Work in Progress: Usable Security vs. Workflow Realities

Jim Blythe
Information Sciences Institute
University of Southern California
blythe@isi.edu

Vijay Kothari
Computer Science Dept
Dartmouth College
vijayk@cs.dartmouth.edu

Sean Smith
Computer Science Dept
Dartmouth College
sws@cs.dartmouth.edu

Ross Koppel
Dept of Sociology
University of Pennsylvania
rkoppel@sas.penn.edu

***Abstract*—The usability of any security measure is often dependent on the environment and context in which it is deployed. A better understanding of context can help avoid a one-size-fits-all approach that can lead to security that is burdensome to use and does not address the most relevant vulnerabilities. A key aspect is a group's workflow — repeated group activities that selectively change the importance of vulnerabilities and that selectively restrict the time and cognitive budget available for security. Here we describe a number of case studies (drawn mainly from our own fieldwork) in which the workflow renders unusable a security approach that may be effective in other environments. We distinguish cases where the problem arises from individual tasks, from multiple paths through a workflow that may be unexpected, from time or cognitive stress introduced by the workflow and from barriers to passing needed information for the organization's mission. We present general approaches to design and improve upon security solutions so that they fit organizational workflow. Moreover, we discuss our ongoing efforts of conducting a broad cross-organization security-workflow oriented survey, cataloging and analyzing a wide-range of security failures and successes (many of which stem from workflow and security interactions), and agent-based simulation efforts.**

I. Introduction

In many cases, an organization's security posture, comprising deployed software, policies and recommendations for its individuals, is developed using a one-size-fits-all approach, combining existing off-the-shelf software with generally reasonable password policies and other regulations. This approach often fails, resulting in software that does not address the most pressing vulnerabilities of the organization and in policies that are hard to follow in practice and engender workarounds. The workflow of the organization is often a major reason for the poor fit. In some cases, repeated group activities have the effect of emphasizing some vulnerabilities over others, and also of selectively restricting the time and cognitive budget available for individuals to follow policies. In others, the workflow may allow for several alternative paths for achieving a task, yet security policy only addresses a subset of them.

Workshop on Usable Security (USEC) 2018
18 February 2018, San Diego, CA, USA
ISBN 1-1891562-53-3
http://dx.doi.org/10.14722/usec.2018.23037
www.ndss-symposium.org

In this paper we present a set of case studies that originate mainly from our interviews and surveys with both security professionals and end users. From this body of information we chose a small number of representative cases where the workflow of the organization interferes with the expected performance of the security tools and policies in perhaps surprising ways. These are examples of *mismorphisms* [26], instances where the beliefs of the security developer and of the users are sufficiently different that the developer's intentions are either unrealistic or not seen as appropriate by the users (and thus not followed). In these specific examples related to workflow, the developer typically does not know the circumstances in which the security apparatus is applied, leading it to be drastically reduced in effectiveness or so counterproductive for the organization that it is not used.

Our first set of case studies has the property that one or more activities are made costly by the new policy in a way that the security implementer had either not foreseen or had discounted against the gain in security. Cost here may be measured in a number of ways, including financial, the time required for individuals or lost opportunity. A knowledge of the enterprise workflow might help the implementer to estimate the cost of the policy changes, but the interaction with workflow is at a high level.

Our second set of case studies hinges on the effect of the multiple paths that may exist within a workflow to achieving the same objective. For example, a task that requires privileged access might usually be performed by regular workers, for whom two-factor authentication has been established as a norm, but the task is sometimes performed by external contractors for whom a simpler process is in place. Policy implementers can often catch vulnerabilities by looking for alternative paths to gaining privileged access that exist in the organization and ensuring that any new policy addresses each path. In our examples, the alternatives are sometimes exploited by attackers and sometimes lead to inadvertent vulnerabilities.

Our third and fourth sets of case studies both focus on specific types of workflow issues that are common enough to examine separately. In the third set we look at timing issues caused in environments with mobile, time-stressed users, who rapidly combine working on computers with moving away from the computers. Timed auto-logout solutions have been introduced in a number of different domains that have faltered on the time cost of frequent re-authentication, whether real or perceived. In the fourth set, information needs to be passed from individuals in one point of the workflow to another, but

security barriers to information flow are perhaps too broad, lacking sensitivity to subtypes of information or contexts where information must be passed rapidly. This typically leads to circumvention by frustrated workers, in turn leading to new vulnerabilities if not defeating the original information barriers.

Finally we present an example of positive interaction between the workflow and security policy.

In addition to the case studies, we discuss, in Section IV, approaches we think will serve useful in designing, adapting, and evaluting security solutions for workflow. We also discuss our ongoing efforts to aid in this pursuit.

II. Related Work

It is widely acknowledged that users circumvent security measures, not because they are evil, but because complying seems time-consuming, annoying, unreasonable, and it arguably gets in the way of more important things [1], [6]. We briefly review the literature on user compliance and workflow.

Many models have been constructed to understand non-compliance. Herley explained user rejection of seemingly sound security advice by using cost-benefit analyses that incorporated the negative externalities of following advice [13]. Florêncio et al. used an optimization model to explain password portfolio management strategies and made policy recommendations based on their findings [10]. Beautement et al. presented a model of non-compliance where users comply with security measures so long as expended effort does not exceed a compliance budget [4], and Anderson et al. created a logic for it [3]. We built a model grounded in semiotic triads for understanding the underlying causes of security problems, many involving workarounds, non-compliance, and workflow factors [26]. Safa et al. provided perspective and conducted a study applying social bond theory and involvement theory to information security compliance [24]. Kolkowska et al. developed a Value-Based Compliance method and design principles to help understand why employees may or may not comply with info-sec policies [17]. Dey et al. used sequential games to evaluate the efficacy of education vs. enforcement in preventing security circumvention, finding that enforcement alone is insufficient [7].

Simulations have also been built: Renaud and Mackenzie created SimPass, an agent-based simulation tool to understand the ramifications of password policy changes [22]. We argue that cognitive and behavioral agent-based simulation is useful in understanding and improving aggregate security [20], and continue to refine a password simulation based on this approach [21]. More broadly, Kavak et al. classify and discuss of a variety of security-focused simulations [14].

There's also been much work on studying the intersection of security and workflow through ground truth data gathered by ethnography, surveys, focus groups, and other means. Becker et al. analyzed survey results to understand how attitudes towards security policies and behavioral types vary across business segments within an organization [5]. Heckle conducted a 15-month ethnographic study of the development of an SSO system within a hospital, exploring the tensions between clinical workflow and security [12]. We documented many workarounds commonly employed in clinical settings to comply with security while getting work done [19]. AlKalbani conducted surveys and analyzed responses to determine the impact of various socio-organizational factors on user compliance [2]. Kirlappos et al. studied shadow security: workarounds to existing security measures that well-intentioned users develop to achieve what they deem a reasonable balance between doing their work and realizing perceived goals behind security measures [15]. They conducted interviews with employees from two large organizations and analyzed the responses to better understand what drove users to develop shadow security measures, to explore the implications of shadow security, and to draw lessons to inform security practice.

Our work here complements studies from the literature. We focus on the tension between workflow and security, conducting case studies across different workflow contexts to explore tensions between security and workflow and the underlying causes of security failures. We also reflect and discuss methods to design, evaluate, and adapt security solutions for workflow contexts, relaying some of our ongoing pursuits.

III. Case Studies

These case studies highlight the problem encountered and the generic security solutions employed, observations about why and how they failed, and suggestions for domain-independent methods to reduce the likelihood of failures of each type.

Many examples discussed in the case studies come from surveys with users, computer personnel and UI developers that we conducted over the past 7 years (*e.g.* [18]. (Survey instruments can be found at shucs.org.) We have 74 completed 8-page surveys from users and security personnel (most via survey monkey); we have over 40 face-to-face interviews with computer security leaders; and we have interviewed over 150 users about log-on and authentication issues–most focusing on the relationship of access to their workflow.

A. Policy interaction with individual workflow tasks

Our first case studies have the property that one or more activities are made costly by a new policy that the security implementer had either not foreseen or had discounted against the gain in security. A knowledge of the enterprise workflow might help the implementer to estimate the cost of the policy changes. In these examples, in contrast with those later in the paper, the cost increase to the individual from following the policy is the sum of the cost increases from each affected activity.

In interviews, we learned that a government establishment had such a thorough (and lengthy) process for approving a new computer account with network access that the time to receive an account was longer than the period of work for most summer interns. However, some of their tasks required access to a network. Interns connected to wi-fi that was available from local businesses such as coffee shops, with no authentication and without encryption, jeopardizing the assets that were intended to be protected by the security policy. A formal analysis of the organization workflow alongside the security policy would indicate that, at some point, a task requiring network access would need to be performed by an individual who had not yet gained authorization. If the authorization

process could not be shortened or started earlier, guest access to a controlled network would reduce the vulnerability, or tasks might be completed under supervision.

Our second example has a very different authority structure, coming from a manufacturer's agreement rather than a contract as a government employee. The license agreement, signed by purchasers of John Deere tractors and backed by embedded software, prohibits repairs not made by dealerships and authorized repair shops [16]. These facilities are typically not close to the farm, however, so this requirement can in practice lead to critical equipment being unusable for days longer than necessary, with high costs in lost productivity at critical times. As a result, a growing number of clients use cracked versions of the software purchased from the Ukraine to effect repairs locally without fear of the manufacturer disabling the equipment. In effect, the continued operation of the equipment is now in the hands of unknown hackers.

In both these cases, policies had a high-cost impact on one or more activities in the workflow that was deemed unacceptable by the user. While blame might be placed on users who take actions that have a clear risk and are in many cases explicitly forbidden, we note that the individuals feel their priority is to their tasks, and view the security consequences as secondary to the task. Information about workflow is often relatively easy to acquire and can be used to design policies that do not encourage workarounds. While these policies may be weaker than those originally intended, they will probably be stronger than the intended policy would be given the human ability to find workarounds [6].

B. Alternative paths through a workflow

Under certain conditions, a lax security protocol may supplant a standard, more stringent one. This may be in accordance with organizational policy or it may conflict with it, instead arising from human considerations; in both instances, it is often the case that these de facto differential protocols emerge from unique workflow considerations. While it may make sense to use a lax protocol under select conditions, it also may increase the attack surface for a skilled adversary who can simulate these conditions to exploit the lax protocol. Below, we discuss a few examples of differential protocols.

One example of alternative authentication protocols involves re-issuing hotel keys. As we have seen first-hand, some hotels issue a new key card based solely on room number. This may not be in accordance with hotel policy, but it often becomes the de facto protocol Indeed, Schneier [25] points out a very real consequence of such differential protocols: rape. The source link provided in the Schneier article is now down; yet, now, almost 10 years later, there continue to be many instances of rape arising from hotel card keys being issued without adequate vetting [23]. The workflow conditions that hotel employees face, deficiencies in the de facto protocol, and insufficient training may be contributors to this problem. Users lose room access all the time: they lose room card keys, they leave them in their rooms, and magnetic stripes fail. Employees may empathize with presumed occupants, try to avoid conflict, or simply desire to be helpful. Understanding this interplay between the organization's policy, human considerations, and workflow factors is key to implementing effective, practical security measures. Glass provides complementary discussion on engineering such an attack [11].

Accommodating the user who experiences an anomalous event is critical, although in select contexts it may paradoxically be routine. For example, many biometrics fail when the user experiences an accident, *e.g.*, a thumbprint reader fails when the user cuts their thumb. So there is often a back-up protocol in place. At one high-profile sporting event, admission to a sensitive area was supposed to be guarded by hand-geometry biometrics. Penetration testers we knew discovered a way to circumvent: wrapping a bandage around the hand that was to be scanned, and then using a fake ID instead.

For another example, password authentication protocols must acknowledge the reality that some users may lose access to their accounts because they cannot recall their passwords. Often, this results in a weak password reset or retrieval mechanism that can be exploited either by an adversary or a user who is infuriated with password management. Indeed this has happened in the past, often through social engineering. One example from informal interviews with Navy and Air Force Intelligence personnel involves a clever workaround to keep passwords in the face of required password resets. Every k days, a password reset would occur, requiring users to reset their passwords. However, if the user called the help desk right after the password reset, stating that they forgot the (new) password, the help desk would allow the user to again reset the password during a ten minute window. That window however induced a temporary amnesia of past passwords. That is, all the users' previous passwords were forgotten and the user was free to choose any password. Using this workaround, the user can indefinitely reuse the same password, despite mandatory password resets.

Heckle [12] documents many ways clinicians employ or even develop alternative workflow paths to do their jobs. Clinicians discovered that a password was required to enter time reports online, but not over the phone. In another example, Heckle mentions that while clinicians should scan a "bar-coded identification bracelet," they copied these bar-codes on a clipboard, scanning them instead when it was hard to scan the bracelet.

The diversity in the context under which alternative workflow paths are employed is illuminating. Sometimes workflow considerations find their way into organizational policy, resulting in two paths to the same end. Other times, the employee develops an alternative path to do their job. This path may be benign, aligning with organizational objectives though not captured in the policy, or it may directly conflict with the policy and its underlying intents of the policy, resulting in disastrous consequences.

C. Time-stressed mobile environments and auto-logout

Automated log-outs based on time away from computers are designed to prevent unauthorized use when a user fails to log out. Choosing the right interval before log outs can be tricky. If too soon, one interferes with work. If too late, one leaves the machine open for misuse. In medical settings especially, timed auto-logouts create workflow breaks that result in patient safety dangers. In the case we examined [6], clinicians frequently must leave the workstation to find

a document or device, or interact with the patient or another clinician. However, policy implementers assumed that physicians and other clinicians were near the computer for the entire time of the interaction with the patient and attempted to stop unwanted timeouts by attaching a proximity indicator to every mobile computer on the hospital floor. As might be expected, the unwanted logouts caused massive delays–requiring relogins, interruption of thought flow, and lost data that were not saved. As a result, the clinicians found that putting an inverted styrofoam cup over each proximity indicator defeated its function and allowed them to continue uninterrupted.

In another timeout case, we found that teams would assign the lowest ranking member to sit by the computer and periodically hit the space bar to defeat the inactivity auto-logout mechanism. Since then, we have seen vendors selling "mouse jigglers" that fool the auto-logout timers by periodically moving the mouse.

A different category of time-related security workarounds is found when workers can't be at the assigned computer or desk to perform a needed action. For example from an interview, when clinicians are in a patient's isolation room (infection control) they can't access computers in other rooms without removing gloves, masks, gowns etc. They thus signal another clinician to insert medical orders in their name.This requires sharing passwords in addition to other artifacts (e.g., phones for two factor authentication, or ID cards, or e-keyfobs with chips). As above, the needs of workflow and the expectations of cybersecurity do not fit together.

A major goal would be to develop a data-informed model of the implementation and optimization process, illustrating/documenting the many vectors/forces involved, the tradeoffs, the constant adjustments, the delays and perhaps even the occasional jumps (as compared to incremental improvements).

Of course, many time-related circumventions do not involve technology. We have seen nurses indicate they are reserving a specific computer by draping a sweater over it; we have had clinicians tell of being logged into one device but needing access to another in a different patient's room—and calling out to a colleague to log them out of the distant computer, because security rules prevent being logged into more than one device at one time.

A non-medical example is seen in the financial industry, where seconds mean millions to traders. In these settings, the computer screens are festooned with yellow stickies each reflecting the day's passwords, and auto-logout policies are strongly condemned by users.

All of these examples illustrate the many tradeoffs between expediency and safety.

D. Bypassing barriers in information flow

Fieldwork by our team – and many others – reveals numerous cases where, to accommodate requirements from workflow, users will send information through channels that completely jump out of the system governed by the organization's security posture.

In one set of examples, users find an information path that bypasses the official controls. In both finance and medical settings, we have heard users brag about how they transmitted data they believed essential by embedding the data in an image (e.g., a screenshot pasted into a PowerPoint presentation) so the content would be hidden from exfiltration guards. In fieldwork in critical infrastructure IT, users expressed that the existence of well-known default passwords were necessary for safety and reliability – in urgent settings, one needs immediate access.

In another set, users find computational means to bypass barriers. A security officer in a multi-level security organization told of his fear that employees would optimize their workflow by bridging the security-critical airgap. In other government settings, workers moved their operations offsite, such as to a coffee shop, to avoid rules that prevented them getting their jobs done. In both finance and medical settings, we have heard IT-savvy users bragging about bypassing the official software (and its controls) by writing their own programs to extract data.

One common set of examples of this pattern are all the inventive ways users come up with to access systems using passwords that do not belong to them. In [19], we cataloged:

> In hospital after hospital and clinic after clinic, we find users write down passwords everywhere. Sticky notes form sticky stalagmites on medical devices and in medication preparation rooms. Weve observed entire hospital units share a password to a medical device, where the password is taped onto the device. We found emergency room supply rooms with locked doors where the lock code was written on the door... One vendor even distributed stickers touting to write your username and password and post on your computer monitor .... Clinicians share passwords with others so that they can read the same patients charts even though they might have access in common.

E. Positive effects for workflow interaction

There are of course workflow and security interactions that are beneficial. For example, the New York Stock Exchange automatically logs out all users and sweeps the system shortly after 5pm. This policy greatly reduces the attack surface for the exchange. In many organizations such a policy would meet resistance from users trying to finish off the final tasks of the day. Since no trading takes place after hours, though, it works as expected.

There is, of necessity, a caution that accompanies beneficial interactions: when adapting a policy that has been successful elsewhere, policy implementers should ensure that an approach that was successful in an original setting was not due to non-generalizable idiosyncrasies.

IV. Discussion: Workflow-Inspired Security Design

We examined a number of cases where security policy does not have the intended effect in the context of the workflow of the organization for which the policy was developed. Many problems can be captured this way, highlighting the value of a workflow-based analysis in designing policies that are both usable and effective with as little readjustment as possible. While workflow has been considered in this context, this work is novel in its analysis of different aspects of this

interaction. A significant advantage of such an analysis is the relative ease of noting features of the workflow that may impact the effectiveness of a policy, compared with the difficulty of extracting generalizable aspects of human behavior in security. We drew out four themes in workflow interactions with security policy. While clearly not exhaustive, they capture many of the examples seen in our field work and in the literature. First, policies may cause certain tasks in the workflow to become unreasonably hard. Second, policies may not reflect multiple paths in the workflow to gain some form of access, and each should be protected without undue burden if possible. Third, in highly time-dependent and particularly mobile environments, policies are likely to falter if they place an undue time commitment at critical moments. Finally, developers must pay attention to when and why users need to share information. Information sharing is often informal, and blanket policies to restrict information based on type and without respect to context can lead to a cottage industry of alternative representations and communications that defeat security and privacy.

We now present methods to design, develop, maintain, and revise security solutions to fit the organizational workflow. We also provide a snapshot of our work in progress in this area.

Ethnography, Surveys, and Focus Groups: These and other approaches provide an understanding of organizational workflow context from the user's perspective that is essential to workflow-inspired security. As mentioned, Heckle [12] demonstrated how an ethnographic field study improved both usability and security of a single sign-on system, leading to the eradication of problems that security practitioners had unintentionally and unknowingly introduced. Results from security-focused surveys [5] or psychometric tests [8], [9] on the target subpopulation may also help predict the efficacy of a security solution. It's also immensely valuable to incorporate users in the design process, as argued by Kirlappos et al. [15]; doing so may avoid issues that commonly arise when designing for the other [27]. There's also great value in establishing a usable, reliable, and meaningful bi-directional communication channel between users and security practitioners.
Our Efforts: As mentioned in the beginning of Section III, we are continuing to conduct surveys, which are available on shucs.org. We seek to establish connections to further the reach of these surveys and also to gather feedback. Assistance in either pursuit would be much appreciated.

Analyzing and Cataloging Security Failures: Many papers catalog instances where security fails, often due to usability problems stemming from workflow and other considerations [1], [12], [26]. There's a wealth of valuable information that decision-makers could operationalize if they only knew about it. Decision-makers contemplating a new security solution would be well-served by descriptions and analyses of deployments of similar security solutions elsewhere, even though workflow and context may differ. These research findings, however, are seldom communicated to practitioners. This is evidenced in part by the recurrence of same security failures time after time. We must be able to communicate this information concisely. Moreover, we must understand practitioner needs so that we can best serve them. To this end, we call on the academic community and practitioners to collaborate in an effort to catalog and categorize the broad set of security failures and successes that occur in practice— and to develop a tool allowing developers to find relevant episodes, perhaps beginning with a simple front-end for keyword matching.
Our Efforts: We are in the early stages of cataloging a corpus of security scenarios to do just this. Our intended approach involves the following intertwined subgoals: (1) cataloging instances of security scenarios; (2) refining the underlying model and data representation for the scenarios; (3) creating a submission interface for researchers and academics to communicate their findings and observations; (4) building a front-end for practitioners to make use of the catalog. We soon hope to discuss this idea with security researchers and practitioners.

Other Predictive Tools: Additionally, the security community continues to develop tools to predict the efficacy of security solutions that may be useful in developing security solutions for workflow contexts. For example, agent-based simulation is well-suited for contexts where we can derive reasonable models of user behavior and workflow interactions, but where the environment is too complex for an analytical prediction of the effectiveness of security solutions [22], [21]. Despite the promise of such solutions, we note that first, it is paramount to identify the assumptions on which the tools rely, and second, while these tools can complement other approaches, they should not replace them. We advocate the development and adoption of tools that accurately model the deployment context; this means there must be sufficient understanding of that context, which often is acquired through alternative methods such as ethnography.
Our Efforts: While predicting a precise circumvention can be difficult, noting that an existing security posture makes it difficult for people to do their jobs is a key factor in predicting low compliance. If employees perceive the cognitive effort or time required to comply with organizational policy to be too expensive, the offending security solutions will probably be jettisoned. Agent-based simulation serves as a useful tool in predicting such organizational stresses. We are expanding upon earlier work, which focused on modeling password management practices [21].

ACKNOWLEDGMENT

This material is based upon work supported by the Maryland Procurement Office under Contract No. H98230-14-C0141. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the Maryland Procurement Office. Koppel's work was supported in part by NSF CNS-1505799 and the Intel-NSF Partnership for CyberPhysical Systems Security and Privacy. Blythe's work was supported in part by the Science and Technology Directorate of the United States Department of Homeland Security under contract number HSHQDC-16-C-00024. The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official policies or endorsements, either expressed or implied, of the Department of Homeland Security or the US Government.

REFERENCES

[1] A. Adams and M. A. Sasse, "Users are not the enemy," *Commun. ACM*, vol. 42, no. 12, pp. 40–46, Dec. 1999. [Online]. Available: http://doi.acm.org/10.1145/322796.322806

[2] A. AlKalbani, H. Deng, and B. Kam, "Investigating the role of socio-organizational factors in the information security compliance in organizations," *arXiv preprint arXiv:1606.00875*, 2016.

[3] G. Anderson, G. McCusker, and D. Pym, "A logic for the compliance budget," in *International Conference on Decision and Game Theory for Security*. Springer, 2016, pp. 370–381.

[4] A. Beautement, M. A. Sasse, and M. Wonham, "The compliance budget: managing security behaviour in organisations," in *Proceedings of the 2008 New Security Paradigms Workshop*. ACM, 2009, pp. 47–58.

[5] I. Becker, S. Parkin, and M. A. Sasse, "Finding security champions in blends of organisational culture," *Proc. USEC*, vol. 11, 2017.

[6] J. Blythe, R. Koppel, and S. W. Smith, "Circumvention of security: Good users do bad things," *IEEE Security & Privacy*, vol. 11, no. 5, pp. 80–83, 2013.

[7] D. Dey, A. Ghoshal, and A. Lahiri, "Security circumvention: To educate or to enforce?" in *Proceedings of the 51st Hawaii International Conference on System Sciences*, 2018.

[8] S. Egelman and E. Peer, "The myth of the average user: Improving privacy and security systems through individualization," in *Proceedings of the 2015 New Security Paradigms Workshop*. ACM, 2015, pp. 16–28.

[9] ——, "Scaling the security wall: Developing a security behavior intentions scale (sebis)," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 2015, pp. 2873–2882.

[10] D. Florêncio, C. Herley, and P. C. Van Oorschot, "Password portfolios and the finite-effort user: Sustainably managing large numbers of accounts." in *USENIX Security Symposium*, 2014, pp. 575–590.

[11] C. Glass, "How to Use Social Engineering to Gain Unauthorized Access to a Hotel Room," *Wonder How To. Null Byte*, 2015. [Online]. Available: https://null-byte.wonderhowto.com/how-to/use-social-engineering-gain-unauthorized-access-hotel-room-0161055/

[12] R. Heckle, "Security dilemma: Healthcare clinicians at work," *IEEE Security & Privacy*, vol. 9, no. 6, pp. 14–19, 2011.

[13] C. Herley, "So long, and no thanks for the externalities: the rational rejection of security advice by users," in *Proceedings of the 2009 workshop on New security paradigms workshop*. ACM, 2009, pp. 133–144.

[14] H. Kavak, J. J. Padilla, and D. Vernon-Bido, "A characterization of cybersecurity simulation scenarios," in *Proceedings of the 19th Communications & Networking Symposium*. Society for Computer Simulation International, 2016, p. 3.

[15] I. Kirlappos, S. Parkin, and M. A. Sasse, "Shadow security as a tool for the learning organization," *ACM SIGCAS Computers and Society*, vol. 45, no. 1, pp. 29–37, 2015.

[16] J. Koebler, "Why American Farmers Are Hacking Their Tractors With Ukrainian Firmware," *Motherboard*, 2017. [Online]. Available: https://motherboard.vice.com/en_us/article/xykkkd/why-american-farmers-are-hacking-their-tractors-with-ukrainian-firmware

[17] E. Kolkowska, F. Karlsson, and K. Hedström, "Towards analysing the rationale of information security non-compliance: Devising a value-based compliance analysis method," *The Journal of Strategic Information Systems*, vol. 26, no. 1, pp. 39–57, 2017.

[18] R. Koppel, J. Blythe, V. Kothari, and S. Smith, "Beliefs about cybersecurity rules and passwords: A comparison of two survey samples of cybersecurity professionals versus regular users," in *Workshop on Security Fatigue, Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. Denver, CO: USENIX Association, 2016. [Online]. Available: https://www.usenix.org/conference/soups2016/workshop-program/wsf/presentation/koppel

[19] R. Koppel, S. W. Smith, J. Blythe, and V. Kothari, "Workarounds to computer access in healthcare organizations: you want my password or a dead patient?" in *ITCH*, 2015, pp. 215–220.

[20] V. Kothari, J. Blythe, S. Smith, and R. Koppel, "Agent-based modeling of user circumvention of security," in *Proceedings of the 1st International Workshop on Agents and CyberSecurity*. ACM, 2014, p. 5.

[21] C. Novak, J. Blythe, R. Koppel, V. Kothari, and S. Smith, "Modeling aggregate security with user agents that employ password memorization techniques," in *Workshop on Adventures in Authentication, Symposium on Usable Privacy and Security (SOUPS)*, 2017.

[22] K. Renaud and L. Mackenzie, "Simpass: Quantifying the impact of password behaviours and policy directives on an organisation's systems," *Journal of Artificial Societies and Social Simulation*, vol. 16, no. 3, p. 3, 2013.

[23] C. Rizzo, "Hotel Safety: 7 Rape Cases in 7 Years After Front Desks Give Away Room Keys," *Travel and Leisure*, 2017. [Online]. Available: https://www.yahoo.com/news/hotel-safety-7-rape-cases-180910464.html

[24] N. S. Safa, R. Von Solms, and S. Furnell, "Information security policy compliance model in organizations," *Computers & Security*, vol. 56, pp. 70–82, 2016.

[25] B. Schneier, "Giving Out Replacement Hotel Keys," *Schneier on Security*, 2008. [Online]. Available: https://www.schneier.com/blog/archives/2008/11/giving_out_repl.html

[26] S. W. Smith, R. Koppel, J. Blythe, and V. Kothari, "Mismorphism: a Semiotic Model of Computer Security Circumvention," in *Proceedings of the 2015 Symposium and Bootcamp on the Science of Security*. ACM, 2015, p. 25.

[27] B. Vandenberghe and K. Slegers, "Designing for others, and the trap of hci methods & practices," in *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 2016, pp. 512–524.