit from the Cybersurvival Task?

It is clear from our expert feedback session that security experts in organisations make assumptions about their institution’s security culture and that these assumptions are not always correct. This means that organisations may not be providing staff with the necessary and/or relevant training programmes. While the Cybersurvival Task does not measure employee compliance – it is possible that employees engage in all behaviours on the list – it can be used to obtain a snapshot of security subcultures within an organisation, and to identify any misinformation that might be circulating in those subcultures. This would allow experts the opportunity to tailor solutions that would help prevent the proliferation of non-compliant security practices.

The Cybersurvival Task can also serve as a sanity check for an organisation’s security policies. During the first step when the organisation’s security experts modify and rank the list of behaviours, they can identify any policy items that may no longer apply, or others that they may not have considered before. Additionally, this process should make experts aware of what the most important message to staff should be. The act of having to rank a particular behaviour as first or second on a list can give pause for thought – how are these important behaviours being

communicated to staff across the organisation? Note, too, that rankings may change in keeping with the dynamic cybersecurity threat landscape.

Lastly, it may be possible to use the Cybersurvival Task as a training tool, exploiting the way it can readily highlight misconceptions and promote discussions about why the experts prioritise certain behaviours and why staff might find these behaviours challenging to execute in their own work contexts. While such an approach would require a greater degree of co-ordination (e.g. scheduling for both employees and experts), the direct outcome with regards to mutual understanding by both parties would seem to be beneficial. The task certainly generated high levels of engagement across all groups – something which is not always said of cybersecurity training material.

5.4 Limitations and Future Work

The most obvious limitation regarding this implementation of the Cybersurvival Task related to the time taken to conduct the workshops and collate and present the findings. Despite our participants finding it an enjoyable task, we do recognise that length could be an issue for both organisations and individuals. In future settings, the individual rankings could be completed online and analysed before the group meeting to discuss differences and agree a consensus ranking (thus speeding up the process). We are hesitant to suggest running the complete task online as it is currently presented, as this would miss out on valuable qualitative data that shows the reasoning behind the rankings and reveals any underlying misconceptions or erroneous mental models that management can then address. However, it may be possible to redesign some of the activities (e.g. the group ranking task) to accommodate digital technologies for carrying out the workshops in a distributed manner and reducing the time taken to complete them.

We also recognise that our deployments have been restricted to academic organisations and so, in future work, we aim to take the tool into other sectors, streamlining some aspects of the data collection process, and exploring the automatic generation of reports.

5.5 Conclusions

In this paper, we have shown that security experts and staff do not always agree on the most important security behaviours and this will be a big concern for organisations. Ideally, all members of an organisation should be working towards the same security goals and should understand their role in achieving those goals, yet we have found that group discussions on cybersecurity behaviours in fact led to more disagreement between staff and expert priorities. We have shown that a simple ranking task, conducted individually and then in groups, can highlight such disagreements and illustrate the different normative beliefs held by specific staff groups as well as illustrating the differing priorities shown by security experts and employees at different levels of the organisation. We believe the Cybersurvival Task would be useful for any CISO seeking to understand the kinds of sub-optimal security subcultures that develop within their organisation.

6. REFERENCES

- [1] A Hacker’s Dream: American Password Reuse Runs Rampant: 2017. <https://www.infosecurity-magazine.com/news/american-password-reuse-runs/>.

- [2] Adams, A. and Sasse, M.A. 1999. Users are not the enemy. *Communications of the ACM*. 42, 12 (Dec. 1999), 40–46.
- [3] Alotaibi, M., Furnell, S. and Clarke, N. 2016. Information Security Policies : A review of Challenges and Influencing Factors. *The 11th International Conference for Internet Technology and Secured Transactions* (Dec. 2016), 352–358.
- [4] Ashburn-Nardo, L. and Johnson, N.J. 2008. Implicit outgroup favoritism and intergroup judgment: The moderating role of stereotypic context. *Social Justice Research*. 21, 4 (Dec. 2008), 490–508.
- [5] Bada, M. and Sasse, M.A. 2014. *Cyber Security Awareness Campaigns Why do they fail to change behaviour ?*
- [6] Baxter, J. 2002. Jokers in the Pack: Why Boys are More Adept than Girls at Speaking in Public Settings. *Language and Education*. 16, 2 (Jun. 2002), 81–96.
- [7] Beautement, A., Becker, I., Parkin, S., Krol, K. and Sasse, M.A. 2016. Productive Security: A scalable methodology for analysing employee security behaviours. *Proceedings of the Symposium On Usable Privacy and Security (SOUPS)* (2016), 1–18.
- [8] Beautement, A., Sasse, M.A. and Wonham, M. 2008. The compliance budget. *Proceedings of the 2008 workshop on New security paradigms - NSPW '08* (New York, New York, USA, 2008), 47.
- [9] Blythe, J.M., Coventry, L. and Little, L. 2015. Unpacking security policy compliance : The motivators and barriers of employees ' security behaviors. *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)* (2015), 103–122.
- [10] Bollmer, J.M., Harris, M.J., Milich, R. and Georgesen, J.C. 2003. Taking Offense: Effects of Personality and Teasing History on Behavioral and Emotional Reactions to Teasing. *Journal of Personality*. 71, 4 (Jun. 2003), 557–603.
- [11] Briggs, P., Jeske, D. and Coventry, L. 2017. The design of messages to improve cybersecurity incident reporting. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* (Jul. 2017), 3–13.
- [12] Bursztein, E., Margolis, D., Archer, A., Pitsillidis, A. and Savage, S. 2014. Handcrafted Fraud and Extortion : Manual Account Hijacking in the Wild. *In Proceedings of the 2014 Conference on Internet Measurement Conference* (2014), 347–358.
- [13] Carlton, M. and Levy, Y. 2015. Expert assessment of the top platform independent cybersecurity skills for non-IT professionals. *Conference Proceedings - IEEE SOUTHEASTCON* (Apr. 2015), 1–6.
- [14] Cheswick, W. 2013. Rethinking Passwords. *Communications of the ACM*. 56, 2 (Feb. 2013), 40–44.
- [15] Cooke, R.A. and Kernaghan, J.A. 1987. Estimating the Difference Between Group Versus Individual Performance on Problem-Solving Tasks. *Group & Organization Management*. 12, 3 (Sep. 1987), 319–342.
- [16] Dembo, M.H. and McAuliffe, T.J. 1987. Effects of perceived ability and grade status on social interaction and influence in cooperative groups. *Journal of Educational Psychology*. 79, 4 (1987), 415–423.
- [17] Dudley, M.G. and Harris, M.J. 2003. To think or not to think: The moderating role of need for cognition in expectancy-consistent impression formation. *Personality and Individual Differences*. 35, 7 (Nov. 2003), 1657–1667.
- [18] Egelman, S., Harbach, M. and Peer, E. 2016. Behavior Ever Follows Intention?: A Validation of the Security Behavior Intentions Scale (SeBIS). *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (2016), 5257–5261.
- [19] Egelman, S. and Peer, E. 2015. Scaling the Security Wall: Developing a Security Behavior Intentions Scale (SeBIS). *Proceedings of the ACM CHI'15 Conference on Human Factors in Computing Systems* (2015), 2873–2882.
- [20] Fogg, B.J. 2002. Persuasive technology. *Ubiquity*. 2002, December (Dec. 2002), 2.
- [21] Furnell, S. and Thomson, K.L. 2009. From culture to disobedience: Recognising the varying user acceptance of IT security. *Computer Fraud and Security*. 2009, 2 (Feb. 2009), 5–10.
- [22] George, J.M. 1992. Extrinsic and Intrinsic Origins of Perceived Social Loafing in Organizations. *Academy of Management Journal*. 35, 1 (Mar. 1992), 191–202.
- [23] Giessner, S.R., van Knippenberg, D., van Ginkel, W. and Sleebos, E. 2013. Team-oriented leadership: The interactive effects of leader group prototypicality, accountability, and team identification. *Journal of Applied Psychology*. 98, 4 (2013), 658–667.
- [24] Giessner, S.R. and Mummendey, A. 2008. United we win, divided we fail? Effects of cognitive merger representations and performance feedback on merging groups. *European Journal of Social Psychology*. 38, 3 (Apr. 2008), 412–435.
- [25] Government Communications Headquarters 2015. *Simplifying Your Approach: Password Guidance*.
- [26] Grawemeyer, B. and Johnson, H. 2011. Using and managing multiple passwords: A week to a view. *Interacting with Computers*. 23, 3 (Mar. 2011), 256–267.
- [27] Guerin, L. 2007. *Smart Policies for Workplace Technologies: Email, Social Media, Cell Phones & More*. Nolo, Berkeley, CA.
- [28] Harris, M.J. and Perkins, R. 1995. Effects of distraction on interpersonal expectancy effects : A social interaction test of the cognitive busyness hypothesis. *Social Cognition*. 13, 2 (Jun. 1995), 163–182.
- [29] Herath, T. and Rao, H.R. 2009. Protection motivation and deterrence: a framework for security policy compliance in organisations. *European Journal of Information Systems*. 18, 2 (Apr. 2009), 106–125.

- [30] Herley, C. 2009. So long, and no thanks for the externalities. *Proceedings of the 2009 workshop on New security paradigms workshop - NSPW '09* (2009), 133.
- [31] Ifinedo, P. 2014. Information systems security policy compliance: An empirical study of the effects of socialisation, influence, and cognition. *Information and Management*. 51, 1 (Jan. 2014), 69–79.
- [32] Ion, L., Reeder, R. and Consolvo, S. 2015. “... no one can hack my mind”: Comparing Expert and Non-Expert Security Practices. *Symposium on Usable Privacy and Security* (2015), 327–346.
- [33] Kiahr, R., Shah, J.N., Sheriffs, P., Rossington, T., Pestell, G., Button, M. and Wang, V. *Cyber Security Breaches Survey 2017*.
- [34] Kirlappos, I., Parkin, S. and Sasse, M. 2015. Shadow security as a tool for the learning organization. *ACM SIGCAS Computers and Society*. (2015).
- [35] Kirlappos, I., Parkin, S. and Sasse, M.A. 2014. Learning from “Shadow Security”: Why understanding non-compliant behaviors provides the basis for effective security. *Usec '14* (2014), 1–10.
- [36] Kirlappos, I., Parkin, S. and Sasse, M.A. 2015. “Shadow security” as a tool for the learning organization. *ACM SIGCAS Computers and Society*. 45, 1 (2015), 29–37.
- [37] Kolkowska, E. and Dhillon, G. 2013. Organizational power and information security rule compliance. *Computers and Security*. 33, (Mar. 2013), 3–11.
- [38] Lafferty, J.C., Eady, P.M. and Elmers, J. 1974. The desert survival problem. *Experimental Learning Methods*. (1974).
- [39] Lee, C., Lee, C.C. and Kim, S. 2016. Understanding information security stress: Focusing on the type of information security compliance activity. *Computers and Security*. 59, (Jun. 2016), 60–70.
- [40] Li, I., Forlizzi, J., Dey, A. and Kiesler, S. 2007. My agent as myself or another. *Proceedings of the 2007 conference on Designing pleasurable products and interfaces - DPPI '07* (New York, New York, USA, Aug. 2007), 194.
- [41] Li, L., Xu, L., He, W., Chen, Y. and Chen, H. 2016. Cyber security awareness and its impact on employee’s behavior. *Lecture Notes in Business Information Processing* (Dec. 2016), 103–111.
- [42] Littlepage, G., Robison, W. and Reddington, K. 1997. Effects of Task Experience and Group Experience on Group Performance, Member Ability, and Recognition of Expertise. *Organizational Behavior and Human Decision Processes*. 69, 2 (Feb. 1997), 133–147.
- [43] McAninch, C.B., Milich, R. and Harris, M.J. 1996. Effects of an Academic Expectancy and Gender on Students’ Interactions. *The Journal of Educational Research*. 89, 3 (Jan. 1996), 146–153.
- [44] Ohtsubo, Y. and Masuchi, A. 2004. Effects of Status Difference and Group Size in Group Decision Making. *Group Processes & Intergroup Relations*. 7, 2 (Apr. 2004), 161–172.
- [45] Ovelgönne, M., Dumitras, T., Prakash, B.A., Subrahmanian, V.S. and Wang, B. 2017. Understanding the Relationship between Human Behavior and Susceptibility to Cyber Attacks. *ACM Transactions on Intelligent Systems and Technology*. 8, 4 (Mar. 2017), 1–25.
- [46] Pahlila, S., Siponen, M. and Mahmood, A. 2007. Employees’ behavior towards IS security policy compliance. *Proceedings of the Annual Hawaii International Conference on System Sciences* (Jan. 2007), 156b–156b.
- [47] Parkin, S. and Krol, K. 2015. Appropriation of security technologies in the workplace. *Presented at Experiences of Technology Appropriation: Unanticipated Users, Usage, Circumstances and Design* (Oslo, Norway, 2015).
- [48] Parsons, K., Calic, D., Pattinson, M., Butavicius, M., McCormac, A. and Zwaans, T. 2017. The Human Aspects of Information Security Questionnaire (HAIS-Q): Two further validation studies. *Computers & Security*. 66, (May 2017), 40–51.
- [49] Pattabiraman, A., Srinivasan, S. and Swaminathan, K. 2018. Fortifying Corporate Human Wall: A Literature Review of Security Awareness and Training. *Information Technology Risk Management and Compliance in Modern Organizations*. (2018), 142–175.
- [50] Phishing Across the Pond: 70% of U.K. Universities Impacted: 2017. <https://duo.com/blog/phishing-across-the-pond-70-percent-of-uk-universities-impacted>.
- [51] PwC 2017. *The Global State of Information Security® Survey 2017*.
- [52] Rus, D., van Knippenberg, D. and Wisse, B. 2010. Leader power and leader self-serving behavior: The role of effective leadership beliefs and performance information. *Journal of Experimental Social Psychology*. 46, 6 (Nov. 2010), 922–933.
- [53] Schneier, B. 2000. *Secrets and Lies*. John Wiley & Sons.
- [54] Sheng, S., Holbrook, M., Kumaraguru, P., Cranor, L.F. and Downs, J. 2010. Who falls for phish? A Demographic Analysis of Phishing Susceptibility and Effectiveness of Interventions. *Proceedings of the 28th international conference on Human factors in computing systems - CHI '10* (2010), 373–382.
- [55] Stewart, G. and Lacey, D. 2012. Death by a thousand facts. *Information Management & Computer Security*. 20, 1 (Mar. 2012), 29–38.
- [56] Stobert, E. and Biddle, R. 2014. The password life cycle: User behaviour in managing passwords. *SOUPS '14: Proceedings of the Tenth Symposium On Usable Privacy and Security* (2014), 243–255.
- [57] Thomas-Hunt, M.C. and Phillips, K.W. 2004. When What You Know Is Not Enough: Expertise and Gender Dynamics in Task Groups. *Personality and Social Psychology Bulletin*. 30, 12 (Dec. 2004), 1585–1598.

- [58] Toward a Group Facilitation Technique for Project Teams: 2007. <http://gpi.sagepub.com/content/10/3/299.full.pdf>. Accessed: 2016-03-15.
- [59] Vaniea, K.E., Rader, E. and Wash, R. 2014. Betrayed by updates. *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14* (New York, New York, USA, 2014), 2671–2674.
- [60] Vitale, F., Mcgreneere, J., Tabard, A., Wendy, M.B., Lyon, U., Paris-saclay, U. and Umr, C. 2017. High Costs and Small Benefits : A Field Study of How Users Experience Operating System Upgrades. *CHI 2017* (New York, New York, USA, 2017), 4242–4253.
- [61] Wash, R., Rader, E. and Fennell, C. 2017. Can People Self-Report Security Accurately? *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17* (2017), 2228–2232.
- [62] Zhang-Kennedy, L., Chiasson, S. and Van Oorschot, P. 2016. Revisiting password rules: Facilitating human management of passwords. *eCrime Researchers Summit, eCrime* (Jun. 2016), 81–90.

APPENDIX

A. Initial set of behaviours as presented to experts in Phase I.

The original list of behaviours was obtained from Ion et al. (2015). Below they are presented unranked as seen by the security experts in Phase I.

Behaviour
Be suspicious of links
Be sceptical of everything
Turn on automatic updates
Save passwords in a file
Clear browser cookies
Use a password manager
Use 2-factor authentication
Check if HTTPS
Look at the URL Bar
Install OS Updates
Don't click on links from unknown people
Use strong passwords
Use unique passwords
Don't write down passwords

Visit only known websites
Don't open email attachments from unknown people
Update applications
Don't enter passwords on links in emails
Use antivirus software

B. Ranked list of behaviours agreed by experts in Phase I.

Our two security experts from Phase I were given the opportunity to add, remove, and rename behaviours from the original list (Appendix A). Below is the final agreed rank list from experts for Phase I.

Ranking	Behaviour
1	Use strong passwords
2	Use antivirus software
3	Turn on auto software updates
4	Check every message is genuine
5	Keep OS up to date
6	Be aware of fake phone calls
7	Use different passwords
8	Be suspicious of links
9	Ask for advice when unsure
10	Check URL bar
11	Check if HTTPS
12	Don't download attachments from unknown senders
13	Don't enter password on website from link
14	Don't click links from unknown senders
15	Update applications
16	Only visit known websites
17	Don't write down passwords
18	Use a password manager
19	Use 2 factor authentication
20	Clear cookies

C. Scatter Plots Comparing Expert and Staff Rankings

Here we present the higher quality versions of the scatter plots from Figure 3. These are omitted from the paper due to space.

Calibration Between Support Staff and Expert Rankings

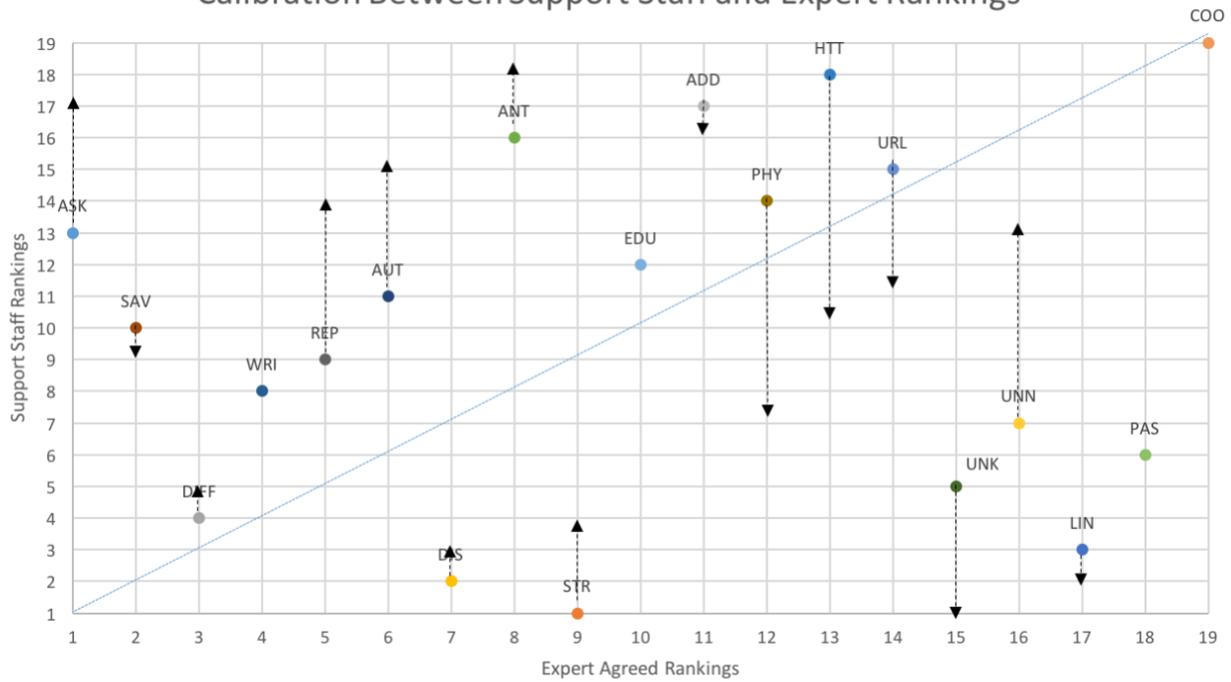


Figure C.1: Scatter Plots comparing the rankings of experts (X-axis) against support staff. Arrows show the shift from mean individual rankings to final group rankings (dots).

Calibration Between Academic Staff and Expert Rankings

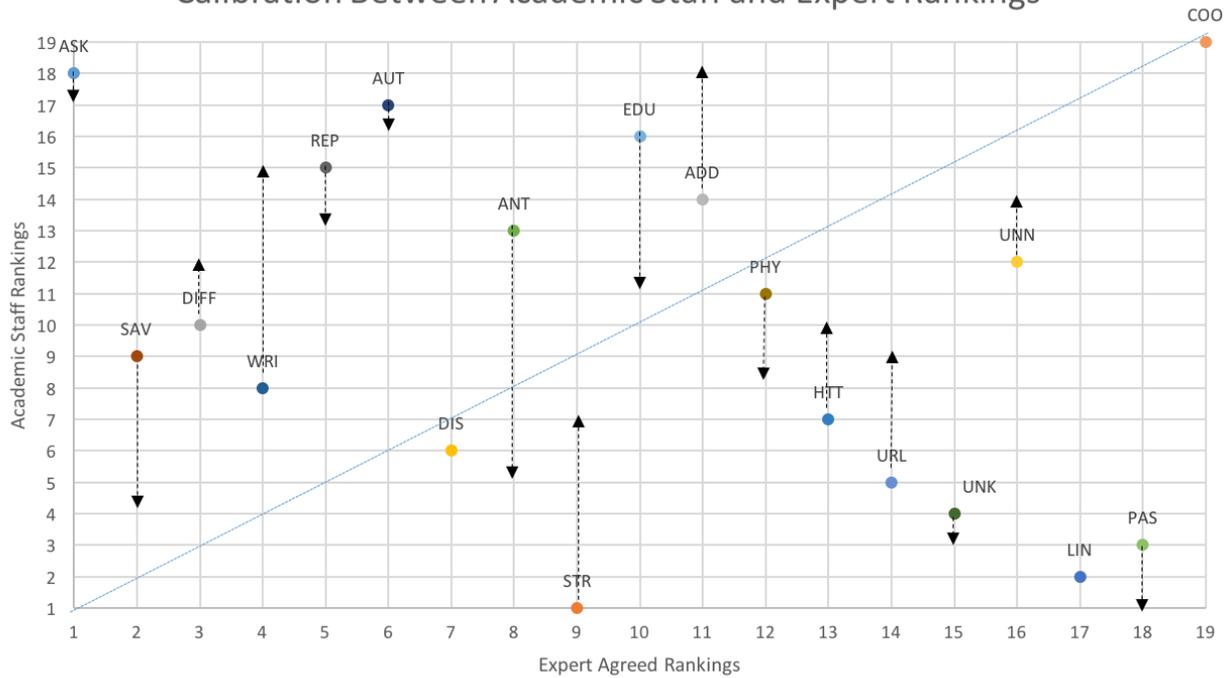


Figure C.1: Scatter Plots comparing the rankings of experts (X-axis) against academic staff. Arrows show the shift from mean individual rankings to final group rankings (dots).