Heuristics and Biases: Implications for Security

Vaibhav Garg
Indiana University

L. Jean Camp
Indiana University

Abstract

How can heuristics and biases improve the design of security technologies to leverage end-user behaviors? This position paper argues both for the importance of this question, and the specific identified examples. We discuss the limitations and criticisms of a heuris- tics and biases approach to understanding human behavior. We introduce some emerging theories in social-psychology that are more quantitative, and thus can be both predictive as well as descriptive. Our work is in response to the use of a rational actor model in computer security, which prescribes normative outcomes to individual preferences. How- ever, observed behaviors consistently deviate from these normative prescriptions. These errors are typically attributed to biases that under-lie a decision-making paradigm based in quick and dirty heuristics compared to an algorithmic and accurate approach. Typical investigations based in bounded rationality have focused on perverse incentives for end users to underinvest in security and privacy. Here we argue for a different perspective to bounded rationality, one that build upon it as a foundational tool rather than an obstacle that must be avoided.[1]

1 Introduction

Anderson [2] demonstrated that security risks are mediated not only through technology but also by the incentives that individuals and institutions have to protect themselves and others (and the lack thereof) [14]. Explorations based in expected utility have been conducted for both security and privacy risks. These have provided valuable insight into areas such as stakeholder investment in secure technologies [35].

The theory of the rational actor, though reasonably successful for studying markets, has been limited in its ability to predict or describe individual decision-making [37]. Conundrums such as the privacy paradox are difficult to address with a model that presupposes a rational actor [15]. Similar dichotomies between attitudes and behaviors have been observed in the physical world and have been better addressed through behavioral models [22]. The ideal rational actor has thus paved for the behavioral concept of bounded rationality [24].

A behavioral approach to security and privacy has previously been explored [1, 10, 19, 36]. Its application underlies but is not the primary contribution of this work. While several of these papers discuss underlying heuristics and biases that facilitate bounded rationality in decision making, they primarily illustrate the fallibility of human judgement. In this paper, we examine how these heuristics and/or biases can enable better decisions online. We expand the discourse by applying the limitations and the criticisms of this approach that have been the subject of several papers in the psychology community.

Section 2 introduces the theory of the rational actor and provides background on the deviations from normative behavior that led to the development of bounded rationality paradigm under prospect theory. Section 3 introduces several heuristics and biases and their application in the security domain. Section 4 discusses the limitations and criticisms of bounded rationality as well as introduces emerging theories in the socio-psychology domain that are more quantitative, i.e. are predictive as opposed to bounded rationality, which is primarily descriptive. Section 5 concludes.

2 Normative Decision Frames

Expected utility theory implies an algorithmic approach to decision-making that presupposes rational choice. However, this idealized model is not always realized by a real world rational actor, as carefully documented by Kahneman and Tversky in their seminal work [11].

[1] Reference as V. Garg, and L. Jean Camp,“Heuristics and Biases: Implications for Security Design”, IEEE Technology & Society, Mar. 2013

Actors' deviations from the normative model limits the descriptive and predictive ability of theories based in a paradigm of rationality. Kahenman et al. [11] approach decision-making as the risk of choosing between two gambles or *prospects*. They reported three scenarios where observed behavior differed from that prescribed by expected utility theory[2]:

1. Certainty Effect: There are two possible outcomes A and B. In A the participant has the opportunity to win a high amount with a small probability, while in B the participant has the opportunity to win a small amount with a high probability.
   A(500, 0.1) B(50, 1.0)
   The expected value from both A and B is quantitatively equal, i.e. a gain of $50. According to expected utility theory neither of the outcomes should be preferred. However, when this choice presented to participants in an experimental setting, most participants prefer B over A. Thus, when presented with a small certain gain as opposed to a large probable gain, participants choose the smaller certain gain with the equal expected value.

2. Reflection Effect: As in the previous case study, there are two possible outcomes A and B. In A the participant can suffer a big loss with a small probability, while in B the participant can suffer a small loss with a high probability.
   A(500, 0.1) B(50, 1.0)
   As in the previous case study the expected value from both A and B is quantitatively equal, i.e. a loss of $50. However, participants overwhelmingly prefer outcome A over B. Thus, when presented with a larger probable loss as opposed to a smaller certain loss, participants choose the larger probable loss with the same expected value.

3. Isolation Effect: Consider a two stage game. In stage 1, the probability of losing the game and winning nothing is 0.75. The probability of moving to stage 2 is 0.25. There are two outcomes possible for stage 2, A1 and B1.
   A1(400, 0.8) B1(300, 1.0)
   Taking into account stage 1, the probability of winning $400 is 0.8*0.25=0.20 and that of $300 is 1.0*0.25= 0.25. Thus, the real choice is between A2 and B2.
   A2(400, 0.2) B2(300, .25)
   If the outcomes are presented as a one stage choice participants prefer the former outcome of winning $400 with probability 0.2. However, if the outcomes are presented as a two stage game participants prefer the latter outcome of winning $300 with probability 0.25. Thus, we witness a preference reversal when the same choice is presented in two different forms. It has been suggested that when the game is presented in two stages participants do not take into account the first stage. This is rationalized since the first stage is common to both outcomes. If this assumption is true, then in the two stage game participants perceive the game as a choice between (400, 0.8) and (300, 1.0). Thus, just like in the first case study, it becomes a choice between a larger probable gain and a smaller certain gain. As noted in the first case study, participants will choose the smaller certain gain over larger probable loss. Notice that utility theory predicts the converse as the expected value is higher.

3 Heuristics and Biases

Given the deviations from expected utility theory, it was posited that people are not rational decision-making machines. Rather, they are bounded in their rationality [25, 26]. Instead of being algorithmic and accurate we use heuristics that are quick and dirty mechanisms of decision-making [31]. While heuristics provide a performance upgrade in terms of being quicker, this comes at the cost of accuracy, i.e. heuristics make us susceptible to making mistakes or biases. Here we outline several heuristics and biases that play a dominant role in our decision-making process:

3.1 Framing

If we combine *certainty effect* with *reflection effect* we get 'framing'. In the example above, the choice or gamble being presented is the same each time, i.e. the expected value of the outcome is $50 each time. However, the probable outcome is preferred over the deterministic version when the prospect is framed as a loss. This preference is reversed when the outcome is framed as a gain. This has significant implications for individual security investment. Security investment for an end user is currently framed as a definite loss, while the risk of not investing in security is a probable loss [19]. Consider the alternative framing. The gain from investing in secure technologies and protecting IT resources is probable, while the gain from investing those same resources in an alternate locale is deterministic. For example, an individual may choose to buy a bigger monitor than buy antivirus. Similarly an organization may want hire an extra IT support than invest in intrusion detection technology. Thus, individual actors will choose **not** to invest in secure technologies under the constraints of bounded rationality.

[2] The notation to denote a choice is presented as A($x, p), where outcome a consists of obtaining $x with probability p (0<1)

This preference reversal based on the decision frame was first noticed by Tversky et al. [32]. This effect has been further investigated by several researchers [18, 21]. Rothman et al. [18] discuss the implications of framing in improving health behaviors. They found that actions which prevent threats were seen as safe, while those that detect threats were seen as risky. In their study safe behaviors were encouraged through a gain frame, while risky behaviors are encouraged though a loss frame. In the security domain, patching prevents malware from exploiting software vulnerabilities. End users should be presented information regarding patching as a gain. Thus, end users can be told that patching would keep their computer running fast and the software cutting edge. On the other hand antivirus detects the presence of malware. Hence, antivirus installation and periodic scans are arguably encouraged by presenting the information in a loss avoidance frame. For example, end users can be told that if that without antivirus they are much more likely to lose their data.

3.2 Assimilation and Contrast

In the previous section we argued that providing the appropriate decision frame is important. Schwarz et al. [21] find that the wrong decision-frame can encourage decision making counter to the intended nature of the communication. Thus, if the information is not presented appropriately it may encourage users to be more risky than they would be otherwise. When an actor acts in valence with the communication, they demonstrate *assimilation* effects [8]. However, if they act counter to the intended goal of the communication, they demonstrate *contrast* effects [5]. In general, if the information provided is used to judge a category it leads to assimilation effects. However, if it is used to judge a specific member of the group it leads to contrast effects. For example, the example of a specific corrupt lawyer may decrease the general opinion of lawyers, but might also lead to an improved opinion of a specific lawyer, such as our personal lawyer. Thus, telling users that clicking on online banner ads might install malware on their systems may make them suspicious of pop-ups in general but might increase their trust in specific banner ads. Users may then be more susceptible to ads that suggest that they have discovered malware on the user's computer, e.g. figure 1[3].

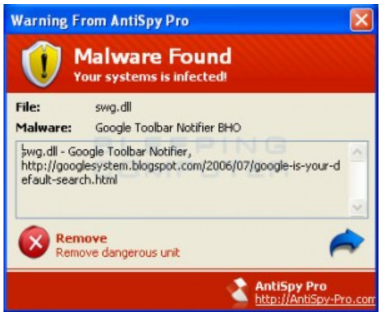
Warning From AntiSpy Pro

Malware Found
Your systems is infected!

File: swg.dll
Malware: Google Toolbar Notifier BHO

swg.dll - Google Toolbar Notifier,
http://googlesystem.blogspot.com/2006/07/google-is-your-default-search.html

Remove
Remove dangerous unit

AntiSpy Pro
http://AntiSpy-Pro.com

Figure 1: Typical Malware Pop-up pretending to be a Malware remover.

Similarly, there may be implications for email-based scams. Herr et al. [8] found that participants, when primed with moderate exemplars and asked to judge ambiguous stimuli demonstrated assimilation effects. However, contrasts effects were seen when unambiguous stimuli were judged or when participants were primed with extreme exemplars. Thus, when participants are shown extreme exemplars of email scams such as 419 scams, they may become more trustworthy of less extreme examples such as generic phishing emails. In general, priming with a stereotype generates assimilation [3], while priming with an exemplar generates contrasts [5]. For example, priming with a professor stereotype led participants to perform better on trivia, while priming with Einstein degraded performance. Thus, priming with the general example of a security guru might lead to safer behaviors online, while priming with a specific example such as Ron Rivest might lead to excessive risk taking. If we account for the decision frame a more effective analogy may be explored. For example, priming the user with a generic unsafe end user might lead to more unsafe behavior, however, priming with an extremely unsafe user (such as a using the example of a specific victim such as Sarah Palin) might lead to much safer behaviors.

3.3 Representativeness

Stereotypes, encouraging assimilation are an illustration of the general power of representativeness. The comprehensive representativeness heuristic is based on the conjunction fallacy [33]. One of the foundational results of probability is that the probability of a conjunction, i.e. P(A&B), can not be more than the probability of either of it's components, P(A) and P(B). However, intuitive decision making leverages the representativeness heuristic that results in decisions and estimates that consistently violate this basic fact. For example, let us consider the case of Linda who is good at math, introvert, diligent, and unimaginative. The probability of Linda being an accountant as well as a historian can not be greater than that her just being a historian. However, people systematically report the former probability to be higher. Representativeness can also get people to disregard base rates. For example, in a class of 100 student there are 30 boys and 70 girls. Student A is a computer nerd, has never had a girlfriend and loves Star Trek. These are attributes may

[3] `http://velocity93.blogspot.com/2010/08/malware.html`

usually be attributed to a nerdy boy in his teens. Thus, even though the probability of Student A being a girl is higher (0.7) than that of being a boy (0.3), individuals will give an estimate that is higher than the base rate on these conjoined probabilities.

Similarly, while the number of legitimate emails asking for personal information is much lower than phishing emails, the former may appear higher depending upon how well the phishing email has been crafted. Considering the example of pop-ups, e.g. fig 1, there are arguably no pop-ups that provide legitimate antivirus services. However, the probability of one being accurate is heightened by how well the pop up resembles real antivirus softwares. This can become a major issue when security indicators are used incorrectly [29]. It may be easier for phishing websites to appear legitimate by misusing security indicators, such as the lock sign (e.g. figure 2).

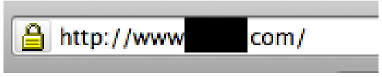

http://www[redacted]com/

Figure 2: A well placed lock sign may inspire a sense of trust and hence security, even absent critical indicators such as https.

3.4 Availability

Like representativeness, the availability heuristic acts upon probability judgements of likelihood of a risk to occur. In general, people tend to rate the probability of a risk to be high if the ease of recalling an instance of that risk is easy and low if it is difficult. For example, Sherman et al. [23] asked participants in their study to rate the ease with which they could imagine contracting a particular disease. They found that the ease of imagination was correlated with participants perceived probability of contracting the disease. Thus, people may be more scared of terrorism than food poisoning, even though statistically people are more likely to die of food poisoning than of terrorism[4]. Thus, heavily publicized risks are more salient in human imagination rather than the more threatening one. This impinges not only on individual decision-making but can also guide public policy and investment.

This has several implications in the security domain. In general, when security works best then nothing happens. The lack of incidence is not a salient event that can be made readily available. This also means that people will find it easier to remember identity theft rather than the mechanisms, e.g. phishing, that underlie identity theft. This heuristic is also leveraged in phishing and other spoofing attacks. In general, phishing websites differ from their legitimate counterparts in small details. However, small details are not as easily available to the decision-maker and hence they may not look for them and become a victim. This is further complicated by *belief perseverance* [17], i.e. if a person believes a hypothesis A, they will continue to hold that hypothesis as true simply because it is salient. When end users go to a banking website, they believe that they are going to the desired location. They might maintain this belief even after encountering several visual cues that the website is indeed fraudulent.

Availability can, however, be used to design better risk communication. Koehler [13] found that if a person is asked to imagine a hypothesis as being true, it increases their confidence in the truthfulness of that hypothesis. In the security domain, many times users are told that the website they are navigating to is insecure and ask the user to reconsider their decision. Thus, instead of simply suggesting that a website is insecure, users might be asked to imagine that the website is insecure and the resulting implications. This would increase the perceived salience of the risk and may lead to safer behaviors.

Availability can also be leveraged by public relations campaigns in the form of public service announcements or media coverage of security incidents. Increased coverage of risks such as Facebook fired[5] would make the risk more easily accessible and thus available to the end user. This might discourage users from sharing information on Facebook. Information sharing on Facebook is further impinged by availability. In general, Facebook or any site that wants the user to share their information advertises the benefits of information sharing. Thus, even though the users may be aware of the risks, they only pay attention to the benefits as these are more salient. Further, risk communication may cause the user to reflect on why they chose to share the information and thus make it appear more beneficial that it truly is, i.e. belief perseverance. However, this can be alleviated by asking the users to generate a counter hypothesis and explanation [28]. Thus, end users should be asked to generate the benefits of sharing less information or the risks of sharing more information. This hypothesis generation and explanation would make the risks of information sharing more salient.

3.5 Affect

The affect heuristic refers to the general feeling that a person may have towards a certain action. It differs from the previously described heuristics in not being cognitive, rather affect deals with emotions [27]. For exam-

[4] This is based on a US only perspective.

[5] http://www.huffingtonpost.com/2010/07/26/fired-over-facebook-posts_n_659170.html#s118586&title=Juror_Dismissed_After

ple, a person may choose to buy one car over another not based on performance or price, but purely because they find it more attractive. Another example is the difference in increased happiness or sadness when a person wins or loses. Say the choice is between winning $50 or losing $50. Since the expected value of the gain or loss is the same, the gain should lead to the same increase in happiness as witnessed by the loss in terms of increase in sadness. However, the affect heuristic cause us to be more wary of losses than accepting of gains. Thus, winning $50 is less joyful than losing the same amount is painful.

Schwarz [20] noted that people may attribute current affective states to an evaluation irrespective of valence. Thus, people may rate their overall satisfaction with their lives higher on sunny days as compared to rainy days. This effect can be alleviated by making the source of affect explicit, e.g. they can be asked about local weather earlier. In general, good moods may lead to positive evaluation, while bad moods may lead to negative evaluation. Affect can impinge availability and vice versa [12]. Keller et al. found that perceived risk was greater when participants were presented with risk information for 30 years as compared to one year. Thus, it may be better to provide aggregate financial loss due to phishing for several years than just one.

Positive affect impinges cognitive flexibility [9]. In general, better moods increase a person's ability to retrieve, store, and process information. Thus, designers of security risk communication should ensure that it does not create undue anxiety or negative affect in the recipient. This can be critical for older adults who have lower cognitive plasticity than younger adults and have lower technical literacy. This reflects the previous discussion of loss and gain in framing from Section 3.1.

4 Limitations and Criticisms

The concept of bounded rationality is based on observed deviations from normative probabilistic prescriptions. This premise was challenged in a paper by Gigerenzer [7]. He argued that the experiments conducted by social psychologists typically evaluate single events. He argues not that particular biases are incorrect but that they can not be subject to measurement. For example, an experiment in the security domain might ask a participant the probability of an email being a phishing email or legitimate. Since this is a single event a *frequentist*[6] would argue that probability and statistics are not applicable. He noted, for example, that the conjunction bias disappeared (or reduced drastically) when the problem was framed in terms of frequency. He also argued that base-rate neglect could be alleviated if the participants were ensured of random sampling. An opposing view to a frequentist outlook to probability is *subjectivist or bayesian*. Gigerenzer [7] states that for subjectivists '*rationality is identified with the internal consistency of subjective probabilities*'. Thus, he argues, the outlook under which the observed discrepancy between stated and prior probabilities can be considered an *error* or deviation from a normative model is very narrow; it naively assumes that good decision-making consists of applying Baye's theorem to real life choices.

Tversky and Kahneman replied to criticism by noting that most heuristics and biases had very little to with probability judgements [34]. They also cite several studies that demonstrate base rate neglect despite frequency based framing, e.g. [30]. They further comment that the difference in preferences for frequency framed vs. probability framed is acknowledged by framing effect. Finally, they argue against the notion of subjective probabilities as learned frequencies; several investigations note errors despite the presence of relevant frequency data.

Heuristics and biases have additional limitations. There is neither a general theory, nor a model that explains the underlying cognitive processes. Thus, it is unclear which heuristics dominate the decision-making process under what context. In general, one can only provide a post-hoc explanation for an observed effect. These limitations are addressed by two emerging theories in social-psychology: (1) Decision Field Theory (DFT) [4], and (2) Quantum Information Processing Theory (QIPT)[16].

Both expected utility theory and prospect theory qualify the preference amongst outcomes. However, neither measures the strength of those outcomes. Thus, they do not account for deliberation time for preference formation. DFT was thus introduced by Busemeyer et al. [4] as a stochastic dynamic theory of decision making under risk. Deliberation process can be important in security and privacy risks as these are usually not the primary goal of the end user. Thus, given two strategies to encourage risk averse behavior in end users, the one more likely to succeed is the one that requires smaller deliberation time. DFTs based modeling provides a framework for such evaluation.

QIPT unlike previous theories, which were based in classic probability theory, leverages quantum probability models. A key property of quantum models is that the probability of a conjunction can be higher than that of its components. Thus, quantum probabilities appropriately model several systematic errors explained by prospect theory. They are limited by requirement for an appropriate Hilbert Space and Hamiltonian [16].

[6] Probability is defined by the frequency of observed outcomes over an infinite number of trials.

5 Conclusion

Research in social-psychology observes and explains systematic errors by decision-makers. These errors have been identified as a problem in security design. We argue that such finds should instead be encouraging insights that inform security designs for risk averse decision making by end user. It is critical to remember that bounded rationality serves well and fails only in specific cases. Thus, bounded rationality should not present the decision maker in a negative light, rather it should be treated as a design constraint when provisioning for security in information systems. There are existing examples of research that demonstrate the usefulness of heuristics based decisions when leveraged appropriately [6]. This can be further informed by emerging theories such as QiFT and DFT that facilitate modeling of end user behavior.

References

[1] ACQUISTI, A., AND GROSSKLAGS, J. What Can Behavioral Economics Teach Us About Privacy? *Digital Privacy: Theory, Technologies, and Practices* (2008), 363–377.

[2] ANDERSON, R. Why information security is hard-an economic perspective. In *Computer Security Applications Conference, 2001. ACSAC 2001. Proceedings 17th Annual* (2001), pp. 358–365.

[3] BARGH, J., CHEN, M., AND BURROWS, L. Automaticity of social behavior: Direct effects of trait construct and stereotype activation on action. *Journal of personality and social psychology 71*, 2 (1996), 230.

[4] BUSEMEYER, J., AND TOWNSEND, J. Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological review 100*, 3 (1993), 432–459.

[5] DIJKSTERHUIS, A., SPEARS, R., POSTMES, T., STAPEL, D., KOOMEN, W., KNIPPENBERG, A., AND SCHEEPERS, D. Seeing one thing and doing another: Contrast effects in automatic behavior. *Journal of Personality and Social Psychology 75*, 4 (1998), 862.

[6] EHRLICH, K., KIRK, S. E., PATTERSON, J. F., RASMUSSEN, J. C., ROSS, S. I., AND GRUEN, D. M. Taking advice from intelligent systems: the double-edged sword of explanations. In *Intelligent User Interfaces* (2011), pp. 125–134.

[7] GIGERENZER, G. How to make cognitive illusions disappear: Beyond heuristics and biases. *European review of social psychology 2*, 1 (1991), 83–115.

[8] HERR, P., SHERMAN, S., AND FAZIO, R. On the consequences of priming: Assimilation and contrast effects* 1. *Journal of Experimental Social Psychology 19*, 4 (1983), 323–340.

[9] ISEN, A., AND LABROO, A. Some Ways in Which Positive Affect Facilitates Decision Making and Judgment. *Emerging perspectives on judgment and decision research* (2003), 365.

[10] JACKSON, J., ALLUM, N., AND GASKELL, G. Perceptions of risk in cyberspace. *Trust and crime in information societies, Edward Elgar Publishing Limited, Cheltenham* (2005), 245–281.

[11] KAHNEMAN, D., AND TVERSKY, A. Prospect theory: An analysis of decision under risk. *Econometrica 47*, 2 (1979), 263–291.

[12] KELLER, C., SIEGRIST, M., AND GUTSCHER, H. The role of the affect and availability heuristics in risk communication. *Risk Analysis 26*, 3 (2006), 631–639.

[13] KOEHLER, D. Explanation, imagination, and confidence in judgment. *Psychological bulletin 110*, 3 (1991), 499.

[14] MOORE, T., CLAYTON, R., AND ANDERSON, R. The economics of online crime. *The Journal of Economic Perspectives 23*, 3 (2009), 3–20.

[15] NORBERG, P., HORNE, D., AND HORNE, D. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs 41*, 1 (2007), 100.

[16] POTHOS, E., AND BUSEMEYER, J. A quantum probability explanation for violations of rationaldecision theory. *Proceedings of the Royal Society B: Biological Sciences 276*, 1665 (2009), 2171.

[17] ROSS, L., AND ANDERSON, C. Shortcomings in the attribution process: On the origins and maintenance of erroneous social assessments. *Judgment under uncertainty: Heuristics and biases* (1982), 129–52.

[18] ROTHMAN, A., MARTINO, S., BEDELL, B., DETWEILER, J., AND SALOVEY, P. The systematic influence of gain-and loss-framed messages on interest in and use of different types of health behavior. *Personality and Social Psychology Bulletin 25*, 11 (1999), 1355.

[19] SCHNEIER, B. The psychology of security. In *Proceedings of the Cryptology in Africa 1st international conference on Progress in cryptology* (2008), Springer-Verlag, pp. 50–79.

[20] SCHWARZ, N. *Feelings as information: Moods influence judgments and processing strategies.* New York: Cambridge University Press, 2002, pp. 534–547.

[21] SCHWARZ, N., AND BLESS, H. *Mental Construal Processes: The Inclusion/Exclusion Model.* New York: Psychology Press, 2007, pp. 119–141.

[22] SHERMAN, S. On the self-erasing nature of errors of prediction. *Journal of Personality and Social Psychology 39*, 2 (1980), 211.

[23] SHERMAN, S., CIALDINI, R., SCHWARTZMAN, D., AND REYNOLDS, K. Imagining can heighten or lower the perceived likelihood of contracting a disease. *Personality and Social Psychology Bulletin 11*, 1 (1985), 118.

[24] SIMON, H. A behavioral model of rational choice. *The quarterly journal of economics 69*, 1 (1955), 99.

[25] SIMON, H. Theories of bounded rationality. *Decision and organization 1* (1972), 161–176.

[26] SIMON, H. *Models of bounded rationality.* MIT Press Cambridge, Mass, 1982.

[27] SLOVIC, P., FINUCANE, M., PETERS, E., AND MACGREGOR, D. The affect heuristic. *European Journal of Operational Research 177*, 3 (2007), 1333–1352.

[28] SLOVIC, P., AND FISCHHOFF, B. On the psychology of experimental surprises. *Journal of Experimental Psychology: Human Perception and Performance 3*, 4 (1977), 544.

[29] STEBILA, D. Reinforcing bad behaviour: the misuse of security indicators on popular websites. In *Proceedings of the 22nd Conference of the Computer-Human Interaction Special Interest Group of Australia on Computer-Human Interaction* (2010), ACM, pp. 248–251.

[30] TVERSKY, A., AND KAHNEMAN, D. Availability: a heuristic for judging frequency and probability. *Cognitive psychology*, 2 (1973), 207–232.

[31] TVERSKY, A., AND KAHNEMAN, D. Judgment under uncertainty: Heuristics and biases. *Science 185*, 4157 (1974), 1124.

[32] TVERSKY, A., AND KAHNEMAN, D. The framing of decisions and the psychology of choice. *Science 211*, 4481 (1981), 453.

[33] Tversky, A., and Kahneman, D. Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological review 90*, 4 (1983), 293.

[34] Tversky, A., and Kahneman, D. On the reality of cognitive illusions. *Psychological Review 103*, 3 (1996), 582–591.

[35] Varian, H. System reliability and free riding. *Economics of Information Security* (2004), 1–15.

[36] West, R. The psychology of security. *Communications of the ACM 51*, 4 (2008), 34–40.

[37] Zeckhauser, R. Comments: Behavioral versus rational economics: What you see is what you conquer. *The Journal of Business 59*, 4 (1986), 435–449.