A Technique for using Employee Perception of Security to Support Usability Diagnostics

Simon Parkin
Department of Computer Science
University College London
London, United Kingdom
s.parkin@ucl.ac.uk

Sanket Epili
Department of Computer Science
University College London
London, United Kingdom
sanket.epili.13@ucl.ac.uk

***Abstract*— Problems of unusable security in organisations are widespread, yet security managers tend not to listen to employees' views on how usable or beneficial security controls are for them in their roles. Here we provide a technique to drive management of security controls using end-user perceptions of security as supporting data. Perception is structured at the point of collection using Analytic Hierarchy Process techniques, where diagnostic rules filter user responses to direct remediation activities, based on recent research in the human factors of information security. The rules can guide user engagement, and support identification of candidate controls to maintain, remove, or learn from. The methodology was incorporated into a prototype dashboard tool, and a preliminary validation conducted through a walk-through consultation with a security manager in a large organisation. It was found that user feedback and suggestions would be useful if they can be structured for review, and that categorising responses would help when revisiting security policies and identifying problem controls.**

Keywords-information security; analytic hierarchy process; security policies; human factors of security

I. INTRODUCTION

Information security managers in organisations introduce security mechanisms as controls that users – typically employees – are expected to comply with (such as password composition rules, access cards to restrict entry to buildings, etc.). Users may be expected to comply with a range of security controls within their roles. These controls may be technical or procedural, but research has shown that they are not always effective [6].

When faced with a security control, users have a choice to comply or not to comply [2]. When deciding whether to comply, users are influenced by the perceived benefit and cost of compliance [3], for instance in terms of the associated effort. Individuals consider the benefit and the cost not just for themselves, but also for the people around them in the organisation [2]. If the effort of compliance is seen as too much, employees may avoid specific security controls altogether (as has been seen for authentication controls [18]). There is then a need for security managers in organisations to understand the costs and benefits of security as perceived by employees.

Recent research [11] has shown that even non-expert users can articulate their perception of good and bad security controls in relation to their role, and that there is scope to incorporate these views into the security design process. Improvements to security design can reduce the perceived workload of security for users in organisations [2]. User experiences are however generally ignored in the design of security, despite the importance of user involvement to system design [20]. Security managers can harness user feedback to improve the security of the organisation [11], by identifying pain-points and targeting interventions accordingly.

Here we present a methodology to structure user perception data at the point of collection, and support security managers in subsequent analysis. Users are prompted to perform pair-wise comparisons of the controls that affect them, comparing usability and usefulness. Follow-on questions further structure user perceptions toward identifying remediation activities, where diagnostic rules map these structured perceptions to recommendations emerging from the literature. This streamlines the use of management resources and introduces human factors of security into routine management practices, as seen in other work [18].

As a preliminary validation of the methodology, we incorporate the approach into a survey tool for end-users (populated with existing user perception data) and a dashboard tool for managers, which we present to a security manager in a large organisation. This also explores the potential of the methodology as part of routine management activities, ahead of any wider deployment of the tool.

Related work is discussed in Section II. We discuss our methodology in Section III. Tools incorporating the methodology are discussed in Section IV. Validation through consultation with a security manager is described in Section V. Discussion follows in Section VI, with conclusions presented in Section VII.

II. RELATED WORK

A number of works characterise user non-compliance with security in organisations. Deciding to comply or not to comply is a cost-benefit analysis by the user, weighing the perceived cost of compliance against the perceived benefit of that compliance [2]. Beautement, Sasse and Wonham [2] provide general rules, based on interviews with employees in

organisations, for improving the experience of security controls for users, where it is suggested that overburdening the user makes compliance less appealing and the management of security more costly. We propose directly leveraging users' perception of security as a metric to inform changes to security provisioning in organisations.

Users may also justify non-compliant behaviours through neutralising arguments (as in [9]), such as denial of responsibility, seeing non-compliant actions as necessary, blaming management and unreasonable policies, etc. Similarly, Kirlappos et al [6] categorised common reasons for non-compliance less in terms of personal justification but through the impact that security and policy have on them. Reasons include perceived high compliance costs, lack of understanding, and unavailability of compliance mechanisms, pointing to issues with how security controls are managed. These works collectively illustrate that the reasons for non-compliance and circumvention of security may not be obvious to security managers, and further that security mechanisms can hamper the primary task if user experience is not considered. Here we propose that user perceptions of security can be structured to assist security managers in crafting security provisions as an enabler of work.

Without engagement with users, security mechanisms in organisations are indifferent to their users [10], and so users themselves will act to align security controls with the primary task [6]. It has also been shown that heavy costs to the end-user encourage habitual, yet rational rejection of security [4]. Conversely, an understanding of compliant individuals can help organisations to improve productivity, and leverage compliant individuals to strengthen information security [3]. Other work suggests ways to consider the user in security management practices, such as by managing policy mandates and demands on user effort with greater care [13]. The work presented here identifies when controls are rejected or adopted by users, incorporating user perspectives into security management processes.

There are a number of works which model user security behaviour in organisations, focusing on security policies such as password use [16][17] and USB data transfer [15]. Model outputs – rather than user reports - inform management decisions, with data that exposes the drivers for visible user behaviour. Models are in some cases informed by engagements with users and discussion of their security behaviours, where here we propose directly eliciting and categorising feedback about specific controls. Other work has shown that security managers are amenable to introducing usability factors into their decision-making processes [18], and similarly that Analytic Hierarchy Process (AHP) can be used to weigh up alternative security solutions (such as a security manager's own view of available solutions) [14].

III. Methodology

Non-compliance is an opportunity for security managers to make security better [12]. Kirlappos et al [6] identified the following main reasons for non-compliance:

- *High Compliance Cost.* Non-compliance is deemed necessary by the user – they made a choice not to comply due to more pressing reasons.
- *Lack of Understanding.* Inaccurate perception of the risks associated with their roles and responsibilities, and misunderstanding of a technology's capability to support secure working.
- *Unavailable Compliance Mechanisms.* Provisioned security solutions do not fit the task.

By capturing the attitude of users towards security mechanisms, a security manager can gain insights into the reasons for both compliance and non-compliance.

Prior work has shown that employees' experience of security differs across groups of users in an organisation (see Section III.A), and that users can articulate views of security controls such as policies and technologies. Here we look at how findings from that prior work and similar research can inform the structuring of user feedback during elicitation, to produce data that the manager can use when looking or ways to improve security provisions (Section III.B). User perceptions of security are captured and categorised according to user-defined weightings of the usability and usefulness of controls. Diagnostic rules use research insights to guide remediation efforts by security managers (Section III.C).

A. Use of Existing Security Perception Data

Associated work analysed a set of interviews conducted with 118 employees holding various positions at a large multinational organisation [11]. Excerpts from these interviews are representative of how security is perceived by users.

The interviews were one-to-one, semi-structured, and approximately 50 minutes in duration, exploring employee interaction with – and sentiment toward – their organisation's security policies and controls. The interviews touched upon security awareness and compliance, including the impact of security on a person's role, views of the organisational support for security, and knowledge of the existence of security policies and mechanisms. The interviews were recorded, transcribed, and analysed by the work's authors, producing a preliminary thematic analysis that informs this paper.

The analysis is relevant to understanding the information employees might provide about security controls if prompted. The interviews also illustrate that employees are able to articulate the impact of security upon their work. For one, some participants elaborated on situations when following prescribed security behaviour resulted in slower completion of primary business tasks, or security mechanisms prevented completion of the primary task altogether.

Analysis suggested a general perception that the organisation did not listen to employee feedback, and did not respond to shortcomings in security adequately or in good time. Developing a capability to listen to employees and act on feedback informs improvements, but also conveys a dedication to maintaining good security for the organisation.

Findings from the interview study and related work have begin to characterise security experiences, and how user feedback may be collected, categorised, and applied to support management decisions in a more structured manner. We used extracts from the interviews as test user data for the methodology and tool (Section IV).

B. Categorising Controls

To categorise users' perception of security, we capture an employee's view of the *Usability* and *Usefulness* of security-related controls (any set of mechanisms, systems, and policies that factor into an employee's working day). Controls are rated against each other, to identify controls which may be a bad fit to the primary task, or their purpose poorly communicated. Controls that are successful can also be identified. The goal is to diagnose problem controls to investigate further, or indeed success stories to learn from.

Analytic Hierarchy Process (AHP) is used in the comparison process. AHP is a technique to make individual or group multi-criteria decisions [8]. Pairwise comparisons are made between choices, to identify the best alternative. Crucially, the AHP technique can be used to measure intangible psychological events, such as the beliefs of an individual [7]. This capacity to build metrics from beliefs is applied for weighing criteria relatively, where criteria are rated according to an individual's personal judgment. Successive pair-wise comparisons populate a pairwise matrix of relative weights, so that preferences can be quantified without having a unit of measurement.

In our method, we have two rounds of pairwise comparisons across the set of controls in an organisation:

- *Usability*: The user is asked to compare controls based on how usable they find them. Inversely this may be the Perceived Pain, as the user will experience competition between primary and security tasks – if security feels effortful, it can become a burden [2].
- *Usefulness*: The user is asked to compare controls based on the Perceived Benefit to them in their role.

Pairs of controls are compared along a bounded scale, where the centre position indicates that two controls are equally preferable, and a selection toward one end of the scale indicates a preference for the control associated with that side. In this work, users are encouraged not to pick the centre value – any distinction between two controls can help the manager focus remediation efforts.

Each control is compared to every other control. The user's choices populate a comparison matrix, and priority vectors of the user are created for both usability and usefulness. The control with the highest weight denotes the most usable (or useful) control, and the one with the lowest weight denotes the least usable (or useful) control.

AHP emphasises consistency in pairwise comparisons – that is, if control A is preferred more than control B, and B is preferred more than control C, then A must be preferred more than C. A consistency ratio is calculated from the values in the matrix. Saaty [8] argues that a matrix with a consistency ratio beyond 10% is not acceptable. Users are afforded some room for inconsistency here, as they are not assumed to be security experts in their own right. Responses lacking consistency are still useful to the security manager: users can be confused by controls or disengaged from them [18], or it may be that they are making guesses about controls that do not apply to them.

After every user completes the comparison ratings, controls are categorised according to the two comparison metrics, and targeted follow-up questions asked to capture comments relating to each control.

Each control is sorted into one of four categories based on their Pain and Benefit weightings, as in Table I. Three of these categories are based on the three non-compliance situations listed in [6] and reported in Section III. The fourth category is considered an ideal, where impact on the user is limited and the benefit of the control to a person's role is perceived to be high. Each category has three associated refinements that further categorise attitudes towards a control, where the refinements are representative of prominent views elicited during the prior in-depth interviews (Section II.A). The user selects a refinement statement and can provide further comments to add more context, should they wish to. The categories and refinements structure feedback for review by the security manager later.

C. Categorising Group Responses

Diagnostic rules were developed which collate individual responses and determine the relative perception of each control in a set, for groups of users and the whole organisation. Users are identified as far as being a member of a department, where it is assumed that a group of users has a similar set of primary tasks, and that it is large enough to ensure anonymity of user feedback. However, it is noted that issues can also develop amongst smaller teams within departments [11].

AHP can also capture group decisions, and so is used to generate overall Pain and Benefit values as a combination of responses from a group of users, not only for the entire organisation but also for each individual department. If the matrix values for a user are very inconsistent there is no concrete action that can be taken from the feedback itself, however inconsistent responses can identify where there is a disconnect between users and security. If this occurred across large groups of users, it would imply a "worst case" for streamlining remediation efforts, where users must be engaged directly and at length with no prior perception data.

We use the diagnostic rules to look for inconsistencies and potential explanations for values across the organisation and within departments, at the level of individual controls. We give corresponding recommendations on how the security manager can craft investigation activities, incorporating findings from prior interview-based research of the human factors of security in organisations [2][6][11][12]. These works illustrate that particular perceptions can indicate specific remediation activities without requiring in-depth interviews, affording a security manager opportunities to start aligning security with the needs of users. Structuring data according to these insights also limits the need for a security manager to be well-versed in the human factors of security. The diagnostic rules are outlined in Table II.

TABLE I. CATEGORIES AND REFINEMENTS

Category	*Pain*	*Benefit*	*Refinement*
Ideal	Low	High	I like it just because it is quite simple to comply.
			I like it because it is very important for security.
			I like it because it is both important for me and easy to use.
Could Comply, But Why Should I?	Low	Low	It is pointless.
			I do not mind using it, although I'm not sure of the point of it.
			It does not bother me that much.
Could Comply, But Cost too High	High	High	I understand why it is there, but it is not integrated well generally.
			The plan is good, but it is not implemented well.
			I am not sure where they went wrong, but it is irritating.
Something's Awry, Just Can't Comply	High	Low	It is useless in general.
			I do not see how it relates to my work.
			I know what can be done instead.

The diagnostic rules identify ways to apply user responses:

- *Inter-Control:* One control is compared to another. This can quickly identify controls with either far better or far worse *Usability* or *Usefulness* than other controls. When all controls are considered equal we cannot be sure if they are all equally usable or unusable, and so need a means to differentiate. In case of a *Usability* deadlock (where all controls have equal values), *Usefulness* is also taken into account. In cases of a *Usefulness* deadlock we compare the self-reported knowledge of controls (also captured).
- *Intra-Control:* Examine differences of perception for an individual control. The majority opinion for a control is compared with that of each distinct user group. There may be *Inconsistent Perception* across the organisation or within a particular department, and there is no clear majority opinion. This may occur where security micro-cultures have developed amongst teams [11]. There may otherwise be *Deviation*, when there is a clear majority opinion, but a specific group of users representing a minority view (e.g., one group may have reached or exceeded their perceived available effort for compliance based upon the controls they are tasked with using [2], or rationalised that controls are not worth the effort [4]).

The following excerpt is an example of an Intra-Control Inconsistency, where perception of USB data transfer policy differs between two members of the same department (from the same interviews that informed [11], as in Section II.A):

P32:
I: "Okay, are you aware of any policies around their use?"
P: "Yes, the major policy … is that they have to be a company [flash drive], they can't be an external [flash drive]."

P34:
I: "do you ever use USB flash drives at all?"
P: "I do."
I: "Do you know if there's any kind of policies around their use?"
P: "No, I don't."

Alongside organisation-wide categorisation, it is also possible to identify intra-control differences in specific departments. These differences may arise because of the work tasks performed within a department, or some other department-specific factor that promotes different attitudes.

The rules can suggest a range of information-gathering and remediation activities. This can include investigation of what it is that users appreciate about a "good" control (as it may be that traits can be copied to another control). If a control is provoking a range of reactions amongst users, direct engagement may be required to understand the cause (such as through in-depth exploratory interviews similar to those described in Section A). Enacting changes in technology or procedures without considering user experiences can perpetuate existing problems or cause further damage [11], and so rather than advocating specific controls, here we help managers explore the fit of controls to users' working lives.

Prior work suggests that reviewing the design of security measures is the most direct way to reduce mental and physical demands upon users [2], and that users are a useful source of information about how to improve those designs [11][12], so design activities and application of feedback drive the majority of diagnostic activities in Table II.

IV. IMPLEMENTATION

The methodology informed a prototype survey and dashboard toolset, built to investigate the viability of incorporating human factors thinking into security management. The tool has two components; the End-user Tool (for collecting user opinions of security controls) and the Manager Tool (which presents analysis and responses).

A. End-user Tool

As a running example, the security manager has chosen four security policies as the controls to gather opinions about – Password Policy, Email Policy, USB Policy and Clear Desk Policy. The following steps make up the process for the user:

1. *Initialising:* The user indicates their department.
2. *Knowledge Phase:* The user rates each control on how well they think they know them. This uses a set of indicative comments produced from the interview analysis described in Section III.A.

TABLE II. INTER-CONTROL AND INTRA-CONTROL DIAGNOSTIC RULES

Classification	*Situation*	*Reason*	*Recommendation*
Inter-Control			
Usability	Pain of Control-1 far higher than Pain of all other Controls	Control-1 has some inherent problem	Control-1 needs to be improved, or removed entirely if unworkable (use User feedback to identify actions)
	Pain of Control-1 far lesser than Pain of all other Controls	Control-1 has some likeable characteristics	Control-1 can be used as an ideal control (use User feedback to understand what is perceived as good)
	Pain of all Controls nearly equal	Controls are all equally usable OR equally unusable	1) A mixture of: a) Users are willing to comply – use control-specific feedback to determine if controls are usable OR unusable b) Integration of the Controls needs to better aligned with the culture of the organisation, examine comments c) There may have been changes in the business goals, talk to team managers 2) Revisit all controls, review all per-control responses
Usefulness	Benefit of Control-1 far higher than Benefit of all other Controls	Control-1 has high perceived benefit	1. The control integrates well with the primary task of users, and is a candidate ideal control, and/or; 2. Users appreciate the need for it, so communication around the control is good
	Benefit of Control-1 far lesser than Benefit of all other Controls	Control-1 has low perceived benefit	There is a need to improve communication or consider removal (review control-specific comments)
	Benefit of all Controls equal, 1. Knowledge of all high 2. Knowledge of all low 3. Knowledge is random	1. Users may understand the control enough for their role, OR not know at all 2. Perceived Benefit low for all 3. There is confusion about the full range of controls	1. Determine if users know the controls, OR if there are widespread misconceptions 2. Controls need to be communicated better 3. Take an altogether new approach to communicating controls, as there is a systemic issue
Intra-Control			
Inconsistent Perception	Two different Pain extremes for Control/Department	Control has different burden on different users	1. Recognise the primary tasks that are affected more than others 2. Revisit the grouping of users based on their primary tasks
	Pain distributed for Control/Department	Control not aligned well to users' primary tasks	Understand varying primary tasks of users, and improve control to accommodate new understanding (start with Department-specific comments)
	Two different Benefit extremes for Control/Department	Goal of the control is not clear	Communicate the goal of the control in a better manner
	Benefits distributed for the Control/Department	1. Control itself not clear 2. Communication is weak	1. Make the control more understandable 2. Revisit how the control is communicated
Deviation	Low Pain for a Department (Group) within a High Pain Control (Department)	Workaround possibly established within Department (Group)	Monitor for possible workarounds for the control, to improve the control
	High Pain for a Department (Group) within a Low Pain Control (Department)	Possibly low technical support for the Department (Group)	Explore views of the entire Department/Group in more detail
	High Benefit for a group within a Low Benefit Department	1. Naturally complying individuals within a Department 2. Good Control	1. Monitor to identify the group, and reward – others will copy good behaviour 2. Replicate the good qualities of the control
	High Benefit for a Department within a Low Benefit Control	1. Department understands the reason for the control better (This Control might be important for their primary task) 2. Department has a good security culture	1. Understand why it is important – control may complement primary tasks 2. Show recognition of the Department, perhaps through rewards

3. *Usability Phase:* Each control is compared to all others in sequence to assess Usability (e.g., Password against Email, Password against USB, Password against Clear Desk) – see Figure 1. When comparing n controls against each other, there must be a total of n(n- 1)/2 comparisons. There will be a total of n-1 such steps to complete all possible comparisons. The number of controls being reviewed then influences the length of the process, and in turn how usable the tools are for security managers themselves.
4. *Usefulness Phase:* Follows the same structure as Step 3, but to inform the perceived Usefulness of controls.
5. *Follow-up Phase:* The controls (or policies, as in our running example) will each have been assigned one of the four categories as per Table I (see Section III.B). In this last phase, further follow-up

questions are asked to narrow the reasoning for the ratings. The user is presented with three refinements corresponding to a control's category. On selecting a refinement, users can provide further free-text feedback. A lack of feedback from a user group could indicate that they are too busy, or that they are deliberately limiting interactions with security.

An example of "ideal" user feedback for Password Policy could be as follows (extracted from the interview set described in Section III.A), where the user discusses password reset time and password complexity in fine detail:

"I wish it was sort of 3 months rather than every month. I think…I'd rather have a longer password …it's still is it 8/9 characters long /yeah/ and a mixture of uppercase [and] lower case and all that sort of stuff. You know, I'd rather have a 10 figure password and change it every 3 months, than an 8 one and change it every month"

B. Manager Tool

The manager can use this dashboard to configure the user tool, analyse user responses, and view recommendations. Navigation of user responses is by Control or by Department.

The manager can choose to view perception data for controls by Usability, Usefulness, or Both. There is also a drop-down list to compare against a particular department. Bar charts are available that show overall weights for the selected criteria, including series comparing overall data with that of a specific department (see Figure 2). This supports use of the data by security managers; summary data can provide a top-level view, and "drilling down" opens up more detailed information.

A 'Show Issues' button highlights control-specific diagnostic recommendations, otherwise organisation-wide recommendations are shown. If a particular department is selected, the recommendations are shown for that department. Recommendations are presented as a Reason-Action pair, where for each recommendation a reason and follow-up activity is suggested, as per Table II.

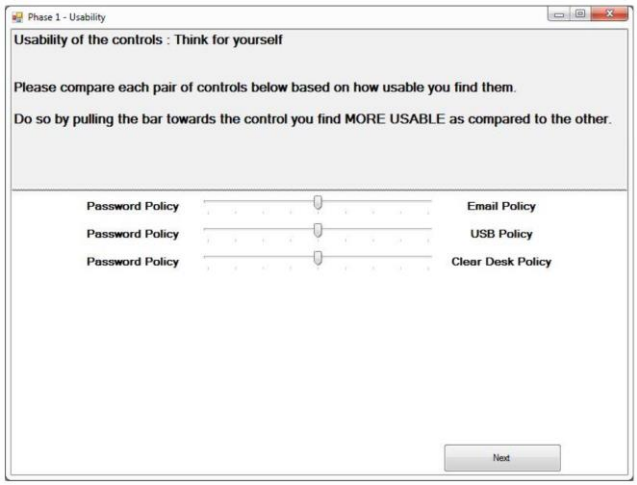

Figure. 1. Pairwise comparison of controls in the End-user Tool.

When a manager selects a Category, a pie-chart is displayed representing the quotient of users that chose each of the refinements in that Category, serving as a high-level indication of users' attitudes towards a specific control. User comments for the selected control are displayed alongside the chart.

V. VALIDATION

A. Validation Approach

As a preliminary validation of the approach, a security manager in a large organisation (many thousands of employees and contractors) was shown the User View (Section IV.A) and Manager View (Section IV.B) - where the tool had been populated with test response data - and asked a series of questions to guide discussion, including:

- Are there any tools or channels through which you already approach the problem?
- Do you believe users have enough knowledge of controls to provide answers that can be useful?
- Could individual recommendations be acted upon?
- How willing are you to go through comments to understand the basis for complaints or suggestions?
- How willing would you be to use feedback and recommendations to tackle compliance issues in your organisation?
- Can the tool link in to any existing (regular/irregular) security management activities?

B. Results

1) Current Practice

The security manager has previously used surveys sent to users to gather information about their views on security, but this happens irregularly, if at all: *"main reason for that is time, no time or resources … but it's definitely something we'd like to do"*. Instead, the manager might take a guess at what to do about security problems, rather than do nothing, though in fact *"it's the taking action bit that takes the time"*.

Referring to the diagnostic rules, and specifically learning from a "good" control: *"you'd have to find out why they liked it, but it's definitely worth doing"*. On maintaining security policies that affect users, *"to go through the whole lot it would take a year, and by the time you've got through them all another year has passed"*.

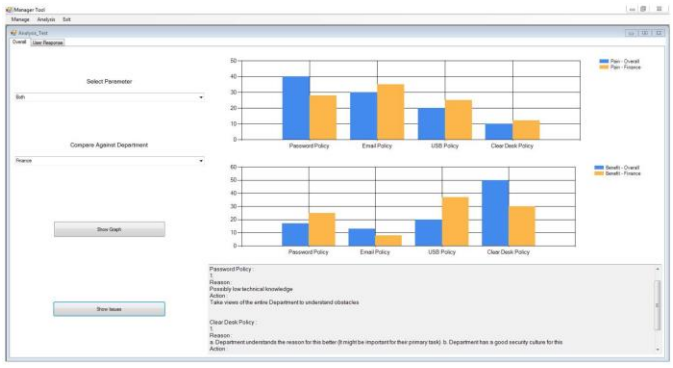

Figure. 2. "Pain" and "Benefit" bar charts in the Manager Tool.

2) Value of User Feedback

The security manager believed it is *"definitely better to ask the users what the problem is, rather than making a guess"*, and that *"at the moment it's quicker and easier to make that guess"*, although *"you can guess, but I'm sure there are other aspects that aren't obvious"*. On the value of user suggestions for improvements, *"they're the ones that have to use it, so in a way they kind of know best really, they give a perspective that chances are we wouldn't have thought of"*, although *"I'm sure there'll be some ideas that aren't doable, but you never know"*. Engagement with users was seen as beneficial for both the security manager and the employees in an organisation: *"Saves us from wasting our time creating something that isn't unusable or won't be used. It saves their time if we can make something easier to use."*

Regarding the visibility of non-compliance, *"it's impossible to know unless someone comes to tell you, so the only way to know is to ask people"*. It may be that *"people don't know [policy is] there or they don't know how to apply it, [so] something like this would help … in improving it"*.

3) Tool Features

The manager appreciated being able to see department responses alongside responses for the whole organisation.

Regarding collection and display of user comments: *"the comments are the most important thing, and you should try and encourage more to put their comments, because especially if you get lots of different people saying the same thing, you know there's clearly something wrong"*. Referring again to the issue of how long it takes to investigate non-compliance *… "it would be difficult to go through them if you get lots and lots but how else do you know how to improve something or why someone is finding something so difficult to use?"*

Categorisation of user feedback and comments was seen as useful, however *"it's worth splitting out on why they find it difficult to use, and their ideas on how to improve it, so … when you want to improve something, you can look at all the 'improvement' comments in one go"*.

4) Application of the Tool

The tool was not expected to see continuous use, by users or managers. Instead, the manager thought the tool (or something similar) could be used in the management of security policy, gathering user feedback *"a certain period of time before [updating policy] so that you could incorporate those findings into the updating process"*. Tying user feedback to the crafting of policy, *"If something happens that shouldn't have, because it was in policy, it would be useful to then see if people didn't know about policy, or if it was too difficult"*. Irregular events may also warrant use of the tool, such as any restructuring within the organisation, or when a number of new staff members join the organisation.

A tool like the prototype could help with understanding the root of specific compliance issues – the example was given of copyright infringement in the workplace: *"You usually only have to tell them once or block them once and they won't do it again. It would be interesting to know what peoples' opinions are on it first."*

The tool might fit into existing processes, rather than being isolated from other decision-support tools: *"[past] surveys were an online thing, whereas we might want to go into departments and talk to people face-to-face, so this could be integrated"*. However, *"we'd still have to go directly to people, things we ask them are a much broader thing, but we wouldn't have gone into this level of detail."*

It was suggested that the methodology and tool be repurposed to manage risks: *"wouldn't have been looking at just certain controls, security in general, data … rather the assets" … "user view on assets is important from a risk perspective … useful to know what they're doing and how we can help them do it better" … "say they're doing something with a certain type of data that they shouldn't be, identifying what they're doing, if it's wrong, and educating them"*.

The security manager also made the point that there may need to be a certain level of maturity in the security management process before using a tool such as this.

VI. DISCUSSION

Although the validation session involved only one security manager, a number of relevant factors were identified. There are costs to resolving issues with security controls, and yet "doing something rather than nothing" without consulting end-users has been a standard approach for many years [1]. Efforts to make user feedback more accessible and cheaper to gather can help to guarantee that remediation actions support users, rather than cause more problems.

Although there is interest in hearing user concerns, they may only be considered during policy review, which can be a lengthy process that does not happen regularly. Other work has proposed that considering the attitudes of users toward security should be considered during policy implementation to reduce non-compliance [2]. Applying the methodology during policy review would fit with interventions suggested elsewhere [2], such as making controls less difficult to use and communicating the applicability of controls to users' roles.

Manager comments supported the diagnostic rules in many ways, indicating that if many users are giving controls poor ratings that this alone is cause for action to be taken. There was support for identifying reasons for non-compliance before acting to fix problems, aligning with the intended use of the diagnostic rules to guide engagement with users. It is of note that some issues – such as copyright infringement – could be resolved once and no longer be a concern, however as was discussed, the hiring of new staff can mean that previously-resolved issues need then to be revisited, perhaps outside of regular policy review.

Potential was seen in using user feedback about security mechanisms, aligning with research showing that where employees may modify improper or inadequate security to produce their own alternative solutions, their rationale for doing so can identify improvements [11][12]. It was implied that there would always be some ideas that are unworkable, suggesting further ways to analyse comments and find value on user feedback.

A. Limitations

One limitation of the approach is that "Pain" may manifest for a user when the cumulative effort of using a set of controls adds together to become burdensome [2]. Free-text comments are then critical for characterising how "bad" each control is.

In a real organisation there may be countless distinct security controls (policies, technologies, etc.). As implied in Section IV.A, each pairwise comparison adds time to the process of gathering feedback, making the process potentially laborious. Managers would then need a rationale for focusing investigation on a limited set of controls, for instance according to risk assessments or asset manifests (as alluded to during validation, Section V). Minimising the effort of using the tool is important, to encourage managers to adopt the methodology rather than enact changes without consulting end-users. To that end, consultation with additional security managers could assess the benefits and costs of current approaches to end-user engagement and remediation of problematic security controls.

VII. Conclusion

Security managers in organisations need to understand how security affects end-users' work, and how users perceive security, toward investigating how to make improvements. We present a technique to enable security managers to capture user perspectives of security in a structured manner, as input to managing improvements. In addition, we present diagnostic rules that interpret collected data to produce guidance on how to proceed in investigation and remediation activities.

We demonstrate the technique as a dashboard tool, for both capturing user perception of security and to help a security manager navigate user responses. As preliminary validation the tool and the supporting diagnostic rules were discussed with a security manager in a large organisation, showing that it has potential to reduce problematic security deployments and align with existing management practices.

In future, our methodology would ideally see wider deployment and validation, to determine its correctness for a large user base and a larger set of security managers respectively. Comparative analysis of the technique alongside existing approaches for user engagement will inform what is gained and lost in data structure and detail, but also in capacity to represent actual behaviours (when compared to e.g., interviews [6] or combined interviews and site observations [21]). Future work will also use perceptions of security to quantify the effort of compliance and to anticipate when users feel over-burdened. Findings from wider research may potentially be incorporated to further structure the diagnostic process and assist security managers in navigating user responses, for instance to identify plausible solutions within user comments.

Acknowledgment

The authors are supported by UK EPSRC and GCHQ, grant nr. EP/K006517/1 ("Productive Security").

References

[1] Adams, A. & Sasse, M. A. 1999. Users are not the enemy. Communications of the ACM, 42, 40-46.

[2] Beautement, A., Sasse, M. A. & Wonham, M. The compliance budget: managing security behaviour in organisations. Proceedings of the 2008 workshop on New security paradigms, 2009. ACM, 47-58.

[3] Bulgurcu, B., Cavusoglu, H. & Benbasat, I. 2010. Information security policy compliance: an empirical study of rationality-based beliefs and information security awareness. MIS quarterly, 34, 523-548.

[4] Herley, C. So long, and no thanks for the externalities: the rational rejection of security advice by users. Proceedings of the 2009 workshop on New security paradigms workshop, 2009. ACM, 133-144.

[5] Kerckhoffs, A. 1883. La Cryptographie Militaire. J. des Sciences Militaires, IX, 5–38.

[6] Kirlappos, I., Beautement, A. & Sasse, M. A. 2013. "Comply or Die" Is Dead: Long live security-aware principal agents. Financial Cryptography and Data Security. Springer.

[7] Saaty, R. W. 1987. The analytic hierarchy process—what it is and how it is used. Mathematical Modelling, 9, 161-176.

[8] Saaty, T. L. 2004. Decision making—the analytic hierarchy and network processes (AHP/ANP). Journal of systems science and systems engineering, 13, 1-35.

[9] Siponen, M. & Vance, A. 2010. Neutralization: new insights into the problem of employee information systems security policy violations. MIS quarterly, 34, 487.

[10] Zurko, M. E. & Simon, R. T. User-centered security. Proceedings of the 1996 workshop on New security paradigms, 1996. ACM, 27-33.

[11] Kirlappos, I., Parkin, S. & Sasse, M.A. Learning from "Shadow Security": Why understanding non-compliance provides the basis for effective security. Workshop on Usable Security (USEC), 2014.

[12] Kirlappos, I., Parkin, S. & Sasse, M. A. Shadow security as a tool for the learning organization. ACM SIGCAS Computers and Society, 45.1 (2015): 29-37.

[13] Herley, C. More Is Not the Answer. IEEE Security & Privacy, 12.1 (2014): 14-19.

[14] Bodin, Lawrence D., Lawrence A. Gordon, and Martin P. Loeb. Evaluating information security investments using the analytic hierarchy process. Communications of the ACM, 48.2 (2005): 78-83.

[15] Beautement, A., et al. Modelling the human and technological costs and benefits of USB memory stick security. Managing Information Risk and the Economics of Security. Springer US, 2009. 141-163.

[16] Arnell, S., Beautement, A., Inglesant, P., Monahan, B., Pym, D., Sasse, A. Systematic decision making in security management modelling password usage and support. International Workshop on Quantitative Aspects in Security Assurance. Pisa, Italy. 2012.

[17] Shay, R. & Bertino, E. A comprehensive simulation tool for the analysis of password policies. International Journal of Information Security. 8.4 (2009): 275-289.

[18] Parkin, S., van Moorsel, A., Inglesant, P., Sasse, A. A stealth approach to usable security: helping IT security managers to identify workable security solutions. Proceedings of the 2010 workshop on New security paradigms. ACM, 2010.

[19] Sasse, M. A., Steves, M., Krol, K., Chisnell, D. The great authentication fatigue - and how to overcome it. HCI International 2014, 6th International Conference on Cross-Cultural Design. Springer International Publishing, 2014. 228-239.

[20] Checkland, P. B. & Poulter, J. Learning for action: a short definitive account of soft systems methodology and its use for practitioners, teachers and students, Wiley, 2006.

[21] Molotch, H. Everyday Security: Default to Decency. IEEE Security & Privacy, vol. 11, no. 6, pp. 84-87, Nov.-Dec., 2013.