

Who's Afraid of Which Bad Wolf? A Survey of IT Security Risk Awareness

Marian Harbach, Sascha Fahl, Matthew Smith
Usable Security and Privacy Lab
Leibniz University Hannover
Hannover, Germany
{harbach,fahl,smith}@dcsec.uni-hannover.de

Abstract—The perception of risk has been established as an important part of the study of human aspects of security research. Similarly, risk awareness is often considered a central precursor for the adoption of security mechanisms and how people use them and interact with them. However, the state of risk awareness in users during their everyday use of the modern Internet has not been studied in detail. While it is well known that users have a limited “budget” for security behavior and that trying to coerce them into considering additional risks does not work well, it remains unclear which risks are on users’ minds and therefore already accounted for in terms of their budget. Hence, assessing which risks and which consequences users currently perceive when using information technology is an important and currently overlooked foundation to shape usability aspects of IT security mechanisms. In this paper, we present a survey of risk and consequence awareness in users, analyze how this may influence the current lack of adoption for improved security measures, and make recommendations how this situation can be alleviated.

I. INTRODUCTION

Research in the area of security and privacy technology faces many challenges. Consequently, in the past decade, a growing amount of research has focused on human factors and usability issues of security mechanisms [1]. Early research of Sasse et al. [2] established that, in contrast to other branches of usability research, usable security technologies are particularly difficult to create, as security is rarely a primary goal or task for users.

Additionally, not only do HCI researchers need to design security and privacy technologies that enable people to remain safe while following their primary task, they also need to create a wish for adoption of these technologies beforehand:

“The challenge is not to enable the individual’s mastery of an application so much as to convince the individual to avoid digital risks by adopting appropriate security tools and application settings, despite the financial and time costs of doing so.” [3]

The problem with convincing users of security risks as well as the interplay with financial and time costs of security measures has been investigated by Beutement et al. [4] and Herley [5]. According to this work, users’ *compliance budgets* are limited and therefore users make a rational choice when they reject (new) security measures if they do not perceive enough benefits. The compliance budget was defined in a corporate context, where security policies are forced on users, who may then choose not to comply.

We extend the reasoning behind the compliance budget to the regular Internet user at home and ask ourselves why users do not adopt security measures that can protect them from the many risks on the Internet. Consequently, the question arises why certain measures are currently within users’ budgets. As there is no mandated policy to comply with, one can think of users spending their compliance budget on security measures as they seem fit. If we assume that users spend their budget on the most important and salient risks more easily, then there may be some risks users could be protected from but do not care enough to take the necessary precautions and some risks users would like to be protected from but aren’t at the moment. Similarly, by looking at salient risks, we may find why some protection measures are currently not within the compliance budget: if a risk is considered unimportant or entirely unknown, users will not have any desire to protect themselves against that risk.

The importance of risk perception for HCI research can also be seen in the large body of HCI security research that analyzed warning messages and attempted to improve risk communication to elicit safe behavior [6]–[9]. Additionally, the factors that govern the perceived severity of a risk have previously been subject to investigations [10]. Researchers have also investigated how users perceive certain threats and risks [10]–[12]. Furthermore, risk plays a major role in technology acceptance models [13], [14], which analyze the human aspects involved in the adoption of new technology. For security technology, previous research in this area indicates that a subject needs to be aware of a threat – or risk – and then come to the conclusion that this threat needs to be dealt with. Only then is a security technology evaluated for its suitability, for example its usability and capability to protect against said threat.

Interestingly, to the best of the authors’ knowledge, little is known about which IT security risks users actually feel exposed to and are aware of during their everyday dealings with today’s Internet. Previous research on risk perception has often focused on describing specific threats before analyzing how users perceive those. Similarly, representative national surveys, such as the Oxford Internet Survey for the U.K. [15] or the Security Report for Germany [16], only ask participants about attitudes towards enumerated risks. While this gives an understanding of how users perceive a specific risk when prompted, these results can not give insight into the risks actually perceived in the users’ everyday lives. As users are no security experts, they can and will only pay attention to

risks they are aware of. Additionally, users may evaluate risks differently from security professionals, for example by considering other threats or reasoning differently about possible consequences. There has, however, not been any recent work addressing this bottom-up view of risk awareness. We believe that this lack of knowledge is detrimental to the foundation on which security and privacy researchers base their work. We argue that it is important to know which risks users are aware of and how they appraise them, since they will only take action based on their own, intrinsic appraisal.

Another factor that influences the perception of risk and therefore IT security technology is its perceived benefit. Beautelement et al. [4] mention avoiding consequences and punishment as key benefits, diminishing the costs a mechanism creates within the compliance budget. Outside a corporate environment (in which Beautelement et al.'s reasoning was situated), negative consequences are the only kind of punishment a users has to fear, as there is no regulatory body punishing misbehavior. Similarly, the technology acceptance model of Herath et al. [14], also features an assessment of effectiveness – i.e. is a technology able to prevent certain consequences from happening – in the appraisal of security technology. Hence, only the prevention of certain relevant consequences using a certain behavior or technology will constitute a benefit for users that can make them accept a certain cost in terms of effort. We therefore hypothesize that without relevant consequences perceived by users, they are unlikely to adopt security technologies or change behavior to guard against risks, even if the risks themselves are known.

To offer a foundation on which to base future developments of security measures, we explore users' risk as well as consequence awareness. Our results can help researchers and developers target the risks actually perceived by users in given situations and also highlight how security technologies need to address the most important risks perceived by users. By looking at the perceived consequences, we can also analyze if there are risks users are aware of, but against which they do not protect themselves, because they do not perceive any relevant benefits. It is our aim to offer researchers a guide as to which risks they can address directly and from which consequences they need to protect users. Overall, we argue that knowing which risks and consequences users want to be protected from is an important precursor for technology adoption or behavior change. A previous investigation of a large-scale security technology roll-out showed that if these human aspects are not adequately addressed, promising security and privacy technologies will not be adopted no matter how good they technically are [17].

In this paper, we give an overview of users' risk awareness while using today's Internet, based on a survey of 210 participants from two different populations. The results suggest that the sporadic and slow adoption of new security technologies, for instance replacing username and password, is not only due to usability problems but is also rooted in these technologies not addressing salient risks. We also find evidence that there are some culture-specific risks, but many risks are also common to all our participants: Malware, hackers and stealing account credentials were the most salient risks and financial losses accounted for the most frequently perceived consequences. We also compare the awareness of salient risks

to agreement with a set of risks commonly warned against: Our results suggest that users are aware of far fewer risks than may currently be believed. Overall, the lesson we learned from this study is that there appear to be two avenues for improving end-user security in the future: The easier one is to directly address those risks which users are already aware of and which are salient to them. The second is trying to support the process of changing risk perception, for example using risk communication and education.

The remainder of the paper is organized as follows: Related work is introduced first, before detailing the research method, the participants and the results. We discuss the implications of and hypothesis arising from our results before we conclude. The appendix presents additional information, comprising the code plan we used to categorize participants' responses as well as the questionnaire contents.

II. BACKGROUND AND RELATED WORK

In 2002, Friedman et al. [18] presented a short paper on a study of users' concerns about risks during the use of "the Web". The focus of their study was to analyze the effect of communities on risk awareness. They interviewed 72 people and found that participants "most often emphasized security, privacy, and threat[s] to computer systems" as potential risks. The more "high-tech" a community was, the more they were concerned about security and privacy. Additionally, users from a suburban community were more concerned about people and their experiences, while a rural community showed significantly less concerns overall. These results were obtained when the Web was relatively young and its use was not as common as it is today. Unfamiliarity and novelty has likely played a role in the participants' views more than 10 years ago and the authors argue that the investigated communities will likely progress towards the views of the high-tech community. Since this study was conducted, interaction with technology and the Internet has changed significantly and it is hence important to draw a current picture of which risks people see for themselves in this changed environment. We extend the study design of Friedman et al. in an effort to create a recent and more detailed understanding of people's general risk awareness when using the modern Internet with a focus on IT security risks.

As mentioned in the introduction, there also are several representative, national surveys compiled on an annual basis about views on IT security and risks (e.g. [15], [16]). These ask their participants about their attitudes towards certain enumerated risks, for example having their computer infected with malware or being under surveillance by governments. This method works well to judge the relative relevance of different sources of risks, but does not address which risks may actually influence day to day behavior. We argue that users will only consider measures against certain risks for their compliance budget if they are intrinsically aware of these risks and hence consider them worth their time.

There also is existing work on risk perception that analyzed users' mental models [12] or how to incorporate these into security solutions [6] with respect to specific threats, such as phishing, hacking or malware. Dhamija et al. [19] have argued that phishing works because users do not sufficiently understand the technology, and therefore its risks. Additionally,

researchers investigated how to communicate a particular risk to a particular user group [20] or which factors influence the perception of specific risks [10], [11]. Further work focused on ways to communicate IT security risks to users more efficiently (e. g. [21]–[23]). These approaches allowed the usable security and privacy community to instill a certain amount of awareness for particular security and privacy measures in users (e. g. anti-phishing measures).

In these papers, however, risks were assumed to arise from specific threats, such as phishing, hackers, or malware, ignoring the fact that users may not be aware of such threats or believe that they do not apply to them. Threats were always simply presented to the users in these studies. For users to adopt security measures protecting them from these threats, they will have to first discover a measure and then decide that it is worth the effort to actually use it. This appraisal, using the model of the compliance budget introduced above, will only have a positive outcome if a benefit, protecting oneself from a risk and its consequence, is perceived. However, previous work has found that participants do often not differentiate threats or risk at all. In a related study [17], participants were confident that using only two different passwords across all their online accounts was safe and saw no problems or risks arising from that practice. Conversely, multiple participants also expressed that they treat the Internet as a generally insecure medium and that they therefore, for example, do not use online banking at all. Participants also expressed doubts that security technologies would actually protect them from the risks they perceive. Among other comments, one participant believed that password managers “*surely could be hacked by someone*”. Another participant said: “*I don’t believe that there will ever be perfect security on the Internet. Whether you use [an alternative mechanism] or continue using passwords [...] there are vulnerabilities everywhere*”. In a study by Fahl et al. [24], many subjects believed that there will be a way to circumvent any security system at some point in time and that virtually anything on the Internet was vulnerable to attack or “hacking”. Similarly, Klasnja et al. [25] found that users “lack understanding of important privacy risks” when connecting to and using Wi-Fi networks. This corroborates the intuition that a specific risk as well as concrete benefits need to be perceived in order for users to consider countermeasures.

A non-tech oriented study by Hogarth et al. [26] investigated everyday risk perception. Participants recorded one risk and the most severe consequence involved in whatever they were doing when receiving a text message from the researchers three times per day on work days over the course of two weeks. They found that the most frequently reported risks were the most salient ones as opposed to the most severe ones. Consequently, they conclude that the risks users are aware of are only a subset of the risks actually faced. Many of the everyday risks commonly studied were also entirely absent from the risks reported by their participants.

To the best of the authors’ knowledge, there has not yet been a recent investigation of the risks and consequences perceived during everyday Internet use without being queried about risks arising due to specific threats. Additionally, this is the first study to analyze risk awareness with a focus on IT security in the modern Internet, providing ideas to facilitate the adoption of security measures.

III. ONLINE RISKS SURVEY

We designed a questionnaire and ran an online survey, aiming to assess the risk and consequence awareness of users. In contrast to Friedman et al. [18], we chose a survey as our research method because surveys can reach people in familiar settings. While it is well known that using surveys to ask people about past behavior causes biases, our survey used scenarios to get people into a certain mindset before eliciting their attitudes within that mindset.

We were concerned that inviting people to interviews may cause biases. For instance, participants who do not see many risks may feel obligated to name more risks in order not to seem careless. It also seemed likely that participants would not share risks they are afraid of because they might feel ashamed. It has been shown that responses are more truthful and open when given in private [27]. Yet, interviews can yield deeper insights into users’ reasoning if care is taken to minimize bias and can therefore be complementary. We thus decided to gain a broader overview first and conduct interviews in future work to explore the results of our survey in more depth.

We also chose a survey over more fine-grained methods, such as the experience sampling used by Hogarth et al. [26], as we believe that it is necessary to gain an initial understanding across a wider range of Internet users. Especially differences in culture and beliefs may influence risk awareness which we would not be able to capture as easily using other methods. To investigate the adoption of security technologies on the Internet’s scale, a wider view is important. We therefore ran the survey on two continents, using a local student population in Germany as well as workers from Amazon’s Mechanical Turk. The questionnaire design and participant demographics are detailed in the following subsections.

A. Questionnaire

The questionnaire was structured to elicit a set of risks to which participants believe to be subject to during their daily Internet conduct. We presented five scenarios in which participants were asked to list which risks they are aware of. These scenarios comprised “using the Internet in general”, “logging in to your social network account”, “shopping online”, “online banking”, and “finding a shared ride using online services”. The reasons for including each of the scenarios are as follows:

- *General Internet use*: This scenario was chosen to induce as little priming as possible to try and capture the base line of risks users are aware of. We hypothesized that users may apply this mindset when considering general protection measures without a specific application. The remaining scenarios include common online use cases.
- *Online shopping and banking* were chosen since they both include obvious financial risks. We chose two financial scenarios so we could examine whether the type of institution influences the risks and consequences users state. While banking constitutes a more severe scenario, it could also be perceived to have less risks, as banks take more precautions to protect their customers.
- *Logging in to a social network site* was chosen because it is a very common activity and is often paid little attention to or even perceived as annoying. It suffers from the common problem of IT security mechanisms, since it is a

barrier keeping users from achieving their primary goal, which in this case is to take part in a social network site. This scenario was chosen, as social networking accounts often hold more sensitive data and therefore potentially have a different protection value than for example credit card details in the shopping scenario.

- *Sharing a ride using online services* was chosen because it includes direct real-world implications, as the user will meet with another person in the physical world. We included this scenario to capture a potential relationship between real-world, physical risks and abstract, technological online risks.

Each scenario was introduced to participants with a short description. The description was brief and simple and aimed to let the participants imagine how they usually interact with such services. The text mentioned a well-known workflow for each scenario and reminded participants to imagine that they were completing this task in a familiar environment. This part of the description aimed to overcome the issue of trust, as they should trust their favorite shopping site, social network or bank similarly. For example, the description of the shopping scenario read:

“Please imagine that you are using the Internet at home as usual. You are visiting your favorite online shopping site and would like to make a purchase. You enter your address and payment information on the site and complete the checkout process. Please answer the following questions in the context of purchasing merchandise in an online shop.”

For each scenario, we first asked participants to state the most severe risk or danger they believe to be subject to within this scenario. We intended that participants state whatever came to their mind and the questionnaire therefore provided text boxes for free-text answers. The text boxes were sized to accommodate approximately one sentence in one line so as not to overwhelm participants. We specifically chose not to define or otherwise explain the concept of risk, as we intended to capture how participants intuitively respond to our questions. We believed that if we asked participants to state a risk according to a certain definition, rationalization would overlay their initial responses as they would think too much about their answers in the context of a given definition. While we acknowledge that this lack of specificity may cause some risks which are technically similar to be described in different ways, this approach prevents that we miss differences in users’ reasoning due to forcing them into a certain definition. Since these differences are one of the main focuses of this work, we chose to accept this limitation and use our results only to derive hypotheses about which differences in risk perception may be observable.

Also, one of the very fundamental techniques taught to security professionals is to evaluate and rank risks by combining the likelihood of a risk and the severity of consequences. Security professionals are trained in differentiating these terms and making decisions based on the technical understanding and hopefully well founded experience. If the general population does not make the same distinction in these terms, their basis for making decisions is different. We believe examining the differences in understanding and perception is a vital founda-

tion to understanding our users and their decisions better. As argued above, we would not have been able to examine this important difference if we had provided participants with expert definitions during the survey.

Next, participants were given the chance to enter three additional risks for each scenario, before being asked when they had last heard about each of these risks from common sources, including friends, family, and media. We then asked participants to rate the completeness of their set of risks and to give an estimate of relative risk arising to their wellbeing in general on a scale from 0 (no risk at all) to 100 (greatest possible risk).

Participants were then requested to state four potential consequences in order of severity for each scenario. As before, we chose open-ended questions with free-text answers so as not to influence their answers. As before, we did not provide a definition or explanation of the term consequence, in order to preserve the participants’ mental models as much as possible. We asked for consequences, as this will allow us to better assess to what extent users conform to the common model of security professionals and also to what extent they are actually ready to do something against a risk or at least to what extent they believe a risk applies to them. Based on previous work [4], [17], we assume that if users see only very improbable or impersonal consequences, they are very unlikely to see a need for measures against the corresponding risks. Therefore, our participants were also asked to judge the relative severity of the most severe consequence in comparison with those arising from other risks and dangers in their life. Additionally, each scenario concluded with a question about the perceived likelihood of the most severe consequence happening to the participant personally. An overview of the questions included in the questionnaire for each scenario can be found in the Appendix.

The questionnaire ended with a block of questions giving the participants a pre-compiled list of 22 common risks users are often warned against or which featured on popular websites about online risks, asking them which they know about and how relevant they consider those on a scale from “not relevant at all” (1) to “very relevant” (4). It included common risks like malware, spam, phishing or online shopping fraud, but also other dangers, such as psychological issues due to exposure to unsuitable content, cybermobbing or Internet addiction. In this part of the questionnaire we intended to compare the open-ended answers from the first part to the risk users would say are relevant or they know about if the risks are presented to them in a list. Users’ views on risks are often collected in this fashion in representative surveys (see above). We believe that this introduces biases that makes results based on enumerated risks less suitable for analyzing why security measures are adopted or not.

The questionnaire also asked participants about their perception of risks and security on the Internet in general, how much which sources of risk information influence their perceptions, and if they had previously been subject to any of the consequences and risks they gave before. Lastly, demographics were collected.

B. Participants

As mentioned above, we recruited students from our university’s study participation mailing list and submitted a task to Amazon’s Mechanical Turk in July 2013. For the university students, the survey was administered in German. The students were offered to enter a raffle of 30 10 Euro Amazon vouchers. On MTurk, we invited only U.S.-based Master workers to our task, offering \$3 for a 20-30 minute survey. Amazon screens Master workers for reliability and they receive a higher compensation per task. This should diminish the impact of workers trying to make as much money with as little effort as possible on the reliability of our results. We still checked all results for irregularities and obvious patterns in the answers and removed one participant who was answering randomly.

Our choice of participant recruitment offers a look at two rather different populations and allows us to spot major differences in risk awareness between these different countries, education and age groups. While the two samples do not represent the countries in general, they do offer a fairly broad view across a diverse set of people.

We received $N_1 = 111$ complete questionnaires from the university students and $N_2 = 99$ complete and valid questionnaires from MTurk. Students spent 24.7 minutes ($sd = 16.5$ min) and Turkers 22.7 minutes ($sd = 11.3$ min) on the questionnaire. Table I gives an overview of participant demographics. MTurk workers were older and comprised more females. Their IT experience was similar while previous experiences with online risks and dangers was reported to be higher by students.

TABLE I. PARTICIPANT DEMOGRAPHICS FOR BOTH SURVEY DEPLOYMENTS.

	Students	Turkers	
N	111	99	
Age Range	18-42	19-66	years
Median Age	23	36	years
Gender	45.0 %	60.6 %	female
	55.0 %	39.4 %	male
Occupation		37.4 %	full-time employee
	100 %	6.1 %	student
		11.1 %	part-time worker
		20.2 %	self-employed
		11.1 %	homemaker
		9.1 %	unemployed
		5.1 %	retiree
IT Experience	21.6 %	18.2 %	is currently or has been working in or studying IT
Risks	59.0 %	35.4 %	previous incidents
	6.7 %	4.0 %	N/A

1) *Differences between Students and Turkers:* As we aim to investigate risk awareness for an as broad as possible population and to keep the results as concise as possible, we combine the two datasets for the analysis we present below. Whenever there were significant differences in the results for a certain aspect, the respective results subsection includes a description of how the two populations differed. Otherwise, the conclusions drawn from the data apply to both populations equally. The differences we found will also be discussed in the Discussion section.

C. Coding

To analyze participants’ responses to the open-ended questions on risks, we used an inductive coding procedure. We chose this method to be able to flexibly represent the responses, as there is, to the best of the authors’ knowledge, no previous research on fine-grained coding and categorizing for user risk awareness concerning the Internet.

We began coding with the list of 22 common risks that was also included in the later part of the questionnaire. One coder went through all 4,200 responses, adding codes whenever an answer did not match an existing code. Codes were also hierarchically refined if a response fitted an existing code but addressed a more specific aspect. Each code could only be assigned once per user and scenario and was otherwise marked as duplicate. We also filtered for responses that were not descriptions of a risk but of a consequence or something else. After this first coding session, codes were refined in a discussion among the authors and a second coder went through the responses again using the refined coding scheme. The same process was applied to the open-ended responses on possible consequences.

Overall, we created 74 risk codes and 38 consequence codes. A table of all risk codes and their hierarchy can be found in the Appendix together with the counts for each code. We did not consolidate the codes further for the purpose of this paper, as we aim to explore risk awareness and its implications for the human factors of the adoption of security technology. A categorization framework is subject to future work.

D. Results

Altogether, 210 participants had the chance to specify 20 risks each. Of these potential 4,200 risks, a total of 1,795 valid responses across the five scenarios were given (median of seven unique risks per participant). The remaining responses were either empty or filtered for several reasons (see below). Figure 1 provides a graphical overview of all mentioned risks, grouped into the categories that emerged during coding. In general, concerns about privacy, account abuse, malware and hackers, as well as financial risks and fraud were most commonly mentioned. There is also a fair amount of miscellaneous risks that were mostly mentioned in the ride-share scenario, as “unreliability of people” was a frequently stated risk in this case. Since Figure 1 also suggests that risk awareness depends essentially on the presented scenario, we look at the scenarios individually below.

Differences Between Populations: We performed a Fisher’s exact test over the data source \times risk-code contingency table. As the table was too large to compute all permutations and had more than 70 % of expected counts at less than five due to the sparse nature of many risks participants specified, we used Monte-Carlo simulations to obtain an approximate p-value. The test indicated a highly significant difference between the two data sources ($p < .0001$, 100,000 replicates). Significant standardized residuals in the contingency table revealed differences within the following codes (all these cells had expected counts larger than 5): Turkers more frequently gave identity theft, abuse of bank details, and theft of physical items as possible risks. The

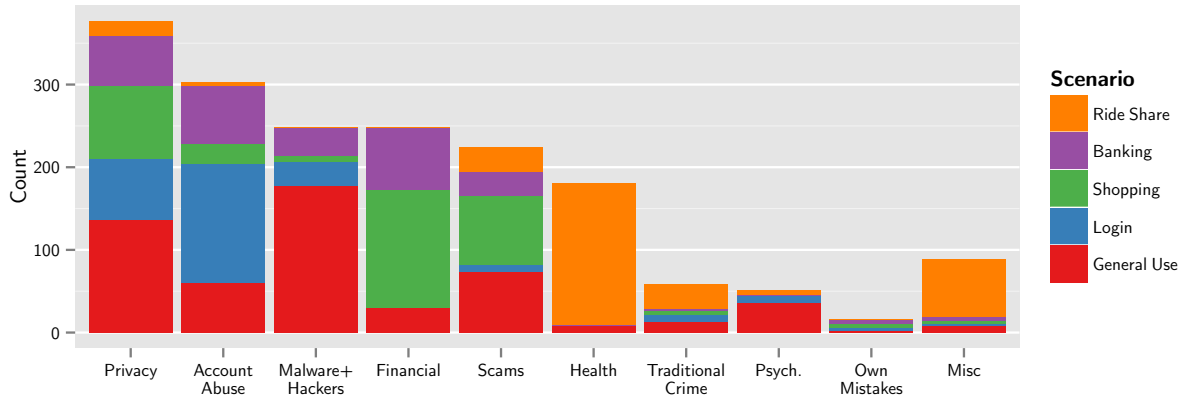


Fig. 1. An overview of all risks mentioned by the respondents grouped by categories and color-coded by scenario (cf. Section III-A). Counts for all risk codes can be found alongside the codeplan in the Appendix. The counts of mentioned risks differed significantly from the expected counts for the categories privacy (students 220/turkers 157), scams (106/118), health (78/102), and traditional crime (12/47) due to the differences in mentioned risks detailed in Section III-D1.

German students more often named fraud, hidden cost, and abuse of personal data (cf. the counts in the Appendix).

1) *Scenarios*: Table II provides an overview of top risks by presented scenario, including these risks where the populations differed. The following results nevertheless largely apply to both populations. In the general use scenario, we find that users are most concerned about malware, identity theft, stolen account credentials and hackers. The lesser concerns mostly comprise privacy risks. The great fear of malware is interesting, as a plentitude of software exists that can help users to protect them from this threat, yet users apparently still feel that they can be subject to adverse effects from it. We suspect that participants may nevertheless be afraid to unknowingly contract malware due to a lack of understanding for technical complexities. The fear of hackers possibly arises from a similar source, as participants may believe that a competent person can break into almost any IT system.

Looking at the top risks mentioned in the four specific scenarios, we find that the more technical risks of the general scenario are superseded by risks that mirror the described scenario. When considering shopping online or online banking, risks pertaining to the respective task, such as abusing account details or fraudulent merchants, were considered to be most important. This confirms the common view that security is a secondary issue, even with respect to risk awareness. Notably, stealing account credentials was more relevant in the Banking scenario, even though banks usually take greater care to protect their customers.

In the scenario that specified an IT-security relevant activity (logging in to a social network), we see that IT security risks were considered most important again. In consequence, this could indicate that users only see a necessity for improved security mechanisms (that target IT security risks) when these risks are obvious in the corresponding situation and are not overlaid by other, more important, non-IT-security concerns.

There also are two risks that occur in most of the specific scenarios: identity theft and stealing private information appear to be concerns that are cross-cutting for most Internet usage

scenarios. While private information often inherently needs to be entrusted to online services if they are to provide a useful service, identity theft could be more difficult if authentication of individuals used appropriate protection measures. This is also a risk that was particularly pronounced for participants from the U.S., as social security numbers are often used for authentication in important and official workflows in this country.

We furthermore compared the provided risks across all responses with the first response participants gave in each scenario. In the questionnaire, a scenario would be introduced and the participants were then asked to first state the greatest risk they thought may arise from this scenario. We found that while the top two or three risks provided based on the greatest risks and on all given risks respectively did not differ, there were some risks that seemed to only occur to participants on second thought (cf. third column in Table II). These included stealing of account credentials using a specific means in the general use scenario, malware in the login scenario, fraud in the shopping scenario, identity theft in online banking, and theft of physical things in the ride share scenario. Hence, if users assess a security measure in a short amount of time, some risks may not be considered. Similarly, future studies that wish to elicit a comprehensive overview of specific risks should therefore plan to include more than one opportunity to specify a risk.

2) *Filtered Responses*: As noted above, participants had a total of 4,200 slots into which they could enter risks. 2,021 slots were left empty. Non-empty slots were filtered further: duplicates, where the same participant stated the same risk twice in one scenario, were found in 201 cases. In thirteen cases, participants stated that a scenario did not apply to them and in 29 cases they stated that they did not see any risk in this scenario. Finally, there were 21 cases where answers were considered off topic by the coders and six cases where a participant was unsure about potential risks for a scenario. Interestingly, in additional 114 cases (74 unique users), we found that users specified things that did not refer to an actual risk. While this can be an effect of fatigue, it may to a certain

TABLE III. USERS' RESPONSES WHEN ASKED HOW COMPLETE THEY ESTIMATE THE SET OF RISKS THEY PROVIDED TO BE. THE TABLE SHOWS PROPORTIONS OF ANSWERS SEPARATED BY SCENARIO AND POPULATION (STUDENTS, $N = 111$, AND MTURK, $N = 99$). AN ASTERISK DENOTES A SIGNIFICANT DIFFERENCE BETWEEN THE TWO POPULATIONS IN THIS SCENARIO USING A FISHER'S EXACT TEST AND STANDARDIZED RESIDUALS GREATER THAN 1.96.

Answer	General Use			Login			Shopping			Banking			Ride Share		
	S	T	Overall	S	T	Overall	S	T	Overall	S	T	Overall	S	T	Overall
all relevant	7.3%	52.0%	28.4%*	22.9%	56.1%	38.9%*	20.0%	59.6%	38.8%*	18.2%	54.6%	35.4%*	13.0%	51.0%	31.1%*
most important	64.5%	27.6%	47.1%*	32.4%	11.2%	22.2%*	50.9%	13.1%	33.0%*	45.5%	18.2%	32.5%*	44.4%	16.3%	31.1%*
all known	19.1%	17.3%	18.3%*	35.2%	29.6%	32.5%*	24.6%	24.2%	24.4%*	28.2%	23.2%	25.8%*	26.9%	25.5%	26.2%*
feel safe	3.6%	3.1%	3.4%*	3.8%	2.0%	3.0%*	1.0%	1.0%	1.0%*	2.7%	4.0%	3.4%*	0.0%	3.0%	1.5%*
space full	3.6%	0.0%	1.9%*	1.0%	1.0%	1.0%*	0.0%	2.0%	1.0%*	1.0%	0.0%	0.5%*	3.7%	3.0%	3.4%*
no answer	1.8%	0.0%	1.0%*	4.8%	0.0%	2.5%*	3.6%	0.0%	1.9%*	4.5%	0.0%	2.4%*	12.0%	1.0%	6.8%*

TABLE II. RISKS MENTIONED BY MORE THAN 5% OF RESPONDENTS GROUPED BY SCENARIO. THE MS COUNT COLUMN LISTS HOW FREQUENTLY THIS RISK WAS STATED AS THE MOST SEVERE RISK (BEING ELICITED FIRST IN THE QUESTIONNAIRE). THE S/T COLUMN LISTS COUNTS FOR THE STUDENT (S) AND MTURK (T) POPULATIONS SEPARATELY IF THE COUNTS DIFFERED SIGNIFICANTLY BETWEEN THEM.

Risk	Count	S/T	MS Count
General Use	548		205
Malware	121		42
Identity theft	55	8/47	28
Stealing account credentials (specific)	46		5
Targeted attacks by third parties ("Hackers")	42		18
Stealing private information	41		22
Loss of privacy in general	26		13
Surveillance	23		12
Abuse of credit card/banking details	21		10
Abuse of personal data	15		7
Stalking	11		1
Shopping	358		200
Abuse of credit card/banking details	125		108
Fraud	35		10
Stealing private information	34		13
Passing private information on to third parties	26		13
Offering non-existent merchandise or services	21		7
Identity theft	16	0/16	6
Stealing account credentials (unspecific)	11		6
Banking	281		184
Abuse of credit card/banking details	58		49
Stealing account credentials (unspecific)	33		28
Stealing account credentials (specific)	29		24
Stealing private information	24		17
Identity theft	23	1/22	7
Targeted attacks by third parties ("Hackers")	22		20
Surveillance	12		7
Malware	12		3
Login	278		184
Stealing account credentials (specific)	67		60
Stealing account credentials (unspecific)	64		59
Stealing private information	19		10
Targeted attacks by third parties ("Hackers")	16		12
Surveillance	12		5
Ride Share	330		195
Risk to health and wellbeing	157		131
Unreliability of other people	69		19
Theft of physical things	19	0/19	0
Fraud	17		15

extent also be indicative of users not thinking about risks like experts. We did not find significant differences in the counts for filtered responses between the student and MTurk population. All in all, we gathered 1,795 valid responses. This suggests that users are only aware of a limited set of risks: only three participants exhausted all 20 fields to enter risks.

3) *Completeness*: When asked to what extent participants believe their answers to be complete for each scenario, they showed high confidence in their answers: In 34.5% of all cases, the participants stated to be sure to have entered all

risks that are relevant for them. In an additional 25.4% of all instances, participants indicated that they did not know about any additional risks. In 2.9% of all cases, participants preferred not to answer this question, in 1.5% of all cases no additional risks were entered because the questionnaire did not provide more space and in 2.4% of the cases no additional risks were entered because participants stated to feel safe on the Internet. In 33.2% of all instances, our participants admitted that they were aware of additional risks, but that they entered the risks most important to them. Table III provides an overview of the answer proportions by scenarios and shows differences between the populations.

It is evident that the student population more frequently admitted that they only stated the most important risks while there are more. This in combination with the fact that almost no participant exhausted all possibilities to enter risks into the survey shows a limitation in awareness combined with a view that there are additional, but currently unknown risks that is more pronounced in our student sample. The majority of participants from MTurk stated that they have entered all risks that are relevant to them.

With respect to differences between the scenarios, our participants more frequently stated that they provided all the risks they know about or consider relevant (71.4% vs. 56.6% overall, Fisher's exact test, $p = .001$) in the more technical login scenario. This indicates that the more abstract nature of this scenario caused participants to be less confident about the completeness of the risks they are aware of.

4) *Prompted vs. Unprompted Risks*: Comparing the risks participants entered in the first, open-ended part of our questionnaire to the risks people stated to know of in the last part confirmed that biases and priming severely impact risk awareness results. While 74.7% of participants stated to know 15 or more of the 22 listed risks, only 22.5% gave six, seven or eight of the risks they indicated to know about in their free-text responses. The remaining participants had less matches and none but one participant had actually previously mentioned all risks that he or she selected from the list.

Among the risks selected from the list, abuse of login credentials was the most well known risk with 96.2% selections and psychological issues due to unsuitable content was the least well known (49.5%). The latter risk was also never stated by any participant without prompting. Note that this and several other risks were not mentioned at all in free-text responses, while all risks on the list were indicated to be known by at least about half of the participants. This highlights

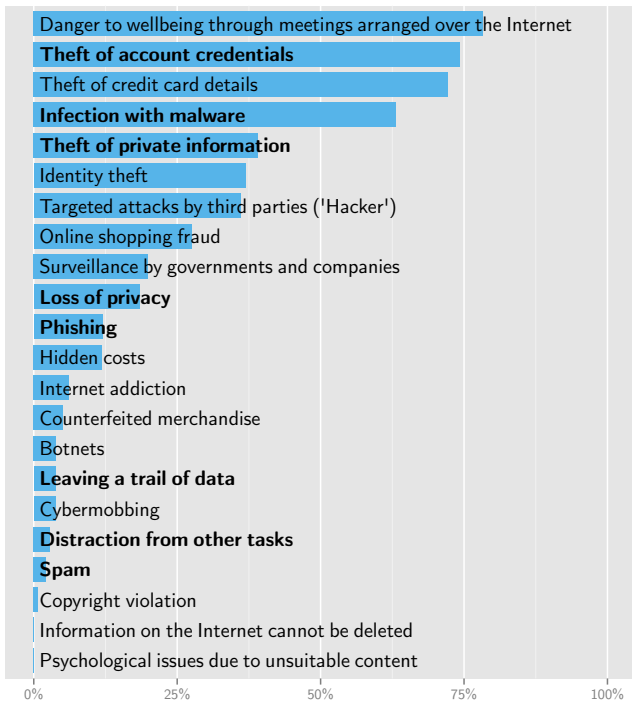


Fig. 2. Percentage of participants who stated a risk without being prompted that they later also selected from a pre-compiled list of risks. Labels printed in bold show the risks that were considered relevant or very relevant by more than two thirds of the users.

the importance of measuring risk awareness without prompting users about specific risks.

Figure 2 gives an overview of the risks our participants were able to select from the list. We also asked our participants to judge the relevancy of each risk to themselves personally as “not relevant at all”, “somewhat relevant”, “relevant” or “very relevant”. The labels printed in bold in Figure 2 show the risks that were considered relevant or very relevant by more than two thirds of the users. The bars in the figure also show how many users provided these risks in the unprompted part of the questionnaire and then also selected them from the list.

In this analysis, the two populations only differed mildly: seven of the items presented in Figure 2 switch places with their neighbors when looking at students only. A notable exception is identity theft, which was considerably less mentioned by students as already discussed in Section III-D1 above.

5) *Last Information on Risk*: When asked when they had last heard about the risks they provided, participants mostly relied on information that was several weeks or older. As Figure 3 suggests, risks for the general use scenario appear to be based on recent information, while participants indicated to have less recent information for the specific scenarios. Considering the results described above, depending on the scenario, participants can more easily rely on available information from friends, family and media. The more specific the scenario, the less information can be obtained from these sources. Furthermore, the additional risks (columns 2-4 in each scenario in Figure 3) appear to be supported by less recent information in several cases, suggesting that risks which participants had recently

heard of came to mind first. This suggests that the availability heuristic of Tversky and Kahneman [28] also plays a role in the appraisal of IT security risks. Furthermore, participants indicated to more frequently never having heard about risks for the banking and ride share scenarios. There were no significant differences between the student and MTurk population in this analysis.

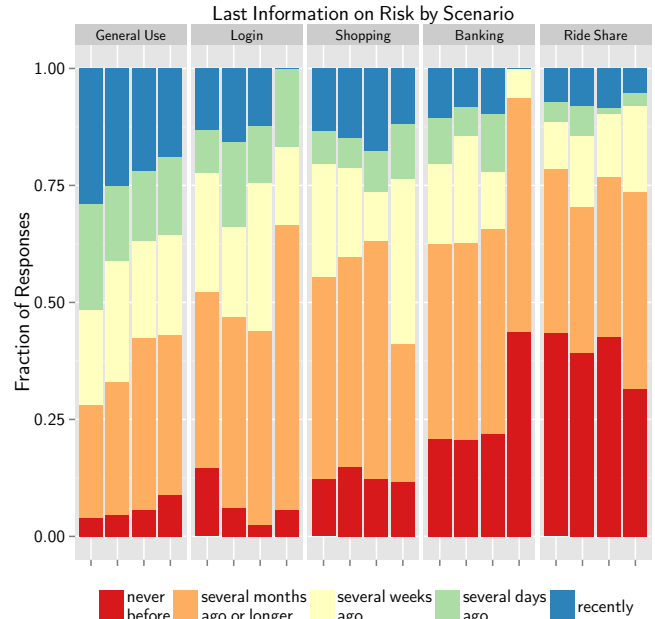


Fig. 3. Overview of participants’ responses to when they had last heard about the risks they provided, grouped by scenario. The leftmost bar in each group corresponds to the most severe risk stated by the users and the remaining three bars to the three additional risks.

6) *Scenario Rating*: Our participants were asked to rate the relative overall risk each scenario poses to their wellbeing in general on a numeric scale from 0 (no risk at all) to 100 (greatest possible risk). Overall, all five scenarios were not perceived as posing a very high risk: The mean rating was between 31.8 and 35.5 (sd between 24.1 and 29.5) for all scenarios, except login, where the mean rating was only 25.1 ($sd = 23.8$). We conducted a $2 \times 2 \times 5$ (data source \times gender \times scenario) mixed ANOVA on the participants’ ratings and found a significant main effect for gender ($F(1, 206) = 6.87, p < .001$) and scenario ($F(4, 824) = 3.12, p = .023$, Greenhouse-Geiser corrected), as well as a significant interaction between scenario and source ($F(4, 824) = 3.30, p = .018$, Greenhouse-Geiser corrected). The main effect of data source ($p = .33$) as well as the remaining interaction effects ($p > .26$) were not significant.

Looking at the effects, female participants rated the overall relative risk arising from our scenarios with 37.3 out of 100 while male participants only gave an average rating of 26.8. Furthermore, Holm-corrected pairwise comparisons revealed that the lower ratings for the login scenario were significant in all cases ($p < .05$) except when compared to ride-sharing, where ratings varied too widely. For the interaction effect, Holm-corrected pairwise comparisons showed that the login and general use scenarios were rated significantly higher by students ($p = .033$ and $p = .018$ respectively). Even though

there were some significant differences between the ratings, the large standard deviations underline that the perception of risk severity varies strongly between individuals.

7) *Risks and Security Measures*: To see to what extent participants could have protected themselves against the risks they currently are aware of, risks were categorized into two classes: Those risks participants gave that can be directly addressed by a particular security technology, for instance malware infection, and those risks that cannot, for instance meeting with a serial killer. We then counted the number of times these risks were stated by participants in the general Internet use scenario. This scenario was chosen since we assume that a user would think about whether or not to adopt another security technology, such as a password manager, within this mindset, as many security measures are considered for all use cases a user may have online. If we then found risks that are currently addressable by security measures in the responses, this can indicate that the user does not feel adequately protected by the currently applied measures and thus possibly desires an improved solution. Also, we believe that the risks a security measure addresses need to be salient enough in their intended audience so that potential users feel a need for improved security through additional measures. Hence, the risks users were aware of and from which they could be protected from are in general good candidates for future developments of improved security measures.

In our sample, we found that 24.6 % of the risks mentioned by participants could be addressed by malware and spam protection, 16.4 % by privacy technology or legislation, 9.5 % by authentication technology, and 8.6 % by end-to-end encryption or access control. The latter proportion was significantly different between students and MTurk participants with 13.1 % of risks mentioned by students addressable by such measures and only 4.0 % of risks mentioned in the MTurk population.

Another possible explanation for this pattern in the data is the influence of advertising and media. The most frequently mentioned risks relating to malware protection are addressed by a flourishing market of protection products that many people use. Similarly, privacy measures have found a large audience in the press and politics, even before the recent reporting on the extent of the NSA's surveillance. Risks commonly addressed by the remaining security technologies – encryption, access control and authentication – however, were not as pronounced in our sample, possibly because they are subtle and not advertised by companies trying to sell a product. This may be an instance of how advertising and media coverage “educate” users to see malware and privacy violations as risks.

Overall, security-measure-related risks comprised 59.1 % of all risks mentioned by participants in the general use scenario (69.0 % in students, 49.3 % in Turkers). This also means that the considerable portion of 40.9 % of risks participants are aware of (31.0 % in students, 50.7 % in Turkers) can currently not be addressed by available security measures or privacy regulations.

8) *Consequences*: Participants mentioned a total of 1,277 consequences across all scenarios. The consequences were elicited in order of severity. Since not all participants provided four valid consequences, 641 most severe consequences were

used in this analysis. Table IV shows an overview of the most frequently mentioned consequences. We find that participants mostly stated financial losses as potential consequences. The third most-frequently mentioned consequence “damage to health” mostly arises from the ride share scenario, where physically meeting with an unknown person was considered very risky by the majority of participants. Inconveniences, annoyances, and loss of time were also frequently mentioned among more serious consequences, including being victim of a crime or losing one's privacy. Again comparing the two populations, damage to health, identity abuse and becoming victim of a crime were more frequently mentioned by participants from MTurk. The risk of identity theft was already more pronounced in turkers as described in Section III-D1. The counts in Table IV additionally provide some evidence that physical consequences are more pronounced in the MTurk population.

Table IV also shows the perceived severity and likelihood of the consequences, based on the instances where the consequence was listed as most severe. Interestingly, losing data was also mentioned frequently, but was not considered to be very severe, especially by MTurk participants. Generally, severe consequences were most frequently mentioned but also perceived as being rather unlikely. Losing privacy was the most likely consequence, according to our participants.

9) *Perception of Security and Information Sources*: We asked our participants for their agreement with three statements about their security when using the Internet on a scale from “do not agree at all” (1) to “completely agree” (7). 38.1 % of our participants indicated that they often worry about risks and dangers in their day to day life at least somewhat (i.e., agreement > 4), 70.7 % of participants stated they feel at least somewhat safe when using the Internet and 31.1 % of participants at least somewhat agreed that they only have little influence on their security on the Internet. The agreement did not differ significantly between students and Turkers (Pearson's χ^2 tests, $p > .08$).

Considering sources of information for risks, participants were asked to indicate their perceived influence on a scale from “no influence at all” (1) to “great influence” (7). Media coverage was perceived to have more than medium influence (rating > 4) by 61.4 % of our participants, 78.6 % stated this for stories told by friends and family, 72.6 % for the influence of information actively sought by the participants themselves and 66.8 % for the influence of own negative experiences. It is noteworthy that the students reported significantly more high-influence ratings in all four cases (Fisher's exact test, $p < .028$). This also supports the hypothesis that advertising may have influenced risk perception to some extent.

IV. DISCUSSION

The results we obtained using our bottom-up survey method show very interesting aspects of risk and consequences awareness in our participants, which we believe to be of interest for future attempts to communicate the benefits of security measures to users. As we mainly provide a broad overview of the saliency of risks in users, we will derive hypotheses that future research should explore in depth.

TABLE IV. TOP 10 CONSEQUENCES MENTIONED BY PARTICIPANTS ACROSS ALL SCENARIOS. MS (MOST SEVERE) COUNT DENOTES THE INCIDENCE OF THIS CONSEQUENCE BEING MENTIONED AS THE MOST SEVERE CONSEQUENCE AND THE CORRESPONDING PERCENTAGE OF VALID CONSEQUENCES. THE TWO ADDITIONAL COLUMNS SHOW WHERE THE PROPORTIONS DIFFERED BETWEEN THE STUDENT AND MTURK POPULATION BASED ON SIGNIFICANT STANDARDIZED RESIDUALS IN THE CORRESPONDING CONTINGENCY TABLE. THE COLUMN UNIQUE COUNT SHOWS COUNTS AND PERCENTAGES OF UNIQUE PARTICIPANTS STATING THIS CONSEQUENCE. THE COLUMN L (LIKELIHOOD) DENOTES THE AVERAGE LIKELIHOOD AND COLUMN S (SEVERITY) THE AVERAGE SEVERITY ASSIGNED TO THIS CONSEQUENCE. BOTH, L AND S, WERE RATED ON A 10-ITEM NUMERIC SCALE FROM “NOT SEVERE/LIKELY AT ALL” (1) TO “VERY SEVERE/LIKELY” (10). RATINGS OF LIKELIHOOD AND SEVERITY WERE COLLECTED FOR THE MOST SEVERE CONSEQUENCES ONLY. ROWS WITH ADDITIONAL VALUES IN BRACES INDICATE VALUES THAT DIFFERED SIGNIFICANTLY BETWEEN THE TWO POPULATIONS (STUDENT’S T-TEST, $p < .05$).

Consequence	Count	Students/MTurk	MS Count	Students/MTurk	Unique Count	L	S
Financial loss	283 (22.2 %)		194 (30.7 %)		145 (69.0 %)	3.78	8.31 (8.1/8.7)
Large financial loss	142 (11.1 %)		122 (19.0 %)		92 (43.8 %)	3.79 (3.1/4.2)	9.01
Damage to health	117 (9.2 %)		112 (17.5 %)	12.8 %/23.0 %	112 (53.3 %)	3.37	9.03
Inconvenience	81 (6.3 %)		17 (2.7 %)		65 (31.0 %)	3.25	6.50
Identity abuse	68 (5.2 %)		31 (4.8 %)	2.3 %/7.8 %	56 (26.7 %)	3.42	7.71
Victim of a crime	66 (5.2 %)	3.0 %/7.3 %	22 (3.4 %)		61 (29.0 %)	3.81	8.52
Loss of privacy	50 (3.9 %)		12 (1.9 %)		44 (21.0 %)	6.00	7.33
Annoyance (e.g. legal)	45 (3.5 %)		6 (0.9 %)		33 (15.7 %)	4.33	8.33
Loss of time	45 (3.5 %)		6 (0.9 %)		50 (19.0 %)	4.00	6.83
Loss of data	41(3.2 %)		5 (0.8 %)		50 (19.0 %)	4.20	3.80 (6.5/2.0)
<i>Invalid: Not a consequence</i>	576	212/364	242	83/159			
<i>Invalid: Duplicate</i>	130	42/88	–				

On a very basic level, our results demonstrate that users differentiate between scenarios when assessing risk for things they do on the Internet. Many of them routinely consider multiple information sources, feel generally at least somewhat safe and believe that they can influence their security on the Internet. However, in contrast to the risks commonly addressed by security mechanisms, the risks our participants are aware of are not very technical but of a more general nature, which needs to be taken into account when security experts try to address them. When other sources of risk in a scenario are more important than IT security (e.g., “Will I be physically harmed doing this?” or “Will the sweater look good on me?”), we hypothesize that people might not think about other, IT-security-related risks. This does not necessarily mean that they are not aware of those risks at all, they are just not aware of them right then.

For example, only a single participant explicitly stated that he was concerned about man-in-the-middle attacks and this participant also self-rated himself as an IT expert. This is a particularly pertinent fact, since it is precisely this risk that the common SSL warning messages attempt to address. Based on our results, we further hypothesize that a major factor is users seeing risks arising from the impersonality of the Internet: being a victim of fraud by unknown merchants or unreliable people, having a “hacker” attack one’s accounts or data, and contracting malware from unknown sources or unknowingly becoming part of a botnet were frequently stated risks. Unknown attackers feature in all those risks and may therefore be believed to be hard to defend against.

Similarly, technical complexity appears to be a cause for concern in users: accounts being hijacked, credentials being “hacked” or stolen, and losing privacy were also commonly mentioned. What it really means to have an account hijacked or how credentials can or cannot get hacked is likely unclear to participants, as their often abstract and unspecific responses suggest. Also, how and to which extent a loss of privacy may occur in our scenarios was often not specified. Related work, for example the study of Rick Wash on users’ mental models of threats [12], found that users often have incomplete knowledge of threats and underestimate the danger for themselves.

Finally, it is important to remember that participants were only able to state a median of seven risks of 20 possible risks. In terms of a compliance or security-measure-adoption budget, this may mean that a new measure either needs to address an already salient risk or find a way to raise awareness for its benefits.

A. Asking Users About Risks

We were able to show that the set of risks users are aware of and that they can readily consider for their decisions is fundamentally different from the risks they indicate to know about given a pre-compiled list. Therefore, simply asking whether or not a user is afraid of a certain risk, scenario, or threat can generate misleading results. This is a very common practice in many studies and the results should be interpreted with care. If we look at the data collected in our study given the precompiled set of risks, many participants would have agreed that phishing, leaving a trail of data, and spam are relevant or very relevant risks. However, few participants actually mentioned these risks in the unprompted part of the survey. Previous work has shown that users don’t readily engage with information in security decision situations and hence would also not be convinced that a particular risk is important in terms of their compliance budget. We therefore postulate that the set of salient and important risks a user is aware of mainly informs decisions in such situations.

B. Risks and Consequences

Considering the consequences participants see for the risks they specified, we find that participants’ reasoning differs from how an expert would evaluate risks and their consequences. Participants often articulated something that is not a true consequence but a risk or simply a state of the world. For example, frequently, the consequence of an account being hacked was “my account is hacked”. Similarly, “my credit card number is given to another person” was given as a consequence of the risk of disclosing one’s credit card details. Thus, we hypothesize that many participants do not evaluate which risks actually have tangible consequences for them and therefore underestimate the impact of a risk for themselves. If they do,

the consequences appear to mainly relate to losing money, damage to their health and inconvenience.

At the same time, the low likelihood ratings indicate that they do not believe this will happen to them personally any time soon. A notable exception in our data is privacy, where the likelihood of having one's privacy compromised was considered fairly high. A potential moderator for these results is users' perceived self-efficacy, meaning how well users think they can protect themselves from a risk and its consequences. Even though we did not explicitly collect information about participants' perceived self-efficacy, it is conceivable that self-efficacy is low for protecting against a loss of privacy and therefore the likelihood of this happening to oneself considered high. Beyond an impact on likelihood ratings, perceived self-efficacy may have caused participants to not state some risks at all, as they feel that there is no threat arising from risks they can cope with by themselves. Future work should explicitly look at self-efficacy as a moderator for risk awareness and hence the adoption of security measures.

During the process of coding the responses, coders noted that many, especially non-financial consequences were phrased in an impersonal way, for example "data stolen", "loss of privacy", "losing friends". While this might very well be a grammatical oversight or abbreviation, the frequency with which we observed a mixing of personal with impersonal statements by the same participant led us to believe that this may have an influence on risk and consequence perception. We thus hypothesize that impersonal consequences cause some risks to be ignored, as the consequences are not perceived to apply to oneself personally and therefore remain abstract. Alternatively, in terms of the compliance budget, the cost of protecting against these risks may not be worth the potential benefits, as the consequences are too abstract to be of sufficient value.

Our results show some important areas where an experts' view and our participants' view of risks and consequences differ considerably. Especially the inability to see personal or any relevant consequences may influence a user's view on the necessity of adopting security technology or behaviors. We believe this represents valuable information, which needs to be taken into account when designing new IT security solutions and risk communication methods for end-users.

C. Differences in Risk Awareness

The choice of scenario influenced the set of perceived risks in our participants. In the banking and shopping scenarios, financial risks were a lot more important, while risks to health and wellbeing overlaid many other risks in the ride share scenario. Our study hence confirms that IT security often plays a secondary role in risk awareness when real-world risks are involved. Also, this means that the usage context of a security measure can influence its appraisal, as some risks become less salient in the light of other problems with regard to a certain task.

Our results also indicate differences in risk awareness between the two populations that were part of our investigation. Workers from Amazon's MTurk, who were all based in the U.S. according to MTurk's filter, reported to be older than our

student sample and included mostly non-students. These participants had a greater fear of identity theft, possibly because of the reliance on social security numbers and credit scores for many important financial aspects, which play a much smaller role in Germany. Similarly, fraud and scams in online shopping as well as hidden costs in services is a common concern for the German student participants. Furthermore, consequences in the real world were also more pronounced in MTurk participants, as they more frequently reported damage to their health and becoming victim of a crime as consequences arising from the risks they provided. We hence hypothesize that taking these kinds of differences in risk awareness into account when designing security mechanisms and developing strategies to deploy or advertise security and privacy measures can have beneficial effects.

Despite the Internet being used by a very diverse population, it seems that many security and privacy mechanisms are currently deployed on a global scale irrespective of culture or background. Future work of the usable security community could examine the benefits of tailoring design, presentation and deployment of security and privacy mechanisms to different cultures. Additionally, we also found that the relative severity of risks arising from our very general and common scenarios varied widely between participants. Female participants also found our scenarios to be more severe than their male counterparts, which may also be a good target for tailored solutions.

D. Awareness of Own Negligence and Mistakes

Another notable result of our survey is the almost total lack of awareness for risks arising due to own mistakes or negligence. Only in eight cases was "leaving an account logged in" or "choosing a weak password" stated as a risk. Particularly weak passwords are a risk which security professionals and researchers have tried to get the general population to take seriously for a very long time. Unfortunately, our participants did not consider these issues to be a major risk. We thus hypothesize that the risk of choosing weak credentials and not logging out of accounts is either unknown to many users, is not important enough to be salient, or users are not aware that they are actually doing these and other security-relevant activities wrong.

E. The Way Forward

Based on our results we postulate that users are only aware of a limited set of risks without being prompted and this set includes many risks that are not addressed by security technologies. We found that the relative importance of risks when using the Internet is perceived to be low in our sample to start with and if a security mechanism only addresses a few of the risks users are aware of, the perceived relevance likely becomes even lower. We therefore hypothesize that, for a security measure to become relevant and be adopted, users need to be aware of a serious risk with personal and immediate consequences, which are addressed by the technology in question.

Alternatively, our results suggest that new security technology can be specifically designed to address the users' greatest existing concerns and therefore more readily find adoption. For example, we believe it is worth researching whether security

measures protecting against man-in-the-middle attacks, such as visual indicators or warnings, might be more readily adopted if users were convinced that they prevent fraud and identity theft. Framing benefits around common scenarios and addressing the risks that are particularly salient in that scenario can help to tip the cost/benefit scales in favor of the security measure.

Another question that arises is whether or not it is possible to create additional awareness in users living in a modern society, with many other concerns competing for their attention. The results from Section III-D7 indicate that it may indeed be possible to raise awareness for particular risks. Malware is possibly the most common and long-standing security threat to end-user IT systems and the installation of anti-virus protection is recommended to most PC users. However, before people started to use information technology on a daily basis, they probably weren't as aware of the malware risk, as the results of Friedman et al. from 2002 suggest [18]. Yet, users may have actually experienced malware on their own device or heard stories from friends or family about such events since then. Additionally, there is considerable advertising for malware protection products that also remind users of the risks and made them learn to be afraid of malware. It needs to be subject of future work to see if and how awareness can also be raised for other IT security risks.

F. Limitations

While two diverse user groups were sampled for our study, especially the incidence of individual risks cannot be generalized. We also aimed to make differences between the two chosen populations clear in the text but also admit that a complete picture can only be painted by redoing this study with a population representative of all Internet users. Similarly, we deliberately chose a particular set of scenarios to test the influence of context on risk awareness. Other scenarios will likely yield different sets of risks, for example when considering the use of different service providers, as a user may trust other services less. However, we believe the patterns found in our results already hold valuable insights concerning the human aspects of IT security research. Future work is needed to look at effects a variation of trust in the scenarios may have. Additionally, we chose to use a survey as a research method, in order to obtain a wide view. It is possible that additional information can be obtained from using an in-depth interviewing technique, which is subject to future work.

Deliberately not providing a definition of the terms risk and consequence also has a potential influence on results. While we believe that this approach biased participants least and provided insights into how users may define the concept of an IT security risk and its consequences for themselves in their everyday conduct, forcing them to adhere to certain definitions could have led to results that are potentially closer to how security experts reason about risks and consequences, as participants may have provided additional risks and consequences they didn't think of or phrased those they provided differently. However, at the time users decide about whether or not to adopt a security measure, there also is no instruction sheet that provides a definition of how risk ought to be appraised before they make a decision.

We used an inductive coding procedure to analyze the risks participants provided in open-ended responses. Analyzing our

data showed that categorizing and coding risks is a task that can be tackled from many different angles. We adopted a pragmatic approach that allowed us to get a general overview of risks. We believe that coding risks in different ways can allow researchers additional insights into particular aspects of risk awareness. To this end we offer to share our complete data and study protocol freely with other researchers, to broaden our understanding of this important topic.

V. CONCLUSION

In this survey of risk awareness during Internet use, we find that users only showed awareness of seven risks on average. While this was to be expected from non-experts, we hypothesize that this also shows that the security community will have a hard time to get new security measures accepted in the general population under the premise that only risks users are aware of are considered in their compliance budget. Furthermore, we present evidence that the overall set of risks users perceive is very diverse and most of these risks were neither very specific nor can users easily protect themselves against these risks by using particular technologies. Additionally, existing security and privacy measures will often only address a small part of the risks users are aware of and focus strongly on the technical risks we found users are generally not too concerned about. This may then create the view that adopting a particular measure will not significantly reduce the risks of being on the Internet and is thus not worth the time, money and effort. We also posit that participants do often not see consequences that apply to them personally, effectively diminishing the benefits in terms of their compliance budget.

The main result of this paper is the hypothesis that users are often not ready to invest effort into changing their behavior or adopting security measures for the above reasons. Our analysis yields new insights into why certain security measures may not be adopted by end-users as well as which factors could influence adoption and hence need to be subject to further research. Security measures that aim to improve end-user security or privacy on the Internet would thus need to be designed to address salient risks and consequences as perceived by their users. The usable security community can support this process by further analyzing the protection needs of individuals, how security mechanisms can be tailored for adoption, as well as investigating possibilities to raise user awareness about important security risks, including their own negligence, effectively. The results presented in this paper can serve as a foundation for this important field of future work.

ACKNOWLEDGEMENTS

We would like to thank the reviewers for their many insightful comments that considerably helped to shape this paper.

APPENDIX

Questionnaire Overview

The questionnaire contained the following questions for each scenario:

- 1) *What do you think is the greatest risk/the greatest danger that arises for you personally from [scenario]?*

- 2) Which additional risks/dangers arising from [scenario] do you know about? You can enter up to three additional risks/dangers.
- 3) When did you last hear about these risks/dangers from others (including media, friends and family)? Answers: never before, a few months ago or longer, a few weeks ago, a few days ago, recently.
- 4) Which of the following statements best describes your listing of risks/dangers arising from [scenario]? Answers:
 - I have entered all risks/dangers that concern me.
 - I have entered the most important risks/dangers, but there are more.
 - I did not enter more risks/dangers, since I don't know about any further risks or dangers.
 - I did not enter more or all risks/dangers, since I feel safe on the Internet.
 - I did not enter more risks/dangers, because all boxes were filled.
 - I don't want to answer this question.
 - Other: [textbox]
- 5) Overall, how high do you believe the risk to your wellbeing from logging in to your social network profile to be? Please enter a number between 0 (no risk) and 100 (very high risk).
- 6) Please enter up to four consequences that may arise from the risks/dangers of [scenario] you provided in the previous question. Please begin with the most severe possible consequence and leave the additional boxes empty if you do not know any further consequences.
- 7) With regard to risks and dangers in other situations of your life, how severe do you consider your most severe consequence "[given consequence]", arising from [scenario], to be? Answers: (1) not severe at all to (10) very severe.
- 8) What do you think is the probability of the most severe consequence "[given consequence]" to happen to you personally? Answers: (1) very improbable to (10) very probable.

Codeplan and Counts

Revised codeplan used for the final round of coding. The numbers next to the items denote the incidence of each code in the students and MTurk deployment respectively. Top level items summarize the counts of all sub-items. Note that these numbers can be higher than the sum of the contained items, as very general responses were counted towards the top-level item.

- Account Abuse – 175/128
 - Stealing credentials (unspecific) – 65/59
 - Stealing credentials (specific) – 92/60
 - Account abuse – 16/6
 - Using account for criminal purposes – 1/0
 - Endangering other accounts – 1/3
- Fraud – 106/118
 - Identity theft – 12/93
 - SSN stolen – 0/2
 - Non-existent merchandise or services – 18/5
 - Low-quality or faked merchandise or services – 8/1
 - Insufficient information on merchant – 2/1
 - Hidden costs – 17/1

- Financial Risks – 110/138
 - Theft/abuse of credit card or banking details (no account access) – 85/127
 - Abuse of online banking (mentioned phishing) – 7/1
 - Abuse of online banking (no phishing) – 8/8
 - Erroneous money transfer – 10/2
- Privacy – 220/157
 - Loss of privacy – 22/22
 - Stealing private information – 67/58
 - Leaving a trail of data – 2/7
 - Personal info stored on third-party server – 0/3
 - Need to give private info to service provider – 0/2
 - Profiling – 10/13
 - Public disclosure of private information – 12/11
 - Passing private information on to third parties – 31/15
 - Information is hard to delete online – 0/0
 - Surveillance – 36/17
 - Government Surveillance – 4/2
 - Companies – 3/0
 - Collection of data in general – 1/0
 - Abuse of personal data – 36/17
 - Abuse by other users of the same service – 1/0
 - Abuse of online photos – 0/2
- Malware and Hackers – 124/125
 - Receiving spam – 0/5
 - Malware infection – 85/61
 - “Drive by Download” – 2/5
 - Abuse of PC for illegal activities by third parties (“Botnets”) – 3/3
 - Targeted attacks from unknown third parties (“Hackers”) – 33/51
 - Abuse of one’s IP-address – 1/0
- Psychological and Societal Risks 25/26
 - Cybermobbing, Bullying – 5/7
 - Psychological issues due to unsuitable content – 0/0
 - Internet addiction – 7/3
 - Being influenced by ads – 1/1
 - Getting depressed – 1/0
 - Account abuse to discredit someone – 0/1
 - Unpleasant social contacts – 3/3
 - Getting distracted from (more important) things – 3/6
 - Being dependent on IT services – 1/0
 - Loosing social contacts – 3/2
 - Influencing politics – 1/0
 - Loss of productivity – 3/0
- Real-world Crime – 14/47
 - Endangering one’s kids – 0/1
 - Copyright violation – 1/0
 - Mixed up in a crime – 1/1
 - Stalking, Internet Predators – 8/13
 - Theft of physical things – 2/22
 - Burglary due to known absence from Internet sources – 2/9
- Health Risks – 78/102
 - Risk to health and wellbeing – 76/83
 - Meeting with serial killer – 1/6
 - Obesity – 1/0

- Health risks because of repetitive motions and sitting – 0/6
- Own Mistakes/Negligence – 5/11
 - Insecure passwords – 1/2
 - Leaving an account logged-in – 0/4
 - Overspending – 3/5
- Misc. – 57/31
 - “General Risk” - 2/0
 - Others changing data – 2/1
 - Unreliability of other people – 44/25
 - Unreliability of services – 4/1
 - Exhausting bandwidth/data plan limit – 0/1
 - Faulty software/programming/services – 5/3
- Negative Codes – 1308/1097
 - N/A – 1115/906
 - This scenario does not apply to me – 10/3
 - Don’t know – 6/0
 - Not a risk – 49/65
 - Off topic – 10/11
 - There is no risk – 18/11
 - Duplicate risk – 100/101

REFERENCES

- [1] L. F. Cranor and S. Garfinkel, *Security and Usability*. O’Reilly, 2008.
- [2] M. A. Sasse, S. Brostoff, and D. Weirich, “Transforming the ‘Weakest Link’ – A Human/Computer Interaction Approach to Usable and Effective Security,” *BT Technology Journal*, vol. 19, no. 3, pp. 122–131, 2001.
- [3] J. Blythe, J. Camp, and V. Garg, “Targeted Risk Communication for Computer Security,” in *Proc. IUI*, 2011.
- [4] A. Beutement, M. A. Sasse, and M. Wonham, “The Compliance Budget: Managing Security Behaviour in Organisations,” in *Proc. New Security Paradigms Workshop (NSPW)*, 2008.
- [5] C. Herley, “So Long, And No Thanks for the Externalities: The Rational Rejection of Security Advice by Users,” in *Proc. New Security Paradigms Workshop (NSPW)*, 2009.
- [6] J. Blythe and L. J. Camp, “Implementing Mental Models,” in *Proc. SPW*. IEEE, 2012.
- [7] C. Bravo-Lillo, L. Cranor, J. Downs, and S. Komanduri, “Bridging the Gap in Computer Security Warnings: A Mental Model Approach,” *Security Privacy, IEEE*, vol. 9, no. 2, pp. 18–26, 2011.
- [8] C. Bravo-Lillo, L. F. Cranor, J. Downs, S. Komanduri, R. W. Reeder, S. Schechter, and M. Sleeper, “Your Attention Please: Designing Security-Decision UIs to Make Genuine Risks Harder to Ignore,” in *Proc. SOUPS*, 2013.
- [9] M. Harbach, S. Fahl, P. Yakovleva, and M. Smith, “Sorry, I Don’t Get It: An Analysis of Warning Message Texts,” in *Proc USEC*, 2013.
- [10] V. Garg and J. Camp, “End User Perception of Online Risk under Uncertainty,” in *Proc. HICSS*, 2012.
- [11] D.-L. Huang, P.-L. P. Rau, and G. Salvendy, “Perception of Information Security,” *Behaviour & Information Technology*, vol. 29, no. 3, pp. 221–232, 2010.
- [12] R. Wash, “Folk Models of Home Computer Security,” in *Proc. SOUPS*. ACM, 2010.
- [13] F. D. Davis, R. P. Bagozzi, and P. R. Warshaw, “User Acceptance of Computer Technology: A Comparison of Two Theoretical Models,” *Management Science*, vol. 35, no. 8, pp. 982–1003, 1989.
- [14] T. Herath, R. Chen, J. Wang, K. Banjara, J. Wilbur, and H. R. Rao, “Security Services as Coping Mechanisms: An Investigation Into User Intention to Adopt an Email Authentication Service,” *Info Systems J.*, 2012.
- [15] W. Dutton, G. Blank, and D. Groselj, “Cultures of the internet: The internet in Britain. Oxford internet survey 2013,” Oxford Internet Institute, University of Oxford, 2013.
- [16] D. Telekom/T-Systems, “Sicherheitsreport 2013,” <http://www.telekom.com/medien/konzern/198366>, Aug 2013.
- [17] M. Harbach, S. Fahl, M. Rieger, and M. Smith, “On the Acceptance of Privacy-Preserving Authentication Technology: The Curious Case of National Identity Cards,” in *Proc. PETS*. Springer, 2013.
- [18] B. Friedman, D. Hurley, D. C. Howe, H. Nissenbaum, and E. Felten, “Users’ Conceptions of Risks and Harms on the Web: A Comparative Study,” *Proc. CHI EA*, 2002.
- [19] R. Dhamija, J. D. Tygar, and M. Hearst, “Why Phishing Works,” in *Proc. CHI*, 2006.
- [20] V. Garg, L. Huber, L. J. Camp, and K. Connelly, “Risk Communication Design for Older Adults,” *Gerontechnology*, vol. 11, no. 2, p. 166, 2012.
- [21] A. De Luca, B. Frauendienst, M.-E. Maurer, J. Seifert, D. Hausen, N. Kammerer, and H. Hussmann, “Does MoodyBoard Make Internet Use More Secure?” in *Proc. CHI*, 2011.
- [22] M.-E. Maurer, A. De Luca, and H. Hussmann, “Data Type Based Security Alert Dialogs,” in *Proc. CHI, Extended Abstract*, 2011.
- [23] F. Raja, K. Hawkey, S. Hsu, K. Wang, and K. Beznosov, “A Brick Wall, a Locked Door, and a Bandit: A Physical Security Metaphor for Firewall Warnings,” in *Proc. SOUPS*, 2011.
- [24] S. Fahl, M. Harbach, T. Muders, and M. Smith, “Helping Johnny 2.0 to Encrypt His Facebook Conversations,” in *Proc. SOUPS*. ACM, 2012.
- [25] P. Klasnja, S. Consolvo, J. Jung, B. M. Greenstein, L. LeGrand, P. Powladge, and D. Wetherall, “‘When I am on Wi-Fi, I am Fearless,’” in *Proc. CHI*, 2009.
- [26] R. M. Hogarth, M. Portell, A. Cuxart, and G. I. Kolev, “Emotion and Reason in Everyday Risk Perception,” *Journal of Behavioral Decision Making*, vol. 24, no. 2, pp. 202–222, 2011.
- [27] N. C. Schaeffer, “Asking Questions About Threatening Topics: A Selective Overview,” in *The Science of Self-report: Implications for Research and Practice*, A. A. Stone, C. A. Bachrach, J. B. Jobe, H. S. Kurtzman, and V. S. Cain, Eds. Psychology Press, 1999, pp. 105–121.
- [28] A. Tversky and D. Kahneman, “Availability: A Heuristic For Judging Frequency and Probability,” *Cognitive Psychology*, 1973.