Social, Behavioral, and Economic Influences in Security Decision Making: Lessons from Early Work in Cybersecurity

Lucien Randazzese, Ph.D. and David Balenson, SRI International (February 15, 2017)

The challenge of cybersecurity is an almost existentially difficult one. The creation and targeted sharing of ever larger pools of data have become indispensable to the foundations of United States social and political stability: improving human health, expanding food production, reducing energy use, enhancing industrial efficiency, and countless other prerequisites for a growing and secure nation.

Paradoxically, the greater interconnectedness essential for progress also makes us more vulnerable. Virtually all information technology assets now share connections and data across multiple individuals and organizations, and often across international borders. Billions of individuals and millions of organizations use the public internet to access countless connected devices and infrastructure including medical devices, vehicles, buildings, airports, power grids, and now entire cities. The cybersecurity of these critical assets represents an essential component of their operational efficiency and protection from harm. As the scope of this cybersecurity challenge has grown, so has the range and diversity of potential solutions. Once viewed as a merely technical problem marked by the seesaw escalation in capabilities by those who would attack and defend the nation's information systems, networks and critical infrastructure, cybersecurity now enjoys research and analytic attention from economists and other social scientists exploring the role of human decision making in effective security. The study of human cognitive biases in particular represents an area of inquiry with great promise to improve security.

The authors have considered this question within the context of cybersecurity research and public policy, and the discussion that follows primarily reflects work on this specific aspect of security [1]. Nevertheless, we believe the implications of our analysis are relevant for the domain of national security generally. At its heart, our work considers how people in sensitive situations – those in charge of valuable data and critical infrastructure – make decisions regarding protection and use of information and assets, what factors influence these decisions, and how policy might be informed by an understanding of human decision making. This focus has obvious relevance to security of almost any variety.

Of particular interest to us has been recent work to understand how human cognitive biases affect decision making. This emerging understanding is being used in the creation of policies in a wide range of areas outside of security including public health, crime prevention, financial decision making, energy efficiency,

and tax collection, to name just a few. The realm of cybersecurity has only recently begun to attract the attention of social, behavioral, and economic scientists. Early findings hint at the potential value of more research of this kind.

Initial work to understand decision making in the context of cybersecurity framed choices in terms of marginal economic costs and benefits analysis, and focused exclusively on investment by firms in cybersecurity measures [2]. Over time, research expanded to include attention to the decisions individuals made with respect to their own data privacy and online conduct [3, 4, 5]. Given its unit of analysis – the individual – this newer work is inherently more behaviorally focused, and serves to highlight the impact of cognitive biases and their influence on behavior.

We know, for example, that people are less likely to behave offensively online when their actual identities are revealed online [3], and are more likely to divulge sensitive information when they believe others have done so [4]. Both findings are examples of a cognitive bias in favor of emulating peers (referred to as peer influence). Other research shows that decisions regarding privacy settings can be influenced by how setting choices are framed within the larger set of personal device settings, highlighting a framing bias [5].

Peer influence and framing are just two examples of a remarkably diverse set of cognitive biases affecting human behavior, which also includes the phenomena of loss aversion, representativeness, and choice cost, among others. Peer Influence results in susceptibility to peer pressure and also leads people to rely on peers as sources of low-cost information about how to choose or behave. Framing bias causes people to assess options depending on how they have been presented. Loss aversion refers to the tendency of people to prefer avoiding losses over acquiring equivalent gains. Representativeness causes people to draw incorrect conclusions about causation and distribution because they assume small sample sizes are representative of system-level phenomena. Choice cost tells us that the process of choosing is difficult, and that people will often make choices in a way that minimizes the effort in making the choice, with little or no consideration of the value of different outcomes.

Our understanding of how these biases work rests on an enormous body of accumulated insight into human behavior gained from a variety of disciplines, including psychology, sociology, and neuroscience, a body of applied research often referred to as *Behavioral Economics*.

One of the things that makes understanding of cognitive biases so promising a tool for informing research and policy in the field of cybersecurity is the magnitude of the effects biases can have on actual behavior and their track record of success in real world applications in other policy areas. Numerous academic

studies of even the most highly-trained specialists have shown both expert and layman susceptibility to systematic failures of human cognition [6], and a growing number of policy changes and interventions are being rolled out based entirely on the specifics of how these cognitive failures affect behavior [7]. For example, recent research has shown that cybersecurity professionals' probability perceptions are as susceptible to anchoring effects as those of the general population [8,9].

Overall, the process of translating the research insights from behavioral research into cybersecurity policy recommendations is still in its infancy. In almost all cases it considers individual choices regarding privacy and how people treat their own personal data. Almost none of this research considers the choices made by people in organizational contexts, either as executive decision makers or as those who serve in operational roles and as stewards of an organization's sensitive data or critical infrastructure.

Neither does behavioral research often focus on organizations as the unit of analysis. Organizations possess unique characteristics that differentiate them from other organizations in their propensity to be secure or unsecure. Research in the area of medical record protection, for example, has shown hospital security outcomes to correlate with specific hospital characteristics [10]. Other work has shown organizations to be subject to the same "psychological" biases as are people. For example, organizations are more likely to reduce spam when spam levels are publicly reported, but tend to do less to reduce spam when the worst-reported offender produces a greater absolute level of unwanted mail, thus making other firms appear less abusive in comparison [11]. Organizations are, in essence, coordinated collections of individuals, and organization decisions are made, ultimately, by people. It should not be surprising then that they "act like people," including in ways that highlight organizational susceptibility to bias. Though we are unaware of any studies that address how even larger collections of individuals such as societies and nations exhibit cognitive-bias like behaviors, we suspect that they do. More work in this area is clearly needed.

To maximize the chance that good decisions are made, it is essential that policy and incentives aimed at influencing cybersecurity decisions appreciate how the information environment can affect decision making. When information is insufficient for good decision making, there is a role for policy in stimulating information creation and dissemination, and possibly in creating and provide missing information directly. Conversely, when the environment is crowded with enormous amounts of contradictory, diffuse, and rapidly changing information (a situation often referred to as information overload), public and quasi-public institutions can encourage better decision making by helping to simplify the information environment, for example through

the creation and promulgation of standards and frameworks, such as the National Institute of Standards and Technology (NIST) Cyber Security Framework [12].

Practical constraints on incentive-oriented policies constitute another reason for use of behavioral based policy. The government is not usually in the positon, legally or practically, to reward or threaten to punish people to a degree sufficient to force compliance with desired cybersecurity practice. Consider the example of password sharing. The government could, in principle, require a mandatory ten-year prison sentence for anyone found guilty of sharing a commercial password. Such a law would dramatically affect people's economic incentives regarding password sharing, causing the frequency of password sharing to drop precipitously. But such a policy is clearly not practical, requiring alternative ways to influence behavior.

So how might policy makers influence cybersecurity decisions in situations for which either too little or too much information is available, or when practical considerations make appeals to cost-benefit calculus unrealistic? The experience in a very wide range of policy areas outside of cybersecurity suggests that the answer lies, at least in part, with greater attention to the effects of cognitive biases in decision making. Many of these areas face challenges and pursue policy goals similar to the challenges and goals of cybersecurity, and to national security more broadly. Some applications of behavior-based policy are quite novel, while others are straightforward, perhaps even obvious. What makes all of them noteworthy is their level of effectiveness in addressing real-world problems.

In work that we have done for the United States Department of Homeland Security Science and Technology Directorate (DHS S&T), we have identified a range of cybersecurity economic incentive research areas for which a behavior focus could add considerable insight. Our proposed research agenda emphasizes the impact of cognitive biases on cybersecurity behavior, and has an applied and empirical focus, directing attention as much as possible toward understanding real behavior in real-world situations.

As highlighted above, most of the current research one may categorize in "cybersecurity decision making" focuses on the decisions individuals make with respect to their own personal data and online activity. Conducting research on how organizations and people behave in their official capacity – e.g. as employees who have access to sensitive data or critical infrastructure controls, or executives making decisions on cybersecurity procedures and investments – is the obvious next step for research.

There are practical challenges to this type of research, but we believe it can be done. Recent experimental research examined how actual information resource owners respond to notification of abuse on their systems, looking at, among other things, how notifier reputation affects efforts to remedy abuse [13]. Other

interesting empirical research into cybercriminal and attacker behavior is further proof that security researchers can be quite creative in their methodological approaches. Some attacker research is made possible by analysis of servers and related assets used in the commission of actual cybercrime [14].

The information tools that are accelerating progress in almost all domains of human life are the same tools would-be antagonists use to threaten national security and well-being. For this reason, comprehensive cybersecurity cannot rely on technical approaches alone. The tools and techniques and analytical methods must include a diverse portfolio of approaches. In this white paper we have made the case for one such approach: understanding the fundamental cognitive biases influencing all human behavior and decision making. We believe this understanding holds considerable potential for the specific domain of cybersecurity, in which we have discussed its application, and also in a host of other national security areas. Advances in this area of research promise to provide tangible insights, new-to-world analytic techniques, and enhanced strategies policy makers can use to materially influence behavior, by good actors and bad. The resulting improvements to our ability to defend the nation's data repositories, information systems, and critical infrastructure will go a long way toward ensuring our national security and stability.

Acknowledgements

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Department of Homeland Security or the United States Government. The work reflected in this paper was funded by the US Department of Homeland Security Science and Technology Directorate (DHS S&T) under Contract No. HSHQDC-10-C-00144. The authors thank the DHS S&T Cyber Security Division for its support and Dr. Joseph Kielman, retired Science Advisor to the DHS Cyber Security Division for his valuable guidance.

For more information on DHS S&T work in the area of cybersecurity and incentives, please see the Cyber Risk Economics program homepage at: https://www.dhs.gov/science-and-technology/csd-cei

References

[1] SRI International, SRI International Work on Cybereconomic Incentives for the Department of Homeland Security Science and Technology Directorate Computer Security Division, January 31, 2015.

[2] Gordon, Lawrence M and Loeb, Martin P. "The Economics of Information Security Investment." *ACM Transactions on Information and System Security* 5.4 (2002): 438-457.

[3] Cho, Daegon. "Real Name Verification Law on the Internet: A Poison or Cure for Privacy?" Tenth Workshop on the Economics of Information Security. George Mason University, Fairfax, VA, USA. 14-15 June 2011.

[4] Acquisti, Alessandro, et al. "The Impact of Relative Standards on the Propensity to Disclose." Eighth Workshop on the Economics of Information Security. University College London, England. 24-25 June 2009.

[5] Adjerid, Idris, et al. "Framing and the Malleability of Privacy Choices," 13th Workshop on the Economics of Information Security. Pennsylvania State University, State College, PA. 23-24 June 2014.

[6] Englich, B., & Mussweiler, T. (2001): "Sentencing under uncertainty: Anchoring effects in the courtroom", Journal of Applied Social Psychology, 31, 1535–1551.

[7] UK Cabinet Office Behavioral Insights Team, "Test, Learn, Adapt: Developing Public Policy with Randomised Controlled Trials," June 14, 2012.

[8] Mersinas, Konstantinos et al. "Experimental Elicitation of Risk Behaviour amongst Information Security Professionals," 14th Workshop on the Economics of Information Security. Delft University of Technology, The Netherland. 22-23 June 2015.

[9] Mersinas, Konstantinos et al. "Measuring Attitude towards Risk Treatment Actions amongst Information Security Professionals: an Experimental Approach." 15th Workshop on the Economics of Information Security. University of California Berkeley, Berkeley CA, USA. 13-14 June 2016.

[10] Kwon, Juhee and Johnson, M. Eric. "Security Resources, Capabilities and Cultural Values: Links To Security Performance And Compliance" Eleventh Workshop on the Economics of Information Security. Berlin, Germany. 25-26 June 2012.

[11] Tang, Qian, et al. “Improving Internet Security through Social Information and Social Comparison: A Field Quasi-Experiment.” Twelfth Workshop on the Economics of Information Security. Georgetown University, Washington, D.C. 11 June 2013.

[12] Framework for Improving Critical Infrastructure Cybersecurity, Version 1.0, National Institute of Standards and Technology, February 12, 2014.

[13] Cetin, Orcun, et. al. “Understanding the Role of Sender Reputation in Abuse Reporting and Cleanup,” 14th Workshop on the Economics of Information Security. Delft University of Technology, The Netherland. 22-23 June 2015.

[14] Stone-Gross, Brett, et al. “The Underground Economy of Fake Antivirus Software.” *Economics of Information Security and Privacy III.* Schneier, B. (Ed.). (2013): 55-78, Springer New York.