Privacy Papers for Policy Makers

2012

The publication of "Privacy Papers for Policy Makers" was supported by AT&T, Microsoft, and GMAC.

November 7, 2012

We are delighted to provide you with FPF's third annual "Privacy Papers for Policy Makers," representing cutting-edge research and analysis on a variety of important privacy issues.

The featured works were selected by members of the Future of Privacy Forum Advisory Board (scholars, privacy advocates, and Chief Privacy Officers) based on criteria emphasizing clarity, practicality and overall utility. Given the excellent submissions we received, choosing was a difficult task. But we believe our Advisory Board has chosen well, and has put together a diverse and thought-provoking collection. Two of the papers were recipients of the IAPP award for best papers presented at the 2012 Privacy Law Scholars Conference.

We hope this relevant and timely scholarship helps inform and stimulate thinking among policy makers and policy influentials in the US and around the world, with whom we are sharing this compilation.

We are delighted to share new ways of thinking about privacy.

We want to thank AT&T, Microsoft, and GMAC for their special support of the "Privacy Papers" project.

Sincerely yours,

Christopher Wolf
Founder and Co-chair

Jules Polonetsky
Director and Co-chair

Future of Privacy Forum Advisory Board

Alessandro Acquisti
Associate Professor of Information Technology and Public Policy at the Heinz College, Carnegie Mellon University

Jim Adler
Chief Privacy Officer & General Manager, Data Systems, Intelius

Ellen Agress
Senior Vice President and Deputy General Counsel, News Corporation

Annie I. Antón
Professor and Chair, Georgia Tech School of Interactive Computing

Stephen Balkam
CEO, Family Online Safety Institute

Kenneth A. Bamberger
Professor of Law, Berkeley School of Law

Elise Berkower
Associate General Counsel, Privacy, The Nielsen Company

Debra Berlyn
President, Consumer Policy Solutions

Joan (Jodie) Z. Bernstein
Counsel, Kelley Drye & Warren, LLP and former director of the Bureau of Consumer Protection at the Federal Trade Commission

Michael Blum
General Counsel, Quantcast

Bruce Boyden
Assistant Professor of Law, Marquette University Law School

Allen Brandt
Corporate Counsel, Data Privacy & Protection, Graduate Management Admission Council (GMAC)

Jim Brock
CEO, PrivacyChoice

Justin Brookman
Director, Consumer Privacy, Center for Democracy & Technology

Kathryn C. Brown
Senior Vice President, Public Policy Development and Corporate Responsibility, Verizon

James M. Byrne
Chief Privacy Officer, Lockheed Martin Corporation

Ryan Calo
Assistant Professor, University of Washington School of Law Affiliate Scholar, Stanford Center for Internet and Society

Dr. Ann Cavoukian
Information and Privacy Commissioner of Ontario

Brian Chase
General Counsel, Foursquare Labs, Inc.

Danielle Citron
Professor of Law, University of Maryland Law School

Maureen Cooney
Senior Counsel and Deputy Chief Privacy Officer, Sprint Nextel

Lorrie Faith Cranor
Associate Professor of Computer Science and Engineering, Carnegie Mellon University

Mary Culnan
Professor Emeritus, Bentley University

Simon Davies
Founder, Privacy International

Michelle De Mooy
Senior Associate, National Priorities, Consumer Action

Elizabeth Denham
Information and Privacy Commissioner for British Columbia

Michelle Dennedy
Chief Privacy Officer, McAfee

Leslie Dunlap
Vice President of Privacy, Policy and Trust, Yahoo! Inc.

Benjamin Edelman
Assistant Professor, Harvard Business School

Erin Egan
Chief Privacy Officer, Policy, Facebook

Keith Enright
Senior Corporate Counsel, Google

Leigh Feldman
SVP, Senior Privacy Executive Global Compliance Risk – Enterprise Privacy, Bank of America

Eric Friedberg
Co-President, Stroz Friedberg

Rip Gerber
President and CEO, Locaid

Scott Goss
Senior Privacy Counsel, Qualcomm

Susan Gindin
Sr. Privacy Manager, Wal-Mart

Jennifer Barrett Glasgow
Chief Privacy Officer, Acxiom

Greg Goeckner
Executive Director, Merchant Risk Council

Kimberly Gray
Chief Privacy Officer, IMS Health

Sean Hanley
Director of Compliance, Zynga Game Network, Inc.

Pamela Jones Harbour
Former Federal Trade Commissioner; Partner, Fulbright & Jaworski LLP

Megan Hertzler
Director of Information Governance, Xcel Energy

Michael Ho
VP Business Development, Bering Media

David Hoffman
Director of Security Policy and Global Privacy Officer, Intel

Marcia Hoffman
Staff Attorney, Electronic Frontier Foundation

Tim Hollenbeck
Director of Global Ethics and Compliance, Procter & Gamble

Chris Hoofnagle
Director, Berkeley Center for Law & Technology's information privacy programs and senior fellow to the Samuelson Law, Technology & Public Policy Clinic

Jane Horvath
Apple, Inc.

Sandra Hughes
Sandra Hughes Strategies, Ltd.

Brian Huseman
Director, Public Policy, Amazon

Jeff Jarvis
Associate Professor; Director of the Interactive Program, Director of the Tow-Knight Center for Entrepreneurial Journalism at the City University of New York

David Kahan
General Counsel, Jumptap

Ian Kerr
Canada Research Chair in Ethics, Law & Technology, University of Ottawa, Faculty of Law

Bill Kerrigan
CEO, Abine, Inc.

Brian Knapp
Chief Privacy Officer and Vice President, Corporate Affairs, Loopt

Jerry Kovach
Senior Vice President, External Affairs, Neustar

John Kropf
Deputy Counsel, Privacy and Information Governance, Reed Elsevier

Fernando Laguarda
Vice President, External Affairs and Policy Counselor, Time Warner Cable

Manuj Lal
Vice President, Legal Affairs, Press Ganey Associates, Inc.

Barbara Lawler
Chief Privacy Officer, Intuit

Peter Lefkowitz
Chief Privacy Officer, Oracle

Adam Lehman
Chief Operating Officer and GM, Lotame Solutions

Gerard Lewis
Senior Counsel and Chief Privacy Officer, Comcast

Chris Libertelli
Head of Global Public Policy, Netflix

Chris Lin
Executive Vice President, General Counsel and Chief Privacy Officer, comScore, Inc.

Brendon Lynch
Chief Privacy Officer, Microsoft

Mark MacCarthy
Vice President of Public Policy, The Software & Information Industry Association

Siobhan M. MacDermott
Chief Privacy Officer, AVG Technologies

Fran Maier
Founder and Board Chair, TRUSTe

Jennifer Mardosz
Chief Privacy Officer, Fox Entertainment Group

William McGeveran
Associate Professor, University of Minnesota Law School

Terry McQuay
President, Nymity

Scott Meyer
CEO, Evidon

Doug Miller
Global Privacy Leader, AOL, Inc.

Saira Nayak
Director of Policy, TRUSTe

Future of Privacy Forum Advisory Board (continued)

Lina Ornelas
General Director for Privacy
Self-Regulation, Federal Institute for Access to Information and Data Protection Mexico

Kimberley Overs
Assistant General Counsel, Pfizer, Inc.

Harriet Pearson
Partner, Hogan and Lovells

George Pappachen
Chief Privacy Officer, Kantar Group

Christina Peters
Senior Counsel, Security and Privacy, IBM

Robert Quinn
Chief Privacy Officer and Senior Vice President for Federal Regulatory, AT&T

Peter Rabinowitz
Chief Privacy Counsel, American Express

MeMe Rasmussen
VP, Chief Privacy Officer, Associate General Counsel, Adobe Systems

Katie Ratté
Executive Counsel, Privacy Policy and Strategy, The Walt Disney Company

Joel R. Reidenberg
Professor of Law, Fordham University School of Law

Neil Richards
Professor of Law, Washington University Law School

Shirley Rooker
President, Call for Action

Mike Sands
President and Chief Executive Officer, BrightTag

Russell Schrader
Chief Privacy Officer and Associate General Counsel – Global Enterprise Risk, Visa Inc.

Paul Schwartz
Professor of Law, University of California-Berkeley School of Law

Cary Sherman
Chairman and CEO, The Recording Industry Association of America

Ho Shin
General Counsel, Millennial Media

Meredith Sidewater
Senior Vice President and General Counsel, Lexis Nexis Risk Solutions

Emery Simon
Counselor, Business Software Alliance

Dale Skivington
Chief Privacy Officer, Dell

Daniel Solove
Professor of Law, George Washington University Law School

Cindy Southworth
Vice President of Development & Innovation, National Network to End Domestic Violence (NNEDV)

JoAnn Stonier
SVP and Global Privacy & Data Protection Officer, MasterCard

Zoe Strickland
VP, Chief Privacy Officer, United Health Group

Greg Stuart
CEO, Mobile Marketing Association

Lior Jacob Strahilevitz
Sidley Austin Professor of Law, University of Chicago Law School

Peter Swire
Professor, Ohio State University Moritz College of Law

Omar Tawakol
CEO, BlueKai

Omer Tene
Associate Professor, College of Management School of Law, Rishon Le Zion, Israel

Owen Tripp
Co-Founder and Chief Operating Officer, Reputation.com

Catherine Tucker
Mark Hyman, Jr. Career Development Professor and Associate Professor of Management Science, Sloan School of Management, MIT

Steven Vine
Chief Privacy Officer, PulsePoint

Hilary Wandall
Chief Privacy Officer, Merck & Co., Inc.

Mark Weinstein
Founder and CEO, Sgrouples

Michael Zimmer
Assistant Professor in the School of Information Studies, University of Wisconsin-Milwaukee

General Electric

Table of Contents

*Recipients of the IAPP award for best papers at the 2012 Privacy Law Scholars Conference

Out of respect for copyright law and for ease of reference, this compilation is a digest of the papers selected by the Future of Privacy Forum Advisory Board and does not contain full text. the selected papers in full text are available through the referenced links.

Bridging the Gap Between Privacy and Design

Deidre Mulligan and Jennifer King

Full paper available at: http://www.futureofprivacy.org/privacy-papers-2012/

Executive Summary

This article explores the gap between privacy and design in the context of "lateral privacy"—privacy issues arising among users of a service rather than from the service provider — on social networking sites (SNSs) and other platforms by analyzing the privacy concerns lodged against the introduction of Facebook's News Feed in 2006. Our analysis reveals that the dominant theory of privacy put forth by regulators, privacy as individual control, offers little insight into the experiences of privacy violation claimed by users. More importantly, we show that this theory is ill equipped to guide the design of SNSs and platforms to avoid similar harms in the future. A rising tide of privacy blunders on social networking sites and platforms drives the search for new regulatory approaches, and privacy regulators across the globe are increasingly demanding that the Fair Information Practice Principles, the embodiment of privacy as individual control, inform the design of technical systems through Privacy By Design. The call for Privacy By Design — the practice of embedding privacy protections into products and services at the design phase, rather than after the fact — connects to growing policymaker recognition of the power of technology to not only implement, but also to settle policy through architecture, configuration, interfaces, and default settings. We argue that regulators would do well to ensure that the concept of privacy they direct companies to embed affords the desirable forms of protection for privacy.

Ideally, there would be a widely used set of methods and tools to aid in translating privacy into design. Today, neither is true. We identify three gaps in the "informational self-determination" approach that limit its responsiveness to lateral privacy design decisions in SNSs and platforms and then explore three alternative theories of privacy that provide compelling explanations of the privacy harms exemplified in platform environments. Based on this descriptive utility, we argue that these theories provide more robust grounding for efforts by SNSs and platform developers to address lateral privacy concerns in the design of technical artifacts. Unlike FIPPs, which can be applied across contexts, these theories require privacy to be discovered, not just implemented. To bridge this discovery gap, we turn to the field of Human Computer Interaction ("HCI") and dip into the related field of Value Sensitive Design ("VSD") to identify tools and methodologies that would aid designers in discovering and ultimately embedding these contextual, socially-oriented understandings of privacy in technical artifacts. Finally, we provide some tentative thoughts on the form and substance of regulations that would prompt corporations to invest in these HCI approaches to privacy.

Authors

Deirdre K. Mulligan is an Assistant Professor at the UC Berkeley School of Information (I School). She came to the I School from the UC Berkeley School of Law (Boalt Hall), where she was a clinical professor of law and the director of the Samuelson Law, Technology & Public Policy Clinic. She served previously as staff counsel at the Center for Democracy & Technology in Washington.

Professor Mulligan's current research agenda focuses on information privacy and security. Current projects include qualitative interviews to understand the institutionalization and management of privacy within corporate America, and role of law in corporate information security policy and practice. Other areas of current research include digital rights management technology and privacy and security issues in sensor networks and visual surveillance systems, and alternative legal strategies to advance network security.

Jennifer King is a Ph.D candidate in Information Science at UC Berkeley's School of Information, where she is advised by Professor Deirdre Mulligan. Ms. King's work uses human-computer interaction methods to examine the "privacy gap" between people's expectations and how technological systems actually function. Her publications include privacy focused investigations into mobile systems, online social networks, radio-frequency identification [RFID], and digital video surveillance. Ms. King holds a professional master's degree in information management and systems also from Berkeley's i-School. Prior to her research career, Ms. King worked in security and product management for several Internet companies, most recently Yahoo!.

'Going Dark' Versus a 'Golden Age of Surveillance'

Peter Swire and Kenesa Ahmad

Full paper available at: http://www.futureofprivacy.org/privacy-papers-2012/

Executive Summary

Law enforcement and national security agencies are worried that they are "going dark" due to new technology. In 2011 testimony about the "going dark" problem, FBI General Counsel Valerie Caproni stated: "As the gap between authority and capability widens, the government is increasingly unable to collect valuable evidence in cases ranging from child exploitation and pornography to organized crime and drug trafficiking to terrorism and espionage – evidence that a court has authorized the government to collect."

"Going dark" is an evocative and compelling image. The phrase invites us to imagine

communications shrouded in darkness, so that the eyes of the agency are blind. Although it did not yet have the name, the "going dark" argument was used extensively in the 1990's, when the FBI and the NSA fought to prevent widespread use of encryption on the Internet. Today, the argument is used to suggest that agencies need a new generation of enforcement powers, such as expansion of CALEA's current requirements that telephones be "wiretap ready." Notably, agencies are seeking to expand the wiretap requirements of CALEA broadly into hardware and software connected to the Internet.

This article explains, however, that "going dark" is the wrong image. Instead, today should be understood as a "golden age of surveillance." Consider three areas where law enforcement has far greater capabilities than ever before:

1. Location information. We are in the first period in history where most people carry a tracking device, the mobile phone.

2. Information about contacts and confederates. Police have ready access to the to/from information for an exploding number of phone calls, texts, e-mails, chat, and social network interactions. Police thus have unprecedented visibility into a suspect's social graph – the list of potential criminal confederates.

3. All the other digital databases. More generally, the amount of information about individuals stored online, and accessible to the police, has risen exponentially – medical records, each purchase by credit or debit card, marketing records, and many more.

A simple test can help the reader decide between the "going dark" and "golden age of surveillance" hypotheses. Suppose the agencies had a choice of a 1990-era package or a 2011-era package. The first package would include the wiretap authorities as they existed pre-encryption, but would lack the new techniques for location tracking, confederate identification, access to multiple databases, and data mining. The second package would match current capabilities: some encryption-related obstacles, but increased use of wiretaps, as well as the capabilities for location tracking, confederate tracking and data mining. The second package is clearly superior - the new surveillance tools assist a vast range of investigations, whereas wiretaps applied only to a small subset of key investigations.

The battle between the images of darkness and light is important. If the overall truth were that agencies are "going dark," then legislatures and agencies would have an important argument for expanding surveillance powers. On the other hand, careful review of the facts shows that we live in a "golden age of surveillance." This article is part of a broader research project on why effective encryption should be encouraged to create security on the Internet. CALEA in the U.S. and other anti-encryption laws in China and India should not create holes in that security. Skepticism should accompany agency requests for new powers.

Authors

Peter P. Swire is the C. William O'Neill Professor of Law at the Moritz College of Law of the Ohio State University. He is a Senior Fellow with the Future of Privacy Forum, and also a fellow with the Center for American Progress and Center for Democracy and Technology. He has been a recognized leader in privacy, cybersecurity, and the law of cyberspace for well over a decade, as a scholar, government official, and participant in numerous policy, public interest, and business settings.From 2009 until August, 2010 Professor Swire was Special Assistant to the President for Economic Policy, serving in the National Economic Council under Lawrence Summers. From 1999 to early 2001 Professor Swire served as the Clinton Administration's Chief Counselor for Privacy, in the U.S. Office of Management and Budget, as the only person to date to have government-wide responsibility for privacy issues. Among his other activities when at OMB, Swire was the White House coordinator for the HIPAA Medical Privacy Rule, and chaired a White House working group on how to update wiretap laws for the Internet Age. In 2012, Professor Swire is lead author for two new books that are the official guides for Certified Information Privacy Professional examinations. Many of his writings appear at www.peterswire.net.

Kenesa Ahmad is an information privacy and cyber security attorney. She received her law degree from the Moritz College of Law of The Ohio State University, where she served as an Articles Editor of the Ohio State Law Journal. She also received her LL.M. from Northwestern University Law School. From 2011–2012 Ahmad completed a legal and policy fellowship with the Future of Privacy Forum. She is now an Associate in the global privacy practice of Promontory Financial Group.

How Come I'm Allowing Strangers to go Through My Phone? Smart Phones and Privacy Expectations

Jennifer King

Full paper available at: http://www.futureofprivacy.org/privacy-papers-2012/

Executive Summary

This study examines the privacy expectations of smartphone users by exploring two specific dimensions to smartphone privacy: participants' concerns with other people accessing the personal data stored on their smartphones, and applications accessing this data via platform APIs. We interviewed 24 Apple iPhone and Google Android users about their smartphone usage, using Altman's theory of boundary regulation and Nissenbaum's theory of contextual integrity to guide our inquiry. We found these theories provided a strong rationale for explaining participants' privacy expectations, but there were discrepancies between their expectations, smartphone usage, and existing platform designs and data access practices by application developers. We conclude by exploring this "privacy gap" and recommending design improvements to both the platforms and applications to address it.

Author

Jennifer King is a Ph.D candidate in Information Science at UC Berkeley's School of Information, where she is advised by Professor Deirdre Mulligan. Ms. King's work uses human-computer interaction methods to examine the "privacy gap" between people's expectations and how technological systems actually function. Her publications include privacy focused investigations into mobile systems, online social networks, radio-frequency identification [RFID], and digital video surveillance. Ms. King holds a professional master's degree in information management and systems also from Berkeley's i-School. Prior to her research career, Ms. King worked in security and product management for several Internet companies, most recently Yahoo!.

Mobile Payments: Consumer Benefits & New Privacy Concerns

Chris Jay Hoofnagle, Jennifer M. Urban, and Su Li

Full paper available at: http://www.futureofprivacy.org/privacy-papers-2012/

Executive Summary

Payment systems that allow people to pay using their mobile phones are promised to reduce transaction fees, increase convenience, and enhance payment security. New mobile payment systems also are likely to make it easier for businesses to identify consumers, to collect more information about consumers, and to share more information about consumers' purchases among more businesses. This is a radical change from the current payment system, which by design and by legal arrangement, limits the ability of participants to fully track consumer purchases. The shift to mobile payments has large implications for consumer tracking and profiling, and because of nuances in existing anti-marketing laws, the shift could mean that individuals will receive much more spam and telemarketing.

While many studies have reported security concerns as a barrier to adoption of mobile payment technologies, the privacy implications of these technologies have been under examined. To better understand Americans' attitudes towards privacy in new transaction systems, we commissioned a nationwide, telephonic (wireline and wireless) survey of 1,200 households, focusing upon the ways that mobile payment systems are likely to share information about consumers' purchases.

We found that Americans overwhelmingly oppose the revelation of contact information (phone number, email address, and home address) to merchants when making purchases with mobile payment systems. Furthermore, an even higher level of opposition exists to systems that track consumers' movements through their mobile phones.

This last result speaks directly to emerging business models that attempt to track individuals uniquely through signals emitted from phones. For instance, Navizon I.T.S. claims that it can track, "any Wi-Fi enabled smart phone or tablet, including iPhones, iPads, Android devices, BlackBerry, Windows Mobile, Symbian and, of course, laptops." As with many other tracking technologies, it seems to be designed to operate without the knowledge of the individual. Navizon claims, "Unobtrusive surveillance / Navizon I.T.S. works in the background, quietly and unobtrusively locating Wi-Fi- enabled devices…No application is needed on the devices to be tracked. The only requirement is that their Wi-Fi radios be turned on, which is the default in most smart phones, tablets and laptops."

In this paper, we explain some advantages of mobile payment systems, some challenges to their adoption in the United States, and then turn to our main finding: Americans overwhelming reject mobile payment systems that track their movements or share identification information with retailers. We then suggest a possible remedy for such information sharing: adapting provisions of California's Song-Beverly Credit Card Act, which prohibits merchants from requesting personal information at the register when a consumer pays with a credit card, to mobile payments systems. Our survey results suggest that consumers would support limitations on information collection and transfer. Song-Beverly could be adopted to accommodate those who wish to share their transaction data.

Authors

Chris Jay Hoofnagle is director of the Berkeley Center for Law & Technology's information privacy programs and senior fellow to the Samuelson Law, Technology & Public Policy Clinic. He is an expert in information privacy law. He teaches computer crime law and a seminar on the Federal Trade Commission and online advertising. Hoofnagle's research focuses on the challenges in aligning consumer privacy preferences with commercial and government uses of personal information.

Jennifer M. Urban is an Assistant Clinical Professor of Law and Director of the Samuelson Law, Technology & Public Policy Clinic at the UC Berkeley School of Law.

Broadly, her research considers how values such as free expression, freedom to innovate and privacy are mediated by technology, the laws that govern technology, and private ordering systems. Her clinic students represent clients in numerous public interest cases and projects at the intersection of societal interests—including civil liberties, innovation, and creative expression—and technological change. Recent Clinic projects include work on individual privacy rights, copyright and free expression, artists' rights, free and open source licensing, the "smart" electricity grid, biometrics, and defensive patent licensing.

Professor Urban comes to Berkeley Law from the University of Southern California's Gould School of Law, where she founded and directed the USC Intellectual Property & Technology Law Clinic. Prior to joining the USC faculty in 2004, she was the Samuelson Clinic's first fellow. Prior to that, she was an attorney with the Venture Law Group in Silicon Valley. She graduated from Cornell University with a B.A. in biological science (concentration in neurobiology and behavior) and from Berkeley Law with a J.D. (intellectual property certificate). She was the Annual Review of Law and Technology editor while a student at Berkeley Law, and received the Berkeley Center for Law and Technology Distinguished Alumni Award in 2003.

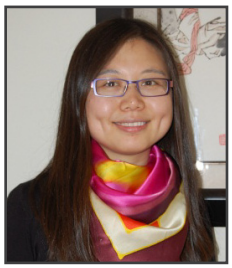

Su Li received a PhD in Sociology in 2006 and a MS in Mathematical Methods for Social Science in 2002, both from Northwestern University. She received additional trainings in quantitative methods from the Stanford Institute for the Quantitative Study of Society (SIQSS) and the Interuniversity Consortium for Political and Social Research (ICPSR) at the University of Michigan. Su Li joined Berkeley Law in January 2010, as the statistical consultant for the school of law. She works with professors, editors of the California Law Review, J.D. and Ph.D. students, as well as affiliated researchers and scholars on research papers/projects, government reports, law suit cases, and dissertations. She provides consultation services on data retrieving, modeling construction, results interpretation and other relevant issues.

Su Li's research interests are on quantitative methods, social network analysis, law and society, gender and social inequality, economic sociology and organizations. Her previous research/publications focus on gender segregation and inequalities in higher education. Su Li's current research involves the development and change of the legal profession, especially in the field of white collar criminal litigation. She also participates in projects on the implications of privacy laws in the context of virtual world and beyond.

Before joining Berkeley Law, Su Li was an assistant professor of Sociology at Wichita State University in Wichita Kansas.

Privacy by Design: A Counterfactual Analysis of Google and Facebook Privacy Incidents

Ira Rubinstein and Nathan Good

Full paper available at: http://www.futureofprivacy.org/privacy-papers-2012/

Executive Summary

Regulators here and abroad have embraced "privacy by design" as a critical element of their ongoing revision of current privacy laws. This idea of building in privacy from the outset is commonsensical, yet it remains rather vague—surely it should enhance consumer privacy but what does it mean and how does it really work? In particular, is privacy by design simply a matter of developing products with greater purpose or intention (as opposed to haphazardly) or of following specific design principles? What do regulators expect to accomplish by encouraging firms to implement privacy by design and how would this benefit consumers? This paper seeks to answer these questions by presenting case studies of high profile privacy incidents involving Google and Facebook and analyzing them against a set of well established design principles.

At a minimum, privacy by design implies building in privacy (in the form of Fair Information Practices or FIPs) when creating software products and services. But FIPs are not self-executing. Rather, privacy by design requires the translation of FIPs into engineering *and* usability principles and practices. The best way to ensure that software includes the broad goals of privacy as described in the FIPs and any related corporate privacy guidelines is by including it in the definition of software "requirements." And a main component of making a specification or requirement for software design is to make it concrete, specific and preferably associated with a metric. Equally important is developing software interfaces and other visual elements that are focused around end-user goals, needs, wants and constraints.

The Article offers the first comprehensive analysis of engineering and usability principles specifically relevant to privacy. Based on the relevant technical literature, it derives a small number of principles and illustrates them by reference to ten recent Google and Facebook privacy incidents. Relying on news accounts, company statements, and detailed regulatory reports, we analyze these privacy incidents in detail to determine whether the two firms might have achieved better privacy results if they had implemented privacy by design. Despite the somewhat speculative nature of this "what if" analysis, we believe that it reveals the strengths and weaknesses of privacy by design and thereby helps inform ongoing regulatory debates. The Article concludes that all ten privacy incidents might have been avoided by the application of these privacy engineering and usability principles. Further, we suggest that the main challenge to effective privacy by design is not the lack of design guidelines. Rather, it is that business concerns often compete with and overshadow privacy concerns. Hence the solution lies in providing firms with much clearer guidance about applicable design principles and how best to incorporate them into their software development processes. Greater guidance is also needed for how to balance privacy with business interests, and there must be oversight mechanisms as well.

This Article has three parts. In Part I, we present a general review of the design principles relevant to privacy. This requires a brief analysis of the strengths and weaknesses of FIPs as a source of privacy design principles. Here we mainly focus on the failure of the notice-and-choice model of FIPs and the shortcomings of all versions of FIPs insofar as they rely primarily on a control conception of privacy. Next, we closely examine what it means to design for privacy, defining "design" in terms of two broad and at times overlapping ideas: back-end software implementations of networking and related

systems infrastructure, which are generally hidden from the user but drive the heart of any system; and front-end, user interfaces, which (in the privacy setting) handle tasks such as notification, consent, access, preference management, and other user experiences. We therefore analyze privacy by design from two complementary perspectives: *privacy engineering*, which refers to the design and implementation of software that facilitates privacy, and *usable privacy design*, which refers to design tasks that focus on human-computer interaction (HCI). The former focuses on building software satisfying the abstract privacy requirements embodied in the FIPs (in some cases overlapping with security engineering), the latter on ensuring that users understand and benefit from well-engineered privacy controls. Our discussion of privacy engineering draws mainly on four key papers in the technical design literature and the works cited therein. In contrast, our discussion of usable privacy design looks at a rather different body of work that finds inspiration in the writings of Irwin Altman, a social psychologist, and Helen Nissenbaum, a philosopher of technology, both of whom analyze privacy in terms of social interaction. In Part II, we offer ten case studies of Google and Facebook privacy incidents and then rely on the principles identified in Part I to discover what went wrong and what the two companies might have done differently to avoid privacy violations and consumer harms. We conclude in Part III by considering what lessons regulators might learn from this counterfactual analysis.

Authors

Ira Rubinstein is a Senior Fellow at the Information Law Institute. His research interests include Internet privacy, electronic surveillance law, online identity, and Internet security. Rubinstein lectures and publishes widely on issues of privacy and security and has testified before Congress on these topics on several occasions. In September 2009, he organized a conference at the law school on Federal Privacy Legislation, and he participated in the December 2009 Federal Trade Commission Roundtable: Exploring Privacy. In July 2010, he testified at a hearing on a new privacy bill, H.R. 5777, the Best Practices Act, before the House Subcommittee on Commerce, Trade, and Consumer Protection. In 2011, he was awarded a research grant to explore regulatory issues related to "privacy by design." In March 2011, he was an invited speaker at a Boalt Hall Law School symposium on "Technology: Transforming the Regulatory Endeavor," where he discussed his paper entitled "Regulating Privacy by Design." This paper has been published in the Symposium Issue of the Berkeley Technology Law Journal (2012). He also recently commented on the White House proposal to encourage a multistakeholder process for developing consumer privacy codes of conduct. In June 2012, he co-authored a paper with Nathan Good entitled "Privacy by Design: A Counterfactual Analysis of Google and Facebook Privacy Incidents," which was chosen for the IAPP Privacy Law Scholars Award at the 5th Annual Privacy Law Scholars Conference. Other recent publications include "Privacy and Regulatory Innovation: Moving Beyond Voluntary Codes," 6 I/S: A Journal of Law and Policy for the Information Society 356 (2011), which was selected by the Future of Privacy Forum in their best "Privacy Papers for Policy Makers" competition, and "Data Mining and Internet Profiling: Emerging Regulatory and Technological Approaches," co-authored with Ron Lee and Paul Schwartz, 75 U. Chi. L. Rev. 261 (2008). Prior to joining the ILI, he spent 17 years in Microsoft's Legal and Corporate Affairs department, most recently as Associate General Counsel in charge of the Regulatory Affairs and Public Policy group. Before coming to Microsoft, he was in private practice in Seattle, specializing in

immigration law. He graduated from Yale Law School in 1985. From 1998-2001, Rubinstein served on the President's Export Council, Subcommittee on Encryption. He has also served on the Editorial Board of the IEEE Security and Privacy Magazine. In 2010, he joined the Board of Directors of the Center for Democracy and Technology.

Dr. Nathan Good is Principal and Chief Scientist of Good Research. A fundamental goal of his work is helping companies create networked systems devices and services that are simple, secure and respectful of people's privacy. He is a co-author of the 2012 web privacy census, and contributing author to books on privacy and security. Prior to

Good Research, Nathan was at PARC, Yahoo and HP research labs. At Berkeley, he worked with TRUST and the Samuelson Law & Technology Clinic and was a member of the 2007 California Secretary of State Top-to-Bottom Review of Electronic Voting Systems. Nathan has published extensively on user experience studies, privacy, and

security related topics and holds patents on software technology for multimedia systems and event analysis. His research has been reported on in the New York Times, CNN and ABC and he has testified on his research before the House, Senate and FTC. Nathan has a Phd in Information Science and a MS in Computer Science from the University of California at Berkeley and was a member of LifeLock's Fraud Advisory Board.

Dr. Nathan Good is Principal and Chief Scientist of Good Research. A fundamental goal of his work is helping companies create networked systems devices and services that are simple, secure and respectful of people's privacy. He is a co-author of the 2012 web privacy census, and contributing author to books on privacy and security. Prior to

Good Research, Nathan was at PARC, Yahoo and HP research labs. At Berkeley, he worked with TRUST and the Samuelson Law & Technology Clinic and was a member of the 2007 California Secretary of State Top-to-Bottom Review of Electronic Voting Systems. Nathan has published extensively on user experience studies, privacy, and

security related topics and holds patents on software technology for multimedia systems and event analysis. His research has been reported on in the New York Times, CNN and ABC and he has testified on his research before the House, Senate and FTC. Nathan has a Phd in Information Science and a MS in Computer Science from the University of California at Berkeley and was a member of LifeLock's Fraud Advisory Board.

The "Re-identification of Governor William Weld's Medical Information: A Critical Re-examination of Health Data Identification Risks and Privacy Protections, Then and Now

Dr. Daniel Barth-Jones

Full paper available at: http://www.futureofprivacy.org/privacy-papers-2012/

Executive Summary

The 1997 re-identification of Massachusetts Governor William Weld's medical data within an insurance data set which had been stripped of direct identifiers has had a profound impact on the development of de-identification provisions within the 2003 Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule. Weld's re-identification, purportedly achieved through the use of a voter registration list from Cambridge, MA is frequently cited as an example that computer scientists can re-identify individuals within de-identified data with "astonishing ease".

However, a careful re-examination of the population demographics in Cambridge indicates that Weld was most likely re-identifiable only because he was a public figure who experienced a highly publicized hospitalization rather than there being any certainty underlying his re-identification using the Cambridge voter data, which had missing data for a large proportion of the population. The Cambridge population was nearly 100,000 and the voter list contained only 54,000 of these residents, so the voter linkage could not provide sufficient evidence to allege any definitive re-identification. The statistics underlying this famous re-identification attack make it clear that the purported method of voter list linkage could not have definitively re-identified Weld. While the odds were somewhat better than a coin-flip, they fell quite short of the certainty that is implied by the term "re-identification".

The complete story of Weld's re-identification exposes an important systemic barrier to accurate re-identification known as ***"the myth of the perfect population register"***. Because the logic underlying re-identification depends critically on being able to demonstrate that a person within a sample data set is the only person in the larger population who has a set of combined characteristics (known as "quasi-identifiers") that could potentially re-identify them, most re-identification attempts face a strong challenge in being able to create such a complete and accurate population register. Importantly, each person missing within an imperfect population register is directly protected from re-identification attempts the register -- but these missing individuals also importantly confound attempts to re-identify others whenever such incomplete registers are used in re-identification attempts. When just a single person sharing the same quasi-identifier characteristics with a purported re-identification victim is missing from the voter register, then the probability of a correct re-identification for this target is only 50%.

This strong limitation not only underlies the entire set of famous Cambridge re-identification results but also impacts much of the existing re-identification research cited by those making claims of easy re-identification. – a fact which must be understood by public policy-makers seeking to realistically assess current privacy risks posed by de-identified data.

Fortunately, HHS responded to the concerns raised by the Weld/Cambridge voter list privacy attack and, through the HIPAA Privacy Rules, acted to help prevent re-identification attempts. Re-Identification risks today under the HIPAA Privacy Rule reveal dramatic reductions (thousands fold) of re-identification risks for de-identified health data as protected by the HIPAA Privacy Rule de-identification provisions since 2003. Available evidence further suggests that re-identification risks under current HIPAA protections are now well-controlled.

- In 2007, the National Committee on Vital and Health Statistics received testimony that 0.04 percent (4 in 10,000) of the individuals in the U.S. population within data sets de-identified using the "Safe Harbor" method could possibly be identified on the basis of their year of birth, gender and three-digit ZIP code.
- A 2010 study estimated re-identification risks under the HIPAA Safe Harbor rule on a state-by-state basis using voter registration data. The percentage of a state's population estimated to be vulnerable (i.e., not definitively re-identified, but potentially re-identifiable) ranged from 0.01 percent to 0.25 percent.
- The Office of the National Coordinator for Health Information Technology conducted a 2011 study examining an attack on HIPAA de-identified data under realistic conditions, testing whether HIPAA Safe Harbor de-identified data could be combined with external data to re-identify patients. The study was performed under practical and plausible conditions and verified the re-identifications against direct identifiers—a crucial step often missing from this sort of study. The study used 15,000 de-identified patient records and showed a match for only two of the fifteen thousand individuals (a re-identification rate of 0.013 percent). Even when maximally strong assumptions were made about the possible knowledge of a hypothetical "data intruder", the re-identification risk (under the questionable assumption that re-identification would even be attempted) was likely to be less than 0.22 percent.

Because a vast array of healthcare improvements and medical research critically depend on de-identified health information, the essential public policy challenge then is to accurately assess the current state of privacy protections for de-identified data, and properly balance both risks and benefits to maximum effect. While one can point to very few, if any, cases of persons who have been harmed by attacks with verified re-identifications, virtually every member of our society has routinely benefited from the use of de-identified health information.

Considerable costs come with incorrectly evaluating the true risks of re-identification under current HIPAA protections. It is essential to understand that de-identification comes at a cost to the scientific accuracy and quality of the healthcare decisions that will be made based on research using de-identified data. Balancing disclosure risks and statistical accuracy is crucial because some popular de-identification methods, such as "k-anonymity methods", can unnecessarily, and often undetectably, degrade the accuracy of de-identified data for multivariate statistical analyses. This problem is well understood by statisticians and computer scientists, but not well-appreciated in the public policy arena. Poorly conducted de-identification and the overuse of de-identification methods in cases where they do not produce real privacy protections can quickly lead to incorrect scientific findings and damaging policy decisions.

De-identified health data is the workhorse that supports numerous healthcare improvements and a wide variety of medical research activities. This critical role that de-identified health information plays in improving healthcare is becoming increasingly more widely recognized, but properly balancing the competing goals of protecting patient privacy while also preserving the accuracy of research requires policy makers to realistically assess both sides of this coin. De-identification policy must achieve an ethical equipoise between potential privacy harms and the very real benefits that result from the advancement of science and healthcare improvements which are accomplished with de-identified data.

The paper also provides recommendations for enhancements to existing HIPAA de-identification policy, such as:

- Prohibiting of the re-identification, or attempted re-identification, of individuals and their relatives, family or household members.
- Requiring parties who wish to link new data elements (which might increase re-identification risks) with de-identified data to confirm that the data remains de-identified.
- Specifying that HIPAA de-identification status would expire if, at any time, the data contains data elements specified within an evolving Safe Harbor list which should be periodically updated by HHS.
- Formally specifying that for statistically de-identified data, anticipated data recipients must always comply with specified time limits, data use restrictions, qualifications or conditions set forth in the statistical de-identification determination associated with the data.
- Requiring those holding and using de-identified data to implement and maintain appropriate data security and privacy policies, procedures and associated physical, technical and administrative safeguards.
- Requiring those transferring de-identified data to third parties to enter into data use agreements which would oblige those receiving the data to also hold to these conditions, thus maintaining an important "chain-of-trust" data stewardship principal accompanying de-identified data.

Conclusion

William Weld's 1997 "re-identification" had an important impact on improving healthcare privacy because it led to regulations that help to importantly protect patients from re-identification risks. But the Weld saga does not reflect the privacy risks that exist under the HIPAA Privacy rules today. We should not let today's minimal re-identification risks cause us to abandon our use of de-identified to protect privacy, save lives and continue to improve our healthcare system.

Author

Daniel C. Barth-Jones, MPH, PhD is a statistical disclosure control researcher and HIV epidemiologist serving as an Assistant Professor of Clinical Epidemiology at the Mailman School of Public Health at Columbia University and an Adjunct Assistant Professor and Epidemiologist at the Wayne State University School of Medicine. Dr. Barth-Jones' work on statistical de-identification science focuses the importance of properly balancing two public policy goals: effectively protecting individual's privacy and preserving the scientific accuracy of statistical analyses conducted with de-identified health data.

Smart, Useful, Scary, Creepy: Perceptions of Online Behavioral Advertising

Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang

Full paper available at: http://www.futureofprivacy.org/privacy-papers-2012/

Executive Summary

In recent years, Internet advertising has become increasingly tailored to individual users. In the simplest case, contextual advertising, advertising networks choose which ads to display on a webpage based on the contents of that page. In the more complex technique of online behavioral advertising (OBA), advertising networks profile a user based on his or her online activities, such as the websites he or she visits over time. Using this profile, advertising networks show ads that are more likely to be of interest to a particular user.

OBA presents both benefits and downsides to users. If their interests have been accurately profiled, users will receive more relevant advertising. However, collecting data about users' online activities can potentially violate their privacy. Previous research has found that users have substantial privacy concerns about OBA, while marketing surveys have found that consumers like OBA and that discomfort with OBA is reduced when users are informed that non-personally identifiable information is used for OBA. Whereas past work employed surveys, which can sample a large number of individuals but are not conducive to open-ended questions exploring attitudes and motivations, we conducted interviews to learn how past experiences, knowledge, and understanding factor into users' attitudes toward online behavioral advertising.

In this paper, we report results of 48 semi-structured interviews that unpack the factors fueling users' attitudes about OBA. Beyond asking participants their opinions, we investigated their knowledge of the current practice of OBA and tools to control it, their understanding of how profiles can be created, and the extent to which the circumstances of data collection and the identity of the advertising network influence their attitudes.

Attitudes About Internet Advertising and OBA

Participants were surprised that OBA currently occurs. While a number of participants believed that browsing history could theoretically be used to target advertising, few were aware that this technique is currently used. In contrast, many participants were familiar with contextual ads on first-party sites, such as Amazon and Facebook.

After learning about OBA, many participants perceived some benefits from behavioral advertising, yet the majority of participants noted that the practice negatively impacted their privacy. Participants mentioned lack of transparency and control as well as discomfort with being monitored. Taken as a whole, participants found OBA smart, useful, scary, and creepy at the same time. They often believed that personal information is collected during OBA, potentially influencing their attitudes toward the practice. Participants varied in the types of browsing situations in which they would like data to be collected for OBA purposes, basing these decisions on both privacy and utility.

Effectiveness of Notice and Choice Mechanisms

Participants' responses suggested that current approaches for providing notice about OBA are ineffective. Only a handful of participants understood the meaning of industry-created icons intended to notify consumers about OBA. Instead, they believed that icons intended to provide notice about OBA would let them express interest in the product being advertised or purchase their own ads. Participants could not accurately determine what information is collected for OBA purposes, or by whom, and they assumed the worst, leading them to oppose a practice they ex¬pected would involve the collection of personally-identifiable and financial information.

Our results also identify disconnects between participants' mental models and current approaches for giving consumers control over OBA. Participants were unaware of existing tools for controlling OBA, and they were unsure where to turn to protect their privacy. To exercise consumer choice, participants expected that they could use familiar tools, such as their web browser's settings, deleting their cookies, or antivirus software suites. However, mechanisms to exercise choice about OBA in browsers are limited and difficult to use. Deleting cookies, participants' most common response, would nullify their opt-outs. A Do Not Track header has been designed to allow users to set a preference in their browser that does not disappear when cookies are deleted. However, efforts to define fully the meaning of Do Not Track are still ongoing in the W3C Tracking Protection Working Group.

Furthermore, existing privacy tools ranging from opt-out pages to browser plug-ins expect consumers to express OBA preferences on a per-company basis. However, participants misunderstood the role of advertising networks in the OBA ecosystem, evaluating companies based solely on activities unrelated to advertising. Participants expressed complex OBA preferences that depended on the context of their browsing, an approach that is unsupported by current mechanisms. Future investigation is needed to test notice and choice mechanisms that better align with users' understanding of OBA, particularly by taking users' mental models of the process into consideration.

Conclusions

Participants found behavioral advertising both useful and privacy-invasive. The majority of participants were either fully or partially opposed to OBA, finding the idea smart but creepy. However, this attitude seemed to be influenced in part by beliefs that more data is collected than actually is. Participants understood neither the roles of different companies involved in OBA, nor the technologies used to profile users, contributing to their misunderstandings.

Given effective notice about the practice of tailoring ads based on users' browsing activities, participants would not need to understand the underlying technologies and business models. However, our research suggests that current notice and choice mechanisms are ineffective. Furthermore, current mechanisms focus on opting out of targeting by particular companies, yet participants displayed faulty reasoning in evaluating companies. In contrast, participants displayed complex preferences about the situations in which their browsing data could be collected; yet they currently cannot exercise these preferences. Our results suggest that rather than tools for opting out of tracking by individual companies, there is a need for easy-to-use tools that allow consumers to opt-out of certain types of tracking or data practices they find objectionable, or to opt-out of tracking on certain types of websites or in certain contexts (e.g., healthcare). In addition, our results suggest a need for more effective communication with users about when and how OBA occurs.

Authors

Blase Ur is a second-year Ph.D. student in the School of Computer Science at Carnegie Mellon University. His research focuses on usable security and privacy, including passwords, online behavioral advertising, and privacy decision making. He received his undergraduate degree in computer science from Harvard University.

Pedro Giovanni Leon is a Ph.D. student in Engineering and Public Policy at Carnegie Mellon University. His research focuses on investigating strategies that protect non-expert users' privacy in today's complex Internet ecosystem. In particular, he is interested in assisting the design and implementation of both regulations and technologies that improve current transparency and control mechanisms in the context of online tracking and behavioral advertising. He received a master's degree in Information Security Technology and Management from Carnegie Mellon University and a bachelor's degree in Telecommunications Engineering from the School of Engineering at the National Autonomous University of Mexico. Before coming to CMU, he worked for the Central Bank of Mexico.

Lorrie Faith Cranor is an Associate Professor of Computer Science and of Engineering and Public Policy at Carnegie Mellon University where she is director of the CyLab Usable Privacy and Security Laboratory (CUPS). She is also a co-founder of Wombat Security Technologies, Inc. She has authored over 100 research papers on online privacy, usable security, phishing, and other topics. She has played a key role in building the usable privacy and security research community, having co-edited the seminal book Security and Usability (O'Reilly 2005) and founded the Symposium On Usable Privacy and Security (SOUPS). She also chaired the Platform for Privacy Preferences Project (P3P) Specification Working Group at the W3C and authored the book Web Privacy with P3P (O'Reilly 2002). She has served on a number of boards, including the Electronic Frontier Foundation Board of Directors, and on the editorial boards of several journals. She was previously a researcher at AT&T-Labs Research and taught in the Stern School of Business at New York University.

Richard Shay is a fourth-year Ph.D. student in the School of Computer Science at Carnegie Mellon University. His research focuses on usable privacy and security, studying online behavioral advertising and password policy. He received an undergraduate degree in computer science and classics from Brown University, and a master's degree in computer science from Purdue University.

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. 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. Wang previously worked at Intel Research, Fuji Xerox Palo Alto Laboratory, and CommerceNet.

Will Johnny Facebook Get a Job? An Experiment in Hiring Discrimination via Online Social Networks

Alessandro Acquisti and Christina Fong

Full paper available at: http://www.futureofprivacy.org/privacy-papers-2012/

Executive Summary

Nowadays, many job seekers publicly disclose online, personal information that is risky for employers to ask in face-to-face interviews or use in the official hiring process. In most of the United States, for instance, an employer who asks a job applicant questions about her religious affiliation, sexual preference, or family status may be sued for discrimination under the Equal Employment Opportunity laws. Thus, even in extensive interviews, much of that information remains frequently private.

Employers' costs of acquiring the same data online, however, are much lower: the information is often a few clicks away, and the risks of detection are substantially lower. With the rise of social networking sites, micro-blogging, and other Web 2.0 services, new opportunities for labor market discrimination have clearly arisen. Anecdotal evidence and self-report surveys suggest that U.S. firms have, in fact, started using various online services to seek information about prospective hires. According to the employers, the information sought online is benign: firms admit to searching blogs or online profiles for evidence of professional or unprofessional behaviors and traits. However, so much more can be gleaned about prospective hires from their online presences. A tweet can reveal a place of worship. A blog post can imply a person's sexual preference. A photo on LinkedIn can show her race. A comment on Facebook - or even just an image chosen as the online profile's background - can indicate her family status.

To date, however, no controlled experiment has investigated the extent to which firms use online resources to find information about job applicants, and how their hiring activities are influenced by the information they find. In particular, no experiment has established whether "protected" information that employers are discouraged from asking during interviews, but which can be found on social networking sites, affects their employment decisions. We used two randomized experiments to investigate the effects of job candidates' personal information, posted on a popular social networking site, on the search activities of employers. The experiments shared a common design: we used data revealed online by actual members of popular social networking sites and job seeking sites to design resumes and online presences of prospective job candidates. We manipulated those candidates' personal information, focusing on traits that U.S. employers may not lawfully consider in the hiring process, and therefore should not inquire about during interviews, and measured individuals' and HR professionals' responses to those profiles. Our current findings suggest that information found online about prospective job candidates can, in fact, be a source of hiring discrimination.

Authors

Alessandro Acquisti is an associate professor at the Heinz College, Carnegie Mellon University, the director of the CMU PeeX (Privacy Economics Experiments) lab, and the co-director of CMU Center for Behavioral and Decision Research (CBDR). Alessandro 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.

Alessandro's research 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. His studies have been published in journals across several disciplines (including the Proceedings of the National Academy of Science, the Journal of Consumer Research, the Journal of Marketing Research, Marketing Science, Information Systems Research, Social Psychological and Personality Science, the Journal of Comparative Economics, and ACM Transactions), as well as edited books, conference proceedings, and numerous keynotes. 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. His research has been supported by awards and grants from the National Science Foundation, the Transcoop Foundation, Microsoft, and Google.

Alessandro has testified before Senate and House committees on issues related to privacy policy and consumer behavior, and participated in policy-finding activities of the Federal Trade Commission, DARPA, the European Network and Information Security Agency, and various national privacy commissioner authorities. In 2009, he was the invited co-chair of the cyber-economics track at the National Cyber Leap Year Summit, as part of the NITRD Program, under guidance from the White House's Office of Science and Technology Policy.

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 (SSNs) 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). Following his study on face recognition and online social networks, in December 2011 Alessandro was invited to participate in the Federal Trade Commission's forum on facial recognition technology.

Alessandro holds a PhD from UC Berkeley, and Master degrees from UC Berkeley, the London School of Economics, and Trinity College Dublin. While at Berkeley, he interned a Xerox PARC and Riacs, NASA Ames.

Christina Fong is a Senior Research Scientist in the Department of Social and Decision Sciences at Carnegie Mellon University. She has a BA in Economics from University of Michigan, Ann Arbor and an MA and Phd in Economics from University of Massachusetts, Amherst. Prior to earning her PhD, she acquired three years of full-time experience in the public sector working for the U.S. Bureau of Labor Statistics, the Massachusetts State Senate, and the Inter-American Development Bank. She has held visiting positions at the Economic Science Laboratory at University of Arizona and the Department of Political Science at Washington University in St. Louis and received research funding from the National Science Foundation, The John D. and Catherine T. MacArthur Foundation, The Russell Sage Foundation, and the NSF funded program Time-Sharing Experiments for Social Scientists (TESS).

Christina's research focuses on behavioral motives in non-market settings and in imperfectly competitive markets. Her most recent research concerns discrimination in a variety of such settings, including labor markets, interpersonal cooperation, and charitable giving. In perfectly competitive labor markets, discrimination purely on the basis of personal characteristics unrelated to job performance should not occur. If some employers discriminate because of distaste for a particular type of person, other employers could profit by hiring members of this less preferred group, driving their wages up to the competitive level. Thus we must ask: when we see income disparities between different types of people, does that represent rational, profit-seeking discrimination on the basis of different levels of human capital, or does it represent inefficient behavior stemming from distaste for different categories of people? Similar arguments apply to charitable giving and preferences for public policy. If an altruistic donor's or voter's goal is to help the poor in general, then racial and ethnic bias based purely on social group membership should not occur. Fong's research suggests that i) there is some racial bias in charitable giving to poor people but ii) it stems not from a simple distaste for people of different races but from beliefs that members of one's own racial group are more morally worthy (e.g., harder working, less eager to take advantage of handouts) than members of a different racial group. In addition to her recent research on discrimination, Fong has a long-term research agenda on the role of fairness in economic behavior. She has shown that even when we subject data to a high degree of statistical rigor, preferences for redistribution of income and wealth stem not only from economic self-interest but also to a large extent from desires for fairness.

Christina is a frequent reviewer for leading academic journals in economics. She speaks to philanthropic and other non-profit groups about practical implications of research on generosity. Her research has been featured in a variety of media outlets including the Financial Times Magazine, The Pittsburgh Post-Gazette, and the Chronicle of Philanthropy.

Privacy Papers of Notable Mention

To View the Following Papers Visit: http://www.futureofprivacy.org/privacy-papers-2012/

"Differential Privacy as a Response to the Reidentification Threat: The Facebook Advertiser Case Study"
By: Andrew Chin and Anne Kleinfelter

"Dutch Treat? Collaborative Dutch Privacy Regulation and the Lessons it Holds for U.S. Privacy Law"
By: Dennis Hirsch

"Internet Advertising After Sorrell V. IMS Health: A Discussion on Data Privacy & The First Amendment"
By: Agatha Cole

"Why Johnny Can't Opt Out: A Usability Evaluation of Tools to Limit Online Behavioral Advertising"
By: Pedro Giovanni Leon, Blase Ur, Rebecca Balebako, Lorrie Faith Cranor, Richard Shay, and Yang Wang


FPF
THE FUTURE OF
PRIVACY FORUM

This report was designed with the environment and cost-effectiveness in mind. It is printed on recovered fiber paper that has no ozone layer threatening emissions and generates no detectable amounts of sulfur, chlorine, nitrogen, or dioxide gases when properly incinerated.

About the Future Privacy Forum

The Future of Privacy Forum (FPF) is a Washington, DC based think tank that seeks to advance responsible data practices. The forum is led by Internet privacy experts Jules Polonetsky and Christopher Wolf and includes an advisory board comprised of leading figures from industry, academia, law and advocacy groups.

To learn more about FPF, please visit www.futureofprivacy.org