Perceptions of Computing Risks

Vaibhav Garg; L Jean Camp
Department of Computer Science; School of Informatics and Computing
Drexel University; Indiana University
Philadelphia, PA; Bloomington, IN 47401
garg@gmail.com, ljcamp@indiana.edu

Abstract

Understanding end-users' perceptions of information security risks is critical for the design of warnings and interactions that inform non-expert behaviors. To the extent that risk decisions are subject to bounded rationality offline, the perceived probability of risk has been judged by its salience and the perceived magnitude is impinged by the perceived benefits of the risky activity. Does this apply online despite the lack of potential physical harm? In this paper, we build on offline physical risk determinants by investigating the underlying determinants of perceived risk online. We evaluate perceptions of thirty technical risks, each grounded in one of six distinct mental model categories. We analyze the determinants of risk independently and then within each mental model. We compare the determinants and the mental models of experts and non-experts. Identity theft was identified as most risky, while severity was the most important determinant of perceived risk.

Security, Risk Perception, Psychology, Behavioral Economics

1 Introduction

Risk perception is a profound determinant of how individuals choose to act (or not act) to secure their online resources [?]. Technology acceptance is a function of a perceived risk of the proposed technology [?]. Mental models of risks determine the strategies that are adopted to ensure security. For example, Wash found that individuals may choose to use anti-virus software or try to avoid 'bad' web sites based on whether they perceive the risk as targeted or ubiquitous. In fact, Wash found that no extant mental models were correlated to or aligned with a decision to patch known vulnerabilities [?]. Given the well-documented impact of vulnerabilities for which patches exist, encouraging such risk-averse behavior is critical [?, ?]. Changing behaviors has been possible in the past, e.g. when risk is publicized or perceived, there are waves of patches installed [?].

Aligning beliefs and perceptions about risk and risk mitigation with objectively rational risk evaluations requires an understanding of the boundedly rational determinants of perceived risk [?, ?]. Perceived risk has been rigorously studied offline [?]. Yet offline risk perception is grounded in a *fear of physical harm.* There has been little repeatable, quantitative research mapping the determinants of offline risks to online risks. The goal of the work presented here is to *take the first step in leveraging risk perception to alter aggregate human behavior and secure the network.* In doing so, we will inherently introduce and translate the experimental framework used in offline risk to information security risks, offer a repeatable experiment, and briefly delineate the foundational research.

An understanding of technology alone is not adequate for defeating worms, viruses and malware [?]. Human incentives have been illustrated to be important, (e.g. [?]), as has usability [?]. However, incentives assume that individuals are rational [?, ?]; implementing a calculus of risk [?]. Usability assumes that the individuals will have an exogenous awareness of and desire to engage with risk reducing technology. The bounded rationality of end-users is being addressed though nudging [?]. Simultaneously, user awareness is being addressed through security education [?]. Integration of the scholarship on perceived risk is a necessary complement to these approaches.

In this work, we begin with a nine dimensional framework of risk perception based in the psychometric paradigm of expressed preferences [?]. This framework has been used to examine perceived risk in a diversity of contexts, e.g. health risks [?], environmental risks [?]. We combine this framework with a mental models approach introduced by Camp for information security risks [?]. We examine a set of thirty technical risks, each corresponding to a distinct mental model of security.

Section 2 is background and related work. Section 3 discusses the methodology, survey design and deployment procedure. Section 4 presents the results. After implications are enumerated in Section 5; Section 6 concludes.

2 Background and Related Work

Computer security risks are often voluntary and thus much scholarship has been grounded in transaction models rather than risk models. A rational actor paradigm posits that risks are weighed objectively as the product of probability with magnitude. However, while the benefits of online risk taking are immediate, the risks are often not salient [?]. Thus, this model of rationality has failed predictably and systematically for both security as well as privacy risks online[?, ?]. To the extent that humans are bounded in their rationality, risks management is not objective; instead balancing perceived risks with perceived benefits. Thus, online risk acceptance may be driven by contextualized risk taking attitudes [?].

Mental models based approaches have been used successfully to communicate risk in other fields, e.g. health risks [?]. Camp states that security experts typically use five mental models to communicate security and privacy risks [?] : physical, medical, criminal, warfare, and economic. These have been explained as: 1) Medical: physical harm but no intent, e.g., infection. 2) Criminal: economic but not physical harm and intent, individual target, e.g. stolen laptop from bar. 3) Physical: physical harm with intent, e.g., one time assault. 4) Warfare: systematic repeated harm from a potent adversary, economic and physical loss. 4) Economic: economic loss but not with intent, system manipulation, e.g. banking fees. These were extracted from naming of risks, practice, and validated for experts in [?]. It is important to understand how mental models inform determinants of perceived risk as well as identify those that make the risks salient.

3 Methodology

In this paper, we examine the impact of Fischhoff's nine dimensional model of perceived on end-users perceptions of technical security risks. Additionally, we examine the differences between the mental models of information security and their ability to impinge perceived risk. Finally, we want to examine the differences between perceived risk for end-users and security experts.

We used expressed preferences to elicit end-user perceptions of technical risks. We conducted a survey using Amazon's Mechanical Turk service. The risk items were based on the five mental models identified by Camp [?]. Additionally, there was a sixth category for risk items that did not fit any of the previously identified categories of mental models. There were four risk items for each of the five mental models as well four that appeared to fit in none of the above. In addition, we had one fake risk item associated with each category[1]. The purpose of these fake risks was to determine if individuals were capable of distinguishing the true risk from the imaginary. There were overall thirty risk items. These are listed below, the items in *italic* are the fake risk items:

1. Medical: virus, worm, blinding attack, infection, *digital pandemic.*
2. Criminal: identity theft, cracking, spoofing, 419 scam, *516 hustle.*
3. Physical: click jacking, system penetration, brute force attack, buffer overflow *information hunt attack.*
4. Warfare: trojan horse, spyware, logic bomb, DoS attack, *drone payload.*
5. Economic: adware, click fraud, stock spam, pump and dump, *fee flipping.*
6. Catch All: sniffing, cookies, zombies, phishing, *vampires.*

3.1 Survey Design

The survey had three components: 1) demographic information, 2) perceived risk, and 3) validation questions for Mechanical Turk. The demographic information consisted of gender, age, income, education, frequency of Internet use, frequency of financial transactions online, and frequency of location disclosure. Additionally we were also interested in technology attitudes and privacy preferences of participants. The former was operationalized using the oft-cited scale created by Spitzberg [?]. We used a subset of Internet Users' Information Privacy Concerns (IUIPC) scale [?] to measure privacy preferences of participants. We specifically chose the subset that correlates with global information privacy concerns (i.e., we did not consider specific factors such as consumer privacy preferences). Participants were also asked if they had a computer science or associated degree or other formal security training.

The next section of the survey constituted the perceived risk of thirty technical risks. This was modeled after the survey instrument used by Fischhoff et al. [?].This design has been replicated over several studies and

[1] Using fake instances of categories being considered is often used in social psychology, e.g. [?]

extensively tested. Thus, this instrument is optimal for online and offline comparisons. In the first question, participants were asked to rank the thirty risk items in the order of riskiness. They were ask to rate the least risky item as 10. An item that was rated 20 was considered to be twice as risky as the least risky item. This is similar to several previous survey-based evaluations of this framework. Thus, being consistent allows us to compare the results of this research with prior work in offline risk perceptions. The thirty items were presented in a randomized order to each participant to avoid primacy or recency effects. Participants were asked to rank only those items with which they were familiar (thus the use of imaginary risks).

Next the participants were asked to rate the risk of each of the thirty risk items on the nine dimensions. The nine dimensions were defined as:

1. Voluntary: Are people subjected to these risks voluntarily? (1 = voluntary; 7 = involuntary)

2. Immediacy: Is the risk from the threat immediate or does it occur at a later time? (1 = immediate; 7 = delayed)

3. Knowledge to exposed: To what extent are the risks associated with these threats known by the people exposed to them? (1 = known precisely; 7= not known)

4. Knowledge to science: To what extent are the risks associated with these threats known to experts, like computer scientists? (1 = known precisely; 7 = not known)

5. Controllability: If you are exposed to this threat, to what extent can you control (or mitigate) the risk, by virtue of personal skill or diligence? (1 = uncontrollable; 7 = controllable)

6. Newness: Are these risks new, novel ones or old, familiar ones? (1 = new; 7 = old)

7. Chronic-catastrophic: Does this risk affect one or multiple systems at a time? (1 = chronic; 7 = catastrophic)

8. Common-Dread: Is this a risk that people have learned to live with and can think about reasonably calmly, or is it one that people dread on the level of a gut reaction? (1 = common; 7 = dread)

9. Severity of Consequences: How severe do you think the consequences would be if this threat were exploited? (1 = not severe; 7 = severe)

Participants were asked to rate only those items with which they were familiar. The thirty risk items were presented in a random order to every participant to account for biases due to ordering effects. The risk items were interspersed with validation questions. These are described in the next paragraphs.

3.2 Mechanical Turk

We used Amazon's Mechanical Turk service to crowd-source the survey participation. Mechanical Turk provides a crowd-sourced labor force to perform human intelligence tasks that are relatively more difficult to automate, e.g. solving CAPTCHAs. Unlike traditional survey environments, the participants do not interact with the researchers. Thus, financial incentives could drive Mechanical Turk workers to select arbitrary responses. This would result in increasingly noisy data, the insights from which would be limited, if any. Several approaches have been proposed to reduce the sources of noisy crowd-sourced data [**?**, **?**, **?**]. We used a combination of the proposed approaches, as described in this section.

Only Mechanical Turk workers with a HIT approval rating of 95% or above were allowed to participate. HIT or Human Intelligence Tasks refer to the tasks completed by the workers. An approval rating of 95% means that 95% of the tasks completed by a worker were approved and thus paid for. The lower the approval rating, the lower the probability there is of a worker taking the task seriously.

Due to the complex nature of the survey questions, participants were required to be proficient in English. The survey solicitation required potential Turkers to be native English speakers. We determined participants' attentiveness as well as there adherence to survey requirements by asking them to state their native language as one of the survey questions. Responses from participants who answered with a language other than English were rejected.

After the participant was done ranking the thirty risk items in order of perceived risk, they were shown a list of five items and were asked to identify which item was not one of the thirty risk items they just ranked. For this survey, that item was 'eggs'. Since the participants had just viewed and ranked the list of thirty items, it was considered a reasonable validation question. Responses from participants who did not answer this question correctly were rejected.

Participants were also asked a general knowledge question. They were given a list of four American politicians, two of which had been presidents and two which had not. Participants were asked to identify those who had not been presidents. They were also given the option of selecting two. Initially, we rejected those participants who didn't select both of the correct options. However, we reviewed the responses of the participants who selected only one of the correct options, and most selected Hillary Clinton. The other correct option was Ted Kennedy. Based on the rest of their answers, it seemed that those participants might have simply confused Ted Kennedy with John F. Kennedy. Therefore, we retained the responses from participants who selected only one of the correct options. Responses from participants who selected either Obama or Lincoln were rejected. Since Obama is the current president and Lincoln was a particularly famous president, this was considered a reasonable validation question, especially since an attentive participant could quickly lookup the answer during the survey.

Finally, we also examined the time spent by the participants on the survey. We rejected the responses of participants who took less than 500 seconds ($\approx$ 8 minutes) to complete the survey. Given the length of the survey, it was highly unlikely that those participants could complete the survey without arbitrarily selecting responses or simply leaving large sections of the survey incomplete. Upon reviewing the responses of participants who took less than 500 seconds to complete the survey, we found supporting evidence. Most of them ($\approx$85 %) had answered at least one of the validation questions incorrectly. Those who answered the validation questions correctly, left most of the remaining survey blank. Participants that took more than 500 seconds had a much lower validation failure rate ($\approx$40 %) and left fewer responses blank.

3.3 Procedure

The survey was deployed using Mechanical Turk. The participants were recruited by posting an advertisement on Mechanical Turk's website. The advertisement noted that the participants were required to be native English speakers (from any country, not just United States) and have an HIT approval rating of more than 95%. The advertisement further stated that the participants would be asked validation questions to ensure the quality of responses. It noted that the participants would not be compensated for their time if the validation questions were not answered correctly. Participants were allowed up to sixty minutes to complete the survey. They were compensated $2 to participate in the survey. The survey was set up so that each Turk account could complete the survey only once.

The survey was was coded in HTML/JavaScript and then uploaded directly to Mechanical Turk. Survey responses were collected by Mechanical Turk. Responses for participants who answered even one of the validation questions incorrectly were rejected.

4 Results

The total number of responses was 843. The survey was online for five days. We ran twelve rounds of approval/rejection. Only 488 participants answered all the validation questions, described in the previous section, correctly. 355 participants answered, at least, one of the validation questions incorrectly. 102 participants incorrectly identified the American politician who was never the president. 74 participants did not select ‘eggs’ as the risk item that was not ranked. 264 participants did not list their native language as English. 33 participants took less than 500 seconds to complete the survey.

We also rejected additional participant responses as they did not follow the survey instructions. Participants were asked to rate the least risky item as 10. However, 142 participants entered values that were less than 10. Participants were also instructed to rate the risk items on the nine dimensions. The rating was done on a seven point scale, however, 14 participants gave higher ratings. All the responses for these participants were rejected. The final number of participants with valid responses was 332. 112 participants listed their gender as female, 105 as male, and 115 chose not to respond. 37 participants had a computer science or related degree, 293 did not and 2 chose not to respond. 14 participants had formal security training, 286 did not and 32 chose not to respond. 113 participants were between the ages 18-25, 82 were between 26-30, 59 between 31-35, 18 between 36-40, 21 between 41-45, 14 between 46-50, 15 between 51-55, 6 between 56-60, 3 between 61-65, and 1 between 66-70. Table 1 gives the median, mean, and maximum rating of perceived risk for each of the thirty items. The highlighted rows refer to fake risk items. Participants were more likely to provide responses for real risk items than the fake ones. On an average 208 participants responded to a question about real risks, while 140 participants responded to a question about fake risks.

Risk Item	Median	Mean	Max	n
Virus	19.0	354.5	100,000	309
Worm	17.0	26.4	200	248
Blinding Attack	15.0	22.76	150	51
Infection	18.0	30.61	300	187
Digital Pandemic	18.0	21.97	175	61
Identity Theft	22	35780	10,000,000	283
Cracking	16.50	23.99	175	116
Spoofing	15.0	21.93	125	122
419 Scam	13.0	18.39	75	77
516 Hustle	13.50	18.80	80	50
Adware	14.0	18.69	200	272
Click Fraud	15.0	20.26	100	117
Stock Spam	12.00	15.93	75	109
Pump and Dump	14.0	23.06	250	53
Fee Flipping	15.0	18.91	100	54
Click Jacking	15.0	23.68	100	87
System Penetration	20.0	44.12	1000	123
Brute Force Attack	17.0	986.7	100,000	105
Buffer Overflow	14.0	19.45	140	86
Information Hunt	17.0	22.76	125	59
Trojan Horse	18.50	67.36	10000	276
Spyware	15.0	24.26	220	292
Logic Bomb	16.0	24.66	110	61
DoS Attack	15.0	28.17	500	124
Drone Payload	14.0	40.4	1000	50
Sniffing	15.0	21.86	110	78
Cookies	10.0	13.56	100	309
Zombies	11.0	24.83	1000	113
Phishing	15.0	22.08	150	263
Vampires	10.0	28.74	1500	102

Table 1: Mean Risk Ratings

4.1 Real vs. Fake Terms

Table 2 notes the correlation between perceived risk and the nine dimensions for real as well as fake risks. We applied linear regression to the model. We note from table 1 that some of the risks were rated extremely high. The associated box plots suggested several outliers. Thus, perceived risk could not be used as a dependent variable in raw form. A second degree log transformation of perceived risk was normally distributed and thus was used as the dependent variable instead[2].

The independent variables consist of the nine dimensions as well as additional demographic variables: age, gender, income, education, frequency of Internet use, frequency of financial transactions online, frequency of location disclosure, attitudes towards privacy, and technology attitudes. We had two dependent variables: 1) perceptions of real security risks, and 2) perceptions of fake security risks. Separate regression models were constructed for each dependent variable.

For the real risk terms, the complete model gave an adjusted R-square value of 0.1515. Voluntariness, knowledge to exposed, knowledge to expert, control, severity, and age were statistically significant for $p<0.001$. Income was statistically significant for $p<0.01$. Newness, common-dread, gender and privacy attitudes were statistically significant for $p< 0.05$. Maximum variance was explained by severity. Removing frequency of location sharing, frequency of financial transactions online, and immediacy increased the adjusted R-square value to 0.1548.

For fake risk terms, the complete model gave an adjusted R-square value of 0.303. Knowledge to the expert was statistically significant for $p<0.001$. Frequency of financial transactions was statistically significant for $p<0.01$. Income, frequency of location exposure, and attitudes towards privacy were statistically significant for $p<0.05$. Maximum variance was explained by frequency of financial transactions online. Removing severity, knowledge to exposed, age, immediacy, and common-dread increased the adjusted R-square value to 0.3317.

4.2 Mental Models: Experts vs. Non-Experts

The participants were divided into two categories: those with some security expertise and non-experts. Experts were defined as participants who stated

[2] A fourth degree log of perceived risk was the closest fit to normal distribution. However, there were no differences in the regression models for a fourth degree transformation vs. a second degree transformation

Perceived Risk	Real Risks		Fake Risks	
Risk Dimension	Cor	p-val	Cor	p-val
Voluntary	0.142	≈ 0	0.127	0.0166
Immediacy	-0.044	0.008	-0.085	0.111
Exposed	0.019	0.26	-0.053	0.32
Expert	0.027	0.107	-0.099	0.064
Control	-0.135	≈ 0	0.043	0.422
Newness	-0.029	0.076	0.026	0.627
Common-Dread	0.246	≈ 0	0.077	0.15
Chronic-Catas	0.152	≈ 0	0.095	0.075
Severity	0.307	≈ 0	0.103	0.057

Table 2: Perceived Risk vs. Nine Dimensions: Correlations

that they had either a computer science or related degree as well as participants who had formal security training. Participants who did not respond to these questions were *not* automatically assumed as non-experts and were excluded from the analyses.

We constructed two 30*6 matrices, one for experts and the other for non-experts. The rows corresponded to the individual risk items, while the columns corresponded to the five mental models, with the sixth column indicating that the item does not fit either of the mental models. A lack of response was *not* assumed to imply that the term does not fit any mental model, but was excluded from the analyses. The results were plotted using classical multidimensional scaling (CMDS), on two dimensions, as shown in figures 1 and 2 for experts and non-experts respectively[3].

A kmeans clustering for both experts as well as non-experts was computed. We began by assuming six clusters, i.e. five mental models identified by Camp [**?**] as well as one catch all category of not applicable. However, the results for kmeans using six clusters varied for every run. The results for kmeans using four clusters were self similar and consistent for different runs for both experts as well as non-experts. The four distinct clusters are represented in figures 1 and 2 by different colors. The clusters for experts were:

1. Virus, Worm, Infection, Digital Pandemic, and Spyware.

[3] Multi dimensional scaling calculates the smallest distance between two vectors defined by mapping the results into an n-dimensional space. It is a distance metric to enumerate the similarities or dissimilarities in a dataset. For details consult `http://forrest.psych.unc.edu/teaching/p208a/mds/mds.html`; Retrieved May 2nd, 2012.

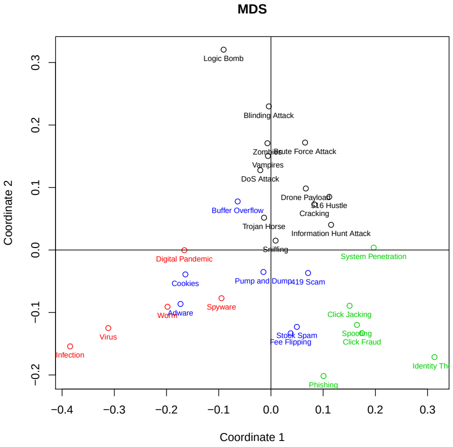

Figure 1: Expert Mental Models: CMDS with kmeans clustering

2. 419 scam, Stock Spam, Pump and Dump, Adware, Fee Flipping, Buffer Overflow, and Cookies.

3. Identity Theft, Spoofing, Click Fraud, Click Jacking, System Penetration and Phishing.

4. Blinding Attack, Cracking, 516 Hustle, Brute Force Attack, Information Hunt Attack, Trojan Horse, Logic Bomb, DoS Attack, Drone Payload, Sniffing, Zombies, and Vampires.

For experts, cluster one corresponds closely to medical mental models. Cluster two, similarly indicates economic mental models. Cluster three was a mix of mental models. The third and fourth clusters seems less determined. We could argue that the items in the third are those that are indiscriminate; with traps laid for anyone to walk into. Cluster 4 is a mix of warfare and physical mental models, consisting of targeted attacks with specific victims in mind. These were similar to those formed by non-experts. A chi-square analysis for the distances for each individual risk item between experts and non-experts was statistically *not* significant, p-value>0.995. The clusters for non-experts were:

1. Virus, Worm, Infection, and Digital Pandemic.

2. 419 Scam, Stock Spam, Identity Theft, Spyware, Adware, Fee Flipping, Click Fraud, and Click Jacking.

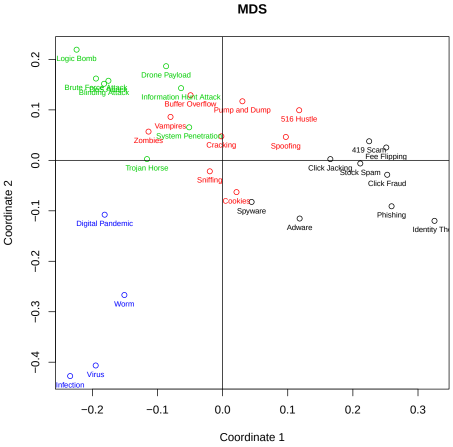

Figure 2: Non-expert Mental Models: CMDS with kmeans clustering

3. Phishing, 516 Hustle, Pump and Dump, Sniffing, Vampires, Cracking, Buffer Overflow, Cookies, Spoofing, and Zombies.
4. System Penetration, Trojan Horse, Drone Payload, Brute Force Attack, Blinding Attack, Logic Bomb, Information Hunt Attack, and DoS Attack.

Cluster one again corresponds to medical mental models. Cluster two is a mix of economic and criminal mental models, where the consequences are financial impact to the victim. Cluster three has several risks that in their offline usage could be confused with naive benign items, e.g. vampires. Cluster four again corresponds to both warfare and physical mental models.

We repeated the kmeans clustering with CMDS without including the fake risk items. The results for kmeans were self similar and consistent for four clusters. The results are plotted in figures 3 and 4, for experts and non-experts respectively. The clusters for experts were:

1. Virus, Worm, Spyware, and Infection.
2. Spoofing, Click Jacking, System Penetration, Phishing, Identity Theft, and Click Fraud.
3. Cracking, Trojan Horse, Blinding Attack, Zombies, Logic Bomb, Brute Force, and DoS Attack.

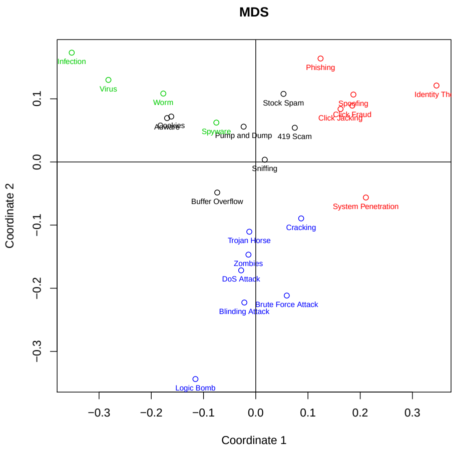

Figure 3: Expert Mental Models, Real Risks: CMDS with kmeans clustering

4. Buffer Overflow, Sniffing, Pump and Dump, Cookies, 419 Scam, and Adware.

For experts, the first cluster corresponds to medical mental model. The second cluster, though constituting several mental models, refers mostly to crimes with financial impact to the victim. The third corresponds mostly to a warfare mental model. The fourth cluster seems to be a catch all category. The clusters for non-experts were similar:

1. Virus, Worm, and Infection.

2. Adware, Spyware, Click Jacking, 419 scam, Stock Spam, Phishing, Identity Theft, and Click Fraud.

3. Trojan Horse, Blinding Attack, System Penetration, Logic Bomb, Brute Force, and DoS Attack.

4. Buffer Overflow, Sniffing, Cracking, Spoofing, Pump and Dump, Cookies, and Zombies.

The first cluster is again corresponds to medical mental model. The second cluster constitutes financial risks. Cluster three corresponds to a warfare and physical mental model. Cluster four is again a catch all category.

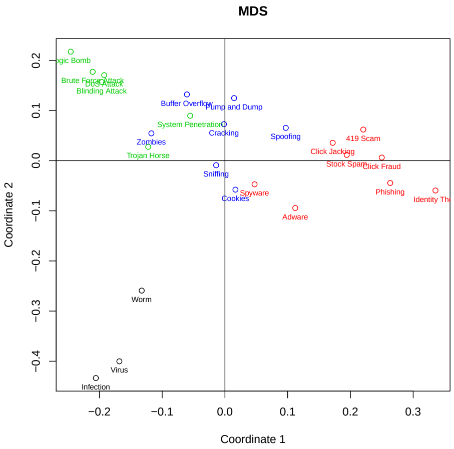

Figure 4: Non-expert Mental Models, Real Risks: CMDS with kmeans clustering

4.3 Cluster Based Analysis

We assumed that distinct clusters refer to specific metal models as expressed by the participants. For each cluster, as expressed separately by experts and non experts, we conducted a regression analysis with perceived risk as the dependent variable.

Cluster 1 for both experts as well as non-experts corresponded to medical mental models. For experts the cluster consisted of virus, worm, infection, digital pandemic, and spyware. For non-experts cluster 1 was virus, worm, infection, and digital pandemic, but without spyware.

We ran a linear regression for experts with the nine dimensions as well as additional variables: age, gender, income, education, frequency of Internet use, frequency of financial transactions online, frequency of location disclosure, attitudes towards privacy, and technology attitudes. The adjusted R-square value was 0.3608. Significant dimensions were knowledge to the expert, control, frequency of location sharing, and privacy attitudes. The best fit for the model was given by knowledge to expert, knowledge to exposed, control, newness, chronic-catastrophic, severity, gender, education, income, frequency of financial transactions, frequency of location sharing, and privacy attitudes; adjusted R-square=0.3998. Knowledge to expert, knowledge to exposed, control, frequency of financial transactions, frequency of location sharing, and privacy attitudes were statistically significant; respective p-values<0.001, 0.05, 0.001, 0.05, 0.01, and 0.01. A similar regression for non-experts with all the independent variables produced an adjusted R-square value of 0.03839. Knowledge to exposed, control, severity and age were statistically significant. The best fit for the model was given by knowledge to the exposed, control, severity, gender, age, income, and frequency of Internet use; R-square=0.06207. Knowledge to exposed, control, severity, and age were statistically significant; respective p-values<0.05, 0.01, 0.01, and 0.05.

Cluster 2 for experts corresponded to economic mental models, while for non-experts it was a mix of economic and criminal mental models. For experts, the cluster consisted of 419 scam, stock spam, pump and dump, adware, fee flipping, buffer overflow, and cookies. For non-experts, the cluster constituted 419 scam, stock spam, identity theft, spyware, adware, fee flipping, click fraud, and click jacking.

A linear regression for experts with the gamut of independent variables gave an adjusted R-square value of 0.3506. Immediacy, common-dread, severity, frequency of Internet use, and privacy attitudes were statistically significant. The best fit was given by immediacy, control, common-dread, severity, gender, age, income, frequency of Internet usage, frequency of financial transactions online, and privacy attitudes; adjusted R-square=0.3854. Immediacy, common-dread, severity, frequency of Internet use, privacy attitudes were statistically significant; respective p-value<0.01, 0.01, 0.05, 0.001, 0.05. Similarly, for non-experts the complete model gave an adjusted R-square value of 0.1578. Knowledge to expert, knowledge to exposed, control, common-dread, severity, age, and education were statistically significant. The best fit was achieved by removing immediacy, chronic-catastrophic, technology attitudes, income and, frequency of location sharing; adjusted R-square=0.1644. Knowledge to expert, knowledge to exposed, control, common-dread, severity, age and education were statistically significant; respective p-values<0.05, 0.05, 0.001, 0.05, 0.001, 0.01, and 0.01.

The discourse on the efficacy of mental models in the third and fourth clusters is less clear.

Cluster 3 for experts could be categorized by global risks, i.e. risks that do no target a specific victim. For non-experts, this cluster constituted terms that in the offline world are benign, i.e. sniffing, vampires, cookies, spoofing, and zombies. However, these were mixed in with phishing, 516 hustle, pump and dump, buffer overflow, and cracking. While the mental model here is more ambiguous, in general these terms were given low risk ratings, table 1.

Linear regression for experts with the complete model gave an adjusted

R-square value=0.2821. Severity was statistically significant. Best fit was given by voluntary, immediacy, knowledge to expert, knowledge to exposed, control, newness, chronic-catastrophic, severity, gender, education, and frequency of Internet use; adjusted R-square=0.3339. Knowledge to expert, severity, and education were statistically significant; respective p-value< 0.05, 0.001, 0.05. Linear regression for non-experts with the complete model gave an adjusted R-square value=0.1479. Knowledge to expert, severity, income and privacy attitudes were statistically significant. The best fit was given by voluntary, knowledge to expert, knowledge to exposed, control, newness, chronic-catastrophic, severity, gender, age, income, frequency of internet use, and privacy attitudes; adjusted R-square value=0.1578. Voluntary, knowledge to expert, severity, income, and privacy attitudes are statistically significant; respective p-value<0.001, 0.001, 0.001, 0.01, 0.01.

Cluster 4 constituted a warfare/physical mental model, with targeted attacks. For experts, the risk items were: blinding attack, cracking, 516 hustle, brute force attack, information hunt attack, trojan horse, logic bomb, DoS attack, drone payload, sniffing, zombies, and vampires. For non-experts, the risk items were: blinding attack, brute force attack, information hunt attack, trojan horse, logic bomb, DoS attack, drone payload, and system penetration. Linear regression for experts with the complete model gave an adjusted R-square value=0.211. Knowledge to expert and age were statistically significant. Best fit was given by immediacy, knowledge to expert, severity, gender, age, income, frequency of Internet use, and technology attitudes; adjusted R-square value=0.263. Knowledge to expert, age, and technology attitudes were statistically significant; respective p-value< 0.001, 0.01, 0.05. Linear regression for non-experts with the complete model gave an adjusted R-square value=0.08105. Knowledge to expert, severity, and age were statistically significant. The best fit for the model was given by knowledge to expert, knowledge to exposed, severity, gender, age, income, frequency of financial transactions online, and privacy attitudes; adjusted R-square value=0.1051. Knowledge to expert, severity, age, and privacy attitudes were statistically significant; respective p-value< 0.001, 0.001, 0.05, and 0.05.

In the later section we discuss future qualitative research which could build upon the questions generated by the more ambiguous groupings in the third and fourth cluster. In general, items in the third cluster seemed not targeted and are perceived as having lower risk.

5 Discussion

In this paper, we investigate the application of Fischhoff's nine dimensional model of perceived risks to technical security risks. Identity theft is clearly identified as the most risky of all the thirty risk items, for either mean, median or maximum values of perceived risk. This finding is a replication of a previous study by Garg et al. [?]. If we look at the maximum value of perceived risk, identity theft is considered several magnitudes more risky than the second most risky activity. On the other hand, the vectors of identity theft, for example phishing, are perceived much less risky. This is again similar to prior research findings [?]. It is grouped, by both experts and non-experts, with other risks that lead to financial loss. However, the other financial risks such as 419 scams are not ranked as high. This implies that identity theft is singular in its ability to inform end-user risk perceptions. Thus, identify theft can be leveraged not only inform risk attitudes but to effect behavior changes. The key would be framing risk communication messages in terms of identity theft and not just financial loss.

The degree to which the nine dimensions inform perceived risk differs significantly for real vs. fake risks. While for real risks the nine dimensions only explain 15.15% of the variance in perceived risk, for fake risks the impact is 30.3%, exactly double. This implies that the less end-users know about the risk, the more they are likely to rely on subjective measures of evaluating risks such as those posited by the nine dimensional framework. Thus, efforts to educate end-users about security risks are not only needed but can in fact have the desired impact.

The framework's ability to explain real online risks, 15.15%, is similar to that seen by Fischhoff et al. for offline risks ($\approx$ 13%) [?]. Thus, people are equally irrational or boundedly rational online as they are offline. Thus, poor risk decisions such as not patching vulnerabilities may be similar to behaviors offline such as smoking or procrastinating on getting a flu shot. The disconnect between attitudes and behaviors might then be driven by context and availability[?, ?]. The design of risk communication should then complement traditional approaches of exaggerating risks, with a contextual input to alleviate benefits. Risk communication cannot then be static or generic, but instead must be dynamic and targeted [?, ?].

The chi-square test for the distances between the different risk items on the nine dimensions was not significant for experts vs. non-experts. Thus, the absolute distance for each risk item for experts vs. non-experts was the same. There relative position on the axis is, however, different (figures 1-4). The x-axis for both experts and non-experts is defined by system vs. individuals. The right hand side of the x-axis is defined by risks that impinge individuals, e.g. identity theft. These risks can be avoided actively by users. Risk avoidance depends on end-user awareness or knowledge. They are voluntary and controllable. The left side, simultaneously, has risks that impact systems. These require experts systems or signals for successful risk mitigation. Thus, risk avoidance depends on knowledge to the expert. The risk items are involuntary and uncontrollable. X-axis constitutes voluntary, knowledge to exposed, knowledge to expert, and control.

The y-axis is defined by targeted vs. generalized attacks. Risks higher on the y-axis are characterized by risks that are targeted or materialize based on specific conditions (other than figure 3 where the y-axis is flipped), e.g. logic bomb, DoS attack etc. These risks are rarely encountered by end-users. For some of these risks, end-users may not be directly impinged thus the impact may appear delayed. Risks lower on the y-axis are characterized by those items that impact all systems/users in general, e.g. virus, phishing. These risks are old and commonly encountered by end-users. They also impinge the end-user directly. Thus, impact is immediate. Y-axis constitutes immediacy, newness, common-dread, and chronic-catastrophic.

We examined a kmeans clustering of risk items. These clusterings were similar between security experts and end-users for real risks. The first cluster corresponded to medical mental models. The second consisted risks with financial impact to the victim by stealing the victim's financial information. Physical and warfare mental models combined to form the third cluster. The fourth cluster was a catch all category with less clear delineation from the third. While for real risks security experts and end-users were similar, there were differences when we also consider fake risks. This is more interesting as it shows how security experts and end-users differ in assimilating new risk information.

Cluster 1 constituted medical mental models. While the nine dimensional model had good explanatory power for experts, it performed poorly for non-experts. A key determinant for experts was knowledge to experts, which was absent for non-experts. Knowledge to exposed was important for both. More knowledge to exposed led to higher perceived risk. For non-experts, severity and control were also important. Thus, simply creating awareness and exaggerating risks might be effective risk behaviors for this mental model. Simultaneously the impact of risk must appear to be uncontrollable.

Cluster 2 constituted economic mental models. Again the model had better explanatory power for experts then for non-experts. For experts, rarely encountered and severe risks were more risky. More Internet use and stricter privacy preferences also led to higher perceived risk. For non-experts most of the variance was explained by common-dread, control, and severity. Rarely encountered, controllable, and severe risks were perceived to be more risky. For this risk, simply creating awareness would not be adequate. While exaggerating the risk would help, the user must also feel that they can control the risk as well as that the risk is not an everyday occurrence. This would be a case where instead of providing average financial losses or providing statistics would not be relevant. Instead providing stories of specific individuals would be more pertinent.

Cluster 3 for experts and non-experts was not similar. For experts it constituted items indicating targeted attack. The key dimension accounting for perceived risk was severity. More severe items were perceived to be more risky. Thus, simply exaggerating risk would allow effective risk communication. For non-experts it was kind of a catch all category. More severe and involuntary items were perceived to be more risky. Perceived risk was also directly correlated with income. Thus, for non-experts when there is no easily identifiable mental model, risk communication should exaggerate the outcomes and stress the involuntary nature of the risk.

Cluster 4 combined warfare and physical mental models, mostly implying targeted attacks. Knowledge to the experts was important for both experts as well as non-experts. Higher knowledge to experts was correlated with lower perceived risks. Age was also a significant determinant, older adults perceived targeted attacks to be more risky. For experts, technology attitudes were important while for non-experts, privacy attitudes were more relevant. This reflects previous findings, where individuals irrationally responded more strongly to a minor change in privacy policy than to data breach [**?**]. Technology attitudes and privacy attitudes are highly correlated (p-value<2.2e-16). Thus, the only difference is severity, which is important for non-experts but not for experts. Thus, exaggerating the risks might be improve risk communication.

While the first and second clusters argue for the salience of mental models in risk classification, the third an fourth clusters argue agains this. Are there reliable internal representation of a external reality of online risks that are based on offline risks? Or do people built entirely new models? The conclusions that the nine dimensions are as significant online and off is the primary contribution of this work. However, the identification of those models that are salient for end users e.g., the medical model of viruses, versus those that leave confusion is also valuable. The contribution and the limitation of mental models argues for more qualitative work in the mode of Wash [**?**].

6 Conclusions

This paper examines the determinants of perceived risk for online risks. We confirmed individuals are as systematically irrational online as off. The lack of the threat of physical harm did not translate in substantive differences in the efficacy of the classic dimensions of risk to predict perception. These findings argue that the scholarship on mental models and risk communication offline can inform online risk communication online. We reference previous under-theorized experimental work that illustrates the efficacy of mental model and described work in progress.

We examined thirty technical risks. We included fake risks to determine if individuals were able to distinguish these. We investigated the relevance of distinct determinants of perceived risk for different mental models. We examined the impact of mental models at making risks salient. Finally, we explored the differences between security experts and naive end-users. Of the nine dimensions, severity was identified as the most significant. Clusters of responses identified the differences between the perceptions of experts and non-experts for the models; and the dimensions of risk in the models.

As would be expected, risk judgments were more subjective for fake risks than for real ones. Risk evaluations for fake risks were based on frequency of technology use and privacy attitudes. Real risks, however, are instead impinged by the nine dimensions of perceived risk given by Fischhoff et al. The ability of the nine dimensional model to explain perceived risk online is similar to that observed for offline risks.

We note that the less familiar the risks (including the fake and thus unfamiliar) the more individuals rely on their subjective measurements for evaluating risk. An additional finding is that technology attitudes dominate risk perception for experts, while risk perceptions for more naive users were dominated by privacy attitudes. This implies that highlighting the connection between security and privacy risks holds potential for motivating behavior changes for users with little expertise.

The design of risk communication must then identify the appropriate mental model that makes risk accessible as well as address the determinants of perceived risk that are most relevant and, thus, most likely to inform behaviors.

In this paper we introduced a classic nine dimensional model of risk perception to evaluate end-user perceptions of security risks online. We find that more knowledge about online risks decreases the affective perception of risks, and encourages more quantified evaluations. That online risks are perceived similar to offline risks allows security researchers and practitioners to leverage existing work in public policy, risk communication, as well as design architectures to facilitate risk averse behaviors online. For example, research in heuristics and biases that impinge decision-making, can be applied online. Research efforts should address contexts where such translation fails. Simultaneously, to the degree that security is a private good with externalities the exhaustion of security commons [**?**] can be addressed, in part, by public awareness campaigns grounded in successful endeavors offline, e.g. Iron Eye Cody's single tear on littering[4]. Since severity is the most significant determinant of risk online, risk communication effort should make it salient, for example demonstrating risk implications not of negligence but of malevolence[**?**].

A

A.1 Scatterplots

The relationship between perceived risk and the nine dimensions was relatively linear, figures 5 and 6. Thus, linear regression was applicable.

[4] `http://youtu.be/j7OHG7tHrNM`, Retrieved May 4th 2012

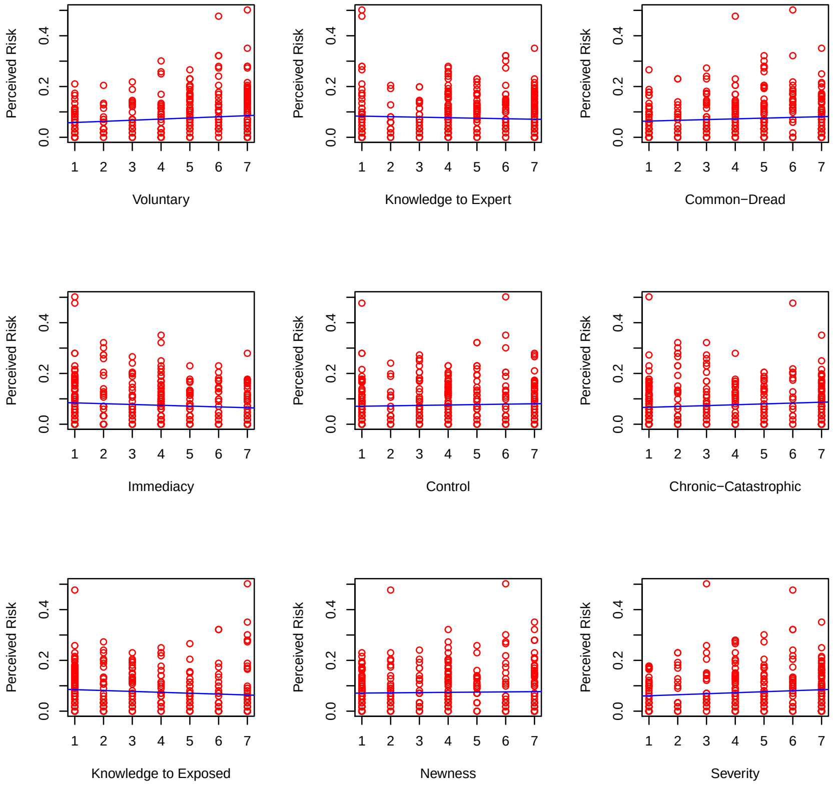

Figure 5: Perceived Risk vs. Nine Dimensions: Fake Risks

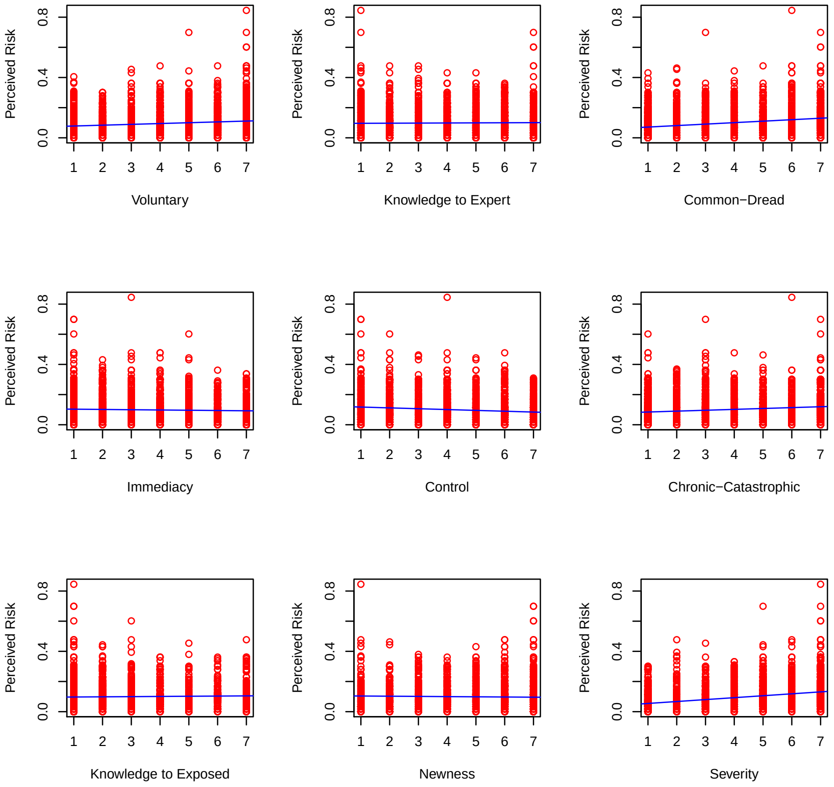

Figure 6: Perceived Risk vs. Nine Dimensions: Real Risks