Cyber Security Policy and Research Institute

THE GEORGE WASHINGTON UNIVERSITY

Social Science, Computer Science, and Cybersecurity

Workshop Summary Report

Lance J. Hoffman

August 22, 2013

Report GW-CSPRI-2013-02

Support for this research was provided through National Science Foundation Division of Computer and Network Systems, Secure and Trustworthy Cyberspace Award CNS-1223630

Social Science, Computer Science, and Cybersecurity
Workshop Summary Report

Lance J. Hoffman

Abstract

This report summarizes a small invitational workshop that brought together computer scientists, social scientists, and others to examine the steps necessarily for integrating social sciences into design of future cybersecurity mechanisms and systems. Its goal was to further the development of communities of researchers who today do not interact, but whose cooperative work is necessary for the development of new and improved cybersecurity systems.

The workshop succeeded in better identifying the greater community of interest, devising a recommended reading list for researchers exploring entry into the field, and arriving at a first approximation of what problems are considered important and identifying where significant progress has been made (and where much remains to be done); similarities and differences between computer scientists and social scientists were noted. Topics were identified that may provide a fertile area for joint research. Mechanisms described in the workshop or a follow-up survey that would facilitate interdisciplinary research in the area included a cybersecurity mentor match system, an open access journal, and a dedicated session at a mainstream cybersecurity conference for social scientists working with computer scientists.

Table of Contents

1. Introduction

Cybersecurity is an important challenge and it is becoming apparent that this complex issue has components based in social science, computer science, and other disciplines. This is the report of a workshop that brought together computer scientists and social scientists in an attempt to start integrating social sciences into the design of future cybersecurity mechanisms and systems. The workshop continued the search for new models of and paradigms for cybersecurity. The hope is that the workshop will lead to the development of communities of researchers that today do not interact, but whose cooperative work is necessary for the development of new and improved cybersecurity systems.

2. Background

2.1. NSF Secure and Trustworthy Cyberspace (SaTC) Program

The Secure and Trustworthy Cyberspace program involves several NSF Directorates: Computer and Information Science and Engineering (CISE), which leads and integrates the program, Social, Behavioral, and Economic Sciences (SBE), Mathematical and Physical Sciences (MPS), Engineering (ENG), and the Office of Cyberinfrastructure (OCI). It encourages research that addresses cybersecurity from one or more of three perspectives: trustworthy computing systems; social, behavioral and economic sciences; and transition to practice, as well as proposals that combine multiple perspectives.

2.2. Relationship of SaTC to NSF and to U. S. Cybersecurity Research

Broad as the SaTC program is, research related to cybersecurity can also be found in a number of other NSF research programs and activities. Outside NSF, many other government agencies conduct research in cybersecurity and information assurance, including the Departments of Defense, Energy, Homeland Security, and others. NSF and other agencies coordinate their research programs through the Cybersecurity and Information Assurance (CSIA) Interagency Working Group under the auspices of the National Information Technology Research and Development (NITRD) program, which reports to the National Science and Technology Council in the Office of Science and Technology Policy (OSTP).

3. Workshop Elements

3.1. Introductions

This report describes a workshop that started on April 23, 2013 at the Hilton Orlando Hotel with a working dinner, followed by brief introductory comments by Peter Muhlberger, Program Director in the Division of Social and Economic Sciences at NSF. Prior to the meeting, each participant had been assigned to a dyad, a two-person group, with a partner. He or she was asked to read the biographical sketch and the white paper of their assigned partner and then prepare a one-minute oral introduction of that person and a two-minute presentation of the highlights or salient points of that person's contributed white paper. The participants delivered these after dinner. Each partner was given the opportunity to respond, expanding on any of the points that were mentioned, adding others, and updating or correcting anything, either about their paper or themselves. Those on the Steering Committee, from NSF, and the facilitator were not exempt – they were also paired and tasked with preparing a one- minute oral introduction of their partner, who was then given one minute to respond to that introduction, in a similar manner.

3.2 Working Sessions

On the following day, the attendees met in plenary session to review the agenda for the day, agree upon facilitation ground rules, and provide further reactions to and comments on the read-ahead material and the comments from the previous evening. (All substantive material from the read-ahead packet (except

initial versions of the white papers if they were revised later) is included in Appendices 1, 2, and 3.) The participants then broke into small working groups after lunch and then reconvened to report their initial observations.
The final day was devoted to finalizing points of agreement on the most productive paths for cybersecurity research that combines efforts of social scientists and computer scientists.

3.3. White Papers

Participants of this workshop were determined by a Steering Committee (see Section 6) that read thirty one-page white papers that had been submitted by potential attendees in response to the Call for Participation in Appendix 1. A dozen of these were selected and their authors invited to attend the workshop. Due to last-minute personal emergencies, two selectees could not attend the workshop, leaving ten of the selectees attending, along with the five steering committee members, two NSF program directors, and a facilitator. Following the workshop, selectees were requested to submit a second white paper that reflected any changes in their views that may have arisen due to the workshop or otherwise since the submission of their first white paper in early 2013. Appendix 2 contains the (final) white papers.

4. Discussion

A number of observations were made in the initial plenary discussion and later in breakout groups. A recurring one was that rather than having social scientists think of cybersecurity as an application, it could instead be considered as a particularization of human interaction through the Internet (e.g. relationship formation). This is also a reason that cybersecurity can be exciting and appealing for social science researchers.

One participant felt that the notion of inviting social scientists in to work on a computer science problem is an insufficient incentive. Instead, he argued that cybersecurity is an instantiation of existing social science problems and research on it allows researchers to advance basic research in their native fields, using already developed social science structures and applying them to cybersecurity.

4.1. Working Groups

After a morning plenary session, the participants broke up into three working groups to tackle different issues and report back to the group.

4.1.1. Working Group 1

Working Group 1 focused on three major areas: metrics and indicators, analysis of the motivations of cybercriminals, and the benefits of data availability.
Metrics and indicators are needed for longitudinal studies on both attackers and users. Some of the major areas of interest include:

- the behaviors and attitudes expressed, and their consequences
- the costs and benefits of focusing research on attackers versus consumers
- how the findings will be able to be translated into actionable data
- repeatability of surveys and experiments on a longitudinal basis

Studies of note included the surveys by Alan Westin from the 1970s until the mid-2000s[1] and the annual Computer Crime and Security surveys[2].

[1] http://repository.cmu.edu/cgi/viewcontent.cgi?article=1857&context=isr
[2] http://gocsi.com/survey

An important issue is the ability to acquire and share datasets. There should be established conditions for access and publication to preserve anonymity. One suggestion was to use a model similar to that of the Census Bureau and procedures similar to its, which allow access by researchers to "semi-public" data.

Cybercrime data appears to be in high demand to help analyze the motivations of cybercriminals. Ways to gather information on this include the use of case studies with potentially an international focus, using both qualitative and quantitative approaches, using researchers on the ground (e.g., Jenna Burrell's work on Nigerian scammers[3]), as well as studying communications networks and analyzing social networks. In connection with this studying of the data, the effects of disruption/poisoning/lemonizing[4] using these databases was also mentioned.

One benefit of more data availability is the ability to develop new and more informed research questions and points of views. Secondary data analysis would be possible and one could study the change over time of important metrics (assuming one has agreed-upon metrics), thus allowing publication at more established venues (e.g., economic journals) with wider readership that are arguably more influential in shaping policy decisions than existing "niche" publications that currently carry this work. Ultimately, informed actions could spring from these studies, including approaches to deter, displace, or convert cybercriminals, and other enforcement actions. Of course, the data itself will have to be protected from unauthorized access and alteration.

4.1.2. Working Group 2

Working Group 2 looked into the cost-benefit analysis of regulations and policy interventions. Emerging from this discussion was the sense of a need to develop indicators that supported an evidence- based approach to cybersecurity. Types of indicators include incidents, organizational performance, user behavior, and security system performance. It will be important to analyze the differences between attitudinal and behavioral indicators and risk perceptions and actual threats. It will also be important when identifying indicators to select those that will stand the test of time. It was also noted that depending on the level of organization being studied (e.g., individual, group, organization, substate, state), different types of social science might be appropriate.

4.1.3. Working Group 3

Working Group 3 considered strategies to disrupt attackers who may have already gained access to a given system through technical means.[5] The group also discussed privacy policies and the chilling effect on privacy-sensitive users caused by the presentation of privacy policies. Perhaps the most interesting conversation related to trust and the difference between interpersonal trust and institutional trust. Between individuals, trust is generally based on past experiences, whereas trust in an institution (for example, a hotel) involves trusting something beyond our control where the individuals that make up that institution (e.g. the housekeepers, bellhops, and managers) will honestly perform their duties even though one may never even meet the individuals.

---

[3] Burrell, Jenna, Problematic Empowerment: West African Internet Scams as Strategic Misrepresentation, http://itidjournal.org/index.php/itid/article/view/308/140.

[4] Vila, Tony, Rachel Greenstadt, and David Molnar. "Why we can't be bothered to read privacy policies: models of privacy economics as a lemons market." *Proceedings of the 5th international conference on Electronic commerce*. ACM, 2003.

[5] Somayaji, A. and Forrest, S. Automated response using system-call delays." Usenix 2000, https://www.usenix.org/legacy/publications/library/proceedings/sec2000/full_papers/somayaji/somayaji.pdf

4.2. Fertile Areas for Social Science Interacting with Computer Science

Reconvening in plenary session, ten areas of interest in cybersecurity were identified as shown in Table 1. Five categories of social science application were also identified as shown in Table 2.

Cybercrime
Terrorism
Espionage
Privacy
Cyberwar
Hackivism
Identity theft
Fraud
Political repression
Distributed denial of service

Table 1: Ten Areas of Interest in Cybersecurity

Attitudes and behaviors
Incentives and constraints
Metrics and indicators
Adaptive adversaries
Trust

Table 2: Five categories of social science application

The development of the social science application in each of the ten areas was then rated in terms of both importance and current progress toward solutions. In doing so, the participants were asked to describe themselves as either "social scientists" or "computer scientists". The results are shown in Figure 1.

	Cybercrime	Terrorism	Espionage	Privacy	Cyberwar	Hacktivism	Identity Theft	Fraud	Political Repression	DDOS
Attitudes and Behaviors	2.5	2.25	2.33	1.4	2	2.33	2	2	2	2
Incentives and Constraints	2.6	3	3	1.67	2.5	2.5	2.5	2.5	2.5	1.33
Metrics/Indicators	1.5	3	3	1.75	3	2	2.25	2.5	1.67	1.33
Adaptive Adversaries	2	3	2.5	2	3	2.5	2	2	2	1.67
Trust	2	3	3	2	3	3	2	2	3	2

(a) Progress (Computer Scientists' View) [1 = making good progress, 2 = limited progress, 3 = no visible progress]

	Cybercrime	Terrorism	Espionage	Privacy	Cyberwar	Hacktivism	Identity Theft	Fraud	Political Repression	DDOS
Attitudes and Behaviors	2	1	3	1.5	3	3	2	2	2	3
Incentives and Constraints	2	1	2	2	2	3	1	1	2	2
Metrics/Indicators	2	2	3	2	3	3	1	1.5	1.5	2
Adaptive Adversaries	2	1	3	2	2	3	2	2	3	1
Trust	1	2	2	2	2	2	1.5	1.5	2	No rating

(b) Progress (Social Scientists' View) [1 = making good progress, 2 = limited progress, 3 = no visible progress]

	Cybercrime	Terrorism	Espionage	Privacy	Cyberwar	Hacktivism	Identity Theft	Fraud	Political Repression	DDOS
Attitudes and Behaviors	A	D		A			A	A	A	
Incentives and Constraints		D		A	A	A	D	D	A	
Metrics/Indicators	A		A	A	A		D		A	
Adaptive Adversaries	A	D	A	A		A	A	A		
Trust				A			A	A		

(c) Progress (Joint View) [A=agree (< 0,5 variance), D= disagree (variance > 1.0)]

	Cybercrime	Terrorism	Espionage	Privacy	Cyberwar	Hacktivism	Identity Theft	Fraud	Political Repression	DDOS
Attitudes and Behaviors	x			x				x		
Incentives and Constraints	x			x			x		x	
Metrics/Indicators	x		x	x		x	x	x		X
Adaptive Adversaries	x		x		x	x		x	x	
Trust	*	*	*	*	*	*	*	*	*	*

(d) Importance (Computer Scientists' View) [X = important, * = Trust was not considered well-defined enough by the computer scientists.]

	Cybercrime	Terrorism	Espionage	Privacy	Cyberwar	Hacktivism	Identity Theft	Fraud	Political Repression	DDOS
Attitudes and Behaviors	x			x			x	x		
Incentives and Constraints	x			x		x		x		
Metrics/Indicators	x		x	x			x	x		
Adaptive Adversaries		x	x		x	x				
Trust				x			x			

(e) Importance (Social Scientists' View) [X = important]

Attitudes and Behaviors	A	A	A	A	A	A	D	A	A	A
Incentives and Constraints	A	A	A	A	A	D	D	D	D	A
Metrics/Indicators	A	A	A	A	A	D	A	A	A	D
Adaptive Adversaries	D	D	A	A	A	D	A	D	D	A
Trust	*	*	*	*	*	*	*	*	*	*

(f) Importance (Joint View) [A=agree (< 0,5 variance), D= disagree (variance > 1.0), * = Trust was not considered well-defined enough by the computer scientists.]

Figure 1. Ratings by Computer Scientists and Social Scientists of Importance of and Progress in Cybersecurity Issues

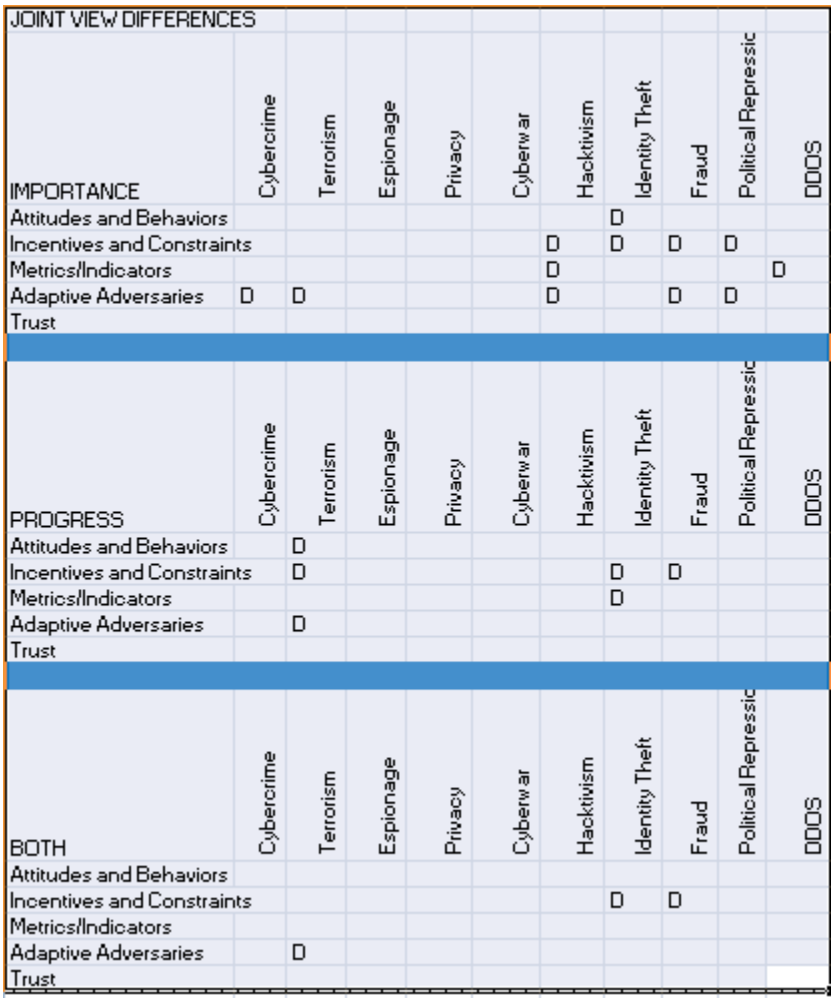

Figure 2. Differences (D) in Joint Views of Computer Scientists and Social Scientists around Cybersecurity Issues

The matrices above provide early indications of directions to pursue that are felt by one or both groups deserving of further research; they also may identify fertile areas for social science interacting with computer science for further research and study. In particular, they highlight some interesting potential similarities and differences in the way computer scientists and social scientists view various aspects of cybersecurity. For example, Figure 1(c) suggests that there is consensus between computer scientists and social scientists on the progress of privacy issues across all categories considered. The most disagreement is over the progress of terrorism issues.

When analyzing importance, as seen in Figure 1(f), consensus was achieved across all categories in espionage, privacy, and cyberwar. The most disagreement was seen with the importance of hacktivism.

We see from Figure 2 that the most stark differences in the very small sample of social and computer scientists at the workshop -- where they disagree on both importance and progress to date -- come in how they view the issues of adaptive adversaries with regard to terrorism and incentives and constraints with regard to hacktivism and identity theft.

4.3. Interdisciplinary Studies in Cybersecurity

Attempts to integrate work in disparate fields can be seen in the research applying social-psychological literature to online behaviors. Also, by grounding this work in more applied research of risk and safety engineering[6] it may be possible to have new solutions to empower non-experts. Similarly, there is some research examining underground markets[7] where cybercriminals exchange information on credit cards, zero day exploits. Our understanding of the development of these markets is limited by the lack of underlying economic theory. A trivial example is the use of the phrase "organized cybercrime". Often it is unclear what either "organized" or "cybercrime" means. While there is an intuitive understanding of the latter, the notion of organization in the context of cybercrime is unexamined. Leveraging existing literature in economics and criminology can provide deeper insights.

4.4. Validation of Interdisciplinary Studies and Support Mentors

Cybersecurity is interdisciplinary and to date its development has not had a lot of theory behind it (except in certain subareas like cryptography). Thus, it can be difficult to achieve recognition of work and research within the field. This is not a problem unique to cybersecurity, nor is it new – it exists to some extent in all research that is interdisciplinary[8,9].

While interdisciplinary work is fascinating, it is too often also high-risk for traditional academic careers due to the focused nature of tenure and promotion decisions in academe. It is often not easy to find "fellow travelers" with an interdisciplinary outlook. Nor is it easy to find successful mentors. Mentorship is important to individual and group success and thus it was suggested that the use of MentorNet10 and similar services might be used as a model for creating a cybersecurity mentor match system.

---

[6] Garg, V., L. Jean Camp. Heuristics and Biases: Implications for Security Design. IEEE Technology and Society, 32(1): 73-79, 2013.

[7] McCoy, Damon and Andreas Pitsillidis, Grant Jordan, Nicholas Weaver, Christian Kreibich, Brian Krebs, Geoffrey M. Voelker, Stefan Savage, and Kirill Levchenko. 2012. PharmaLeaks: understanding the business of online pharmaceutical affiliate programs. In *Proceedings of the 21st USENIX conference on Security symposium* (Security'12). USENIX Association, Berkeley, CA, USA, 1-1.

[8] Salter, Liora and Alison Hearn, Outside the Lines: Issues in Interdisciplinary Research, McGill-Queen's University Press, 1996.

[9] Rhoten, Diana. "Interdisciplinary research: Trend or transition." *Items and Issues* 5.1-2 (2004): 6-11.

[10] http://www.mentornet.net

4.5. Community

Some discussion tried to identify the "greater community" in cybersecurity research. One way to do this is to identify relevant associations, conferences, and publications. Initial ones identified included those in Tables 3, 4, and 5. This is not a complete list; in particular, this does not include several conferences in cybersecurity (such as the IEEE Symposium on Security and Privacy, Usenix Security, RSA, and others) where there is usually little specific attention paid to interdisciplinary research.

Computer Science Faculty
Social Science Faculty

Table 3. Informal Groups

Association for Computing Machinery
Institute for Electrical and Electronics Engineers
American Sociological Association (Section on Communication and Information Technologies)
American Political Science Association (Information Technology and Politics Section)

Table 4. Associations

Privacy Law Scholars Conference*
Workshop on the Economics of Information Security
Workshop on Security and Human Behavior*
IEEE International Symposium on Technology and Society
New Security Paradigms Workshop
APWG eCrime Researcher's Summit
Workshop on Information Systems Economics (WISE)
Symposium on Usable Privacy and Security (SOUPS)
Usable Security Workshop (USEC) at Financial Cryptography

Table 5. Conferences

* = Invitational Event

Emerging from the workshop and also examining results of a post-workshop survey administered to the attendees, one gets the sense that it is still relatively difficult to bring together researchers in computer science and social science to work together on cybersecurity because of a dearth of common conferences and other venues to meet at and no common publication typically read by both groups. One solution proposed was an open access journal, where the workshop participants (and others) could serve on the editorial board. While everyone appeared to enjoy the interaction and meeting new people and researchers, one event is not enough – a continuing venue appears necessary to obtain significant interdisciplinary cooperation. One possible approach would be to convince more mainstream cybersecurity conferences (e.g., USENIX Security) to have a track set aside for social scientists working with computer scientists. One attendee suggested expanding this to include not only social scientists such as anthropologists but also ethicists, telecommunications researchers, consumer affairs and advertising researchers, and lawyers.

5. Additional References

Participants also discussed and compiled a list of important interdisciplinary papers and resources in cybersecurity. Some participants felt that those individuals wishing to do interdisciplinary cybersecurity doctoral research should have read, or at least skimmed, most or all of these.

General

Brashears, M. E. (2013). Humans use Compression Heuristics to Improve the Recall of Social Networks. Sci. Rep., 3. doi:10.1038/srep01513

Cheshire, Coye & Cook, Karen S. (2004). The Emergence of Trust Networks under Uncertainty: Implications for Internet Interactions. Analyse Kritik 26/2004 p. 220-240. http://people.ischool.berkeley.edu/~coye/papers/Cheshire-Cook.pdf

Dasgupta, A., Punera, K., Rao, J., & Wang, X. (n.d.). Impact of Spam Exposure on User Engagement. http://www.justinmrao.com/spamfx.pdf

Dourish, P., & Anderson, K. (2006). Collective Information Practice: Exploring Pr ivacy and Security as Social and Cultural Phenomena. HUMAN-COMPUTER INTERACTION, 21, 319–342. http://www.douri.sh/publications/2006/DourishAnderson-InfoPractices-HCIJ.pdf

Morselli, Carlo, Cynthia Giguère, and Katia Petit. (2007). The efficiency/security trade-off in criminal networks. Social Networks 29(1):143-153.

Nissenbaum, H. (2011). "A Contextual Approach to Privacy Online," Daedalus 140 (4), Fall 2011: 32-48. https://www.amacad.org/publications/daedalus/11_fall_cover.pdf

Rose, Jeremy & Jones, Matthew (2005) "The Double Dance of Agency: A Social Theoretical Account of How Machines and People Interact." Signals, Signs Action (1:1). 19-37.

Stajano, F., & Wilson, P. (2009). Understanding the psychology of scam victims: seven principles for systems security. http://www.cl.cam.ac.uk/~fms27/scams/

van Eeten, M., et al. (2010), "The Role of Internet Service Providers in Botnet Mitigation: An Empirical Analysis Based on Spam Data", WEIS 2010

Varian, Hal (1996), Economic Aspects of Personal Privacy. http://people.ischool.berkeley.edu/~hal/Papers/privacy/

NITRD web page on Cybersecurity and Information Assurance, http://www.nitrd.gov/nitrdgroups/index.php?title=Cyber_Security_and_Information_Assurance_(CSIA)#title

Economics of Security

1. Ross Anderson, Economics and Security Resource Page, http://www.cl.cam.ac.uk/~rja14/econsec.html

2. Ross Anderson and Tyler Moore, The Economics of Information Security, Science 27 October 2006: Vol. 314 no. 5799 pp. 610-61, http://www.sciencemag.org/content/314/5799/610.abstract

3. Jean Camp, Information Security Economics Bibliography, http://infosecon.net/workshop/bibliography.php

Economics of Privacy

1. Alessandro Acquisti, Resources on the economics of privacy, financial privacy, and the economics of anonymity, http://www.heinz.cmu.edu/~acquisti/economics-privacy.htm

2. Alessandro Acquisti, "The Economics of Personal Data and the Economics of Privacy", Background Paper #3, Roundtable on The Economics of Personal Data and Privacy: http://www.oecd.org/sti/ieconomy/46968784.pdf

3. 30 Years after the OECD Privacy Guidelines, http://www.oecd.org/sti/ieconomy/46968784.pdf

Behavioral economics/psychology in security and privacy

1. Ross Anderson, Psychology and Security Resource Page, http://www.cl.cam.ac.uk/~rja14/psysec.html

2. Alessandro Acquisti and Jens Grossklags, "What Can Behavioral Economics Teach Us About Privacy?", http://www.heinz.cmu.edu/~acquisti/papers/Acquisti-Grossklags-Chapter-Etrics.pdf

3. Bruce Schneier, The Psychology of Security, http://www.schneier.com/essay-155.html

4. Adele Howe et al., The Psychology of Security for the Home Computer User, 2012 IEEE Symp. on Security and Privacy, http://www.ieee-security.org/TC/SP2012/papers/4681a209.pdf

6. Acknowledgements

This event could not have taken place without the hard work of many individuals. They are acknowledged below.

Principal Investigator

Prof. Lance. J. Hoffman, Cybersecurity Policy and Research Institute, The George Washington University

Program Director at National Science Foundation

Peter Muhlberger, Program Director in the Division of Social and Economic Sciences, National Science Foundation

Steering Committee

Prof. Denise Anthony, Sociology, Dartmouth University
Prof. Alessandro Acquisti, Information Systems and Public Policy, Carnegie Mellon University
Prof. Eugene Spafford, Computer Science, Purdue University
Prof. Jeff Hancock, Communication, Cornell University

Support Personnel

Regina Elwell, Graduate Student, Cybersecurity Policy and Research Institute, The George Washington University
Barbara Robinson, Facilitator, Robinson & Associates

Support for this research was provided through National Science Foundation Division of Computer and Network Systems Award CNS-1223630.

APPENDICES

Appendix 1 – Call for Participation

CALL FOR PARTICIPATION IN NSF-SPONSORED WORKSHOP TO EXPLORE SOCIAL SCIENCE CONTRIBUTIONS TO UNDERSTANDING CYBERSECURITY
April 23-25, 2013

We seek social scientists with expertise in, or willingness to explore, the behavioral, economic, political, psychological or sociological aspects of cybersecurity to develop new scholarship in the social sciences. The goals of this new scholarship are two-fold: 1) to stimulate and advance our understanding of the behavioral and analytic aspects of cybersecurity and 2) to inform cybersecurity policy and development. This small, invitation-only workshop will be held April 23-25, 2013 to foster the development of new models of and paradigms for cybersecurity by developing communities of researchers between social science and computer and information systems experts, who today may not interact, but whose cooperative work is necessary for the development of cybersecurity mechanisms and systems. The workshop will also produce a research agenda in the social sciences related to cybersecurity that appropriately addresses user, economic, and sociopolitical realities related to the topic.

This project complements the burgeoning panorama of interdisciplinary initiatives in this area (such as WEIS, SHB, and SOUPS) by connecting their results to others from social science and by leveraging the growing body of interdisciplinary cybersecurity research to produce a research agenda in the social sciences that appropriately addresses user, economic, and sociopolitical realities related to cybersecurity.

By highlighting where current marketplace, policy, organizational, and other incentive mechanisms fall short of creating strong cybersecurity, and building on the tools from behavioral, economic and other social sciences to address these issues, we hope this work will lead to the development of actual, working mechanisms and incentives for building more secure and stable information systems - rather than the much less effective yet more prevalent practice of adding security after the fact to existing systems. Collaborative social science and computer science research can inform designers of systems in mobile, desktop, and network environments to be able to develop more effective mechanisms to solve or mitigate existing security problems. It can also inform users (individual and organizational) and regulators about how to better create secure information environments and systems.

The research agenda that results from the workshop will inform investments in research and development for a wide range of stakeholders including university researchers, government, commercial firms, and nonprofits. Clearer understanding of incentive mechanisms to design cybersecurity systems will encourage the production of more secure systems in the future, thus promoting the progress of science and advancing national defense and international welfare by having more secure systems in place.

STEERING COMMITTEE:
Alessandro Acquisti, Carnegie Mellon University, acquisti@andrew.cmu.edu
Denise L. Anthony, Dartmouth College, denise.anthony@dartmouth.edu
Jeff Hancock, Cornell University, jth34@cornell.edu
Eugene H. Spafford, Purdue University, gene@spaf.us
Lance J. Hoffman, George Washington University, lanceh@gwu.edu

Travel support will be provided by the National Science Foundation. Attendees will be expected to submit a short paper with their thoughts on the topic and what they can contribute by the deadline below, and then to submit another paper within a week of the end of event that presents their changed thinking or

their reasons that their thoughts on the topic have not changed. This second paper can be more than one page. Both of these papers will appear as appendices in the project final report that will be a public document.

Workshop dates: April 23-25, 2013. Researchers interested in attending should submit a one-page white paper on what they would contribute (limited to between 300 and 600 words on one page, at least 11 point type size) along with their name, affiliation, and contact information to cspriaa@gwu.edu by February 1, 2013. The subject line should read ―CYBERSECURITY WORKSHOP FOR SOCIAL SCIENCE RESEARCHERS—WHITE PAPER. Selected invitees will be notified by February 15, 2013.

Appendix 2 – White Papers

In this appendix are the white papers selected for the workshop. Those with an asterisk (*) next to their titles are the original submitted papers, unrevised after the workshop. If an author could not attend at the last minute due to a family emergency or illness, there is a double asterisk (**) next to their name. In the case of multiple authors, the attending author's name is in boldface.

Interdicting Cyber Social Attacks: Agenda for a New Cybersecurity (Matthew E. Brashears)

Matthew E. Brashears
Cornell University, Department of Sociology
Meb299@cornell.edu, 607-255-4925

The greatest threat to cybersecurity is social, rather than technological, in nature. All electronic systems contain "exploits" that can permit an unauthorized user to gain access to or control over the equipment. The resulting damage ranges from defacing a website to rendering the compromised hardware inoperative (e.g., by "phlashing" the firmware, or re-writing it with corrupt data). Traditionally, the often-superior resources of defenders have counter-balanced the unpredictability of attackers, but several key changes have upset this balance. First, computers, including smartphones, have become cheap enough that substantial amounts of computational power are now widely distributed. Second, the merging of the internet with pervasive cellular data coverage allows large numbers of devices to interact simultaneously. As a result, attackers can use massive numbers of networked devices to overpower their targets (e.g., Distributed Denial of Service, or DDoS, attacks), ensuring aggregate computational and bandwidth superiority (e.g., Anonymous' "Low Orbit Ion Cannon"). Eliminating exploits will be a part of any cybersecurity plan, but large amounts of distributed computing power represents the central cybersecurity threat of the future.

How can we reduce the likelihood of these assaults? First, DDoS attacks often use computers that have been compromised by malicious software. As a result, users must be educated about appropriate security measures and motivated to use them. Success in this area will reduce the total resources available to adversaries and force foreign attackers to rely on assets located at greater distances from U.S. servers, thereby increasing their latency.

Second, many individuals participate in attacks willingly, either because they agree with the goals or for entertainment (e.g., "for the lulz"). Steps must be taken to identify the types of individuals most likely to participate in attacks and to develop methods of interdiction. The former will require accurate behavioral modeling that should use inexpensive and non-confidential data. The latter encompasses a variety of approaches ranging from persuasion to disrupting adversary groups.

Finally, it will be necessary to develop early warning systems for cyber attacks. These systems could use the enormous amount of data available on twitter, 4chan, and other elements of the social web. In all of the above cases, however, the central problems are social rather than strictly technological: how can we inform and motivate friendly parties to deny their assets to attackers; how can we identify and defuse potential participants; and how can we detect attacks that are in preparation?

Cybersecurity planning must also attend to the diverse threat environment. Most often security is compromised by cybercriminals, who seek to use electronic systems to achieve financial gain. These entities seek profit and thus will not expend more resources than they believe they can gain back from their successes. In contrast, terrorist groups and nation-states may be willing to expend resources without the promise of direct profit so long as it offers the opportunity to compromise the target (e.g., by destroying infrastructure). Thus, while cybercriminals will be a frequent and pervasive challenge, they will also exhibit some restraint, whereas terrorists and nation-states will be less frequent but more acute threats. Cybersecurity efforts must be flexible enough to deal with both threats, particularly given that terrorists and nation-state actors may rely on cybercrime markets for both technologies and proxies. Finally, cybersecurity is not a battle to be won but an environment of competition; a successful policy is

one that manages the cybersecurity environment so that losses are tolerable, rather than impossible. This management must be adaptive such that serious threats to infrastructure are resisted aggressively, while less severe threats are held to a more practical standard.

Research on Cybersecurity as an Organizational Phenomenon (Anthony M. Cresswell and Theresa A. Pardo)

Anthony M. Cresswell & Theresa A. Pardo, Center for Technology in Government
University at Albany-SUNY; tpardo@ctg.albany.edu and tcresswell@ctg.albany.edu 518-442-3892

There is growing awareness that cybersecurity effectiveness depends as much on the capability of the organizational setting as on technology, but this view is not adequately reflected in the bulk of cybersecurity research. This awareness is reflected in some expert opinion, such as the Congressional testimony of the GAO's Director of Information Security Issues: "\It is also important to bear in mind the limitations of some cybersecurity technologies and to be aware that their capabilities should not be overstated. "Technologies do not work in isolation" (Wilshusen, 2012, p. 10). In addition, there has been some increased research attention to the organizational side of the problem (e.g., Carol Hsu, Jae-Nam Lee, & Straub, 2012). However the volume of related research is small. This lack of attention to cybersecurity appears to be due in part to a focus on the behavior of individual users, more or less in isolation. Thus security experts typically acknowledge the importance of the social aspects of cybersecurity in terms of this focus: "It is clear to me that computer security is not a problem that technology can solve. Security solutions have a technological component but security is fundamentally a people problem" (Schneier, 2004, p. 2). Indeed there is a large and growing body of cybersecurity research based on the user as the unit of analysis. While this individual level view is critical, it is far too narrow to support the organization-level action that is necessary to sustain cybersecurity effectiveness. Therefore increased attention to organizational science-based research on cybersecurity capability seems valuable and appropriate direction for new research.

Capability focused research has consider able promise in this regard. It can reveal how technical and organizational factors interact in shaping the ways organizations to respond to constantly emerging threats to information assets. These threats are serious, sustained, and constantly escalating, in particular for governments, which are stewards of highly sensitive and valuable information and are obligated to secure many different forms of the public's information. Government data ranges across personal and organizational financial and health records, vital records (land title, birth, marriage, etc.), public safety and law enforcement data, and legal records, among others. In addition to the complex security demands on governments related to the information assets, government capability may be compromised by complex policy and legal requirements, dwindling financial resources, tangled jurisdictional relationships. A focus of research on government cybersecurity, therefore, can offer an important and rich venue for capability-focused cybersecurity research.

The capability for innovation and adaptability is central. To protect their information assets in a dynamic threat environment, governments must continually adapt and innovate both the technical and organizational components of cybersecurity. Technically focused research will not yield a sufficiently complete or nuanced picture of the demands cybersecurity places on organizational capabilities. Similarly, organizational capability studies that do not focus on the challenges of ongoing technical change are insufficient, as are studies that focus on organizational capabilities and resources without attention to how they interact with technical systems. Studies employing the necessary holistic perspective are largely lacking, particularly when government is the focus.

One benefit of a more holistic perspective on cybersecurity is a critical shift in thinking on the role of IT systems in organizational capability, namely that they can be as much a liability as an asset. The conventional view is that IT systems and the information assets they contain enhance capability, though it may not be clear exactly how the enhancement works (Nevo & Wade, 2010). In the words of Gruber et

al., "… current theory is not sufficiently clear on how different kinds of resources and capabilities contribute to performance, nor does it clarify how firms can combine different resources and capabilities to achieve superior performance outcomes" (Gruber, Heinemann, Brettel, & Hungeling, 2010, p. 1337).

A more robust approach to research on the capability for cybersecurity could address basic theoretical questions in both organizational and technical terms. However such a holistic approach will likely require new models. It is not clear, for example, how well current models of organizational capability fit cybersecurity, in which social and technical phenomena are highly entangled in IT systems. Similarly, current organizational capability models do not adequately address IT as a source of either vulnerability or competitive advantage (Bharadwaj, 2000; Gruber et al., 2010; Otim, Dow, Grover, & Wong, 2012). Because of the central role of technical components, cybersecurity systems provide an excellent venue for exploring the sociomaterial nature of technology at work in complex organizations. This research can explore the processes of interaction and governance that establish the workings and failings of cybersecurity in practice. The insights gained can inform the growing body of theory and empirical work and contribute to a more holistic understanding of how technical and organizational capabilities combine and interact to promote and sustain government cybersecurity.

References

Bharadwaj, A. S. (2000). A Resource-Based Perspective on Information Technology Capability and Firm Performance: An Empirical Investigation. MIS Quarterly, 24(1), 169–196.

Carol Hsu, Jae-Nam Lee, & Straub, D. W. (2012). Institutional Influences on Information Systems Security Innovations. Information Systems Research, 23(3), 918–939.

Gruber, M., Heinemann, F., Brettel, M., & Hungeling, S. (2010). Configurations of resources and capabilities and their performance implications: an exploratory study on technology ventures. Strategic Management Journal, 31(12), 1337–1356.

Nevo, S., & Wade, M. R. (2010). The Formation and Value Of IT-Enabled Resources: Antecedents and Consequences of Synergistic Relationships. MIS Quarterly, 34(1), 163–183.

Otim, S., Dow, K. E., Grover, V., & Wong, J. A. (2012). The Impact of Information Technology Investments on Downside Risk of the Firm: Alternative Measurement of the Business Value of IT. Journal of Management Information Systems, 29(1), 159–194.

Ransbotham, S., & Mitra, S. (2009). Choice and Chance: A Conceptual Model of Paths to Information Security Compromise. Information Systems Research, 20(1), 121–139.
Schneier, B. (2004). Secrets and Lies: Digital Security in A Networked World. Indianapolis, IN: Wiley Publications, Inc.

Siponen, M., Pahnila, S., & Mahmood, M. A. (2010). Compliance with Information Security Policies: An Empirical Investigation. Computer, (February), 64–71.

Wilshusen, G. C. Information Security: Cyber Threats Facilitate Ability to Commit Economic Espionage. , Pub. L. No. GAO-12-876T (2012). Washington, DC: US Government Accountability Office. Retrieved from http://www.gao.gov/assets/600/592008.pdf

Beyond increasing awareness* (Neil Gandal and Sonia Roccas)

Professor **Neil Gandal**
Chair, Berglas School of Economics
Tel Aviv University – Gandal@post.tau.ac.il

Professor Sonia Roccas
The Open University of Israel
e-mail: soniaro@openu.ac.il

Cybersecurity is to a large extent determined by the behavior of the end-users. The integrity the network depends on the willingness of the users to adhere to security guidelines. Increasing awareness to security threats and the steps that should be taken to avoid them is an important factor. However, increased awareness is not sufficient to ensure safe behavior. People often behave in ways that expose them to the risk of undesirable consequences even when they are well aware of these consequences (unhealthy eating habits, unsafe driving practices, etc.). Thus, a different approach is needed to identify factors that increase willingness of end-users to adopt safe behavior.

We propose to use apply the vast knowledge accrued on personal values for this purpose. Values are abstract desirable goals that serve as guiding principles in people's lives. They are a core aspect of people's identity, and serve as standards or criteria that provide social justification for choices and behaviors across situations. Unlike needs and motives, which may be unconscious, values are represented cognitively in ways that enable people to think and communicate about them (Schwartz, 1992). Values are ordered by subjective importance, and thus form a hierarchy of value priorities. The relative importance of different values affects perception and interpretation, emotions, daily actions and long term behavior (e.g., Bardi & Schwartz, 2003; Maio, Olson, Allen, & Bernard, 2001; Roccas & Sagiv, 2010). Values are especially good predictors of behaviors over which individuals have some cognitive control or choice (Roccas et al, 2002).

The Schwartz value theory is particularly useful for studying behavior related to cybersecurity because it seeks to provide a comprehensive mapping of values according to the motivations that underlie them (Schwartz, 1992). Schwartz distinguishes between ten motivational goals that are organized in a circular order according to two basic conflicts. Self-enhancement values emphasize pursuit of self-interests, even at the expense of others. They conflict with self-transcendence values which emphasize concern for the welfare and interests of others, close and distant. (See Gandal et al 2006 for details.)

Behavior that is inconsistent with cybersecurity is usually not malicious, and not random. It serves to attain important motivations. Informing users that a specific practice is unsafe, is only useful to the extent to which the primary motivation of the users is to obtain safety. But safety is rarely the core motivations of end-users. We reason that a first step in a program of research aimed at changing behavior of end users is to identify the values that are attained by behavior that breaches security: for example, sharing intimate information is consistent with social connectedness (the motivation captured by self-transcendence values), sharing information about one's success is consistent with to self enhancement values. Effective intervention should target these motivations, and allow their attainment through safe means. We reason that two types of interventions should be particularly effective: 1. Designing IT systems that ensure that the primary motivation of the users is promoted when they implement safe behavior. 2. Providing the users with information that addresses their core motivations. We plan to pursue a research agenda based on this "white paper."

Cybercrime: Lessons from Criminology* (Vaibhav Garg)

Vaibhav Garg
Drexel University

Current research on cybercrime focuses on describing the underground markets, the different stakeholders, and respective transactions. There is limited investigation of how these markets emerged and why they persist. The presumption is that since crime online is a high profit and low risk activity, it is obviously rational for individuals to be thus engaged. There are several limitations to this extant paradigm.

First, the only solution to cybercrime under the current paradigm is deterrence, i.e. increase the cost of criminal engagement and decrease the profits; e.g. prosecution, more secure technologies. However, even this solution is limited by a lack of theoretical grounding. Becker has argued that as prosecution increases, crime becomes organized, and enforcement is constrained by corruption. We have seen this online as well; as deterrence based approaches have become successful criminals have become organized. From a local governance perspective, such crime brings in much needed cash flow to what might be a stagnant economy, thereby increasing net social welfare. Thus, it is not that cybercriminals come from regions with no cyber laws, but that the misalignment of economic incentives prevents from such laws being framed in specific jurisdictions.

Second, cost-benefit optimization is not the only imperative that drives individual criminal behavior. Presumably, most academics attending this workshop do not engage in cybercrime. Given that the profits from cybercrime are noted to be in millions and median computer science professors' salaries tend to be around $100,000, why do these professors not become cybercriminals? It would certainly make sense from a narrow rational choice perspective. Arguably, they have the required skillset. Criminological theories then argue for other economic, structural, and cultural factors that may be equally relevant if not more. For example, legitimate employment is either not (locally) available, or that entry costs are prohibitively high. It is then critical to address opportunity cost of cybercrime.

Third, deterrence based approaches are expensive. Often they cost more than the financial loss due to crime. Under a prosecution only regime, it would be prohibitively expensive to throw the book at every individual on just one underground forum, let alone all of them. Technological solutions cannot be a 100% effective and suffer from lack of adoption. Shouldn't we then explore non-deterrence based solutions?

A broader understanding of the limitations of deterrence-based solutions is then critical to effective long-term policy and technical solutions that are sustainable. Simultaneously, alternative explanations to crime, as provided by criminological research, must be considered and their relevance examined online. My research begins such examination, by considering the macroeconomic, structural, and cultural variables that correlate with the existence of cybercrime. Specifically, I leverage cross-country differences in cybercrime to engender insights for public policy and technical design. Such investigations are much needed for successful envisioning of cyber-security initiatives.

Towards Effective Behavioral and Economic Security Engineering* (Jens Grossklags)

Jens Grossklags
College of Information Sciences and Technology, The Pennsylvania State University

Misaligned economic incentives, perilous user behavior and a constantly evolving threat landscape are some of the reasons why securing systems is hard. However, addressing these problem dimensions in isolation is like trying to pick up mercury with your hands. It does not work particularly well, and it is unsafe.

The objective of my research is to effectively (re)design real-world choice architectures to improve end-user security with combined measures of behavioral and economic design. While it is challenging to address a user's behavioral limitations in an environment with interdependencies and externalities caused by other users and attackers, making progress along these dimensions is of significant societal importance.

To the workshop, I can contribute a combination of an analytic and behavioral/experimental perspective on cybersecurity. In addition, I have gained experience in the translation of my work to the policy world with participation in activities organized by the Federal Trade Commission, the Defense Advanced Research Projects Agency ISAT, the World Wide Web Consortium, the European Network and Information Security Agency and several other institutions.

Descriptive science: Studying the messiness of security decision-making with theory and experiments

My modeling work understands security as a multifaceted good that can be provided in different ways: users can attempt to prevent security breaches, focus on mitigation and recovery methods, and/or rely on risk transfer (e.g., cyber-insurance). Theoretical predictions are derived for different types of network interdependencies and distinctions about whether security is organized as a public or private good.

Under experimental conditions, I am studying the impact of these different risk management options and environmental factors on users' behaviors. In my experiments conducted in economic laboratories, groups of users defend networks with interdependent security consequences. I observe that individuals experiment frequently with the different security options but with mixed results. In contrast to non-probabilistic and non-interdependent environments, convergence to individually rational strategies is slow, and optimal outcomes on a group level are rare (even in the long run).

Prescriptive science: Designing effective behavioral and economic interventions

The design of choice architectures for security needs to address the willingness to secure a resource on an individual level, but should also contribute to better overall security outcomes in the presence of externalities and interdependencies.

Descriptive models and experimental evidence help with the design in several ways. First, they provide an experimental economic testbed and natural benchmark for the impact of interventions. Second, obstacles apparent from the model (e.g., tradeoffs between different kinds of security technologies) and evidence from the experiments (e.g., learning behavior) guide the development of metrics to evaluate the performance of incentive schemes.

In my work, I systematically study the effectiveness of different hard economic (i.e., carrot-and-stick approach) and soft behavioral incentives in security settings. To design effective educational

programs and intervention mechanisms for security, we need to better understand probabilistic and interdependent environments, in theory and practice. My research agenda provides a pathway to derive generalizable principles to achieve this goal.

I would be delighted to have the opportunity to contribute to the workshop and to continue building a community around social, economic and technical aspects of cybersecurity.

Privacy Norms and Cybersecurity* (Christine Horne)

Christine Horne**
Department of Sociology
State University Pullman, WA

My view of cybersecurity is informed by my research on social norms in the lab, as well as by my new investigations into privacy norms in the context of the Smart Grid. I see the social/behavioral issues related to cybersecurity as substantially overlapping with privacy. Privacy norms regulate information flow – what information goes to whom and for what purpose (Nissenbaum 2010). Security breaches are problematic because they result in the wrong information going to the wrong people, who may in turn use that information in ways that harm individuals or collectivities. That is, security breaches can be conceptualized as particularly problematic invasions of privacy – invasions that often have financial implications.

If so, then understanding privacy norms may contribute to understanding of the human component of cybersecurity. Norms researchers often rely on understanding of the collective goals of a group and information about how behaviors affect those goals to make predictions about norm enforcement and emergence (Coleman 1990; Horne 2009). Typically, they look at situations in which all group members have the same interests. But, of course, interests may diverge. In order to understand privacy norms more specifically, it may make sense to identify the goals of the parties to an interaction (which may differ), and analyze the costs and benefits of behaviors in relation to those goals (see, e.g., Horne et al. 2013). This approach may help us to identify what systems designers/managers and users see as problematic security issues (and where their perceptions diverge), as well as design interventions to increase good security behaviors. In the context of the Smart Grid, for example, this approach would lead us to expect that designers/managers of the Grid would be concerned about security (theft, etc.) while residential end-users would be more worried about the (legal) knowledge that a utility company might gain about them. These concerns have implications for attentiveness to security issues.

References

Coleman, James S. 1990. *Foundations of Social Theory.* Cambridge, MA: Belknap Press.

Horne, Christine. 2009. *The Rewards of Punishment.* Stanford: Stanford University Press. Horne, Christine, Brice Darras, Scott Frickel, and Anurag Srivistava. 2013. ―Privacy

Norms and the Smart Grid. Unpublished manuscript.

Nissenbaum, Helen. 2010. *Privacy in Context.* Stanford: Stanford University Press.

Accountability as an Interface between Cybersecurity and Social Science∗ (Joan Feigenbaum, Aaron D. Jaggard, and Rebecca N. Wright)

Joan Feigenbaum† **Aaron D. Jaggard‡** Rebecca N. Wright§

Overview "Accountability" is generally agreed to be important, although this term is used to mean many different things. Across these various uses, it is related to the idea of deterrence instead of prevention. Computer-science-based work on accountability is providing models to formalize accountability and disambiguate it from related notions. This work also raises many questions that need to be informed by social science, suggesting the potential for fruitful interaction at this cross-disciplinary interface.

The benefit of computer-science perspectives We have proposed a formal view of "accountability" in trace-based and game-theoretic terms. This framework helps to make precise different notions related to "accountability" and to distinguish between them. This and other frameworks also open up the possibility of proving formal relationships (e.g., implications and tradeoffs) involving different accountability-related notions. Interaction with the social sciences will enrich formal models of accountability and make them more realistic. In turn, these enhanced models will help answer real-world social-science questions related to accountability.

The need for social-science perspectives Considering formal models of accountability in computer science, we may identify issues for which we expect social-science perspectives to make significant contributions. Some of these are below. However, increased interaction between computer and social scientists may change the framing of these questions. More fundamentally, interactions between computer and social scientists would better identify the types of questions that should be studied at this cross-disciplinary interface.

In studying deterrence, the question of what constitutes "effective deterrence" is an important one: bad behavior should actually be deterred, which may be somewhat independent of whether any particular technical definition is satisfied. An understanding of human responses to a range of incentives, both positive and negative (such as payments, incarceration, shame, and praise) would inform both more complete models for further study as well as more effective systems for real-world use.

Relatedly, the utility functions studied in connection with accountability may differ dramatically between people. One utility function may be most typical in a population, another one slightly less typical, and so on. Effectively deterring undesired behavior by a "typical" individual may be relatively easy using socially acceptable means. By contrast, deterring even the sociopaths, without knowing in advance who they are, might require measures that would be draconian if applied uniformly. This leads naturally to the example of "three strikes" laws, which suggests that it would be beneficial to further explore formal frameworks for accountability in connection with a broader and deeper knowledge of criminology, sociology, law, and other disciplines.

Accountability often makes use of causality. Our model uses causality to connect punishments to violations; as noted by various people, causality also arises in treating "blame" as something other than a black box. Additionally, identifying violations and assigning blame often make use of some sort of evidence. Causality and evidence have been studied, e.g., at the boundary of computer science and philosophy. Further work on these concepts in the context of accountability would help ensure that accountability systems are viewed as legitimate (for example that punishment is meted out only when it is sufficiently justified).

"Accountability," "deterrence," and related terms have various connotations in colloquial usage. For example, our candidate definition of accountability is in terms of punishment and not "calling to account," etc. Should we instead use a term like "deterrence" for what is captured by our candidate definition of "accountability?" Our formal framework also intentionally allows for a violator to be automatically punished, without necessarily identifying the violator or even revealing that a violation occurred. As Weitzner has asked, is it better to reserve the use of "accountability" for cases in which someone knows that a violation occurred? Even once accountability-related terms are sufficiently disambiguated for interdisciplinary communication, properly framing them will be important for communicating to broader audiences (such as end users or consumers) what accountability systems are intended to do as well as what they do not do.

∗ Partially supported by NSF grants CNS-1016875 and CNS-1018557.
† Department of Computer Science, Yale University. joan.feigenbaum@yale.edu
‡ Formal Methods Section (Code 5543), U.S. Naval Research Lab oratory. aaron.jaggard@nrl.navy.mil
§DIMACS and Department of Computer Science, Rutgers University. rebecca.wright@rutgers.edu

Heuristics and Biases in Cybersecurity Dilemmas (Richard S. John, Heather Rosoff, and Tracy Cui)

Richard S. John, Heather Rosoff, Tracy Cui
University of Southern California

Cybersecurity often depends on decisions made by human operators, who are often thought of as a major cause of security failures – "the weakest link in the chain" (Schneier 2008). Most security software is designed to give the user multiple options when a potential threat is detected. Alternative responses by the user often range in terms of risk and convenience. The most expedient, convenient option is usually the most risky, while the safest option is often more time consuming and requires more effort on the part of the user. Human operators of computer systems are often required to make security related decisions, e.g., whether to install software to protect internet usage, whether to download desired files to a local storage device, whether to submit personal information for identity purpose, etc. In such cases, the human operator is often required to make a trade-off between risk and convenience in responding to cybersecurity threats.

We conducted 2 experiments, with over 500 respondents, to explore whether and how cybersecurity decision making responses depend on gain-loss framing and prior near-miss experiences. Tversky and Kahneman (1981) first reported the gain-loss framing heuristic using the "Asian disease" decision problem. In that study, respondents were more likely to select the sure thing over a (risky) gamble when the outcomes were framed as gains, but were more likely to select a (risky) gamble when the outcomes were framed as losses. As defined by Dillon and Tinsley (2005), "An event is considered a "near-miss" if the outcome is non-hazardous, but if a hazardous or fatal outcome could have occurred." A near-miss occurs when an event (such as a computer virus), which had some nontrivial probability of ending in disaster (loss of data), does not because good fortune intervenes. Dillon and Tinsley (2005) report an empirical study suggesting that individuals tend to evaluate near-misses as a type of success.

In experiment I, we employed a 2x2 within-subjects factorial design, manipulating the frame (gain vs. loss) and the presence vs. absence of a near miss experience. A total of 266 college students responded to all four cybersecurity scenarios, and indicated the extent to which they would recommend a risky response vs. a safer response to a close friend. Over all four scenarios and all four treatment conditions (framing and near miss experience), respondents were more likely to endorse the safer action. Results suggest that the experience of a near miss significantly increases respondents' endorsement of safer response options. There were no significant differences for framing; contrary to hypothesized effect, subjects were no more likely to endorse the risky option in the loss frame than in the gain frame. These results were not moderated by respondent sex or previous cyber victimization.

Experiment II followed the same general paradigm as the first experiment, with the following modifications. The framing manipulation was dropped, and the near miss manipulation were revised to include 3 different types of past experiences: false alarm, near miss, and a hit involving a loss of data. In addition, respondents were given a screen shot for 3 different cybersecurity scenarios. Order and pairing with near miss condition were counterbalanced. A diverse sample of 247 respondents were recruited from Amazon Mechanical Turk. Results indicate that the experience of a hit significantly increases respondents' endorsement of safer response options relative to the near miss past experience. In addition, the experience of a false alarm significantly decreased respondents' likelihood of endorsing safer response options, compared to the near miss past experience. These results were not moderated by respondent sex, age, level of education, income level, self-reported computer knowledge, whether the

respondent had protection software installed, and whether the respondent had been victim of a computer virus.

Results indicate that respondents tend to be more protective given the knowledge of past hits and near-misses. In particular, when respondents were "reminded" of a past loss or near escape from a cyber threat, they are significantly more likely to advise a best friend to avoid the risk. Instead of interpreting the recovery from a danger, e.g., computer froze but functioned normally after restarting, as good luck and a sign of vulnerability, people are more likely to perceive it as an indication of resilience. In contrast to the result reported by Dillon, Tinsley and Cronin (2011) for hurricane insurance purchase, where people believe their house could survive even if a hurricane hit the area, this study indicates that respondents are more concerned about the consequences of a cyber threat following recall of a near miss experience. We suspect that respondents feel more vulnerable following a near miss event because of direct prior experienced bad consequences from cyber threats, e.g., loss of data, loss of privacy, and financial loss from online transactions. Therefore, in the near-miss conditions used in this study, people perceive the near-misses as "near failure" rather than as a "difficult success."

Characterizing and Mitigating Risks in Long-Lived Personal Data Archives (Chris Kanich)

Chris Kanich
Computer Science Department
University of Illinois at Chicago

Although cybersecurity is popularly thought of as a moving target, concepts like the principle of least privilege and public key cryptography have been well understood for decades. The explosive growth of computer and Internet use as integral parts of society at all levels is forcing us to rethink our understanding of secure system design. The fact that humans from several different walks of life are interacting with these systems on a daily basis has prompted a paradigm shift: rather than designing secure systems with arbitrarily defined use models, we must design secure systems with use models informed by how people interact with each other, computers, and information. This security paradigm necessitates a close collaboration between technical and social scientists so that the design of secure systems incorporates an understanding of the needs and capabilities of the billions of people that will rely on them.

My interest in cybersecurity centers on securing the Internet by leveraging an understanding of the human-level motivations of the attackers. I am eager to use my technical expertise as part of a collaboration with social scientists to improve our understanding of the way that users—both legitimate and illegitimate—interact with systems online so that we can design systems that prevent harm to legitimate users.

One project I am particularly interested in pursuing relates to characterizing users' understandings of privacy within long-lived digital storage. Currently, cloud services like replicated online backup and web-based email allow users to keep information intact and globally accessible for decades. These storage mediums make it simple to accumulate gigabytes of information without paying a single cent. These services are no doubt a boon to users, but exhaustive, decades-long digital archives can also become an even greater liability if a cybercriminal were to gain access to them. I wish to explore several angles of this situation: what information is lucrative to a cybercriminal? How much financial value could be extracted from this information? Would end users recognize this liability and erase old information, or information known to be valuable to a cybercriminal? Can we devise systems that maintain guarantees about availability while presenting a more difficult avenue to monetization for cybercriminals who do gain access? Each of these questions hinges on a full understanding of the technical aspects of security alongside an understanding of the humans who will be using these systems on a daily basis.

Statement on Interdisciplinary Cybersecurity Research (Tyler Moore)

Tyler Moore, Computer Science and Engineering Department
Southern Methodist University
http://lyle.smu.edu/~tylerm/

I am interested in conducting research that fundamentally improves information security. I have learned a great deal from my technical experience identifying attacks and designing countermeasures for a broad range of systems, from SS7 signaling protocols used in telephone switches to wireless sensor networks and phishing attacks on the Internet. But technical solutions alone are not enough to make systems secure. In my view, economic analysis should complement technical design in order to better understand, evaluate and manage threats. In my research, I combine practical skills in secure system design with relevant tools from economics, notably incentive analysis, modeling and econometrics.

My primary research method is to leverage empirical observation to improve our understanding of how cybercriminals operate, as well as identify ways in which defenders can strengthen their information security posture. Existing security mechanisms have not kept up with the rapidly evolving strategies of adversaries. By deriving attacker behavior from what can be directly monitored, we can better explain what happens and why. For example, I have studied phishing scams to peer into the world of online crime, detailing the attackers' strategies and the defenders' responses. Empirical analysis offers our best hope of designing realistic threat models that stay on top of attacks. Conducting this empirical analysis requires a combination of expertise from computer science and the social sciences, particularly econometrics.

A related research topic that follows on from conducting empirical analysis is the development of indicators that track the prevalence of incidents and cybersecurity levels over time. Economists have developed a broad set of indicators that track the health of economies, such as GDP, market indices, trade flows, etc. Because there is some consensus over what should be measured, there is now a rich history of time-series data on these indicators available for many jurisdictions. This has enabled specialization within the economics discipline – some organizations focus on collecting data to construct reliable indicators, allowing other researchers to take the indicators as input to their own work.

In the nascent field of security economics, researchers are burdened with the dual task of collecting and analyzing the data. This has raised the barriers to entry, particularly for social scientists lacking the technical skills to collect and store data, as well as decide what data is worth keeping. As a security economics researcher coming from a technical background, I am keen to help design the indicators and implement the infrastructure necessary to collect the requisite data that can be used by social scientists.

The barriers to carrying out effective interdisciplinary research go both ways, of course. Computer scientists are often unfamiliar with methods of quantitative analysis, particularly concerning experimental design. While cybersecurity experts trained in computer science should not seek to become experts in social sciences (and vice versa), there is considerable value in each "side" acquiring basic proficiency in terminology and methods of the other. Without it, those working together on interdisciplinary approaches to cybersecurity will likely talk past each other instead of solving the most pressing problems. Consequently, any research agenda on interdisciplinary approaches to cybersecurity should also emphasize the need to develop curriculum to bring experts quickly up to speed on key results across disciplines.

A Policy Analysis Framework for Cybersecurity Operations* (Amy Silva and Richard Andres)

Amy Sliva**
College of Computer Science &
Department of Political Science
Northeastern University
asliva@ccs.neu.edu

Richard Andres
National War College & Institute for
National Strategic Studies National Defense
University
rich.andres@gc.ndu.edu.

Over the last few years, states and associated cyber militias have demonstrated increasing willingness to use cyber tools to harm each other 's public and private infrastructure. Stuxnet, Flame, Duqu, Red October, and Iran's ongoing attacks against the U.S. banking system have brought the phenomenon into the nightly news, but the pattern of attacks includes numerous older and less discussed programs such as Titan Rain, Ghostnet, and the attacks on Estonia and Georgia. The White House and National Security Agency have estimated that cyber related offensive actions cost the global economy hundreds of billions each year, and Secretary Clinton has recently warned China to desist from its ongoing campaign to steal U.S. intellectual property and infiltrate critical U.S. infrastructure.

Unfortunately, the increase in harm caused by cyber weapons has not been accompanied by an equal increase in research on how to think about and classify various destructive cyber related actions and effects. When should we consider an attack a criminal matter? What qualifies an attack as espionage? When should policymakers discuss a cyber attack in the same terms they usually reserve for a kinetic military action? These questions are particularly urgent because militaries around the world are beginning to play a larger role in cyber conflict, and countries regularly talk about cyber quarrels in military terms. U.S. Secretary of Defense Leon Panetta has publically described America's willingness to consider kinetic responses to cyber attacks and compared cyber incidents with traditional conventional operations. If states now see some cyber incidents as rising to the level of war, how should policymakers decide which attacks cross this line or determine which agencies or departments should have responsibility for particular types of attacks?

Previous work on this subject has attempted to classify the gravity of cyber incidents in relation to international law (see particularly Schmitt 1999). However, this approach has limited practical utility because major cyber attacks, like incidents related to more traditional forms of espionage and military violence, do not easily lend themselves to legal frameworks. Generally, the law is silent on cases involving espionage, and whether a state chooses to classify an act as warlike is a political, rather than a legal question. Even when the law is likely to be clear on an incident, dynamics surrounding putative anonymity in cyberspace make finding and presenting legal evidence difficult.

In this paper, we will present a policy analysis framework to aid in thinking about and classifying cyber incidents as crime, espionage, or military actions. Our goal is to provide a practical and versatile analytic guide for understanding cyber incidents. This framework will consider the interplay between the directly observable elements of an attack, such as damage done, and those factors that can only be ascribed indirectly, such as intentions or norms. We seek to identify which attributes are most crucial for decision-making in various situations and how different classification schemes will impact the practical implementations of cyber policy and national diplomacy. We illustrate our framework by applying it to several recent cases including Stuxnet, Titan Rain, and the non-use of cyber operations by NATO against Libya in Operation Odyssey Dawn.

References

Bharadwaj, A. S. (2000). A Resource-Based Perspective on Information Technology Capability and Firm Performance: An Empirical Investigation. MIS Quarterly, 24(1), 169–196.

Carol Hsu, Jae-Nam Lee, & Straub, D. W. (2012). Institutional Influences on Information Systems Security Innovations. Information Systems Research, 23(3), 918–939.

Gruber, M., Heinemann, F., Brettel, M., & Hungeling, S. (2010). Configurations of resources and capabilities and their performance implications: an exploratory study on technology ventures. Strategic Management Journal, 31(12), 1337–1356.

Nevo, S., & Wade, M. R. (2010). The Formation And Value Of It-Enabled Resources: Antecedents And Consequences Of Synergistic Relationships. MIS Quarterly, 34(1), 163–183.

Otim, S., Dow, K. E., Grover, V., & Wong, J. A. (2012). The Impact of Information Technology Investments on Downside Risk of the Firm: Alternative Measurement of the Business Value of IT. Journal of Management Information Systems, 29(1), 159–194.

Ransbotham, S., & Mitra, S. (2009). Choice and Chance: A Conceptual Model of Paths to Information Security Compromise. Information Systems Research, 20(1), 121–139. Schneier, B. (2004). Secrets and Lies: Digital Security in A Networked World. Indianapolis, IN: Wiley Publications, Inc.

Siponen, M., Pahnila, S., & Mahmood, M. A. (2010). Compliance with Information Security Policies: An Empirical Investigation. Computer, (February), 64–71.

Wilshusen, G. C. Information Security: Cyber Threats Facilitate Ability to Commit Economic Espionage. , Pub. L. No. GAO-12-876T (2012). Washington, DC: US Government Accountability Office. Retrieved from http://www.gao.gov/assets/600/592008.pdf

Cybersecurity Research Collaboration Between Computer Scientists and Social Scientists* (Yang Wang)

Yang Wang
School of Information Studies, Syracuse University

While I am a computer scientist by training, I have a keen interest in understanding the human aspect of cybersecurity. I believe that cybersecurity is inherently a socio-technical challenge that requires inter/multi-disciplinary solutions. I am thrilled to learn about this workshop because I have benefited from collaborating with social scientists on cybersecurity research. While this interdisciplinary, collaborative approach is promising, it's also risky. A senior researcher at the workshop rightly pointed out that interdisciplinary collaboration is hard and particularly risky for junior faculty members who have not had tenure yet.

However, like a few other junior faculty members at the workshop, I am enthusiastic about interdisciplinary research on cybersecurity. Part of my training is in Human-Computer Interaction (HCI) which is itself an interdisciplinary field that draws from computer science as well as social sciences such as psychology and anthropology. My primary research area is usable privacy and security which is primarily concerned with designing systems that ordinary people can understand and use to protect their privacy and security. I have drawn from social science theories and methodologies in studying people's privacy perceptions and behavior, and designing privacy- enhancing technologies.

I have had the opportunities to collaborate with social scientists on a number of cybersecurity projects. My latest collaboration with behavioral economists has been particularly productive. We drew from literature on human cognitive and behavioral biases as well as soft paternalism in designing user interfaces that nudge people to be more thoughtful about their information disclosure decisions. But not all my collaboration experiences were as productive as I would like. I think we need to pay attention to factors that may promote or perturb this kind of collaboration. Based on my own experience, I think there are a number of factors that could strain the collaboration such as misaligned incentives, different disciplinary cultures or norms, and lack of bridge persons or common ground.

It's very encouraging to see many positive signs for building the common ground for computer and social scientists at the workshop. The occurrence of the workshop itself is a good sign. One of my favorite activities at the workshop was that the computer scientists teamed up and so did the social scientists, and the two groups independently rated every cell of a matrix, where each column represents a cybersecurity problem (e.g., privacy) and each row describes an aspect of the problem (e.g., user perception and behavior), on (a) how much progress have been made in their own fields, and (b) how important is that aspect of the cybersecurity problem that we should invest more efforts. This activity yielded a good start of a cybersecurity research agenda. It's also striking to observe that the two groups largely agreed with each other---another good sign of common ground. At the workshop, we have also discussed other ways to help build the interdisciplinary cybersecurity research community such as establishing mentorship, compiling a list of survey papers from different disciplines about cybersecurity (each workshop participant already recommended one of their favorite papers), and starting a new interdisciplinary conference or journal on cybersecurity. The NSF SaTC program also explicitly supports collaborations between computer scientists and social scientists on cybersecurity research. The recent SaTC EAGER program even requires such a collaboration.

While we need to be mindful about interdisciplinary collaboration, it's not a bad time to explore such collaboration opportunities for cybersecurity research!

Appendix 3 – Participant Biographies

RESEARCHERS

Alessandro Acquisti

Alessandro Acquisti is an associate professor at the Heinz College, Carnegie Mellon University (CMU) and the co-director of CMU Center for Behavioral and Decision Research. He investigates the economics of privacy. His studies have spearheaded the application of behavioral economics to the analysis of privacy and information security decision making, and the analysis of privacy and disclosure behavior in online social networks. Alessandro has been the recipient of the PET Award for Outstanding Research in Privacy Enhancing Technologies, the IBM Best Academic Privacy Faculty Award, multiple Best Paper awards, and the Heinz College School of Information's Teaching Excellence Award. He has testified before the U.S. Senate and House committees on issues related to privacy policy and consumer behavior. Alessandro's findings have been featured in national and international media outlets, including the Economist, the New York Times, the Wall Street Journal, the Washington Post, the Financial Times, Wired.com, NPR, and CNN. His 2009 study on the predictability of Social Security numbers was featured in the ―Year in Ideas‖ issue of the NYT Magazine (the SSNs assignment scheme was changed by the US Social Security Administration in 2011). Alessandro holds a PhD from UC Berkeley, and Master degrees from UC Berkeley, the London School of Economics, and Trinity College Dublin. He has held visiting positions at the Universities of Rome, Paris, and Freiburg (visiting professor); Harvard University (visiting scholar); University of Chicago (visiting fellow); Microsoft Research (visiting researcher); and Google (visiting scientist). He has been a member of the National Academies' Committee on public response to alerts and warnings using social media.

Denise Anthony

Denise Anthony is Associate Professor and past-Chair (2007-11) in the Department of Sociology at Dartmouth College. She is also Research Director of the Institute for Security, Technology, and Society (ISTS) at Dartmouth, Adjunct Associate Professor in the Department of Community and Family Medicine at Geisel School of Medicine, and a faculty affiliate at the Center for Health Policy Research at The Dartmouth Institute for Health Policy and Clinical Practice. Dr. Anthony's works in a number of theoretical areas to explore issues of group dynamics, organizational behavior, and institutional change. She studies cooperation, trust, and social capital in a variety of settings, from micro-credit borrowing groups to online groups such as Wikipedia and Prosper.com. In health care, she has studied organizational and institutional variation in managed care practices, physician referral behavior, and patient preferences for care. More recently her work examines the use and implications of information technology in health care, including effects on quality, as well as the implications for the privacy and security of protected health information in health care delivery. Her multi- disciplinary research has been published in journals in sociology as well as in health policy and computer science, including among others the American Sociological Review, Social Science and Medicine, Journal of the American Medical Association, Health Affairs, and IEEE Pervasive Computing. She has received grants from the National Science Foundation, and the Department of Health and Human Services Office of the National Coordinator for Health IT SHARP program, among others.

Matthew E. Brashears

Matthew E. Brashears is an Assistant Professor of Sociology at Cornell University, specializing in social network analysis. He is the sole P.I. on a Defense Threat Reduction Agency grant intended to develop new methods of identifying terrorist groups preparing chemical, biological, radiological, and nuclear (CBRN) attacks. He is also sole P.I. on a National Science Foundation grant exploring the connections between elements of cognition and social network structure. He has been published in a number of outlets, including Nature Scientific Reports, the American Sociological Review, Social Networks, and Social Psychology Quarterly.

Anthony Cresswell

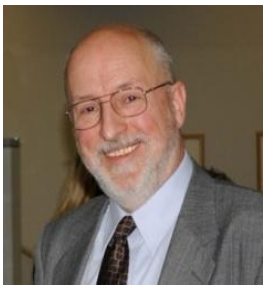

Dr. Cresswell is a Senior Fellow at the Center for Technology in Government, University at Albany. He works with government, corporate, and university partners to conduct applied research on the policy, management, and technology issues of government IT innovation. Dr. Cresswell joined CTG as a senior research fellow in 1994 and served as Deputy Director and later as interim director in 2008-09. His studies include the public value of investment in government IT, and problems of interorganizational information sharing, organizational capability, and IT impacts on practice. Dr. Cresswell joined the University at Albany in1979. He holds faculty appointments in Educational Administration and Information Science. He previously served on the faculties of Northwestern University and Carnegie-Mellon University, and as Faculty Advisor in the US Office of Management and Budget. His international experience includes information system and policy analysis projects in Africa, Asia, Europe, the Middle East, and Caribbean. He holds a doctorate from Columbia University.

Jeremy Epstein

Jeremy Epstein is Program Director of the NSF Secure and Trustworthy Cyberspace (SaTC) program. He is on loan to NSF from SRI International, where his research focused on voting systems security, an inherently interdisciplinary field that has tightly coupled computer science and social science issues. He holds an BS in Computer Science from New Mexico Tech, an MS in Computer Sciences from Purdue University, and is ABD from George Mason University.

Neil Gandal

Neil Gandal is Professor of Economics and Head of the Berglas School of Economics at Tel Aviv University. He received his B.A. and B.S. degrees from Miami University (Ohio) in 1979, his M.S. degree from the University of Wisconsin in 1981, and his Ph.D. from the University of California-Berkeley in 1989. Professor Gandal is a research fellow at the Centre for Economic Policy Research. He was the managing editor of the International Journal of Industrial Organization from 2005-2012. Professor Gandal has published numerous papers in industrial organization, the economics of information technology, the economics of the software Internet industries, and the Economics of Information Security.

Vaibhav Garg

Vaibhav Garg is a post doctoral researcher in the computer science department at Drexel University. His research investigates the cross- section of security and human behavior. It combines elements of social psychology, behavioral economics, and risk communication. He received his PhD from Indiana University in Security Informatics. His dissertation research examined the determinants of perceived risk online and their ability to inform both non-expert behaviors and attitudes. Garg is also interested in the emerging field of cybercrime science. Specifically, he examines the macro level economic, structural, and cultural variables that explain the geographic concentration of cybercrime and resulting victimization. The first paper in this won the best paper award at APWG's eCrime Researcher's Summit in 2011. He is also interested in information ethics, eGovernance, eHealth, and policy.

Jens Grossklags

Dr. Grossklags is an Assistant Professor and holds the endowed Haile Family Early Career Professorship at the College of Information Sciences and Technology at the Pennsylvania State University. Previously, he served as a Postdoctoral Research Associate at the Center for Information Technology Policy, and as a Lecturer of Computer Science at Princeton University. In 2009, he completed his doctoral dissertation at UC Berkeley's School of Information. While at UC Berkeley, he also obtained master's degrees in Computer Science, and Information Management and Systems. He is studying information privacy, security, technology policy and networked interactions from a theoretical and practical perspective. Specifically, Dr. Grossklags is motivated to contribute to a better understanding of the current and future marketplace for personal and corporate information, and improved designs of the underlying evolving security infrastructure. His academic work is very cross-disciplinary and utilizes analytic, empirical and experimental methodologies

Jeff Hancock

Jeff Hancock is an Associate Professor in the Departments of Communication and Information Science, where he is co-Chair, and he is the co-Director of Cognitive Science at Cornell University. He is also the Associate Editor of the journal Discourse Processes. His work is concerned with the psychological and interpersonal dynamics of social media, with a particular emphasis on language use and deception. His research is supported by funding from the National Science Foundation and the Department of Defense, and his work on lying online has been featured frequently in the media, including New York Times, CNN, and TED. Dr. Hancock earned his PhD in cognitive psychology at Dalhousie University, Canada, and joined Cornell in 2002.

Lance Hoffman

Lance J. Hoffman, educator and researcher, is Distinguished Research Professor of Computer Science and Director of the Cybersecurity Policy and Research Institute at The George Washington University in Washington, D. C. Professor Hoffman developed the first regularly offered course on computer security at the University of California, Berkeley in 1970 after serving on the Advisory Committee to the California Assembly Committee on Statewide Information Policy. His second book, Modern Methods for Computer Security and Privacy, published in 1977, was a standard textbook in the few computer security courses offered at the time around the world. A Fellow of the Association for Computing Machinery (ACM), Dr. Hoffman institutionalized the ACM Conference on Computers, Freedom, and Privacy. He has served

on a number of Advisory Committees including those of Federal Trade Commission, the Department of Homeland Security, the Center for Democracy and Technology, and IBM. He has chaired the Information Security Subcommittee of the IEEE Committee on Communications and Information Policy and is a Member of the Subcommittees on Law, and Security and Privacy of the U. S. Public Policy Council of the ACM

Christine Horne

Christine Horne is associate professor of sociology at Washington State University. Her research focuses on social norms, in particular their emergence and enforcement. She tests theory in the lab and applies theoretical insights to explain substantive norms including, for example, privacy norms in the context of smart meters, international human rights norms, and norms regulating gender relations in Africa. She is author of The Rewards of Punishment (Stanford University Press) and editor (with Michael Hechter) of Theories of Social Order (Stanford Social Science).

Aaron D. Jaggard

Aaron D. Jaggard is a mathematician in the Formal Methods Section (Code 5543) at the U.S. Naval Research Laboratory. His current work draws on techniques ranging from game theory to formal methods in order to study various aspects of trustworthy network-mediated interactions. This work includes proving guarantees and fundamental tradeoffs in security and privacy, formalizing and reasoning about accountability and related properties, and proving convergence properties and guarantees of reliable behavior in dynamics that arise in a wide variety of distributed-computing settings. Jaggard took his Ph.D. in mathematics from the University of Pennsylvania; before joining NRL, he held positions at Tulane (Mathematics), Rutgers (DIMACS), and Colgate (CS).

Richard John

Richard John is an associate professor of psychology at USC's College of Letters, Art, and Science, as well as the Director of Undergraduate Studies for the psychology department. In addition to having taught at USC for the past 27 years, Professor John has published a myriad of research articles throughout his career. Articles in referred journals include, ―Cognitive Function in Asymptomatic HIV Infection‖ (1997), Reference Effects: A Sheep in Wolf 's Clothing (1980), and Co-parenting: A link between marital conflict and parenting in two parent families (2001).

Chris Kanich

Chris is an Assistant Professor in the Department of Computer Science at the University of Illinois at Chicago. Chris Earned his Ph.D. in Computer Science and Engineering from UC San Diego in 2012, and his B.S. in Mathematics and Computer Science from Purdue University in 2005. His research centers around Internet security and Internet measurement, with a particular focus on fully characterizing attackers' motivations, capabilities, and strategies.

Tyler Moore

Tyler Moore is an Assistant Professor of Computer Science and Engineering at Southern Methodist University. His research interests include the economics of information security, the study of electronic crime, and the development of policy for strengthening security. Moore holds BS degrees in Computer Science and Applied Mathematics from the University of Tulsa, and a PhD in Computer Science from the University of Cambridge, He is a Director and Vice President of the International Financial Cryptography Assocation (IFCA) and Vice Chair of the IFIP 11.10 Working Group on Critical Infrastructure Protection. Previously, Moore was a postdoctoral fellow at Harvard University's Center for Research on Computation and Society, and the Norma Wilentz Hess Visiting Assistant Professor of Computer Science at Wellesley College. He is a 2004 Marshall Scholar.

Peter Muhlberger

Peter Muhlberger is a National Science Foundation (NSF) Program Director in topical areas such as political science, cybersecurity, and data-intensive research. He is a member of the NSF Cyberinfrastructure Framework for 21st Century Science and Engineering Strategic Leadership Group and the Expeditions in Education Working Group. He is currently on leave from his position as the Director of the Center for Communication Research in the College of Media and Communication at Texas Tech University. He received his Ph.D. in political science from the University of Michigan. Dr. Muhlberger has published in such journals as Political Psychology, Political Communication, the Journal of Information Technology and Politics, and Information Polity. He designed and directed research on Carnegie Mellon University's Virtual Agora Project, a NSF-funded grant project investigating the political, social, and psychological effects of computer-mediated political engagement. He was also principal investigator on the Deliberative E-Rulemaking Project, a NSF-funded project to apply natural language processing and multi-level deliberation to federal agency online rulemaking.

Amy Sliva

Amy Sliva is an Assistant Professor of Computer Science and Political Science at Northeastern University. With this interdisciplinary appointment, she is researching new artificial intelligence models and large-scale data analytics for understanding, forecasting, and responding to behavioral dynamics in intergroup conflict, security policy, and international development. She is currently collaborating with the National Defense University on a policy analysis framework for cyber warfare and is developing computational models of the strategic and behavioral components of cybersecurity to aid policy makers in real-time decision-making. Amy previously worked for the University of Maryland Laboratory for Computational Cultural Dynamics, where she developed decision-support tools for the National Security and Intelligence Communities for counterterrorism analysis, and created similar behavioral modeling technologies at the World Bank for education development in Nigeria. Amy received her Ph.D. in Computer Science from the University of Maryland in 2011. She also has a B.S. in Computer Science from Georgetown University (2005), an M.S. in Computer Science from the University of Maryland (2007), and a Master of Public Policy (M.P.P.) in International Security and Economic Policy from the University of Maryland (2010).

Eugene Spafford

Eugene H. Spafford is a professor of Computer Sciences at Purdue University. He is also a professor (courtesy) of each of Electrical and Computer Engineering, Philosophy, Political Science, and Communication. He is also the founder and Executive Director of the Center for Education and Research in Information Assurance and Security, a campus-wide multi- disciplinary Center with a broadly- focused mission to explore issues related to protecting information and information resources. Spafford has been in computing for over 30 years. Some of his work is at the foundation of current security practice, including intrusion detection, firewalls, and whitelisting; his most recent work has been in cybersecurity policy, forensics, cyber conflict, and future threats. His interests range over these and many other topics, and this has been one of the reasons why he is considered by many to be a polymathic futurist, although some view him as simply an iconoclastic crank. Professor Spafford is a Fellow of the AAAS, ACM, IEEE, (ICS)2, and a Distinguished Fellow of the ISSA. Among many other activities he is currently the chair of the Public Policy Council of ACM (USACM), and is editor-in-chief of the journal Computers Security.

Yang Wang

Yang Wang is an assistant professor in the School of Information Studies at Syracuse University. His research is centered around privacy and security, and social computing. He was a research scientist at CyLab in Carnegie Mellon University. There, he collaborated with Bell Labs on privacy enhancing technologies, and researched privacy issues in online behavioral advertising and privacy concerns of online social networks across different cultures. He has also been working on studies, models and preventive systems related to regrettable behavior in social media. His work has won Best Paper Honorable Mention at the ACM CHI Conference and Future of Privacy Forum's annual―Privacy Papers for Policy Makers‖. His work has also appeared in popular media such as New York Times, Wall Street Journal, and BusinessWeek. He received his Ph.D. in information and computer sciences from University of California, Irvine. In his thesis work, he built a privacy enhancing personalization system that takes into consideration privacy regulations and individuals' privacy preferences. Additionally, Wang worked at several industry research labs such as Intel Research, Fuji Xerox Palo Alto Laboratory, and CommerceNet.

PROJECT SUPPORT PERSONNEL

Regina Elwell

Regina Elwell is a Research Assistant for the Cybersecurity Policy and Research Institute at The George Washington University. She holds a BS in Biology from The University of North Carolina at Chapel Hill and a MS in High Technology Crime Investigation from The George Washington University. While attending The George Washington University, Regina was a CyberCorps student and Information Assurance Scholarship Program Participant. She focused her studies on computer forensics, network intrusion investigation, and information assurance.

Barbara Robinson

Barbara Robinson is the Principal, Robinson Associates, a management-consulting firm based in Washington, DC., which offers a range of services, including: strategic planning, executive coaching, change management, team building, conflict resolution and facilitation. She also coaches individuals to enhance their job satisfaction and productivity at work, as well as to explore career transition options. Barbara has been honing her facilitation skills since the 1980s when she was trained by Interaction Associates, a leader in the field. She has facilitated both small and large groups including retreats, board meetings, and complex neighborhood meetings in the public and private sectors. Her client base spans the public, private, and non-profit sectors in North America and Europe. Barbara has been an adjunct faculty member in the School of Library and Information Science at Catholic University, and has conducted training workshops across North America on strategic budgeting and question handling. She has served on a number of boards of for-profit and non-profit organizations. Barbara has a B. A. from Mount Holyoke College and an M.L.S. from Simmons School of Library and Information Science.