The Security Cost of Cheap User Interaction

Rainer Böhme
University of Münster
Leonardo-Campus 3
48149 Münster, Germany
rainer.boehme@uni-muenster.de

Jens Grossklags
Pennsylvania State University
329A Information Sciences & Technology Bldg
University Park, PA 16802
jensg@ist.psu.edu

ABSTRACT

Human attention is a scarce resource, and lack thereof can cause severe security breaches. As most security techniques rely on considerate human intervention in one way or another, this resource should be consumed economically. In this context, we postulate the view that every false alarm or unnecessary user interaction imposes a negative externality on all other potential consumers of this chunk of attention. The paper identifies incentive problems that stimulate overconsumption of human attention in security applications. It further outlines a lump-of-attention model, devised against the backdrop of established theories in the behavioral sciences, and discusses incentive mechanisms to fix the misallocation problem in security notification, for instance the idea of a Pigovian tax on attention consumption.

Categories and Subject Descriptors

H.1.2 [**Models and Principles**]: Human/Machine Systems—*human factors, human information processing*; C.2.0 [**Computer Communication Networks**]: General—*security and protection*; K.6.0 [**General**]: Economics

General Terms

Security, Human Factors, Economics

Keywords

Interdisciplinary Security and Privacy, Attention Economics, Usable Security, Bounded Rationality, Security Warnings, Notice and Consent, HCI, Security Economics, Policy

1. MOTIVATION

"Security is determined the weakest link. And the weakest link is most likely the user." This mantra is sounding from thousands of security awareness trainings around the globe. Many protection mechanisms are not purely implemented by means of technology, but are only complete if potential security violations can be escalated to the level of user interaction. In principle, it is not a bad idea to let the user know if a remote server's secure shell identity has changed, a Transport Layer Security (TLS) handshake has failed, or potential malware is about to be executed. Humans often posses more contextual knowledge and better capabilities to extract operable conclusions from it than machines—sufficient security knowledge provided [4]. A typical implementation of such user interaction consists of a dialog awaiting a decision from the user on how to proceed [37]. In theory, of course, this dialog would rarely occur. In practice, the average user makes several dozens of decisions per day in response to interception dialogs, which interrupt the user's primary task.

Sometimes these decisions may have substantial economic, social, or legal consequences. So considerable attention and cognitive effort should be devoted to finding the right response. Yet, the averse circumstances of an interception dialog already hamper an elaborate decision. And the main problem is that too many of these decisions are requested in error. In the long run, users get habituated to taking meaningless decisions [33]. As a consequence, the few really meaningful decisions might escape the user's attention.

Two approaches are conceivable in principle to overcome this dilemma: first, getting the user out of the loop. This might be a way forward in certain situations, but it seems unlikely to be feasible in all cases. Hence, in this paper we will elaborate on the second approach, that is to economize user interactions. We argue that user attention is an extremely scarce resource, which should be best allocated to the primary task and the decisions that really matter. One of our main contributions is to interpret unnecessary user interactions as inflicting *negative externalities* on other, possibly more relevant, decisions.

Understanding user attention as a public good may sound exaggerated at the first glance, but it is only a logical consequence in a succession of resources that appeared abundant until people realized their rivalrous nature. In the 18th century, pasture seemed abundant in most places of the world, yet population growth and urbanization led to "tragedy of the commons" in its literal meaning [49]. In the 19th century, industrialization brought pollution and the need to fix the externalities in the consumption of clean environment, a public good that was previously believed to be abundant [50]. Until the late 1980s, adding free computing resources to a network would have been considered as a charitable act, and only few might have realized the negative externalities emerging from unsecured programmable nodes in a network [111]. In all these cases, policies have been established—or

NSPW'11, September 12–15, 2011, Marin County, California, USA.

are under discussion [110]—to internalize these externalities. User attention might just be the resource to be overconsumed and thus exhausted in the 21st century [101]. So it is time to reason about ways to fix the allocation problem.

Our resource of interest is human attention, more specifically the capability to take decisions in response to questions posed by a system. This perspective is general enough to unify topics that have previously been studied rather independently. Concrete security warnings with actionable options are only one example. Our arguments extend to topics of privacy regulation, notably notice and consent, and computer-mediated contract negotiations [61]. The specifics of these situations can be reflected in our model by adjusting its parameters. There may be further applications without direct security or privacy implications, which thus are beyond the scope of this paper.

Section 2 presents the problem in its various instances—spanning security warnings, privacy notices, and click-wrap licenses—by borrowing from social psychology and cognitive theories of responding to closed-form questions. We will argue, and support with evidence, that optimizations of the user interface design are likely to be only marginally effective in the long run (if at all). Therefore the only way forward is economizing the use of human cycles for decisions. Section 3 outlines an economic model of user attention as scarce resource and discusses strategies to incentivize sparsity in user interactions by internalizing the negative externalities. This section also enumerates challenges that prevent our "security paradigm" from being a silver bullet. Section 4 explains similarities and differences of our approach to a range of theories in social psychology and economics, for instance the economics of attention and the paradox of choice. Section 5 concludes this paper with a summary of implications.

2. PROBLEM STATEMENT

The proliferation of computers to serve as interaction partners with human beings in everyday life continues to improve efficiency and productivity [17]. This development goes along with a shift in the *relative distribution* of transaction costs between the involved parties. In human–computer interactions, automation has substantially reduced the cost of asking standardized questions to many human beings. Conversely, the cost for every individual to provide an elaborate response remains constant at best. In the presence of frictions, which prevail in most markets for information goods [99], the side represented by a computer is likely to oversupply questions, thereby dumping the burden to answer on the human part. The software industry's habit to include lengthy End User Licenses Agreements (EULAs) drafted in obscure legalese is just one example of this shift in the relative transaction costs.

The transaction costs involved in answering questions is paid from a special budget, namely users' attention. Social psychologists use dual-path models as tools to relate the attention devoted to a task to its likely outcome. Prominent examples are the Elaboration Likelihood Model (ELM) [88] and the Heuristic–Systematic Information Processing Model (HSM) [24]. Both models are designed to explain attitude changes and have been examined in numerous empirical studies [25]. As our interest is less on attitude change than on the decision process when reacting to questions, we use dual-path models inspired by the ELM and HSM, but adapted specifically to the process of answering (standardized) questions. These models evolved in the 1980s and 1990s in the context of survey research, where scholars tried to understand the cognitive roots of response errors with the intention to minimize or correct them [109].

2.1 Optimizing versus Satisficing

Figure 1 visualizes the cognitive processes involved in answering questions according to Krosnick and Alwin's dual-path model [65, 64]. This model distinguishes two idealized response strategies, *optimizing* and *satisficing*. Users who choose to optimize first read and interpret the question, then they retrieve all relevant information from the long-term memory and possibly additional sources, then form a judgment by weighting the information according to predispositions and beliefs (e. g., by comparing possible outcomes of their decision), and ultimately express their response by picking the alternative that fits their conclusions best. Translated to the example of a browser warning, this means a user reads the full warning message including the underlying certificate details, then recalls what she knows about web safety, the TLS protocol, and her organization's security policy, then weighs contextual information (How likely and how severe would a man-in-the-middle attack be? How urgently do I need to access this website?), and finally expresses her decision by clicking the appropriate button.

Satisficing, by contrast, implies that neither the question is properly read nor understood, information retrieval is bound to salient cues residing in the short-term memory, and judgment is made by simple heuristic rules. Some of these steps might even be skipped. The result is a suboptimal response that can be found with minimal effort. In our example, this corresponds to the type of user who clicks warnings away almost instinctively.

A number of empirical studies, primarily in the area of survey research, investigated determinants for the choice of the optimizing, respectively satisficing path [65, 64, 109]. However, we are not aware of any studies of browser warnings building on the dual path model (in our previous work, we conducted a study on consent dialogs [11]). Determinants supporting optimizing include

- High interest in the topic of the question,
- Very knowledgeable about subject matter,
- Positive expectations about the outcome of the decision, and
- High need-for-cognition, a personality trait [18].

Determinants for satisficing include

- Low motivation,
- High difficulty of the question,
- Absence of economic or behavioral incentives, and
- Monotonous repetition.

As is apparent from this list, security-related interception dialogs have a very dim chance of being processed in the systematic path, in particular, if they appear excessively without observable consequences.

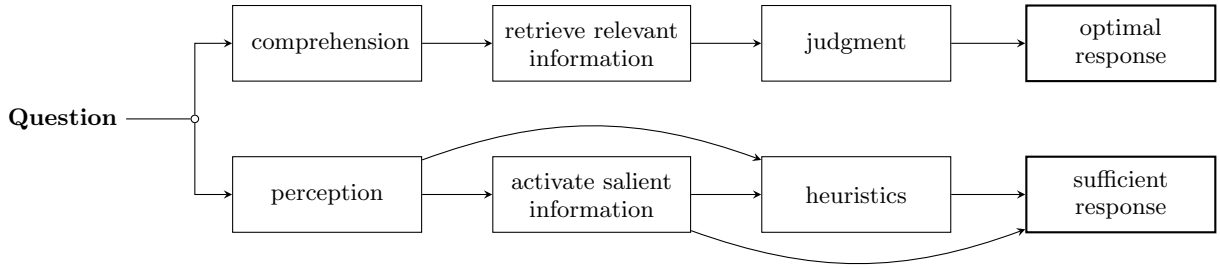

Figure 1: Dual-path model of response process: *optimizing* (upper path) and *satisficing* (lower path) [65, 64]

2.2 Dead Ends

In the following, we are interested in considering a number of solution approaches that have resulted in dead ends.

First, some of the determinants of optimizing and satisficing behavior can be positively influenced with interface design techniques. A number of user warning approaches are based (at least partially) on more intrusive interruptions of the desktop experience. A popular technique is to display a highlighted warning message that necessitates a user decision before the activity can be resumed. This approach can be implemented in degrees of different harshness. For example, the User Account Control (UAC) system in Windows Vista/7 disallows any (even entirely unrelated) desktop activity until the user makes a decision about the warning message [92]. In a recent small-scale user study, however, 20 % of the participants entirely disabled the UAC [83]. Another small-scale user study tested even more intrusive methods to catch the users' attention when opening a potentially dangerous email attachment [60]. One interface required the user to do some proof of work (e. g., retyping a file name). Another interface combined surveillance with the threat of consequences by asking for a justification to open the attachment, suggesting that the stated reason would be stored and audited. Unlike in [83], [60] report that the interfaces led to more secure behavior and positive subjective usability evaluations. But objective measures, such as response latency deteriorated. Results from experimental studies on user interface design might even be too optimistic, as every new mode of presentation is likely to receive more attention while it is novel. But its eye-catching effect might go away once the users become habituated to it [37, 107, 11].

Second, in order to counteract the desire of users to switch off notifications altogether, the frequency of warning messages can be adjusted (i. e., typically reduced) by the system designer. Likewise, the user may be able to access an options menu to selectively disable certain notification types. On the one hand, this trend suggests that there may be room for flexibility with respect to the need for notification. In contrast, some commentators feel that software companies bow to the pressure of consumers (who feel overburdened by too many warning messages) and as an imperfect response leave the software vulnerable to external intrusions [42].

A related notice approach are polymorphic advice messages [16, 60, 115]. The main idea is to create context-specific dialogs asking the user questions about a security-related incident (e. g., does the user expect the message containing an attachment). With the gathered information, the system can then set a suitable course of action (e. g., blocking the attachment). To dissuade the user from giving inaccurate information or to cheat the system in order to circumvent the security mechanism, an auditor can be involved to administer rewards or penalties (e. g., shutdown of email client for a period of time). This approach spreads the attention-consuming actions to additional parties (e. g., the auditor). Further, context comprehension of the security situation is a difficult problem for the general case of all potential user actions (e. g., if irrelevant options are shown to the user trust in the system will diminish).

A third approach is to lower the stakes for security compromises. The principle of least privilege as implemented, for example, with Low-privileged User Accounts (LUA), exemplifies the idea that users operate within a more restricted space, but also better contain the damage that an outside attack may be able to cause. According to a report issued by Microsoft back in 2005, about 85 % of its corporate users executed their everyday business using administrator privileges [89]. In their recent laboratory study of college employees and students, Motiee *et al.* found that all participants utilized high privilege user accounts, and that knowledge about the benefits of restricting access rights was largely absent [83]. This conflict is unlikely to be resolved at the user notification level, but is perhaps best addressed in the program development phase. That is, the tradeoff between convenience and security does not solely originate from the user, but programmers need to understand the interdependency between simplicity and realistic threat modeling [108].

Fourth, content providers and desktop software companies may rely on third party security software to be installed by the user to prevent and mitigate problems, and essentially to sidestep the notification process. However, the reports published by the National Cyber Security Alliance have shown that the penetration of end user systems with security technologies is unsatisfactory. In their 2010 study, only 37 % of the researched user population had, in fact, a full software security suite installed on their systems [84]. Moreover, a recent research study indicates that investments in security technology might be partially offset by users' increased risk-seeking behavior [26]. This observation can be explained with Peltzman's theory of risk compensation stating that individuals typically make a joint assessment for the demand for safety and usage intensity (e. g., faster driving on a highway or more adventurous browsing behavior on the Internet) [87]. Further, users are overly optimistic that they can handle online security risks or that they will not be affected by attacks [20].

A related fifth strategy is to leverage user attitudes. Countless studies have shown that users are concerned about their privacy [1, 2, 66, 104] and security [45, 67, 84]. Common

sense would dictate that users who are strongly concerned about security and privacy would also exhibit more careful and diligent behaviors. However, Spiekermann *et al.* demonstrated that actual behaviors do not correspond strongly with privacy preferences [104], and Acquisti and Grossklags published similar findings extending also to the domain of security [2]. The lack of fit in these simple correlation analyses calls for the adoption of more elaborate models to explain behavior from reported intention, such as the theory of planned behavior [5], in the context of privacy and security. Moreover, purely survey-based studies should be treated with caution when it comes to explaining behavior (see, for example, [70]).

A sixth approach is to rely on social influence from peers, friends and family members or expert influence by administrators under the assumption that group decisions or more proficient external advice are likely to result in a more robust security outcome. From a theoretical perspective, Wilson's model of cognitive authority suggests that individuals rely on first-hand experience, but also second-hand narratives to construct knowledge [119]. The argument is that users rely on information from other individuals to navigate situations outside of their own expertise. It is unclear how suitable this model describes security-relevant interactions. For example, a number of technical proposals rely on improvements of security configurations based on data by peers (see, for example, [116]). Human influence, however, may be subject to conflicting advice and herding effects. For example, Salganik and Watts showed that weak signals can easily be overwhelmed by group behaviors, while only strong signals might survive initially conflicting herding processes [93]. Security notifications likely fall into the former category. Another concern is that peer groups lack precise boundaries which might enable attackers to infiltrate them.

This list is likely not exhaustive (e. g., one could benefit from user experience whether real or simulated with attacks), but should serve as introductory evidence for the complexity of the problem space.

2.3 Instances

An abundance of research and industry studies indicates the current ineffectiveness of the involvement of the user in security decision-making. Below, we summarize a number of key studies sorted into different topical areas with the consideration of security warnings being of primary relevance for our argument.

2.3.1 Security Warnings

Users regularly miss deployed security indicators [46, 95, 107]. These findings are robust even in contexts where users should have particularly strong prevention incentives, e. g., in the context of phishing [33, 37, 121] or fake/malicious warning messages [100, 105]. Some proposed design improvements yield more encouraging results, but rarely affect the overwhelming majority of experimental subjects [103], or utilize more involved interaction techniques (such as two-factor authentication or customized security indicators) [32].

The effectiveness of warnings has been studied in a number of different context, e. g., nutrition and health. Argo and Main conducted a meta-analysis of several impact dimensions. Reading and comprehension as well as recall of the warning were shown to be not affected by moderating variables that included vividness, location, familiarity, product type etc. Attention was impacted by the former three moderating variables. Finally, cost of compliance was found to impact behavioral compliance with the warning [6].

Several other factors may impact the significance of warning messages. For example, there is a discrepancy between individuals' ability to absorb different kinds of information. A recent study has shown that users can express a robust opinion about the visual appeal of a website in as little as 50 ms [72], while absorbing the content of a warning message obviously takes much longer. Further, there is typically a positive correlation between trusting the informational content of a site and its visual appeal [35]. It follows that trust judgments about a site may impact the incentives to scrutinize a warning message.

Felten's exclamation that "given a choice between dancing pigs and security, users will pick dancing pigs every time" [80] serves as a reminder that handling security dialogs is typically not the primary objective of a user. Security warnings that are perceived as interruptive are likely dismissed or, if possible, ignored [46]. A further implication is that curiosity frequently trumps cautiousness.

However, frequent interruptions are not only undesirable but actually harmful. Recent research shows that, in particular, older users suffer from interruption recovery failures, i. e., an inability to dynamically switch between activities [28]. As a consequence a more significant effort is needed to refocus on the primary task.

2.3.2 Online Privacy Notices

In contrast to warning messages, privacy notices are usually not presented to the user in a vivid and engaging manner. Instead, the terms are available in a privacy statement accessible via a link from the main page of a website or installation dialog of a program. In addition, either the user is presented, similar to browse-wrap or click-wrap agreements, with an agree/disagree choice or agreement is implicitly assumed when utilizing a product or site. This distinction is related to the opt-in versus opt-out debate, however, many differences in implementation details exist and would require a more thorough treatment. Essentially, consumers face two types of cost. First, they have to make an assessment on whether they might like the product or service. Second, they have to decide whether to invest attention in the analysis of the associated privacy consequences. For opt-in this decision is made up-front. For opt-out, it may occur at a later time. The relative desirability of these two regimes may depend on many factors.[1]

Vila *et al.* dismiss the usefulness of state-of-the-art privacy notices [114]. They argue that privacy claims are too uninformative to overcome information asymmetries between consumers and service providers. Longitudinal studies have shown that accessibility, writing style and content quality

[1] A previous analysis established a cost–benefit argument with the view of opt-out being more consumer-friendly than opt-in [74]. This thesis is likely difficult to support under more general assumptions. In contrast, Bouckaert and Degryse develop an economic model showing that if consumers can exercise their option at no cost, opt-in and opt-out are equivalent. If exercising the option is costly for consumers, opt-out results in the lowest degree of privacy protection as personal information is marketed, opt-in offers an intermediate degree of protection as personal information flows only within a firm, and anonymity obviously provides the highest degree of protection [13].

are insufficient and do not evolve satisfactorily [59]. In fact, another study evidenced that for average notices length increased and readability declined over time [81]. Readability studies have also been undertaken for health and financial privacy notices [15, 52, 53, 54]. Typically, the results show that notices require at least a college level reading proficiency which is not appropriate for the generic consumer with highly diverse educational backgrounds. A point estimate of the aggregated opportunity cost of reading privacy policies in the United States' online economy gives an order of magnitude of US$ 780 billion per year, or about 35 times the value of all online advertising [79].

Modifying the privacy notice experience can be helpful [58]. In experimental settings, embedding privacy indicators close to search results was shown to impact consumer choice even if being associated with a small price premium [38]. Similarly, privacy warnings increase the awareness of users about privacy risks related to information disclosures [69]. In contrast, the evidence for seals is mixed. Mai *et al.* report positive price premiums after controlling for vendor characteristics [75] whereas Larose and Rifon find no effect on the disclosure intentions of subjects in a convenience sample [69]. Edelman explains such null results with the lack of substantial verification of trust seal approvals and the underlying incentives suggesting that less trustworthy companies are particularly interested in seeking accreditation [36]. Design proposals increasing the transparency of privacy and security actions and consequences may be helpful to overcome some of these unwanted marketplace outcomes [43].

Other related notice contexts might benefit from research efforts on privacy notices. See, for example, the development of financial or health privacy notices in response to the Gramm–Leach–Bliley Act [62, 71] and the Health Insurance Portability and Accountability Act [23], respectively. Further, online sales and membership programs that piggyback on user-initiated purchases at a wide range of merchant sites have led to a recent outpour of consumer complaints arising from the self-stated lack of awareness by consumers about having ever enrolled in such programs [30]. The recently passed Restore Online Shoppers' Confidence Act addresses some structural problems of these practices (e. g., the hidden transfer of payment credentials to a third party without consumer knowledge).

The improved notice regulatory requirements for the aforementioned acts all rely on the FTC Fair Information Practice Principles [41]. The key aspects are a reliance on disclosure of substantial terms and the opportunity for consumers to consent to these terms. As our discussion suggests, the principles can guarantee only a limited effectiveness from the user perspective.

2.3.3 End User License Agreements

Allegedly, Judge Richard Posner of the United States Court of Appeals for the Seventh Circuit and Chief Judge John Roberts of the Supreme Court of the United States admitted to not reading boilerplate notices [77]. However, such agreements might contain unfavorable terms that may also contain provisions impacting users' privacy and security [47]. The latter finding might initially appear surprising. But End User License Agreements (EULA), in particular, need to address modern business models for consumer programs that include data collection, targeted advertisements, etc. More generally, experienced contract lawyers argue that "the typical vendor software license has much less to do with the licensing of technology than it does with the creation of multiple revenue streams flowing from the user to the vendor and the elimination or minimization of most forms of accountability from the vendor to the user" [86].

The question is why do apparently consumer-unfriendly practices survive in the open marketplace [94]? Partly, increased concentration in some business segments might allow vendors to dictate strict terms.[2] However, recent empirical research showed that while price for software programs appears to be indirectly correlated with increased competitiveness of a market segment, the same does not apply to the harshness of terms [76]. Beales *et al.* argue that market concentration will predictably impact attributes that are readily observable to the consumer, i. e., price, but will have only a vague or exploitative influence less visible factors [7]. Our discussion of the lack of readability and usability of notices directly applies also to the case of EULAs indicating that privacy and security terms are much less likely apparent for the average consumer [22, 48]. Moreover, any rephrasing of boilerplate language is affected by practical limits of how much complex legal concepts can be simplified [78].

2.4 Summary

We have discussed three relevant instances where software developers tend to delegate relevant security decisions to the user and we have cited a range of empirical results on how users fail to react in the intended way. This mismatch and the inadequacy of simple apparent solutions, such as user interface optimizations, can be explained with established dual-path models of cognitive processes for responding to closed-form questions (Sect. 2.1). According to this theory, a common weak spot of the solution approaches reviewed in Sect. 2.2 is that they do not scale. Therefore, the most promising way forward is to ration user interactions. This leads to questions of which interactions are dispensable and how to get software developers to implement fewer interactions. In the following, we present an economic model, informed by the dual-path model, to tackle the allocation of user attention.

3. LUMP OF ATTENTION MODEL

In this section, we devise a stylized economic model that describes user attention as a public good. Unlike other public goods, user attention is not a common pool resource, but distributed over many users. This imposes additional constraints on the allocation problem: a system (or society) is at risk not only if the aggregate attention is exhausted, but already if the attention budgets of a significant amount of users are maxed out. So, in principle, the objective function should minimize the consumption of the aggregate attention budget while balancing the distribution of resource consumption over all users.

3.1 Setup

Our model distinguishes between *users* and *defenders*. The latter are responsible for the design of security mecha-

[2]Some legal scholars believe that one-sided contracts are not consumer-unfriendly. In fact, they argue that consumers (due to an absence of reputation concerns) can act opportunistically in their transactions with a company. The vendor, therefore, offers more lopsided contract terms to prevent such exploitation [9].

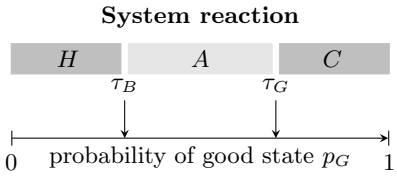

Figure 2: Defenders' choice: thresholds τ_B, τ_G

nisms whereas the former contribute to security by making security decisions if asked. Software vendors, banks, and website operators are the real-life equivalent of the defenders in our model. As a starting point, we assume two defenders and only one user who is endowed with an attention budget $a = 1$. Each defender designs one system to be used by the users. During operation, these systems will face situations undecidable by the machine, for instance whether a remote host is authentic despite an invalid TLS certificate. In the simplest case, there are two states, *good* and *bad*, and we assume that the system knows the probability distribution $(p_G, p_B = 1 - p_G)$. The defender has three options to design the system reaction:

1. **Continue operation (C):** This is the right choice if the true state is good. In this case the user extracts utility u and the vendor benefits (indirectly by prolonged service contracts) with utility γu. Otherwise, a major security problem causes costs to the user and the defender (through reputation loss and compensation to the user).

2. **Halt (H):** This is the appropriate choice in the bad state. It avoids the costs of a security problem s. Independent of the state, the user incurs opportunity costs from the missed utility (which she prefers in the bad state nonetheless).

3. **Ask (A):** Interrupt the user with an interception dialog. The user will identify the state and choose the right option (C if in state G, or H if in state B) if her attention budget is not exhausted ($a > 0$). This optimal decision will consume a part of her attention budget and decrease a by Δa. Otherwise, if a is already exhausted, the user decides along the heuristic path and chooses without further reflection the option she has more often chosen in the past.

Table 1 shows representative payoffs for users and defenders associated with each combination of system reaction and true state.

Some additional assumptions stand to reason. In particular, we assume that defenders and users suffer equally from a security breach. This avoids assumptions about the settlement of damages. Observe from Table 1 that the cost of a security breach are not passed on to the defender if the user has been asked. The defender might decline compensation and avoid reputation loss if responsibility can be credibly dumped on the user. We also assume that $u < s$, normalize $u = 1$, and restrict our analysis to the case of $0 < \gamma \leq 1$, hence $\gamma \leq 1 < s$. A final simplifying assumption is that $\Delta a = 1$. This means only the first decision is systematic and every following decision repeats the outcome of the initial systematic decision.

Table 1: Payoff matrix by state and system reaction

		True state	
		G (good)	B (bad)
User's payoff			
C (continue)		u	$-s$
H (halt)		0	0
A (ask)	$a > 0$	u	0
	$a = 0$ & C	u	$-2s$
	$a = 0$ & H	0	0
Defender's payoff			
C (continue)		γu	$-s$
H (halt)		0	0
A (ask)	$a > 0$	γu	0
	$a = 0$ & C	γu	0
	$a = 0$ & H	0	0

The derivation of (conditional) probabilities for the events (H, A, C) and (G, B) are given in the appendix. Since the model does not terminate, we compare asymptotic intertemporal payoffs with discounting ($r = 5\,\%$ per round) and search for optima on numerical grids.

3.2 Equilibrium Attention Consumption

Each defender sets two choice variables to define the thresholds for the system reaction as a function of p_G: the system reacts with H for $0 \leq p_G \leq \tau_B$, with A in the interval $\tau_B < p_G < \tau_G$, and with C for $\tau_G \leq p_G \leq 1$ (see Figure 2). We will now analyze the resulting equilibrium attention consumptions for a single defender and for different modes of interactions between multiple defenders.

3.2.1 Single Defender's Optimization Problem

Figure 3 displays a single selfish defender's optimal choices for (τ_B, τ_G) as a function of his opportunity-to-risk ratio, the only free parameter determining the defender's payoff. Quite naturally, the defender is very conservative (i. e., never allows C) if he does not participate in the utility but bears the risk of a false decision. His strategy changes slightly if his incentives become aligned with the usefulness of the system. Nevertheless, the gap between τ_B and τ_G remains wide open, indicating that users are requested to judge between 50 % (rightmost) and 70 % (leftmost) of undecidable security decisions. So defenders who optimize selfishly impose a substantial burden on users. We obtain this result because the negative outcomes of attention loss and subsequent decision errors by the users are not internalized in the defender's decision function.

Note that everything that holds for a single defender is also true for a setup with multiple homogeneous defenders who tap the same attention budget.

3.2.2 Defenders as Players

Negative externalities of a different nature emerge if two heterogeneous defenders are assumed. For example, an online bank and an online flower shop may have different exposures to risk, as materialized by $s_{\text{bank}} \gg s_{\text{shop}}$. Yet, they compete for the same attention budget with regard to a users' reaction to a TLS warning. If both defenders optimize

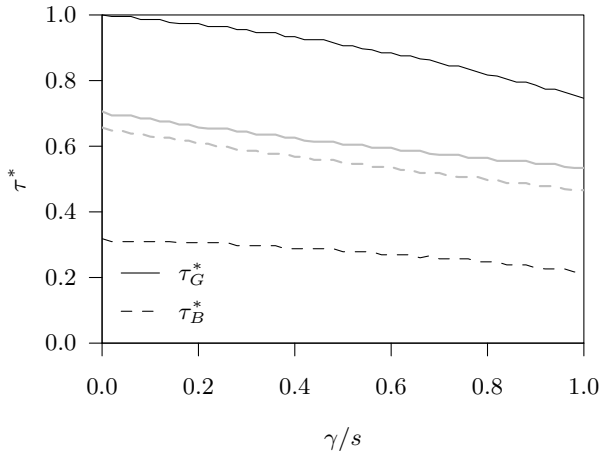

Figure 3: Single defender: optimal threshold settings as a function of the defender's opportunity-to-risk ratio. The black (outer) lines refer to a selfish optimization problem (Sect. 3.2.1). The gray (inner) lines are the solution if the user's payoff is included in the objective function (Sect. 3.3).

locally, they might come to different solutions for (τ_B, τ_G). While the user would be best off to spend her valuable attention on security decisions related to the bank's website, this does not stop the flower shop from demanding attention when subjectively determined to be optimal. Hence, this lax threshold imposes negative externalities not only on the user, but also on the other defender. The habit to confront users with coercive choice dialogs, such as EULAs and other take-it-or-leave-it decisions, can be interpreted as an extreme case of wasting the attention budget. The mechanics between two or more heterogeneous defenders can be expressed in a game-theoretic formulation where defenders are players. We defer its formalization to future work.

3.2.3 The Attacker as Player

Once in the game-theoretic domain, it is a logical next step to include the attacker as a strategic player [44, 82]. Two variants are conceivable. First, an attacker who acts like a defender and intentionally spoils the attention budget. This is akin to penetrating a network to increase the signal-to-noise ratio of an intrusion detection system such that a later real attack is less easily reacted upon. The second category includes fake security messages to capture the user's attention and exploit her naive reaction. Interestingly, while web banners mimicking Windows warning dialogs were prevalent in the early days of the commercial Internet, their proportion has faded. One can speculate if this is caused by the heightened professionalism of the online ad industry, or rather because of fading click-through rates resulting from users' increasing indifference to warning messages.

3.3 Optimal Allocation

In our model, users are no strategic players. They are best off if they are never asked, because responding once opens the path to cost $2s$ in the case of $(A \to C, G)$. The probability of this payoff is zero only if $\tau_B = \tau_G$ (cf. Eq. (17) in the appendix). All other outcomes are identical whether the system chooses automatically or asks the user via A.

As a result, the optimal allocation in the simple optimization problem (i. e., one defender) is to aggregate the payoff of the user and the defender to a measure of social welfare. In the absence of any further guiding assumptions both components may account with equal weight. The introduction of application-dependent weights is certainly possible and smooths the outcome between the two corner cases (see, results presented in Sect. 3.2.1 and $\tau_B = \tau_G$).

The gray lines in Figure 3 show a single defender's optimal choice for (τ_B, τ_G) if the user's cost is included for $s \gtrapprox u$. Even under these conservative parameters, the gap between τ_G and τ_B narrows substantially. This means in the socially optimal situation, users see interception dialogs much less frequently than in the equilibrium solution.

3.4 Fixing Misallocation

Once the existence of externalities has been determined by comparing the social to the selfish equilibrium outcome, one can reason about ways to ameliorate the allocation problem. In this section, we will discuss several approaches—formally when applicable, otherwise informally.

3.4.1 Changing the Norm

First, we will consider benevolent defenders and what they can do to ration user attention. Many software developers might not even perform the above cost–benefit calculation and just implement what is common industry practice. Similarly, software designers have to address a large number of security problems and usability aspects, and, therefore, do not fully appreciate the complexity of the intersection of these two areas [90]. As a result, the selection of security problems that are escalated to the user level is frequently inadequate and the design of the dialogs is not conducive to effective consumer decision-making [122].[3]

If such descriptive social norms govern the standard of security user interfaces, a first step could be to change the norms. For example, Microsoft recently proposed the NEAT guidelines considering the practical particularities of notice design and frequency of user interaction. NEAT stands for "Necessary, Explained, Actionable, Tested" [90]. Software designers are informed about the best practices through a variety of teaching tools such as information cards and hourly seminars that include specific examples.[4] This attempt to

[3]Zurko and colleagues described the flawed escalation process when a system detects something dangerous as follows: "It can ignore the danger and proceed, it can disallow the action silently, it can disallow the action and report the failure [...], or it can ask the user what to do [122]." They concluded that "it would be highly unusual for the user to have enough information to determine whether the action is proper or not. The warning box might as well say: Do you feel lucky?"

[4]NEAT specifically targets user warnings and notices. Broader definitions of effective usable security have been provided in the past. E. g., Whitten and Tygar consider security software as usable if the people who are expected to use it: "1. are reliably made aware of the security tasks they need to perform; 2. are able to figure out how to successfully perform those tasks; 3. don't make dangerous errors; and 4. are sufficiently comfortable with the interface to continue using it [117]." Even broader requirement catalogues distinguishing "polite" from "selfish" software include, inter alia, the principles of respectfulness ("software [...] does not preempt user choices"), helpfulness ("helps users make informed choices"), and personalization ("polite software remembers

Table 2: Modified payoff to find fair attention tax t

		True state	
		G (good)	B (bad)
Defender's payoff			
C (continue)		γu	$-s$
H (halt)		0	0
A (ask)	$a > 0$	$\gamma u - t$	$-t$
	$a = 0$ & C	$\gamma u - t$	$-t$
	$a = 0$ & H	$-t$	$-t$

establish a new injunctive norm is a step in the right direction and complementary to our systematic analysis of previous work and economic interpretation of relevant externalities.

Incidentally, our model suggests a rational explanation why big players in the industry are more likely to take the initiative. Defenders with a large "user surface" internalize more of the negative externalities of attention consumption than small specialized defenders. For example, security warnings of Internet Explorer compete for the same unit of attention with the Windows OS, so Microsoft as a whole should act with caution. By contrast, the installer of a rare software tool has little to lose and might demand user attention more recklessly. Another argument for the same hypothesis can be made on the difference in potential damage to the defender's reputation.

3.4.2 Feedback

Software developers might not be aware of how much user attention they consume. If there is a problem of overconsumption, one should be able to measure it. A simple solution would be to install feedback channels which count the number of occurrences and report it—aggregated and anonymized at the collection stage to protect user privacy—to the responsible parties. This could inform a benevolent defender to improve his strategy. If the defender is not benevolent, such aggregated statistics can be used by a regulator or independent watchdog to discipline the market participants' behavior. There is a range of options from simple blame-and-shame (coupled with the hope of anticipatory obedience), via fixed fines, to a tax regime specifically designed to internalize the externalities as much as possible.

3.4.3 Attention Tax

We elaborate the idea of an attention tax by recapitulating our formal model. A Pigovian tax is levied to internalize the disutility of a market activity generating negative externality with the tax rate. In our case, this could compensate for the relative shift in the distribution of transaction costs in human–computer interactions (see Sect. 2). However, our model measures attention and payoff on two different scales. Therefore, we have to solve the model to find the monetary equivalent of one unit of attention. This defines the attention tax rate t to be charged from every user interaction, as specified in the modified payoff structure of Table 2. Note that we do not consider tax differentiation between the six outcomes of state A. This is reasonable because neither the attention budget nor the true state is readily observable in practice.

its past interactions with the user, and carries forward past choices") [118].

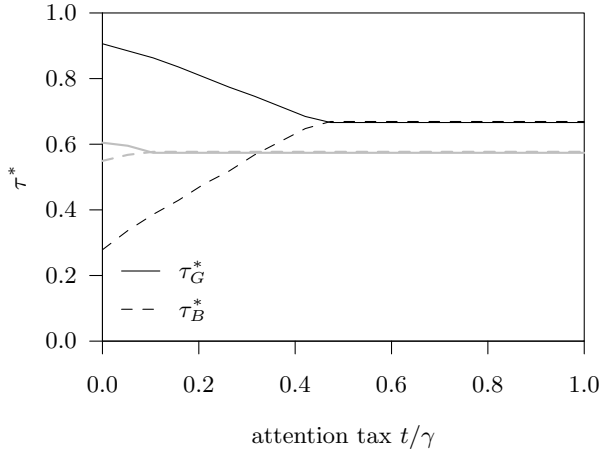

Figure 4: Attention tax rations user interaction ($\gamma/s = 0.5$, color coding as in Fig. 3)

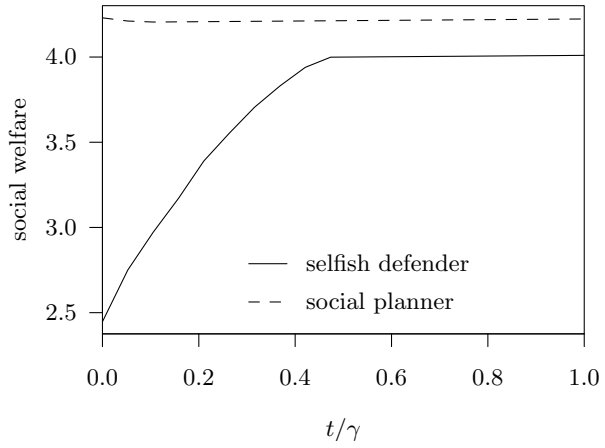

Figure 5: Welfare effect of attention tax ($\gamma/s = 0.5$)

Figure 4 shows a single defender's optimal thresholds as a function of the tax rate t standardized by his share of u in the good state. The opportunity-to-risk ratio has been fixed at $\gamma/s = 0.5$, all other parameters are similar to Figure 3. Observe that the gap between τ_B and τ_G closes steadily until user interactions disappear at a tax rate of slightly below 50 % of the defender's operative revenue (from undecidable situations). The next question is whether such a drastic measure is truly necessary? We reply affirmatively. As can be seen in Figure 5, social welfare increases with the tax rate until $\tau_B = \tau_G$. Increasing the rate further is not beneficial as the social optimum will never be reached. The reason for this discrepancy is the different operating point $\tau_B = \tau_G$ obtained from the private versus the social optimization problem. Note that without constraints on s and γ, it is impossible to design a Pigovian tax that reaches the social optimum unless the realization of the true state G or B is known.

Depending on the frequency of undecidable situations, the tax rate of 50 % can be prohibitive. It is presumably higher than most ad-financed businesses can afford. Yet, this incen-

tivizes both the rationing of user attention and maybe the reduction of undecidable situations in the first place. Both outcomes are good for security.

This analysis is not a full equilibrium closed economic model, as we concentrate on the incentive effect of the attention tax and ignore the revenue effect. Further refinements could compare situations where the revenue is socialized or paid directly to the users in compensation for their contribution. Both ways come with some caveats which need to be studied separately.

A final concern regarding the practical feasibility of a tax collection scheme can be addressed by referring to the possibility of automation. If fully-fledged auctions can be executed to decide individual ad placement, then the technical challenges of an attention tax mechanism appear manageable. Besides, tax schemes have been proposed for other externalities of information technology, such as security vulnerabilities [19].

3.4.4 Cap and Trade

Attention taxes promise to fix the allocation problem in theory, yet in practice it remains difficult to find the right rate. In such situations, market mechanism are the jackknife solution. For example, we could seek inspiration from the carbon emission markets and implement a cap and trade scheme. If a revenue effect is politically feasible, the problematic allocation of initial rights can be evaded. The choice of the cap level can depend on the overall network security situation.

This completes our discussion of incentive measures.

3.4.5 User-side Automation

Finally, a technical solution is conceivable as alternative or complementary measure (e. g., financed by tax savings). The idea is to reduce user interactions by automatizing responses as much as possible, either through intelligent systems that make use of contextual knowledge and policies, or by generalizing from one (hopefully considered) user decision to many similar cases. In this sense, proposed schemes for security information sharing in a web of trust [113] can be interpreted as means to realize economies of scale for user decisions. Examples for policy-based automation exist in the realm of privacy-enhanced identity management [29] and machine-readable privacy policies [31, 97]. An example for reputation-based mechanisms include Microsoft's download evaluation mechanism included in IE SmartScreen.[5]

Another way to frame this idea is the adaptation of the social navigation paradigm [55] to guide user decisions on security [34] and privacy [10] options. Social navigation refers to the principle of communicating (aggregate) behavior of other users in order to facilitate choice in otherwise unmanageably large decision spaces. From the point of view of social psychology, this corresponds to supporting the formation of descriptive social norms [27]. These norms are then easy to follow even by individuals who tend to satisfice. If this reduces the cost of decision-making while maintaining the decision quality, a gain in decision efficiency can be achieved. This would limit the impact of negative externalities even though the problem is not tackled at its very source. Empirical results from a pilot lab experiment on personal data disclosure decisions suggest that social navigation may influence behavior if the visual cues to signal the norms are strong enough [10]. However, this study is inconclusive about the overall decision efficiency because the decision effort, measured in time to completion, increased with the strong cue, and the quality of the cue was not endogenized in this setup. The very fact that users follow other users when uncertain about privacy-relevant actions is also supported with empirical evidence from a larger field study in the context of online social lending [12].

[5] `http://blogs.msdn.com/b/ie/archive/2011/05/17/smartscreen-174-application-reputation-in-ie9.aspx` (not to be confused with an identically-named blacklist)

3.5 Practical Problems

Our idea is certainly not without problems. In this section we collect reasons why it is too early to expect a silver bullet.

3.5.1 Quantification

The main obstacle is to measure the true values for the cost and benefit parameters in the model. This is crucial because of the risk that wrong estimates discourage user interaction in situations where it is utmost needed.

Another quantification problem is to determine the probability of the good state p_G, which our model assumes to be known by the system. In practice, the system has to make a decision based on noisy proxies for p_G and decision errors due to this measurement noise have to be accounted for.

3.5.2 Attribution

The technical details of an undecidable situation are usually not as straightforward as in our model. Most often multiple parties are involved. For example, a TLS warning could be caused by deficiencies of the website, any of the certification authorities in the chain of trust, the proxy server, the web browser, the operating system, or a real attacker. Now, if a security warning is displayed to the user, who of the involved parties should pay the attention tax? This highlights the problem of attributing responsibility for security actions on a very microscopic level.

3.5.3 Individual Differences

Users are not homogeneous. They may differ in cognitive capacity depending on personality traits [18] and situational states. This exacerbates the quantification problem and it is an open research question if an incentive mechanism exists that is invariant to individual differences, or at least one that reacts gently to minor mistakes in the underlying assumptions. Clearly all measures aiming at user education and awareness training directly or indirectly affect the specific capacity and efficiency in dealing with security decisions and thus the severity of externalities from user interactions.

4. RELATED THEORIES

We conducted a broad literature review but could not find any works that take a formal economic approach in studying externalities on human attention consumption, not to mention in a security or privacy context. Nevertheless, we identified a number of related fields. In this section, we briefly review the most relevant directions and their relation to this work.

4.1 Economics of Attention

Economics of attention can be understood as the flip side of information economics. The prevalence of information systems has created an abundance of information so that

human attention and processing capability have become relatively scarce resources. This coined the distinction between information-rich and information-poor environments [101]. Attention economics seeks to optimize the allocation of a design space capable of attracting human attention by treating it as virtual real estate. Examples include positions in search engine results [56] or brand recognition in name spaces.

A number of recent economic studies investigate different strategies how websites can gain the attention of individuals and eventually a larger group [40, 120]. Likewise, significant effort is invested into advertisement effectiveness studies to translate exposure to ads into click-through and eventual purchasing actions [3]. Our work differs because in the user notification context indiscriminate attraction of attention is not necessary and frequently unwanted. Users want to avoid wasting attention on activities lacking any form of benefit. A primary example is unsolicited communication, e. g., spam [73, 63], or spyware [46]. However, similar to the context in our paper, certain email messages as well as consumer programs with problematic features (e. g., a tracking software for children utilized by parents) cannot be easily filtered by an automated agent without user involvement.

More technically oriented researchers have developed a proposal to appropriately manage personal communication availability in a network of human agents [91]. Their goal is to protect privacy and to prevent unwanted interruptions for individuals themselves and their communication partners.

4.2 Rational Inattention

In general, inattention can be rational if the cost of acquiring and processing information exceeds its expected utility, e. g., as extracted from better-informed decisions. Rational inattention has become a label for a stream of research in finance and monetary economics which has brought in Shannon's information theory and assigns costs to the information processing of market participants, including organizations (bearing so-called wiring costs) and human decision makers (facing cognitive bounds) [102]. Researchers in this tradition derive abstract stochastic models to explain deviations of observable market prices from the predicted equilibria under perfect information. We understand that this research is mainly descriptive and independent of a specific cognitive model of human behavior. So far, the models are interpreted to guide central bank communication, suggesting that a heavily filtered view of its own assessment avoids overburdening the market participant's information processing capabilities and reduces the risk of overreaction. This line of work is less concerned about finding the socially optimal allocation of information processing capability on the available information *between* individuals.

4.3 Dual-Path Models in Survey Research

Dual-path models have a long tradition in social psychology. Their main application is the domain of persuasion research, where the models help to explain attitude change [25]. The processes involved in answering (closed-form) questions have been studied intensively in the 1980s. While some—notably untested—theories of rational choice have been proposed [39], the dominant direction of the field was behavioral [109]. Dual-path models include Krosnick and Alwin's (see above, [65]), and further models specifically designed for attitude questions, which are therefore less adaptable to our research question [21, 106, 96].

4.4 Paradox of Choice

Perfectly rational economic actors will always benefit from additional options. However, in practice, humans suffer from the availability of two many choices because of additional cognitive requirements and imperfect trade-off calculations (e. g., concerning opportunity costs) [57, 98]. Similarly, it is likely incorrect to claim that more information is always better. Instead, it merely allows for more efficient decision-making [112]. Further, in the long run, an abundance of related options and too much information may negatively impact individuals' psychological well-being [98]. In a recent study, it was shown that this paradox of choice is directly applicable to the Internet search context. Presented with fewer search results (i. e., a lower visualized recall), individuals reported a higher subjective satisfaction with their eventual choice and greater confidence in its correctness compared to an alternative scenario with a more customary number of search results [85].

While clearly related, the paradox of choice is distinct from our primary research question because it focuses on an excessive *number of alternatives* (beyond a handful) rather than on the *number of occasions* to choose (between few options). Moreover, regret and anticipated regret have been identified as important factors behind this paradox. Both are likely weaker in our scenario, where consequences are not directly observable or comparable. Lastly, the negative impact on individuals psychological well-being is stronger for individuals who try to optimize, whereas the context of our work is defined by the transition from optimizing to satisficing. Merging research on this paradox with security warnings nevertheless appears to be an interesting direction for more theoretically-founded empirical work.

4.5 Illusion of Control

Unrealistic perceived control over a situation has been demonstrated in countless laboratory and field studies starting with the work by Langer [68]. In the security context, individuals might consider the lack of adverse outcomes in essentially undecidable choice situations to be a function of their personal skill and knowledge. This effect is further amplified due to the typically delayed consequences of privacy intrusions and security breaches [2].

In the privacy context, researchers report that users confuse control over publication of private information (e. g., on a social networking site) with control over accessibility and use of that information by unrelated third parties. For example, the independent posting by another person of previously revealed information in a different context will be considered as unwanted intrusion and loss of control [14].

4.6 Security Paradigms

Several papers in earlier editions of the *New Security Paradigms Workshop* have theorized on the right amount of user involvement in security decisions. Here we point to two relevant approaches and highlight differences to our approach.

4.6.1 The Compliance Budget

Beautement, Sasse and Wonham [8] interview employees about their compliance with their organizations' security policy and find that the willingness to comply is not unlimited. The authors coin the notion of a "compliance budget" and suggest that it should be managed like any other budget, gently reminding decision makers that too many rules

can be counter-productive. Their and our work share the idea of a limited human ability to contribute to security, yet many differences prevail. Our approach seeks to quantify and formalize analytically, whereas theirs is qualitative and empirical. We focus on the discrepancy of transaction costs specifically in human–computer interactions, whereas they take a more general perspective including all kinds of compliance with policies in an organizational context. Consequently, our work targets software developers, theirs managers. Most importantly, the compliance budget accounts for the cost of taking a more cumbersome of multiple options for the sake of security (i. e., the result of a decision), whereas our attention budget accounts for the *cost of making a security decision* in the first place. Our model does not presume that acting securely has a higher total cost (no matter if in the organization's or one's own interest), but it accounts for the cost of making wrong security decisions.

4.6.2 Rethinking User Advice

Herley [51] uses an interdisciplinary approach similar to ours to tackle a related question: users' reactions to externalities generated as a consequence of their security decisions. The differences between his work and ours are that he focuses on user advice, typically given in awareness trainings independently of a specific situation. By contrast, we study user interactions with interception dialogs in concrete and specific situations. Another difference is in the perspective on the roots of decision-making. Herley's work emphasizes the rational user who often ignores current advice due to cost–benefit calculations. Our approach stresses the partly irrational aspects of decision-making against the backdrop of behavioral research. Both his and our work concur in the recommendation to economize the use of users time, respectively attention; both seems less abundant than often believed by developers. The attempt to reason about ways of internalizing the externalities of attention consumption is unique to our work.

5. IMPLICATIONS

Our new security paradigm merges attention economics with usable security. In an increasingly automated world, human attention has become the scarce resource tapped by many competing players, thereby creating a new tragedy of the commons. In human–computer interactions, automation allows the developer of the computer interface to realize economies of scale for transaction costs, whereas the side represented by a human being cannot realize equal savings. This stimulates over-consumption of human attention, because the developer who decides if a (security) question shall be asked does not pay the full cost of finding a response. Specifically in the domain of security, this over-consumption may exhaust the users' attention budget quicker than necessary and hence deprive users of the ability to defend against significant risks, thereby imposing considerable social cost.

We have argued that optimizing user interface design is not an adequate solution to this problem, as witnessed by empirical studies which account for habituation effects. Instead, user interactions should be rationed in order to save the scarce resource of attention for the most important decisions. Our stylized analytical model shows, in principle, how this can be implemented and how developers might react. One idea is a Pigovian tax on attention consumption.

The paradigm of rationing user interaction is incompatible with policy initiatives calling for more user involvement, for instance in terms of mandatory notice and consent. The Gramm–Leach–Bliley Act, the Health Insurance Portability and Accountability Act, or the Restore Online Shoppers' Confidence Act are current examples for policy containing such provisions. Mandatory warnings and notices are detrimental to our medium-term vision of finding and establishing viable mechanisms to fix the misallocation of user attention in security decisions along the lines sketched in this paper.

Acknowledgments

We thank the anonymous reviewers and the workshop participants for valuable suggestions to improve this paper.

6. REFERENCES

[1] M. Ackerman, L. Cranor, and J. Reagle. Privacy in e-commerce: Examining user scenarios and privacy preferences. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 1–8, Denver, CO, Nov. 1999.

[2] A. Acquisti and J. Grossklags. Privacy and rationality in individual decision making. *IEEE Security & Privacy*, 3(1):26–33, Jan.–Feb. 2005.

[3] A. Acquisti and S. Spiekermann. Do interruptions pay off? Effects of interruptive ads on consumers' willingness to pay. *Journal of Interactive Marketing*, 25(4):226–240, Nov. 2011.

[4] A. Adams and A. Sasse. Users are not the enemy. *Communications of the ACM*, 42(12):41–46, Dec. 1999.

[5] I. Ajzen. The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50(2):179–211, Dec. 1991.

[6] J. Argo and K. Main. Meta-analyses of the effectiveness of warning labels. *Journal of Public Policy & Marketing*, 23(2):193–208, Fall 2004.

[7] H. Beales, R. Craswell, and S. Salop. Efficient regulation of consumer information. *Journal of Law & Economics*, 24(3):491–539, Dec. 1981.

[8] A. Beautement, A. Sasse, and M. Wonham. The compliance budget: Managing security behaviour in organisations. In *Proceedings of the New Security Paradigms Workshop (NSPW)*, Lake Tahoe, CA, Sept. 2008.

[9] L. Bebchuk and R. Posner. One-sided contracts in competitive consumer markets. *Michigan Law Review*, 104(5):827–835, Mar. 2006.

[10] A. Besmer, J. Watson, and H. R. Lipford. The impact of social navigation on privacy policy configuration. In *Proceedings of the Symposium On Usable Privacy and Security (SOUPS)*, page article 7, Redmond, WA, July 2010.

[11] R. Böhme and S. Köpsell. Trained to accept? A field experiment on consent dialogs. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'10)*, pages 2403–2406, Atlanta, GA, Apr. 2010.

[12] R. Böhme and S. Pötzsch. Collective exposure: Peer effects in voluntary disclosure of personal data. In *Proceedings of the International Conference on*

Financial Cryptography and Data Security (FC'11), Rodney Bay, St. Lucia, Feb.–Mar. 2011.

[13] J. Bouckaert and H. Degryse. Opt in versus opt out: A free-entry analysis of privacy policies. In *Proceedings of the Workshop on the Economics of Information Security (WEIS)*, Cambridge, UK, June 2006.

[14] L. Brandimarte, A. Acquisti, and G. Loewenstein. Privacy concerns and information disclosure: An illusion of control hypothesis. Under submission, 2010.

[15] P. Breese and W. Burman. Readability of notice of privacy forms used by major health care institutions. *Journal of the American Medical Association*, 293(13):1593–1594, Apr. 2005.

[16] J. Brustoloni and R. Villamarín-Salomón. Improving security decisions with polymorphic and audited dialogs. In *Proceedings of the Symposium On Usable Privacy and Security (SOUPS)*, pages 76–87, Pittsburgh, PA, July 2007.

[17] E. Brynjolfsson and L. Hitt. Computing productivity: Firm-level evidence. *The Review of Economics and Statistics*, 85(4):793–808, Nov. 2003.

[18] J. Cacioppo and R. Petty. The need for cognition. *Journal of Personality and Social Psychology*, 42(1):116–131, Jan. 1982.

[19] J. Camp and C. Wolfram. Pricing security. In *Proceedings of the CERT Information Survivability Workshop*, pages 31–39, Boston, MA, Oct. 2000.

[20] J. Campbell, N. Greenauer, K. Macaluso, and C. End. Unrealistic optimism in internet events. *Computers in Human Behavior*, 23(3):1273–1284, May 2007.

[21] C. Cannell, P. Miller, and L. Oksenberg. Research on interviewing techniques. *Sociological Methodology*, 16:389–437, 1981.

[22] R. Casamiquela. Contractual assent and enforceability in cyberspace. *Berkeley Technology Law Journal*, 17(1):475–495, Fall 2002.

[23] Center for Information Policy Leadership at Hunton & Williams. HIPAA privacy notice highlights template, February 2003. `http://www.hunton.com/files/tbl_s10News\%5CFileUpload44\%5C10102\%5CHIPAA_Template.pdf`.

[24] S. Chaiken. Heuristic versus systematic information processing and the use of source versus message cues in persuasion. *Journal of Personality and Social Psychology*, 39(5):752–766, Nov. 1980.

[25] S. Chaiken and T. Yaacov. *Dual-process Theories in Social Psychology*. Guilford Press, New York, 1999.

[26] N. Christin, S. Egelman, T. Vidas, and J. Grossklags. It's all about the Benjamins: An empirical study on incentivizing users to ignore security advice. In *Proceedings of the Fifteenth International Conference on Financial Cryptography and Data Security (FC'11)*, Rodney Bay, St. Lucia, Feb.–Mar. 2011.

[27] R. Cialdini, C. Kallgren, and R. Reno. A focus theory of normative conduct: A theoretical refinement and reevaluation of the role of norms in human behavior. *Advances in Experimental Social Psychology*, 24:201–234, 1991.

[28] W. Clapp, M. Rubens, J. Sabharwal, and A. Gazzaley. Deficit in switching between functional brain networks underlies the impact of multitasking on working memory in older adults. *Proceedings of the National Academy of Sciences*, 108(17):7212–7217, Apr. 2011.

[29] S. Clauß and M. Köhntopp. Identity management and its support of multilateral security. *Computer Networks*, 37(2):204–219, Oct. 2001.

[30] Committee on Commerce, Science, and Transportation, Staff Report for Chairman Rockefeller. Aggressive sales tactics on the internet and their impact on american consumers, November 2009. `http://www.hunton.com/files/tbl_s10News\%5CFileUpload44\%5C10102\%5CHIPAA_Template.pdf`.

[31] L. F. Cranor. P3P : Making privacy policies more useful. *IEEE Security & Privacy*, 1(6):50–55, Nov.-Dec. 2003.

[32] R. Dhamija and J. Tygar. The battle against phishing: Dynamic security skins. In *Proceedings of the Symposium On Usable Privacy and Security (SOUPS)*, pages 77–88, Pittsburgh, PA, July 2005.

[33] R. Dhamija, J. Tygar, and M. Hearst. Why phishing works. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'06)*, pages 581–590, Montreal, Canada, Apr. 2006.

[34] P. DiGioia and P. Dourish. Social navigation as a model for usable security. In *Proceedings of the Symposium On Usable Privacy and Security (SOUPS)*, pages 101–108, Pittsburgh, PA, July 2005.

[35] S. Djamasbi, M. Siegel, T. Tullis, and R. Dai. Efficiency, trust, and visual appeal: Usability testing through eye tracking. In *Proceedings of the 43rd Hawaii International Conference on System Sciences (HICSS)*, Kauai, HI, Jan. 2010.

[36] B. Edelman. Adverse selection in online "trust" certifications and search results. *Electronic Commerce Research and Applications*, 10(1):17–25, Jan.–Feb. 2011.

[37] S. Egelman, L. Cranor, and J. Hong. You've been warned: An empirical study of the effectiveness of web browser phishing warnings. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'08)*, pages 1065–1074, Florence, Italy, Apr. 2008.

[38] S. Egelman, J. Tsai, L. Cranor, and A. Acquisti. Timing is everything? The effects of timing and placement of online privacy indicators. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'09)*, pages 319–328, Boston, MA, Apr. 2009.

[39] H. Esser. Können Befragte lügen? Zum Konzept des "wahren Wertes" im Rahmen der handlungstheoretischen Erklärung von Situationseinflüssen bei der Befragung. *Kölner Zeitschrift für Soziologie und Sozialpsychologie*, 38:314–336, 1986.

[40] J. Falkinger. Attention economies. *Journal of Economic Theory*, 133(1):266–294, Mar. 2007.

[41] Federal Trade Commission. Fair information practice

principles. http://www.ftc.gov/reports/privacy3/fairinfo.shtm.

[42] I. Fried. Criticism mounting over Windows 7 security, February 2009. http://news.cnet.com/8301-13860_3-10156617-56.html.

[43] B. Friedman, D. Howe, and E. Felten. Informed consent in the Mozilla browser: Implementing value sensitive design. In *Proceedings of the 35th Hawaii International Conference on System Sciences (HICSS)*, Big Island, HI, Jan. 2002.

[44] N. Fultz and J. Grossklags. Blue versus red: Towards a model of distributed security attacks. In R. Dingledine and P. Golle, editors, *Financial Cryptography*, volume 5628 of *Lecture Notes in Computer Science*, pages 167–183, Berlin Heidelberg, 2009. Springer.

[45] S. Furnell, P. Bryant, and A. Phippen. Assessing the security perceptions of personal Internet users. *Computers & Security*, 26(5):410–417, Aug. 2007.

[46] N. Good, R. Dhamija, J. Grossklags, S. Aronovitz, D. Thaw, D. Mulligan, and J. Konstan. Stopping spyware at the gate: A user study of privacy, notice and spyware. In *Proceedings of the Symposium On Usable Privacy and Security (SOUPS)*, pages 43–52, Pittsburgh, PA, July 2005.

[47] N. Good, J. Grossklags, D. Mulligan, and J. Konstan. Noticing notice: A largescale experiment on the timing of software license agreements. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2007)*, pages 607–616, San Jose, CA, Apr.–May 2007.

[48] J. Grossklags and N. Good. Empirical studies on software notices to inform policy makers and usability designers. In *Proceedings of the International Workshop on Usable Security (USEC'07)*, pages 341–355, Scarborough, Trinidad and Tobago, Feb. 2007.

[49] G. Hardin. The tragedy of the commons. *Science*, 162(3859):1243–1248, Dec. 1968.

[50] R. Haveman. Common property, congestion, and environmental pollution. *Quarterly Journal of Economics*, 87(2):278–287, May 1973.

[51] C. Herley. So long, and no thanks for the externalities: The rational rejection of security advice by users. In *Proceedings of the New Security Paradigms Workshop (NSPW)*, Oxford, UK, Sept. 2009.

[52] M. Hochhauser. Lost in the fine print: Readability of financial privacy notices, July 2001. http://www.privacyrights.org/ar/GLB-Reading.htm.

[53] M. Hochhauser. Compliance v communication. *Clarity: Journal of the International Movement to simplify legal language*, 50:11–19, Nov. 2003.

[54] M. Hochhauser. Readability of HIPAA privacy notices, March 2008. http://benefitslink.com/articles/hipaareadability.pdf.

[55] K. Höök, D. Benyon, and A. J. Munro, editors. *Designing Information Spaces: The Social Navigation Approach*, Berlin Heidelberg, 2003. Springer-Verlag.

[56] B. Huberman and F. Wu. The economics of attention: Maximizing user value in information-rich environments. *Advances in Complex Systems*, 11(4):487–496, Aug. 2008.

[57] S. Iyengar and M. Lepper. When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology*, 79(6):995–1006, Dec. 2000.

[58] C. Jensen and C. Potts. Privacy policies as decision-making tools: An evaluation on online privacy notices. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'04)*, pages 471–478, Vienna, Austria, Apr. 2004.

[59] C. Jensen, C. Potts, and C. Jensen. Privacy practices of internet users: Self-reports versus observed behavior. *International Journal of Human-Computer Studies*, 63(1-2):203–227, July 2005.

[60] F. Keukelaere, S. Yoshihama, S. Trent, Y. Zhang, L. Luo, and M. Zurko. Adaptive security dialogs for improved security behavior of users. In T. Gross et al., editor, *Proceedings of the International Conference on Human-Computer Interaction (INTERACT 2009)*, volume 5726 of *Lecture Notes in Computer Science*, pages 510–523, Berlin Heidelberg, 2009. Springer.

[61] N. Kim. 'Wrap contracts and privacy. In M. Genesereth, R. Vogl, and M.-A. Williams, editors, *AAAI Spring Symposium on Intelligent Privacy Management*, pages 110–112, Menlo Park, CA, Mar. 2010.

[62] Kleimann Communication Group. Evolution of a prototype financial privacy notice: A report on the form development project, February 2006.

[63] R. Kraut, S. Sunder, R. Telang, and J. Morris. Pricing electronic mail to solve the problem of spam. *Human-Computer Interaction*, 20(1):195–223, June 2005.

[64] J. Krosnick. Response strategies for coping with the cognitive demands of attitude measures in surveys. *Applied Cognitive Psychology*, 5(3):213–236, May–Jun. 1991.

[65] J. Krosnick and D. Alwin. An evaluation of a cognitive theory of response-order effects in survey measurement. *Public Opinion Quarterly*, 51(2):201–219, Summer 1987.

[66] P. Kumaraguru and L. Cranor. Privacy indexes: A survey of Westin's studies. Available as ISRI Technical Report CMU-ISRI-05-138, 2005.

[67] S. Laforet and X. Li. Consumers' attitudes towards online and mobile banking in China. *The International Journal of Bank Marketing*, 23(5):362–380, May 2005.

[68] E. Langer. The illusion of control. *Journal of Personality and Social Psychology*, 32(2):311–328, Aug. 1975.

[69] R. Larose and N. Rifon. Promoting i-safety: Effects of privacy warnings and privacy seals on risk assessment and online privacy behavior. *The Journal of Consumer Affairs*, 41(1):127–149, Summer 2007.

[70] Y. Lee and K. Kozar. Investigating factors affecting the adoption of anti-spyware systems. *Communications of the ACM*, 48:72–77, Aug. 2005.

[71] A. Levy and M. Hastak. Consumer comprehension of financial privacy notices: A report on the results of

the quantitative testing, December 2008. http://www.ftc.gov/privacy/privacyinitiatives/Levy-Hastak-Report.pdf.

[72] G. Lindgaard, G. Fernandes, C. Dudek, and J. Brown. Attention web designers: You have 50 milliseconds to make a good first impression! *Behaviour & Information Technology*, 25(2):115–126, Mar.–Apr. 2006.

[73] T. Loder, M. van Alstyne, and R. Wash. An economic answer to unsolicited communication. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 40–50, New York, NY, May 2004.

[74] N. Lundblad and B. Masiello. Opt-in dystopias. *SCRIPTed*, 7(1):155–165, Apr. 2010.

[75] B. Mai, N. Menon, and S. Sarkar. No free lunch: Price premium for privacy seal-bearing vendors. *Journal of Management Information Systems*, 27(2):189–212, Fall 2010.

[76] F. Marotta-Wurgler. Competition and the quality of standard form contracts: An empirical analysis of software license agreements. *Journal of Empirical Legal Studies*, 5(3):447–475, Sept. 2008.

[77] M. Masnick. Supreme court chief justice admits he doesn't read online EULAs or other 'fine print', October 2010.

[78] M. Masson and M. Waldron. Comprehension of legal contracts by nonexperts: Effectiveness of plain language redrafting. *Applied Cognitive Psychology*, 8:67–85, Feb. 1994.

[79] A. McDonald and L. Cranor. The cost of reading privacy policies. *I/S: A Journal of Law and Policy for the Information Society*, 4(3), 2008.

[80] G. McGraw and E. Felten. *Securing Java: Getting down to business with mobile code*. John Wiley & Sons, New York, NY, 1999.

[81] G. Milne, M. Culnan, and H. Greene. A longitudinal assessment of online privacy notice readability. *Journal of Public Policy & Marketing*, 25(2):238–249, Fall 2006.

[82] T. Moore, A. Friedman, and A. Procaccia. Would a 'cyber warrior' protect us? Exploring trade-offs between attack and defense of information systems. In *Proceedings of the New Security Paradigms Workshop (NSPW)*, Concord, MA, Sept. 2010.

[83] S. Motiee, K. Hawkey, and K. Beznosov. Do Windows users follow the principle of least privilege? Investigating user account control practices. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS)*, pages 1–13, Redmond, WA, July 2010.

[84] National Cyber Security Alliance, Norton by Symantec, and Zogby International. 2010 NCSA / Norton by Symantec Online Safety Study, October 2010.

[85] A. Oulasvirta, J. Hukkinen, and B. Schwartz. When more is less: The paradox of choice in search engine use. In *Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 516–523, Boston, MA, July 2009.

[86] M. Overly and J. Kalyvas. *Software Agreements Line by Line: A Detailed Look at Software Contracts and Licenses & How to Change Them to Fit Your Needs*. Aspatore Books, Boston, MA, 2004.

[87] S. Peltzman. The effects of automobile safety regulation. *Journal of Political Economy*, 83(4):677–726, Aug. 1975.

[88] R. Petty and J. Cacioppo. *Communication and Persuasion: Central and Peripheral Routes to Attitude Change*. Springer, New York, 1986.

[89] W. Poole. Financial analyst meeting 2005, July 2005. http://www.microsoft.com/msft/speech/FY05/PooleFAM2005.mspx.

[90] R. Reeder, E. Kowalczyk, and A. Shostack. Helping engineers design NEAT security warnings. In *Proceedings of the Symposium On Usable Privacy and Security (SOUPS)*, Pittsburgh, PA, July 2011.

[91] M. Reichenbach, H. Damker, H. Federrath, and K. Rannenberg. Individual management of personal reachability in mobile communication. In L. Yngström and J. Carlsen, editors, *Information Security in Research and Business (IFIP SEC '97)*, volume 92 of *IFIP Conference Proceedings*, pages 164–174, 1997.

[92] M. Russinovich. PsExec, User Account Control and security boundaries, February 2007. http://blogs.technet.com/b/markrussinovich/archive/2007/02/12/638372.aspx.

[93] M. Salganik and D. Watts. Leading the herd astray: An experimental study of self-fulfilling prophecies in an artificial cultural market. *Social Psychology Quarterly*, 71(4):338–355, Dec. 2008.

[94] R. Schechter. The unfairness of click-on software licenses. *Wayne Law Review*, 4(46):1735–1803, Winter 2000.

[95] S. Schechter, R. Dhamija, A. Ozment, and I. Fischer. The emperor's new security indicators: An evaluation of website authentication and the effect of role playing on usability studies. In *Proceedings of the 2007 IEEE Symposium on Security and Privacy*, pages 51–65, Oakland, CA, May 2007.

[96] S. Schuman and S. Presser. *Questions and Answers in Attitude Surveys*. Sage, Thousand Oaks, 1996.

[97] M. Schunter and C. Powers. The enterprise privacy authorization language (EPAL 1.1), 2003. http://www.zurich.ibm.com/security/enterprise-privacy/epal/.

[98] B. Schwartz. *The Paradox of Choice: Why More Is Less*. Ecco Press (HarperCollins Publishers), New York, NY, 2004.

[99] C. Shapiro and H. R. Varian. *Information Rules. A Strategic Guide to the Network Economy*. Harvard Business School Press, 1998.

[100] D. Sharek, C. Swofford, and M. Wogalter. Failure to recognize fake internet popup warning messages. In *Proceedings of the 52nd Annual Meeting of the Human Factors and Ergonomics Society*, pages 557–560, New York, NY, Sept. 2008.

[101] H. Simon. Designing organizations for an information rich world. In M. Greenberger, editor, *Computers, Communications and the Public Interest*, pages 38–52. John Hopkins Press, 1971.

[102] C. Sims. Rational inattention: A research agenda.

http://sims.princeton.edu/yftp/RIplus/RatInattPlus.pdf, 2006.

[103] J. Sobey, R. Biddle, P. van Oorschot, and A. Patrick. Exploring user reactions to new browser cues for extended validation certificates. In *Proceedings of the 13th European Symposium on Research in Computer Security (ESORICS)*, pages 411–427, Malaga, Spain, Oct. 2008.

[104] S. Spiekermann, J. Grossklags, and B. Berendt. E-privacy in 2nd generation e-commerce: Privacy preferences versus actual behavior. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 38–47, Tampa, FL, Oct. 2001.

[105] B. Stone-Gross, R. Abman, R. Kemmerer, C. Kruegel, and D. Steigerwald. The underground economy of fake antivirus software. In *Workshop on the Economics of Information Security (WEIS)*, Fairfax, VA, June 2011.

[106] F. Strack and L. Martin. Thinking, judging, and communicating: A process account of context effects in attitude surveys. In H.-J. Hippler, editor, *Social Information Processing and Survey Methodology*, pages 123–148. Springer, New York, 2. edition, 1988.

[107] J. Sunshine, S. Egelman, H. Almuhimedi, N. Atri, and L. Cranor. Crying wolf: An empirical study of ssl warning effectiveness. In *Proceedings of the 18th USENIX Security Symposium*, pages 399–416, Montreal, Canada, Aug. 2009.

[108] P. Torr. Demystifying the threat modeling process. *IEEE Security & Privacy*, 3(5):66–70, Sep.–Oct. 2005.

[109] R. Tourangeau, L. Rips, and K. Rasinski. *The Psychology of Survey Response*. University Press, Cambridge, 2000.

[110] M. van Eeten and J. Bauer. *Economics of Malware: Security Decisions, Incentives and Externalities*. STI Working Paper 2008/1, Information and Communication Technologies. OECD, 2008.

[111] H. Varian. System reliability and free riding. In L. Camp and S. Lewis, editors, *Economics of Information Security (Advances in Information Security, Volume 12)*, pages 1–15. Kluwer Academic Publishers, Dordrecht, The Netherlands, 2004.

[112] H. Varian. Computer mediated transactions. *American Economic Review*, 100(2):1–10, May 2010.

[113] C. Viecco, A. Tsow, and J. Camp. A privacy-aware architecture for a web rating system. *IBM Journal of Research and Development*, 53(2):7:1–7:16, Mar. 2009.

[114] T. Vila, R. Greenstadt, and D. Molnar. Why we can't be bothered to read privacy policies: Models of privacy economics as a lemons market. In L. Camp and S. Lewis, editors, *Economics of Information Security (Advances in Information Security, Volume 12)*, pages 143–153. Kluwer Academic Publishers, Dordrecht, The Netherlands, 2004.

[115] R. Villamarín-Salomón and J. Brustoloni. Using reinforcement to strengthen users' secure behaviors. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'10)*, pages 363–372, Atlanta, GA, Apr. 2010.

[116] H. Wang, Y. Hu, C. Yuan, Z. Zhang, and Y. Wang. Friends troubleshooting network: Towards privacy-preserving, automatic troubleshooting. In *Proceedings of the 3rd International Workshop on Peer-to-Peer Systems (IPTPS)*, pages 184–194, San Diego, CA, Feb. 2004.

[117] A. Whitten and D. Tygar. Why Johnny can't encrypt: A usability evaluation of PGP 5.0. In *Proceedings of the 8th USENIX Security Symposium*, pages 169–184, Washington, DC, Aug. 1999.

[118] B. Whitworth. Politeness as a social software requirement. *International Journal of Virtual Communities and Social Networking*, 1(2):65–84, 2009.

[119] P. Wilson. *Second-hand knowledge: An inquiry into cognitive authority*. Greenwood, Westport, CT, 1983.

[120] F. Wu and B. Huberman. Novelty and collective attention. *Proceedings of the National Academy of Sciences*, 104(45):17599–17601, Nov. 2007.

[121] M. Wu, R. Miller, and S. Garfinkel. Do security toolbars actually prevent phishing attacks? In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'06)*, pages 601–610, Montreal, Canada, Apr. 2006.

[122] M. Zurko, C. Kaufman, K. Spanbauer, and C. Bassett. Did you ever have to make up your mind? What Notes users do when faced with a security decision. In *Proceedings of the 18th Annual Computer Security Applications Conference (ACSAC)*, pages 371–381, Las Vegas, NV, Dec. 2002.

APPENDIX

A. DERIVATIONS

Probability of Occurrence.

Assume p_G is drawn uniformly and independently from $[0, 1]$ in each round.

Then we have the probabilities of system reactions as a function of the thresholds,

$$P(H) = \tau_B, \tag{1}$$
$$P(A) = \tau_G - \tau_B, \tag{2}$$
$$P(C) = 1 - \tau_G, \tag{3}$$

and the probability of the good state conditional to the system reaction,

$$P(G \mid H) = \frac{1}{2}\tau_B, \tag{4}$$
$$P(G \mid A) = \frac{\tau_G + \tau_B}{2}, \tag{5}$$
$$P(G \mid C) = \frac{1 + \tau_G}{2}. \tag{6}$$

Let $k = \{1, 2, \dots\}$ be the round number. The the probability of optimizing and satisficing conditional to A and the round number k are given by:

$$\begin{aligned} P(a > 0 \mid A, k) &= (1 - P(A))^{(k-1)} \\ &= (1 - \tau_G + \tau_B)^{(k-1)} \end{aligned} \tag{7}$$
$$\begin{aligned} P(a = 0 \mid A, k) &= 1 - P(a > 0 \mid A, k) \\ &= 1 - (1 - \tau_G + \tau_B)^{(k-1)} \end{aligned} \tag{8}$$

Defender's Utility.

Defender's utility can be calculated by inserting the valued from the payoff matrix.

$$D(H) = 0 \tag{9}$$
$$D(A) = (1 - \tau_G + \tau_B)^{(k-1)} \cdot \frac{\tau_G + \tau_B}{2} \cdot \gamma u \tag{10}$$
$$+ \left(1 - (1 - \tau_G + \tau_B)^{(k-1)}\right) \cdot \left(\frac{\tau_G + \tau_B}{2}\right)^2 \cdot \gamma u \tag{11}$$
$$D(C) = \frac{1 + \tau_G}{2} \cdot \gamma u - \frac{1 - \tau_G}{2} \cdot s \tag{12}$$
$$D(k; \tau_B, \tau_G) = P(A)D(A) + P(C)D(C) \tag{13}$$

We discount future payoffs with factor $r \in (0, 1)$.

$$D(\tau_B, \tau_G) = \sum_{k=1}^{\infty} D(k; \tau_B, \tau_G)(1 - r)^{(k-1)} \tag{14}$$

This is the final expression used in the analysis and as objective function of the defender's optimization problem.

User's Utility.

User's utility can be calculated by inserting the valued from the payoff matrix.

$$U(H) = 0 \tag{15}$$
$$U(A) = (1 - \tau_G + \tau_B)^{(k-1)} \cdot \frac{\tau_G + \tau_B}{2} \cdot u \tag{16}$$
$$+ \left(1 - (1 - \tau_G + \tau_B)^{(k-1)}\right) \tag{17}$$
$$\cdot \left[\left(\frac{\tau_G + \tau_B}{2}\right)^2 \cdot u - 2s\left(\frac{\tau_G + \tau_B}{2}\right)\left(1 - \frac{\tau_G + \tau_B}{2}\right)\right] \tag{18}$$
$$U(C) = \frac{1 + \tau_G}{2} \cdot u - \frac{1 - \tau_G}{2} \cdot s \tag{19}$$
$$U(k; \tau_B, \tau_G) = P(A)U(A) + P(C)U(C) \tag{20}$$
$$U(\tau_B, \tau_G) = \sum_{k=1}^{\infty} U(k; \tau_B, \tau_G)(1 - r)^{(k-1)} \tag{21}$$

This is the final expression used in the analysis.

Attention Tax.

The above analysis can be repeated with the updated payoff matrix and then solving for t such that

$$\frac{P(A)\left[D(A) - t\right] + P(C)D(C)}{P(A)U(A) + P(C)U(C)} = \text{const.} \tag{22}$$