Request for Information (RFI) – National Privacy Research Strategy

Shaping our National Privacy Research Strategy: A Multi-Disciplinary Perspective

Submitted on October 16, 2014 to nprs@nitrd.gov by the following faculty at

Carnegie Mellon University

Alessandro Acquisti, Yuvraj Agarwal, Lujo Bauer, Avrim Blum, Travis Breaux, Lorrie Cranor, Anupam Datta, Steve Fienberg , Christina Fong, Farnam Jahanian, Limin Jia, Jon Peha, Tuomas Sandholm, Norman Sadeh, and Doug Sicker *(in alphabetical order)*

Background

The President's Council of Advisors on Science and Technology (PCAST) has identified "challenges to personal privacy in the digital era as a significant impairment undermining societal benefits from large-scale deployments of networking and IT systems" [NITRD14]. More specifically, in its 2013 report, PCAST noted that:

"*There continues to be no focused coordinated R&D effort in privacy. Important issues to be addressed include how to realize the benefits of collective personal information without compromising the privacy of individuals, how to achieve cybersecurity and security more broadly without unnecessary disclosure of individual information, how to design systems to avoid unintended personal disclosure, how to empower individuals to assert their identity and also make informed decisions about voluntary disclosure, and how to use the science of privacy protection to inform policy decisions." [PCAST13]*

The present document is submitted by a multi-disciplinary group of faculty at Carnegie Mellon University in response to a Request for Information on a *"National Privacy Research Strategy" issued on September 18, 2014* by the Cyber Security and Information Assurance Research and Development Senior Steering Group on behalf of the Agencies of the Federal Networking and Information Technology Research and Development (NITRD) Program [FR14].

The authors are recognized scholars engaged in privacy research. Collectively, their expertise spans computer science, engineering, human and social sciences, economics, and public policy. Their views have been shaped by their ongoing research as well as their many interactions with industry and government in the US and abroad. URLs linking to bios of the authors are included at the end of this document.

This document is organized in accordance with the guidelines provided in the government's request for information. It successively reviews objectives for privacy research and relevant capabilities in different disciplines. This is followed by a discussion of the importance of taking a multidisciplinary approach to research in this area, drawing in part on research currently under way at Carnegie Mellon University. The document finally discusses examples of architectures that could help organize the research, design and implementation of privacy preserving systems.

As this document will hopefully make apparent, privacy gives rise to a complex set of interconnected challenges that span many disciplines. With innovation in many different sectors of the economy and in government relying on the development and introduction of increasingly data-centric products, services and processes, **these challenges will continue to evolve and grow**

in complexity over the foreseeable future. This in turn **calls for a sustained commitment to multi-disciplinary research funding in this area, if we, as a society, are to reap the full benefits of new and emerging technologies and ensure the safety and security of our nation and its people in a manner that is consistent with our expectations of privacy**. The authors hope that their views, as expressed in this document, will be helpful in informing the design of a national strategy for privacy research and remain available to provide further input or feedback in this area.

I. Privacy Objectives

We describe a set of critical privacy issues illustrated using scenarios. These issues span privacy implications of big data analytics, empowering users to make informed privacy decisions, encouraging and enabling organizations to responsibly use unregulated data, reconciling privacy with surveillance, and privacy challenges from the technology trends of modern app ecosystems, internet of things, and smart infrastructure.

1. Privacy Implications of Big Data Analytics:

Web advertisers use complex behavioral profiles and big data analytics to select the ads that they show to website visitors. Similarly data brokers collect detailed profiles of hundreds of millions of individuals and subject this data to detailed analytics. Increasingly large datasets of health data (e.g., genomics data) and education data are being collected to enable analytics that yield insights on how to improve health and education outcomes. These big data sets and their uses raise significant concerns about privacy and related values. In particular, there are concerns of **discrimination** (e.g., targeting credit cards with higher interest rates towards disadvantaged groups, denying employment or insurance to certain protected groups), **lack of transparency** (e.g., inadequate access to individuals to inspect and correct errors in their profiles, lack of explanations for automated classification decisions made by data analytics programs), and **lack of confidentiality** (e.g., as sensitive personal attributes get inferred from other collected attributes and used in the decision-making process). In order to address these concerns, we need new organizational processes and computational tools that can be used to provide oversight of and accountability for the use of big data analytics in decision making processes. We will also need advances in computer science and statistics to create big data analytics algorithms and systems that are non-discriminatory, transparent, and protect confidentiality.

A related question is whether higher availability of consumer data will result in higher economic efficiency or lead to further discrimination. For the past 30 years (and increasingly so in the last 10), economists have been interested in understanding the economic trade-offs associated with privacy and information sharing, Perhaps surprisingly, a review of this literature shows that the microeconomic theory of privacy has brought forward arguments both supporting the view that privacy protection may increase economic efficiency in a marketplace, and decrease it: personal information, when shared, can become a public good whose analysis can reduce inefficiencies and increase economic welfare; when abused, it can lead to transfer of economic wealth from data subjects to data holders. Similarly, empirical evidence has been offered of both the benefits and costs, for data subjects and data holders alike, of privacy protection. It is unlikely that economics can answer questions such as what is the “optimal” amount of privacy and disclosure for an individual and for society. But it can help us think about the trade-offs associated with personal information. Investigating such trade-offs is of particular importance due to the increasing sophistication of, and reliance upon, business analytics and data mining technologies. The conventional wisdom appears to be that the advent of big data will lead to economic efficiencies and increase economic growth. However, what is not clear - and what we should investigate - is

the extent to which the wealth of consumer data available to firms will lead to an overall growth of societal welfare, and the extent to which, instead, the economic effect of such data will mainly have a reallocation effect by transferring the distribution of resources and wealth from data subjects to data holders, or between data subjects. Questions to investigate under this framework may include: will more sophisticated analysis of consumer data lead to economic efficiency or in fact increased discrimination? Will higher availability of consumer data reduce information asymmetries, or will it in fact exacerbate them in that data holders will know so much about data subjects, but not vice versa? Arriving at answers to these questions will require research that cuts across economics and computer science.

2. Empowering Users to Make Informed Privacy Decisions:

In today's data-centric economy, every new product, process, or service seems to rely on the collection of sensitive data. Traditionally, we have relied on the principle of "**Notice and Choice**" to inform data subjects about the collection, use and disclosure of their data. Yet, multiple studies have shown that this approach, **at least in its current form, has little practical value**. On the Web for instance it has been shown that users would have to spend an inordinate amount of time if they were to read the privacy policies of every website they visit (e.g. [McDonald et al. 2008]). To make matters worse, this same research has shown that even after reading these policies, many people struggle to answer basic questions about what they have read [McDonald et al. 2008, Reidenberg et al. 2015]. With smartphones, the emergence of the Internet of Things, and increasing use of data mining, the problem is further exacerbated. In general, while research has shown that people care about their privacy, they also feel **helpless in light of the variety and complexity of processes that collect, use and disclose their information**. Research has shown for instance that most smartphone users have little understanding about the data collected by their devices and the value chains across which collected information flows (e.g. [Lin et al. 2012, Lin et al 2014]). When systems expose settings for users to control these information flows, **the number of such settings often grows beyond anything that a regular user can be expected to manage** and many settings often prove difficult for users to understand (e.g. [Benisch et al. 2011, Kelley et al. 2012, Lin et al. 2014]). A major research challenge in this regard is to develop better models of people's privacy preferences, more practical interface technologies that empower users to effectively manage these settings and maintain sufficient awareness about those data collection, use and disclosure practices they most likely care about (e.g. [Sadeh et al. 2009, Benisch et al. 2011, Sadeh et al. 2013, Lin et al. 2014]. As is discussed later in this document, addressing this fundamental challenge cuts across many disciplines, from modeling people's mental models and their preferences, to the exploration of intelligent interfaces capable of **selectively interacting with users, learning their preferences, informing them about practices that might possibly surprise them**, and occasionally **nudging them** or motivating them in some manner to re-examine their preferences and settings (e.g. [Balebako et al. 211, Lin et al. 2014, Liu et al. 2014a]). It also requires understanding how far the development of such technologies can go in effectively empowering data subjects to regain control over their information and to what extent the limitations of such technologies may not warrant modifications to our existing legal and regulatory landscape.

3. Responsible Use of Unregulated Data:

Increasingly, companies find themselves handling a mix of regulated and unregulated data. This challenge is so daunting that some companies assume their data is unregulated. For example, Google Health initially assumed their online personal health record was not covered by HIPAA, because the record was initiated by individuals as Google users as opposed to being initiated by doctors in a patient-care scenario. Similarly, Facebook assumes that their users are individuals

over the age of 13 years old to avoid being regulated by the Children's Online Privacy Protection Act (COPPA). Thus, if a minor under 13 years of age registers with Facebook or a parent registers their newborn, then these users are violating the company's terms of use. U.S. privacy laws are driven by social norms for handling sensitive data types, such as health and financial data. These laws typically assume that data collectors play professional roles (e.g., hospitals or credit bureaus). Because laws trail behind evolving technology and laws are generally updated only after technology or business practices become particularly egregious, **companies need ways to track privacy preferences when they enter into unregulated spaces** [Hoofnagle and Honig 2005]. Currently, harms that arise from inferring health conditions from unregulated data are similar to the kinds of harms that the Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule as intended to minimize. Individuals seeking help for mental illness or drug and alcohol problems may **avoid such help if it would leave evidence in unregulated datasets that could be shared and repurposed without restrictions**. A responsible use framework should empower companies to innovate while tracking which laws cover their practices as they evolve and face the realities of diverse privacy preferences and risk thresholds among those about whom they hold sensitive data.

4. Privacy Concerns in Modern App Ecosystems:

A defining characteristic of modern personal computing is the trend towards extensible platforms (e.g., smartphones, web browsers, activity bracelets) that can run a large number of specialized applications, many of uncertain quality or provenance. This business model enables platforms to recruit potentially large numbers of third party app developers, who in turn can deliver added functionality to consumers.. The cost of this third party utility, however, is that these apps can potentially access vast amounts of sensitive data, as well as various sensors that can be used to further infringe on users' privacy. Sensitive data includes contact information (e.g., phone number, home address), emails, and unique device identifiers; sensors include microphones, accelerometers, GPS, and cameras, all of which can be leveraged by apps to infringe on users' privacy.

To guard against privacy violations, modern app ecosystems typically attempt to isolate applications from each other and from the underlying platform, and use some variant of a permission system to specify when an application is allowed to access other applications' data and resources. Unfortunately, this approach fails on two fronts: (1) as shown by multiple studies, users are unable to correctly set or understand the policy even for individual applications (e.g. [Kelley et al. 2012, Lin et al. 2012]); (2) commonly afforded protections are much too coarse-grained to meaningfully specify or enforce how or for what purpose users' sensitive information may be used, as evidenced by many discoveries of privilege-escalation or application-collusion attacks, or apps that abuse their ability to access private information by, e.g., additionally using it for advertising or communicating it to a third party (e.g. Lin et al. 2014]). Given the distributed nature of App developers and the sheer number of Apps (over 1Million in both Android and iOS) and frequent updates, protecting user privacy in these ecosystems is a major challenge.

Improving the status quo will require **innovation both in infrastructure—designing mechanisms that are capable of specifying, tracking, and preventing undesired uses of private or sensitive information more precisely than we can today** (see [Jia et al.2013]) —and in the interface between consumers and the platforms—**creating tools and policy abstractions that allow users to understand and control how apps use private or sensitive data**.

5. Internet of Things and Smart Infrastructures:

The recent growth of interest in the internet of things and smart infrastructure – smart buildings, smart homes, smart cities – will bring about several new and unprecedented privacy challenges. On one hand, a smarter infrastructure enables scenarios that improve living conditions, enhances productivity and quality of life. A smart home that can detect how many occupants (or even their identity) are in a home, learn their schedules and requirements, and then combine that information with electrical grid properties like prices and smart meters to provide automatic control of appliances is very promising. However, the same information can be used to track when individuals are home, who they are, even what TV programs they watch and what websites they visit for advertising, or going a step further to determine when to break into their home. Understanding these privacy and functionality tradeoffs can be challenging for end-users.. Extrapolating this scenario to a world where every device is on the network – coffee makers, fridges, garage door openers, lights, doors – is even more daunting from the perspective of user privacy. Furthermore, these privacy leaks do not even need to be direct given **how much can be inferred from seemingly innocuous sensors**. As an example, researchers have been able to understand household demographics, number of occupants and behavioral parameters just from a single smart meter sending coarse grained energy data. This domain will require careful study with previously described challenges related to big data analytics, empowering users to make privacy decisions, and responsible use of unregulated data all coming together under one umbrella.

6. Reconciling Privacy and Surveillance:

Privacy protection can be particularly challenging and controversial when we consider its relation to government-sanctioned surveillance, because there is potential tension between two societal objectives. Government surveillance of our computers and communications systems serves important purposes. Law enforcement agencies use surveillance of information systems to locate suspected criminals, and gather evidence of crimes. Intelligence agencies use surveillance to track agents of hostile foreign powers and terrorist organizations, and to collect information that may be used to prevent or mitigate attacks. Even public health agencies may use surveillance. For example, location information from cellular phones has been used to observe human travel patterns as a means of seeing where malaria is most likely to spread, and therefore where anti-malaria programs would do the most good [Wesolowski et al. 2012, Wesolowski 2014]. Nevertheless, there are reasons to limit these forms of surveillance. This nation was founded on the principle that government's ability to conduct surveillance on its people should be limited. Moreover, some of the features that enable or facilitate useful government surveillance also leave vulnerabilities that can be exploited by dangerous non-government actors. As discussed further in [Peha 2013, Peha et al. 2014], establishing international standards with weaker security will facilitate surveillance by both government agencies and criminal organizations, as will placing back doors in information systems. For example, features built into the cellular phone system that were intended for authorized surveillance by government law enforcement agencies ended up being used in Greece by criminals to conduct their own surveillance on top government officials [Prevelakis and Spinellis 2012].

To understand the many complex trade-offs, interdisciplinary research is needed that includes technology, public policy, economics, and law. In some cases, purely technical innovations may advance both surveillance and privacy protection. For example, **technology may allow specific law enforcement agents that have authority to collect information on specific targets while making it difficult for those agents to collect information on other targets and difficult for non-authorized actors to collect information on anyone**. In other cases, it is the interplay between technology and policy or law. Even if it were clear how to balance our surveillance

objectives with our privacy objectives in the past, new questions emerge every year as technology changes. This was apparent more than a decade ago when surveillance policies designed around wired telephones were confronted with the mobility of cellular systems. Today, surveillance policies based on whether information is stored in a legally-protected place such as a home are confronted with systems that store information "in the cloud" for exclusive use by the data owner, policies based on the distinction between domestic surveillance and foreign surveillance are confronted with telecommunications networks information that do not take national borders into account, and policies based on the distinction between collection and use are confronted with big data systems that make the labels difficult to apply. **Research is needed to understand how previously clear legal distinctions can and should be applied with new technology. Good research should also consider the economic incentives of the actors**. Markets can give operators of information systems strong incentive to protect the privacy of their users, but only when users can discern privacy threats. These issues become even more complex when privacy protection is mixed with government surveillance, since government surveillance by its nature requires that information be withheld, at least from those parties that are the subjects of surveillance.

II. Assessment Capabilities

We discuss concepts, methods, and constructs needed to support and assess privacy; discuss capabilities and models that can express privacy requirements, assess and quantify risks/benefits to privacy, evaluate effects of privacy risk mitigation, and determine the fulfillment of privacy requirements. We organize this section along disciplines that have methods to offer in developing privacy enhancing technologies. The next section illustrates the need for and value of multidisciplinary efforts.

1. Programming Language and Formal Methods:

Privacy properties may be formally viewed as imposing restrictions on personal information flows. Information flow types encompass context-specific direct flows (e.g., transfer of health information from a hospital to an insurance company) [Barth et al. 2006, DeYoung et al. 2010, 4], implicit flows (e.g., use of users' location in a web advertising system) [Sen et al. 2014], and flows of noisy statistics from databases of personal information (e.g., use of customers' ratings to recommend movies) [Dwork 2006]. The restrictions on these types of information flow include role-based restrictions (e.g., permitting certain types of flows between a patient and a doctor) [Barth et al. 2006, DeYoung et al. 2010, May et al. 2006], temporal restrictions (e.g., permitting flows if the data subject consents or has been given notice) [Barth et al. 2006, May et al. 2006], and purpose restrictions (e.g., requiring that information be used only for certain purposes like treatment) [Tschantz 2012, Breaux et al. 2014]. Considerable technical challenge remains in empirically translating privacy properties and requirements from laws and policies to formalism in light of natural language ambiguity [Breaux & Anton, 2008, Reidenberg et al. 2014], which is further compounded when organizations consider multiple legal jurisdictions [Gordon & Breaux, 2013]. If models misrepresent or misinterpret legal primitives, then the conclusions drawn from those methods will be inaccurate or incomplete. Evidence suggests that models of privacy requirements must be robust under multiple, possibly conflicting interpretations of stakeholders and experts [Gordon & Breaux, 2014]. Advances in these techniques are needed to address the several privacy issues highlighted in the previous section including responsible use of unregulated data, empowering users to make informed privacy decisions, privacy challenges in app ecosystems, smart infrastructure, and internet of things, as well as in ensuring the correctness of big data analytics programs.

Checking Audit Logs for Compliance with Privacy Policies:

Motivated by concerns raised by the first healthcare privacy scenario on regulated data, recent work produced the first complete **logical specification** of two US **privacy laws** -- the HIPAA Privacy Rule for healthcare organizations and the Gramm-Leach-Bliley Act for financial institutions [DeYoung et al. 2010]. This work sets the stage for a deeper investigation of general methods for specifying privacy policies and **computational accountability** mechanisms based on **formal methods** that aid in demonstrating that programs and people are compliant with privacy properties, detecting and explaining violations, and designing interventions (e.g., fixing program bugs or punishing human violators). Significant results include audit algorithms that check logs for compliance with temporal [Barth et al. 2006, Garg et al. 2011, 10] and purpose restrictions [Tschantz 2012] on flows, and provide explanations for detected violations.

At the same time, many difficult questions in this space remain unresolved. First, the emergence of health information exchanges is imminent. This development will greatly enhance health information flows across jurisdictional boundaries, raising significant questions about privacy compliance at scale with a large number of state laws, in addition to federal laws. Second, increasingly large volumes of personal health information are being collected by unregulated vendors of healthcare apps and wearable devices. It is important to set policy for the use and sharing of this information and develop computational tools that can be used to provide oversight of their practices as highlighted earlier in this article. Third, much of the audit activity in the healthcare industry has focused on access log analytics to detect suspicious accesses by insiders. It is enormously important to **track disclosures and ensure compliance with applicable policy**.

Checking Source Code of Software Systems for Privacy Properties:

A complementary setting is one in which the data holder would like to use a computational tool to **check that the source code** of its software systems **use personal information in ways that respect privacy expectations**. This problem reduces to program analysis for various kinds of information flow properties. Recent work develops a methodology and tool chain for checking software systems written in big data programming languages (e.g., Scope, Hive, Dremel) for compliance with a class of privacy policies [Sen et al. 2014]. The privacy policies restrict direct and implicit information flows based on role, purpose, and other considerations. The tool chain has been applied to check over a million lines of source code in Microsoft Bing's data analytics pipeline for compliance with its privacy policies.

This work identifies two central challenges in making privacy compliance tools practical. First, privacy policies are often crafted by legal teams while software that has to respect these policies is written by developers. An important challenge is thus to design privacy policy languages that are usable by legal privacy teams, yet have precise operational meaning (semantics) that software developers can use to restrict how their code operates on personal information of users. Second, software systems that perform data analytics over personal information of users are often written without a technical connection to the privacy policies that they are meant to respect. Tens of millions of lines of such code are already in place in companies like Facebook, Google, and Microsoft. An important challenge is thus to bootstrap existing software with policy annotations to enable privacy compliance.

In addition, it is important to design programming languages with appropriate information flow type systems to enable the development of software systems that are compliant by construction.

This problem becomes significantly harder because a natural class of privacy properties is statistical in nature (see [Reed and Pierce 2010] for a representative result).

Checking Black Box Software Systems for Statistical Privacy Properties:

Web advertisers, such as Google's DoubleClick, the Microsoft Media Network, or the Yahoo Ad Exchange, use complex behavioral profiles to select the ads that they show to website visitors. These profiles and their uses raise privacy concerns about the data used. An important set of scientific questions arise in studying these systems like black-boxes, i.e., without access to the source code and data models internally used by the systems. This is a useful setting since web users, privacy advocacy groups, or government regulators are interested in understanding how these systems use personal information to serve advertisements, so that they can hold the companies accountable for inappropriate uses.

For example, a user may be interested in understanding whether gender has a significant effect on job-related online advertisements (a finding of discrimination) or whether browsing substance-abuse related web sites results in related advertisements (a finding suggestive of a healthcare privacy concern). These questions are all instances of **detecting statistical personal information flow in a black-box setting**. Recent work on information flow experiments [Tschantz et al. 2014] provides a precise definition of this problem, proves that it is equivalent to a problem of causal inference, and presents a methodology based on experimental science and statistical analysis for detecting such flows. But much work remains to be done to produce foundational science and sound, scalable tools that can be used to provide oversight of the big data analytics systems that increasingly underlie numerous decision making systems.

2. Usable Privacy

Usability issues limit the effectiveness of many tools that have been introduced to help individuals protect their digital privacy [Cranor et al. 2006]. Users often lack awareness about how and by whom their data is being collected, as well as the privacy consequences of their actions [Ur et al. 2012]. Without **understanding the privacy risks**, uninformed users do not know when they should take steps to protect their privacy. Users who have some understanding of the risks often feel helpless to address them. Even when users are aware of privacy tools, they find them difficult to install and configure, and cumbersome to use. Unlike physical-world privacy tools, such as window shades and doors, online privacy tools tend to be unintuitive. They often lack feedback to indicate whether or not privacy is being protected. Users may mistakenly believe their privacy is protected when it is not [Leon et al. 2012, Whitten and Tygar 1999, Garfinkel and Miller 2005]. In addition, when tools are difficult or time-consuming to use, users will often decide not to use them [Gaw and Felten 2006].

There is increasing realization that it is not enough to develop tools that provide strong technical privacy protections if users find it difficult to use these tools correctly. Despite a growing stream of research into usable privacy and security over the past decade [Cranor and Garfinkel 2005, Garfinkel and Lipford 2014], we have seen only modest improvements in the usability of commercially-available privacy tools. At a day-long workshop at the 2014 Symposium On Usable Privacy and Security, participants reviewed end-to-end email and instant messaging encryption tools, and discussed ways to increase their usability as well as metrics for evaluating usable privacy and security. Most participants agreed that adoption of these tools was hindered by usability problems, and in order to make progress, experts in usability, human-computer

interaction, and design should be involved early on in the software development process [Bonneau2014].

To improve the usability of privacy tools requires gaining better **understandings of the mental models, workflows, and privacy needs of the users** of these tools (e.g. [Lin et al 2012, Liu et al. 2014, Lin et al 2014]. This also includes studying the **complexity and diversity of people's privacy preferences and concerns** (e.g. [Benisch et al. 2011], what people are capable of, **how different privacy interfaces and technologies impact their decisions and behaviors** (e.g. [Wilson et al. 2013, Kelley et al. 2013]). Research studies that focus on how users understand and conceptualize risks, will lead to a better understanding both of where user education may be needed, as well as of how to design interfaces that effectively communicate about privacy risks to users in terms they understand. By understanding users' privacy needs and workflows, tool designers can gain insights into the best place within a user experience to situate privacy information or configuration options [Egelman et al.2009]. Systematic analysis of user interactions with privacy tools using approaches such as the Human-in-the-loop Model [Cranor2008] can reveal gaps that result in failures of both existing and proposed tools. These failures may be mitigated through a combination of user interface changes as well as by automating some of the tasks critical to effective tool use.

Tool configuration is one example of a usability challenge associated with many privacy tools. Recent research has focused on a number of areas that show promise for reducing user burden. For example, user-oriented machine learning techniques have been successfully used to derive privacy profiles and help users refine their privacy settings, as further discussed below under "Artificial Intelligence and Machine Learning." In addition, well-timed privacy "nudges" can remind users to consider whether the current tool configuration matches their privacy preferences for their current task and motivate them to adjust their settings, if necessary[Wang et al. 2014].

3. Artificial Intelligence and Machine Learning:

Artificial Intelligence and Machine Learning are fueling economic growth and have led to advances in areas as diverse as medicine, robotics, Internet search and machine translation, to name just a few. Many of these advances have also given rise to new privacy challenges as they often rely on the collection, processing and sharing of vast amounts of sensitive data. Yet these very same techniques also offer the promise of more advanced functionality that could help users regain control over their information. **Machine learning techniques have been used to help derive privacy profiles that can be used to significantly reduce the number of privacy settings users have to manually configure** (e.g. see the development of privacy profiles to simplify the configuration of location sharing privacy preferences [Ravinchandran et al. 2009] or to reduce the number of mobile app permissions users have to manually configure [Liu et al. 2014a, Lin et al. 2014]). **User-oriented machine learning techniques can also be developed to help users refine their privacy settings**, leveraging user feedback to suggest modifications to these settings (e.g. [Kelley et al. 2008, Cranshaw et al. 2011, Mugan et al. 2011]).

In a somewhat different context, **machine learning, natural language processing, and crowdsourcing** are being used to develop techniques aimed at **semi-automatically understanding website privacy policies**. Such techniques offer the prospect of automatically (or semi-automatically) understanding key aspects of a natural language privacy policy and summarizing its most salient elements to users - possibly in a personalized fashion. The output of such functionality could also be used to populate formal models of privacy policies, which in turn could be used to verify compliance with relevant laws and regulations and/or identify

inconsistencies with a site's actual practices (e.g. presence of specific website trackers while the site's policy states that no such trackers are being used) [Sadeh et al. 2013].

4. Algorithms and Statistics:

New algorithmic and statistical techniques are key to enabling privacy-related challenges in big data analytics as described under "privacy objectives" earlier in the document.

Advances in big data analytics that protect confidentiality are critical. One of the big challenges is development of effective methods that given a large sensitive dataset can produce an alternative synthetic dataset that (a) in a rigorous way preserves the privacy of all members of the true data set, and yet (b) approximately matches the original dataset in terms of wide ranges of properties and statistics, including those that nobody has even thought to measure yet. Such a capability would allow researchers to do all sorts of exploratory analyses and investigations, different kinds of back-of-the-envelope calculations, etc., without needing to access the real data. Approaches such as multiple imputation [Reiter 2009], post-randomization, and data swapping [Fienberg and McIntyre 2005] have been developed with these in mind and implemented in part with government statistical databases. The problem with these methods is the lack of formal criteria for assessing protection. There has been some exciting "proof of concept" type results in this direction showing that this may indeed be possible to do for large tabular datasets [Blum et al. 2013], and perhaps even for social-network data [Blocki et al. 2012], under the rigorous notion of differential privacy [Dwork 2006]. But, none of these results are truly practical yet. The difficulty comes in the scaling of the methods for the release of large amounts of information that can be shared in a form that enables reanalysis of data.[Yang et al. 2012]. Differential privacy offers strong formal guarantees but often at the expense of the utility of the resulting data releases. Development of practical methods for this task would be a big breakthrough. Another set of advances will be required to create big data analytics algorithms and systems that are non-discriminatory and transparent, and also allow for high quality data analytics on the resulting released data. Relaxed privacy criteria may be required. Protection of network data poses special challenges [Fienberg 2013].

5. Economics:

An example of economic trade-offs arising from the increasing availability of personal data is offered by the job market. Just as the Internet has moved markets for many goods online, it has created a new arena for labor market search activities: job candidates look for and apply to job online; companies use online tools (search engines, online social networks, microblogging platforms) both to advertise openings and to search and screen job candidates. Along with opportunities for more efficient job market matching, however, recent research has highlighted how online tools have created the potential for personal information to increase labor market discrimination [Acquisti and Fong 2014]. This raises a number of research questions that combine economics, decision research, and the design of online platforms. To the extent that online markets facilitate not just matching but also discrimination, people in disadvantaged categories may face a dilemma: if they censor their online activity, they may be able to protect their privacy, but a limited online presence might, by itself, signal that there is something to hide, or that something is missing. Imagine a job candidate who has some traits that may disadvantage her in labor markets (such belonging to a disadvantaged racial, religious, or social category), and other traits that are professionally relevant and desirable to employers, such as a strong professional network. If so, how should such a candidate protect information about her professionally disadvantageous traits while still promoting her professionally advantageous traits?

If she were to reduce her online presence to hide her disadvantageous traits, at least three problems arise. The first is a "lemons" effect: people may assume that those with a restricted online presence belong to less preferred categories. The second is what we may term an "online network exclusion effect": by protecting her privacy, the candidate is deprived of the opportunity to present professionally relevant positive traits about herself. The third effect is a bundled traits effect: online mechanisms sometimes bundle together information about undesirable and desirable traits. For the job candidate described above, information about personal traits that are potentially less desirable to employers may be bundled with information about traits that are more desirable to employers, including the strength of her professionally relevant social networks. This presents additional challenges for disadvantaged candidates who may wish to protect information about traits that are disadvantageous in the labor market while promoting their positively viewed traits. This dilemma may motivate design of mechanisms that facilitate communication only about job-relevant characteristics. Thus, a few multi-faceted questions for future research arise. These questions could be asked descriptively: Does the amount of online self-disclosure affect labor market success? (For instance, in an age of self-disclosure, is the absence of social media presence a neutral, positive, or even negative signal?) Will market mechanisms develop to solve some of these problems? It may also be interesting to ask these questions normatively: How should online self-disclosure affect labor market success? How should markets be structured to address these problems?

6. Mechanism Design

Privacy concerns arise in settings such as voting and auctions, and when *designing interaction mechanisms* for such settings, one needs to **supplement the traditional tools of mechanism design from game theory with notions and tools from privacy** – an issue traditionally not considered in mechanism design.

Over the last few years, researchers have started combining these two sets of questions and techniques. They proposed models with and without a mediator that guarantee correctness and preserve the privacy of preferences based on diverse assumptions such as the existence of secure communication channels or the hardness of certain computations (Brandt and Sandholm 2004). Other work investigated the possibility of obtaining information-theoretic bid privacy in sealed-bid auctions without a trusted auctioneer (Brandt and Sandholm 2008) and proved this is possible in first-price sealed-bid auctions but impossible in second-price sealed-bid (Vickrey) auctions. Other results include impossibility of unconditional full privacy in voting, representative distributed protocols that privately compute the outcome of common voting schemes while only revealing a limited amount of information (Brandt and Sandholm 2005a), and positive results for a class of multi-unit auctions that can be executed in a distributed way without a mediator in a small constant number of rounds (Brandt and Sandholm 2005b) while ensuring that no subset of computationally bounded colluding participants is capable of uncovering private information.

Much work in the combination of mechanism design and privacy still lies ahead. Applications include not only voting and auctions, but the same questions arise in essentially all mechanism design settings.

7. Cognitive and Behavioral Analysis:

Online privacy and security decisions are also difficult, for several reasons. First, technologies and threats are constantly evolving, leaving users in a condition of **incomplete and asymmetric information** [Akerlof 1970] regarding how much of their data may be gathered and how it may

be used, or what security vulnerabilities their systems may have. Second, the trade-offs associated with privacy decisions (e.g., sharing or hiding data) or security decisions (e.g., choosing the ``right'' degree and type of protection for a computer system), are often complex and nuanced: security and privacy are rarely end-users' primary tasks, and users have limited mental resources to evaluate all possible options and consequences of their actions --- a phenomenon termed as bounded rationality [Simon 1957]. Both of these problems are exacerbated by the inherent amount of uncertainty, and sometimes ambiguity, associated with the trade-offs involved in privacy decisions. Third, decisions involving the disclosure of information or protection of information systems are particularly prone to being influenced by cognitive and behavioral biases --- systematic deviations in judgments and behaviors from the theoretical choices of a utility-maximizing decision maker [Camerer et al 2003]. In daily interactions, people make privacy decisions often based on heuristics, shortcuts, feelings, and emotions. In particular, simple heuristics may guide privacy decision making of bounded rational users [Simon 1982]. Such heuristics can prove quite successful most of the time --- but can also lead to suboptimal behaviors or regrettable mistakes ranging from online over sharing [Wang et al. 2008] to exposing one's system to cyber-attacks. In fact, interfaces of popular online services and systems sometimes exploit (whether intentionally or not) these heuristics and biases in ways that may nudge users to act in ways that are not always collinear with their actual intentions or desires. For instance, in a 6-year longitudinal study conducted with 5,076 Facebook users, it was found that changes in default visibility settings led users who exhibited privacy-seeking behaviors to significantly increase the amount of public disclosures [Stutzman et al 2013].

In recent years, various streams of research have attempted to understand, and possibly assist, individuals' privacy and security behaviors. Usability research has attempted to help users by making security and privacy interfaces more usable [Sasse et al. 2001, Cranor and Garfinkel 2005]. Behavioral economics and decision research have analyzed which decision making hurdles individuals face when making privacy or security decisions online [Acquisti 2004] or how emotions and cognition can influence disclosure behavior Li et al. 2008, Li et al 2011]. More recently, a growing body of work has started examining how interfaces and technology could be designed to counteract biases responsible for regrettable security and privacy decisions [Wang et al. 2014, Chiasson et al 2008]. For instance, the field of asymmetric, weak, or `libertarian' paternalism [Thaler 2008] has applied lessons from behavioral research to aid consumers' behavior by designing policies and systems that nudge individuals, influencing their choice without restricting it. Further research is required that extends this body of research: better understanding of cognitive and behavioral biases, and developing tools and approaches that take into consideration, or even counter, those biases, in order to assist and ameliorate privacy decision making.

III. Multi-Disciplinary Approach

While privacy stands to benefit from advances in a variety of disciplines (e.g., social and behavioral sciences, software engineering, formal methods, algorithms, statistics, cryptography), significant progress often requires taking a multi-disciplinary approach. This is because progress often needs to be evaluated from multiple perspectives (e.g., human, legal, social, technical, ethical perspectives). This is illustrated below using a few examples.

1. Privacy as Control Over One's Information

Information privacy is often defined as the ability of data subjects to control how information about them can be used. Under particularly restrictive legal and regulatory regimes, data collection and usage practices could be severely limited, thereby reducing the need for people to understand how data about them can be collected and used and reducing the need for them to make potentially complex decisions. Obviously such regimes may also severely limit innovation

as well as the ability of government to ensure the safety and security of its citizens. Privacy under looser legal and regulatory regimes effectively transfers control to data collectors, with the data collectors in turn responsible for determining how much information (about their practices) and control to further transfer to data subjects. Providing too little control to data subjects can lead to a loss of privacy with information being used in ways that are inconsistent with their expectations. Providing a lot of control to people, while in principle a better solution, can lead to data subjects being overwhelmed by a large set of complex decisions, including decisions that may not necessarily matter to them. Tensions such as these require multi-disciplinary approaches to research in privacy. Development of novel interface technologies and better models of people's privacy preferences can help **reconcile tensions between control and user burden and ultimately inform the development and refinement of legal and regulatory regimes** (e.g. [Sadeh et al. 2009, Benisch et al. 2011, Lin et al. 2014]. It could help ensure that these regimes have realistic expectations about what people are able to accomplish and take advantage of ways in which new technologies can help users manage potentially complex and diverse collections of privacy decisions. Such research has to draw on disciplines that include human computer interaction, user-centered design but also artificial intelligence and machine learning. In addition, it also needs to draw on behavioral economics and cognitive psychology to address cognitive and **behavioral biases that have been shown to impede people's ability to make sound privacy decisions** (e.g. [Acquisti 2004, Acquisti et al. 2005]). For instance, interface technologies that appear to empower users to better control their privacy may in fact give people a false sense of confidence and result in exactly the opposite effect – a phenomenon sometimes referred to as "illusion of control". Research on developing better models of people's privacy preferences, more powerful privacy interfaces and a deeper understanding of how cognitive and behavioral biases affect privacy decisions has seen substantial progress over the past ten years. Yet it is also increasingly clear that the design space in this area is extremely large and continues to grow, and that we have only started to "scratch the surface". Development of models of people's privacy preferences and expectations is to be viewed as a never-ending process that has to reflect the emergence of new technologies, products and business models as well as changing perceptions of what is acceptable.

2. Privacy through Accountability

Privacy through accountability refers to the principle that entities that hold personal information about individuals are accountable for adopting measures that protect the privacy of the data subjects, in particular, ensuring responsible use of such information. This is an inherently multidisciplinary research area where philosophy, law, and the social sciences meet computer science and engineering (see [Datta et al. 2014a] for an overview). Over the past few years, a body of work has emerged on translating privacy expectations, laws, and policies into computer languages with precise semantics (e.g., see [Barth et al. 2006, DeYoung et al. 2010], [Sen et al. 2014]) and designing computational accountability mechanisms for checking audit logs [Garg et al. 2011] and software systems for compliance with these policies [Sen et al. 2014]. These results demonstrate the importance and value of interactions between philosophers, legal scholars, and computer scientists to address privacy challenges. Within computer science, the work relies on and advances the state-of-the-art in programming language and formal methods, as well as human-computer interaction.

At the same time, they leave open a number of directions for further work. In particular, while they tackle accountability in the setting of "small data" (i.e., individual flows of information), designing accountability mechanisms to provide oversight of big data analytics systems remains

wide open. A representative early result in this area is recent work on information flow experiments [Tschantz et al. 2014] and its application to study web advertising systems [Datta et al. 2014]. This research direction will require significant interaction among researchers with expertise in machine learning, security and privacy, programming languages and formal methods, and human factors, in order to produce effective accountability tools.

The following two research projects, which are currently under way at Carnegie Mellon University, further illustrate how progress in this area often requires taking the type of multi-disciplinary approaches outlined above.

NSF SaTC Frontier Project on Usable Privacy Policies: Privacy policies are verbose, difficult to understand, take too long to read, and may be the least-read items on most websites even as users express growing concerns about information collection practices. For all their faults, though, privacy policies remain the single most important source of information for users to attempt to learn how companies collect, use, and share data. Likewise, these policies form the basis for the self-regulatory notice and choice framework that is designed and promoted as a replacement for regulation. This project aims to better understand and hopefully overcome the limitations of website privacy policies [Sadeh et al. 2013, UPPP2014]. The approach taken in this research is one where crowdsourcing is being combined with natural language processing and machine learning to understand key elements of privacy policy. The objective is to accurately answer critical questions about statements (or lack thereof) in privacy policies and use this information to build succinct, personalized policy summaries (e.g. in the form of browser plug-ins) that inform users about those privacy issues they most likely care about. This project brings together faculty with expertise in privacy law, human computer interaction, artificial intelligence, machine learning, crowdsourcing, formal methods and behavioral economics. Research on modeling people's privacy preferences and concerns informs the design of questions to be answered as part of the crowdsourcing process as well as the design of succinct and intuitive privacy interfaces to summarize those answers. Research on natural language processing and machine learning helps scale the crowdsourcing process with the objective of eventually automating as much of it as possible, and as an intermediary step help identify paragraphs that are most relevant to answering different questions (e.g. [Liu et al. 2014b]). Output from the crowdsourcing process also help populate formal models, which can in turn be used in support of research to check for compliance with relevant regulations or for inconsistencies in policies. Human subject studies also help evaluate the understandability of policies, while research focusing on both legal and statistical analyses look for areas where privacy policies are ambiguous and compare privacy guarantees (or lack thereof) across different sectors (e.g. [Reidenberg et al. 2015]).

NSF SaTC Privacy Nudging Project

'Privacy nudges' are about using behavioral economics and soft paternalistic strategies to assist and ameliorate privacy (as well as security) decision making. Privacy decision making is becoming increasingly complex, as information systems have vastly expanded our ability to permanently share with others information about ourselves. The complexity is such that human judgments in this area are prone to errors, stemming from lack of information, insight, or computational ability; or from problems of self-control and limited self-insight uncovered by research in behavioral economics and decision research. This project has been investigating and designing systems that anticipate, and sometimes even exploit, those cognitive and behavioral biases that hamper users' privacy decision making. The research team includes computer scientists, behavioral decision researchers, and economists, working together to study systems that may "nudge" users towards certain behaviors that the users themselves have claimed to prefer, or which sound empirical evidence has demonstrated to be preferable, from a privacy

perspective - an approach inspired by the growing body of behavioral research on "soft" paternalism. This has included work on:

a) Behavioral analysis of privacy decision making (aimed in particular at extending the body of knowledge on behavioral and cognitive biases that hamper privacy decision making)
b) Regrets in Web 2.0 (aimed at understanding what types of behaviors and actions individuals particularly regret on Web 2.0 services, such as - for instance - Facebook)
c) Nudging interventions in Web 2.0 (aimed at developing tools for service such as Twitter and Facebook that help users understand the scope and implications of their online disclosures)
d) Nudging interventions in mobile applications aimed to motivating users to review and refine their mobile app privacy settings

IV. Privacy Architectures

Privacy architectures are one way of articulating a vision for privacy research.. At a high-level one can envision several complementary architectures:

- Architectures that **empower data subjects** to effectively manage their privacy
- Architectures to **empower data collectors** to responsibly collect, use and share data they collect
- Architectures that help data collectors and processors across the value chain articulate and exchange privacy practice commitments and enable data collectors to make **end-to-end privacy commitments** to data subjects, namely commitments that are not limited to their own local practices but extend across entire information flows.
- Architectures that **reconcile surveillance and privacy**
- Architectures to **reduce information disclosure**

The following briefly outlines some such architectures.

1. Architectures to empower data subjects to effectively manage their privacy decisions – Towards "Personalized Privacy Assistants"

A major challenge faced by data subjects today has to do with the myriad of actors that collect information about them – from interactions they have through their web browsers, to information collected at points of sale, to applications they download on their smartphones, information they post on social networks, information collected by their cars. While many (though not all) of these contexts are subject to privacy statements and some also offer users ways of configuring some elements of processes involved in collecting and using their information (e.g. browser settings, opt-outs, app permissions), **information and decisions are exposed to users in an inconsistent and piecemeal fashion**. An architecture could be designed that would provide data subjects with **a single point of interaction** responsible for helping them review relevant information and make associated decisions **across all the services, products and processes with which they interact**. This type of architecture could aim at personalizing the way in which decisions are presented to users by building models of their preferences (e.g. [Lin et al 2014]) and exploiting correlations between such models (e.g. privacy profiles that extend across one's web browser, Facebook account, and mobile app privacy permissions). It could be implemented in the form of "personalized privacy assistants" endowed with dialogue, machine learning and even nudging functionality to semi-automate many decisions and very selectively determine when to interact with users (e.g. to just present them with information they care about, tell them about privacy practices they may not expect, ask them about decisions that are difficult to predict based on existing models of their preferences, or motivate them to reconsider some of their existing decisions or preferences) – e.g. [Liu et al. 2014, Mugan et al. 2011, Balebaco et al. 2011]

2. Architectures to limit the disclosure of information

Enhancing network privacy will require the use of multiple complementary architectural approaches, including 1) providing tools for users to manage their data; 2) enhancing and adopting appropriate network privacy protocols and services; and 3) establishing privacy policies and agreements among the stakeholders. However, a precursor and complement to these approaches is to simply **minimize the amount of personal data that is ever released**. The analogy has been made between the vast amounts of digital data that our society creates and the pollution that we pour into our environment [Schneier 2008]. We produce data in massive volumes, it is often burdensome and pricey to control after it is released, and it can have detrimental societal impact. Of course, it is also useful to understand where the analogy breaks down, and in the case of data pollution several important distinctions exist – data replicates easily, it is possible to eliminate it completely, and there are individuals and organizations that actually seek to possess it. These distinctions make it both easier and harder to contain data pollution, but still the analogy gives us a framing for thinking about the technical, economic and policy challenges. In minimizing data pollution, one might seek to prevent information from ever existing on, or leaving, a personal device. For data at rest, can we feel confident that privacy data will not leak out of its containment and create additional pollution? What methods and tools exist to help users manage this data (as described in the section above)? The IETF has explored work on attribute and assertion mechanisms to minimize the release of personal information between trusted devices and/or domains [IETF]. Could this approach be extended across the Internet? For data that has left the device, can we design protocols that ensure acceptable levels of privacy (noting that we've been unsuccessful for decades in adopting meaningful end-to-end encryption protocols)? Can we design architectures that make violations to our privacy so expensive that those seeking to breach privacy protections are economically thwarted? Can we design architectures that help track and mitigate the replication of personal data? Can a data sequestration architecture be created that captures data in a way that contains it and possibly supports a simple "forget" function [CA Law]? Can data architectures be extended to meaningfully enable control back to the owner of that data or some appropriate custodian? To enhance user privacy we might first ask consider how to align the incentives of the legitimate players, while discouraging others. Mapping out such incentives could be a useful exercise as a prelude to any architectural considerations.

3. Privacy Architectures for Maintaining Compliance

In industry, today, leading companies rely on technical area experts to track privacy law and policy to ensure their organizations react appropriately to emerging privacy norms and preferences. This includes filtering news articles to discover privacy harms and monitoring proposed and enacted changes in privacy law that affect their products and services. These individuals typically have technical expertise in product design or project management, and they interface with in-house legal to ensure the project teams are developing products and services in a consistent manner with the how legal experts interpret privacy requirements. In government agencies, this responsibility typically falls under an office of civil rights. Future privacy architectures must address the many challenges of managing information policy, including: (1) how to **collect, synthesize and reconcile policy-relevant information into a stable, comprehensive taxonomy or ontology** that allows designers to recognize when their IT systems exceed, meet or fall below expectations? This includes **mapping privacy laws and corporate rules into actionable IT requirements that are re-usable across systems**. (2) How to **trace design decisions** throughout an organization so that, when a policy, law or legal interpretation changes, appropriate action can be taken to re-design the IT system to return to a state of compliance? (3) How to demonstrate to auditors and regulators that IT and business practices are

consistent with emerging privacy standards and stakeholder privacy preferences. We believe some of these questions can be addressed by a mix of natural language processing, formal methods, crowdsourcing that collect data on both privacy requirements and privacy preferences of stakeholders, including users and those data subjects affected by systems, but who never interact with the system, directly.

References

[Acquisti 2004] A. Acquisti, "Privacy in electronic commerce and the economics of immediate gratification," in Proc. of the 5th ACM Conference on Electronic Commerce, pp. 21-29, ACM, 2004

[Acquisti and Grossklags 2005] A. Acquisti and J. Grossklags, "Privacy and rationality in individual decision making", IEEE Security and Privacy, Vo. 2, No. 1, 2005.

[Acquisti 2014] Acquisti, Alessandro. "The economics of personal data and the economics of privacy." Background Paper for OECD Joint WPISP-WPIE Roundtable 1 (2010).

[Acquisti and Fong 2014] Acquisti, Alessandro, and Christina M. Fong. "An experiment in hiring discrimination via online social networks." Available at SSRN 2031979 (2013).

[Akerlof 1970] Akerlof, George A. "The market for" lemons": Quality uncertainty and the market mechanism." *The quarterly journal of economics* (1970): 488-500.

[Balebako et al. 2011] R. Balebako, P.G. Leon, J. Mugan, A. Acquisti, L.F. Cranor, N. Sadeh, "Nudging Users Towards Privacy on Mobile Devices", In Proceedings of the Second International Workshop on "Persuasion, Influence, Nudge & Coercion through Mobile Devices" (PNC2011), Co-located with the 2011 ACM Conference on "Human Factors in Computing Systems" (CHI 2011), May 2011

[Barth et al. 2006] Adam Barth, Anupam Datta, John C. Mitchell, Helen Nissenbaum: Privacy and Contextual Integrity: Framework and Applications. IEEE Symposium on Security and Privacy 2006: 184-198

[Benisch et al. 2011] M. Benisch, P.G. Kelley, N. Sadeh, and L.F. Cranor, "Capturing Location-Privacy Preferences: Quantifying Accuracy and User-Burden Tradeoffs", *Journal of Personal and Ubiquitous Computing*. Volume 15 Issue 7, pp. 679-694, October 2011

[Blocki et al. 2012] Jeremiah Blocki, Avrim Blum, Anupam Datta, Or Sheffet: The Johnson-Lindenstrauss Transform Itself Preserves Differential Privacy. FOCS 2012: 410-419

[Blum et al. 2013] Avrim Blum, Katrina Ligett, Aaron Roth: A learning theory approach to noninteractive database privacy. J. ACM 60(2): 12 (2013)

[Bonneau 2014] Joseph Bonneau. A Recap of the First EFF CUP Workshop. August 8, 2014.

[Brandt, F. and Sandholm, T. 2008] On the Existence of Unconditionally Privacy-Preserving Auction Protocols. ACM Transactions on Information and System Security, 11(2), Article 10, 21 pages. Conference version in AAMAS-04.

[Brandt, F. and Sandholm, T. 2005a] Decentralized Voting with Unconditional Privacy. In Proceedings of the International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS).

[Brandt, F. and Sandholm, T. 2005b] Efficient Privacy-Preserving Protocols for Multi-Unit Auctions. In Proceedings of the International Conference on Financial Cryptography and Data Security (FC), published in Springer Lecture Notes in Computer Science LNCS 3570.

[Brandt, F. and Sandholm, T. 2004]. On Correctness and Privacy in Distributed Mechanisms. In Proceedings of the Agent-Mediated Electronic Commerce (AMEC) workshop, Springer LNAI 3937.

[CA Law] California Senate approves new protections for minors using Internet, Patrick McGreevyhttp://articles.latimes.com/2013/aug/30/local/la-me-pc-california-senate-approves-new-protections-for-minors-using-internet-20130830, August 30, 2013.

[Camerer et al. 2003] Camerer, Colin, et al. "Regulation for Conservatives: Behavioral Economics and the Case for" Asymmetric Paternalism"." *University of Pennsylvania Law Review* (2003): 1211-1254.

[Chiasson et al. 2008] Chiasson, Sonia, et al. "Influencing users towards better passwords: persuasive cued click-points." *Proceedings of the 22nd British HCI Group Annual Conference on People and Computers: Culture, Creativity, Interaction-Volume 1*. British Computer Society, 2008.

[Cranor et al. 2006] Lorrie Faith Cranor, Praveen Guduru and Manjula Arjula. User Interfaces for Privacy Agents. *ACM Transactions on Computer Human Interaction (ToCHI),* 13(2), June 2006, 135-178.

[Cranor and Garfinkel 2005] Lorrie Cranor and Simson Garfinkel. 2005. *Security and Usability*. O'Reilly Media, Inc.

[Cranor 2008] Lorrie Faith Cranor. A Framework for Reasoning About the Human in the Loop. Usability, Psychology, and Security 2008.
[Cranshaw et al. 2011] Justin Cranshaw, Jonathan Mugan, Norman Sadeh, "User-Controllable Learning of Location Privacy Policies with Gaussian Mixture Models", Proceedings of the 25th AAAI Conference on Artificial Intelligence, AAAI-11, August 2011.
[Datta et al. 2014] A. Datta, M. C. Tschantz, A. Datta, Automated Experiments on Ad Privacy Settings: A Tale of Opacity, Choice, and Discrimination, arXiv:1408.6491, August 2014.
[Datta et al. 2014a] Anupam Datta: Privacy through Accountability: A Computer Science Perspective. ICDCIT 2014: 43-49
[DeYoung et al. 2010] Henry DeYoung, Deepak Garg, Limin Jia, Dilsun Kirli Kaynar, Anupam Datta: Experiences in the logical specification of the HIPAA and GLBA privacy laws. WPES 2010: 73-82
[Dwork 2006] Cynthia Dwork: Differential Privacy. ICALP (2) 2006: 1-12
[Egelman et al.2009] S. Egelman, J.Y. Tsai, L.F. Cranor, and A. Acquisti. Timing is Everything? Effects of Timing and Placement of Online Privacy Indicators. In Proc. of the 2009 Conference on Human Factors in Computing Systems. Boston, MA, April 4-9, 2009. (CHI'2009)
[Fienberg 2013] S.E. Fienberg, Is the Privacy of Network Data an Oxymoron?, *Journal of Privacy and Confidentiality*: 2013 Vol. 4: Iss. 2, Article 1.
[Fienberg and McIntyre 2005] S.E. Fienberg and J. McIntyre, Data Swapping: Variations on a Theme by Dalenius and Reiss. Journal of Official Statistics, 2005: 21, 309-323.
[Garfinkel and Miller 2005] Simson L. Garfinkel and Robert C. Miller. 2005. Johnny 2: a user test of key continuity management with S/MIME and Outlook Express. In *Proceedings of the 2005 symposium on Usable privacy and security* (SOUPS '05). ACM, New York, NY, USA, 13-24.
[Garg et al. 2011] Deepak Garg, Limin Jia, Anupam Datta: Policy auditing over incomplete logs: theory, implementation and applications. ACM Conference on Computer and Communications Security 2011: 151-162
[Gaw and Felten 2006] Shirley Gaw and Edward W. Felten. 2006. Password management strategies for online accounts. In *Proceedings of the second symposium on Usable privacy and security* (SOUPS '06). ACM, New York, NY, USA, 44-55.
[IETF] RFC 4484, Trait-Based Authorization Requirements for the Session Initiation Protocol (SIP), Peterson, Sicker, Polk and Tschofenig, 2006.
[Lane et al. 2014] Lane, Julia, Victoria Stodden, Stefan Bender, and Helen Nissenbaum, eds. Privacy, Big Data, and the Public Good: Frameworks for Engagement. Cambridge University Press, 2014.
[Leon et al. 2012] Pedro G. Leon, Blase Ur, Rebecca Balebako, Lorrie Faith Cranor, Richard Shay, and Yang Wang. Why Johnny Can't Opt Out: A Usability Evaluation of Tools to Limit Online Behavioral Advertising. CHI 2012, May 2012.
[FR14] Request for Information (RFI) – National Privacy Research Strategy, Federal Register, September 18, 2014. https://www.federalregister.gov/articles/2014/09/18/2014-22239/request-for-information-rfi-national-privacy-research-strategy
[Garfinkel and Lipford 2014] Simson Garfinkel and Heather Richter Lipford. 2014. Usable Security: History, Themes, and Challenges. Synthesis Lectures on Information Security, Privacy, and Trust. Morgan & Claypool.
[Hoofnagle and Honig 2005] C.J. Hoofnagle, E. Honig. "Victoria's Secret and Financial Privacy," January 25, 2005. http://www.epic.org/privacy/glba/victoriassecret.html
[Kelley et al. 2008] P.G. Kelley, P. Hankes Drielsma, N. Sadeh, and L.F. Cranor, "User-Controllable Learning of Security and Privacy Policies", First ACM Workshop on AISec (AISec'08), ACM CCS 2008 Conference. Oct. 2008
[Kelley et al. 2012] P. Gage Kelley, S. Consolvo, L. Cranor, J. Jung, N. Sadeh, D. Wetherall, "A Conundrum of Permissions: Installing Applications on an Android Smartphone", Proc. Of Workshop on Usable Security (USEC2012), collocated with the 16th International Conference on Financial Cryptography and Data Security, March 2012
[Kelley et al. 2013] P. Gage Kelley, L. Cranor, N. Sadeh, "Privacy as Part of the App Decision-Making Process" in Proceedings of the 31st annual SIGCHI Conference on Human Factors in Computing Systems, CHI2013, May 2013
[Jia et al. 2013] Limin Jia, Jassim Aljuraidan, Elli Fragkaki, Lujo Bauer, Michael Stroucken, Kazuhide Fukushima, Shinsaku Kiyomoto, Yutaka Miyake: Run-Time Enforcement of Information-Flow Properties on Android - (Extended Abstract). ESORICS 2013: 775-792

[Li et al. 2008] Li, Han, Rathindra Sarathy, and Jie Zhang. "The Role of Emotions in Shaping Consumers' Privacy Beliefs about Unfamiliar Online Vendors."*Journal of Information Privacy and Security* (2014): 36-62.
[Li et al. 2011] Li, Han, Rathindra Sarathy, and Heng Xu. "The role of affect and cognition on online consumers' decision to disclose personal information to unfamiliar online vendors." *Decision Support Systems* 51.3 (2011): 434-445.
[Lin et al. 2014] J. Lin, B. Liu, N. Sadeh, and J.I. Hong, "Modeling Users' Mobile App Privacy Preferences: Restoring Usability in a Sea of Permission Settings", 2014 ACM Symposium on Usable Security and Privacy (SOUPS 2014), July 2014.
[Lin et al. 2012] J. Lin, S. Amini, J. Hong, N. Sadeh, J. Lindqvist, J. Zhang, "Expectation and Purpose: Understanding Users' Mental Models of Mobile App Privacy through Crowdsourcing", Proc. of the 14th ACM International Conference on Ubiquitous Computing, pp. 501-510, Pittsburgh, USA, Sept. 2012
[Liu et al. 2014a] B. Liu, J. Lin, N. Sadeh, "Reconciling Mobile App Privacy and Usability on Smartphones: Could User Privacy Profiles Help?", Proceedings of the 23rd International World Wide Web Conference (WWW2014). April 2014
[Liu et al. 2014b] F. Liu, R. Ramanath, N. Sadeh, and N.A. Smith, "A Step Towards Usable Privacy Policy: Automatic Alignment of Privacy Statements", in Proc. of the 25th International Conference on Computational Linguistics, Dublin, August 2014.
[May et al. 2006] Michael J. May, Carl A. Gunter, Insup Lee: Privacy APIs: Access Control Techniques to Analyze and Verify Legal Privacy Policies. CSFW 2006: 85-97
[McDonald et al. 2008] McDOnald, A.M. and Cranor,L.F., "The cost of reading privacy policies", I/S A Journal of Law and Policy for the Information Society, 4(3), 2008.
http://moritzlaw.osu.edu/students/groups/is/files/2012/02/Cranor_Formatted_Final.pdf
[Mugan et al. 2011] J. Mugan, T. Sharma, N. Sadeh, "Understandable Learning of Privacy Preferences Through Default Personas and Suggestions", Carnegie Mellon University's School of Computer Science Technical Report CMU-ISR-11-112, August 2011.http://reports-archive.adm.cs.cmu.edu/anon/isr2011/CMU-ISR-11-112.pdf
[NITRD14] Report on Privacy Research within NITRD, National Coordination Office for NITRD, April 2014. https://www.nitrd.gov/Pubs/Report_on_Privacy_Research_within_NITRD.pdf
[PCAST13] "Designing a Digital Future: Federally Funded Research and Development in Networking and Information Technology," January 2013.
http://www.whitehouse.gov/sites/default/files/microsites/ostp/pcast-nitrd2013.pdf
[Peha 2013] J. M. Peha, "The Dangerous Policy of Weakening Security to Facilitate Surveillance," Comments to the U.S. Director of National Intelligence, *SSRN,* Oct. 4, 2013.
[Peha et al. 2013] J. M. Peha, T. Davis, E. Burger, J. Camp, D. Lubar, *Risking It All: Unlocking the Backdoor to the Nation's Cybersecurity*, IEEE-USA White Paper, 2014.
[Prevelakis and Spinellis 2012] V. Prevelakis, D. Spinellis. "The Athens Affair." *IEEE Spectrum,* vol. 44, no. 7, 2007, pp. 26-33.
[Reed and Pierce 2010] Jason Reed, Benjamin C. Pierce: Distance makes the types grow stronger: a calculus for differential privacy. ICFP 2010: 157-168
[Reidenberg et al. 2015] J. Reidenberg, T.D. Breaux, L.F. Cranor, B. French, A. Grannis, J.T. Graves, F. Liu, A.M. McDonald, T.B. Norton, R. Ramanath, N.C. Russell, N. Sadeh, F. Schaub, "Disagreeable Privacy Policies: Mismatches between Meaning and users' Understanding", Berkeley Law Technology Journal, to appear (2015). An earlier version of this article was presented at the 42nd Research Conference on Communication, Information, and Internet Policy (TPRC'14 – Arlington, VA – Sept. 2014) - http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2418297
[Reiter 2009] Reiter, J. P. (2009), Using multiple imputation to integrate and disseminate confidential microdata, *International Statistical Review*, 77, 179 - 195.
[Sadeh et al. 2009] N. Sadeh, J. Hong, L. Cranor, I. Fette, P. Kelley, M. Prabaker, and J. Rao, "Understanding and Capturing People's Privacy Policies in a Mobile Social Networking Application", *Journal of Personal and Ubiquitous Computing*, Vol. 13, No. 6, August 2009.
[Sadeh et al. 2013] N. Sadeh, A. Acquisti, T.D. Breaux, L.F. Cranor, A.M. McDonald, J. Reidenberg, N.A. Smith, F. Liu, N.C. Russell, F. Schaub, S. Wilson "The Usable Privacy Policy Project: Combining Crowdsourcing, Machine Learning and Natural Language Processing to Semi-Automatically Answer Those Privacy Questions Users Care About." Tech. report CMU-ISR-13-119, December 2013
[Sasse et al. 2001] Sasse, Martina Angela, Sacha Brostoff, and Dirk Weirich. "Transforming the 'weakest

link'—a human/computer interaction approach to usable and effective security." *BT technology journal* 19.3 (2001): 122-131.
[Schneier 2008] Schneier on Security, at https://www.schneier.com/blog/archives/2008/01/data_as_polluti.html.
[Sen et al. 2014] Shayak Sen, Saikat Guha, Anupam Datta, Sriram Rajamani, Janice Tsai, Jeannette M. Wing: Bootstrapping Privacy Compliance in Big Data Systems. IEEE Symposium on Security and Privacy 2014.
[Simon 1957] Simon, Herbert A. "Models of man; social and rational." (1957).
[Simon 1982] Simon, Herbert Alexander. Models of bounded rationality: Empirically grounded economic reason. Vol. 3. MIT press, 1982.
[Stutzman et al. 2013] Stutzman, Fred, Ralph Gross, and Alessandro Acquisti. "Silent listeners: The evolution of privacy and disclosure on facebook." *Journal of Privacy and Confidentiality* 4.2 (2013): 2.
[Thaler 2008] Thaler, Richard H., and Cass R. Sunstein. *Nudge: Improving Decisions about Health, Wealth, and Happiness*. New Haven, Conn.: Yale UP, 2008.
[Tschantz 2012] Michael Carl Tschantz: Formalizing and Enforcing Purpose Restrictions. PhD thesis, Computer Science Department, Carnegie Mellon University, Technical Report CMU-CS-12-117, May 2012.
[Tschantz et al. 2014] Michael Carl Tschantz, Amit Datta, Anupam Datta, Jeannette M. Wing: A Methodology for Information Flow Experiments. CoRR abs/1405.2376 (2014)
[Ur et al. 2012] B. Ur, P.G. Leon, L.F. Cranor, R. Shay, and Y. Wang. Smart, Useful, Scary, Creepy: Perceptions of Online Behavioral Advertising, SOUPS 2012.
[UPPP2014] Usable Privacy Policy Project Website – Publications List - http://www.usableprivacy.org/relevant-publications
[Wang et al. 2008] Wang, Yang, et al. "I regretted the minute I pressed share: A qualitative study of regrets on Facebook." *Proceedings of the Seventh Symposium on Usable Privacy and Security*. ACM, 2011.
[Wesolowski et al. 2012] A. Wesolowski et al. "Quantifying the Impact of Human Mobility on Malaria." *Science*, vol. 338, no. 6104, Oct. 2012, pp. 267-70.
[Wesolowski 2014] A. Wesolowski, "Quantifying Human Movement Patterns for Public Health." Ph.D. Dissertation, Carnegie Mellon University, 2014.
[Whitten and Tygar 1999] Alma Whitten , J. D. Tygar, Why Johnny can't encrypt: a usability evaluation of PGP 5.0, Proceedings of the 8th conference on USENIX Security Symposium, p.14-14, August 23-26, 1999, Washington, D.C.
[WH12] Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy, White House, February 2012. http://www.whitehouse.gov/sites/default/files/privacy-final.pdf
[Wang et al. 2014] Yang Wang, Pedro G. Leon, Alessandro Acquisti, Lorrie F. Cranor, Alain Forget, and Norman Sadeh. A field trial of privacy nudges for Facebook. CHI 2014: Conference on Human Factors in Computing Systems. Toronto, CA, April 2014.
[Yang et al. 2012] X. Yang, S.E. Fienberg, A. and Rinaldo, Differential Privacy for Protecting Multi-dimensional Contingency Table Data: Extensions and Applications, *Journal of Privacy and Confidentiality*: (2012) 4: Iss. 1, Article 5.

URLs with Authors' Bios

A. Acquisti (Heinz) http://www.heinz.cmu.edu/~acquisti/bio.htm
Y. Agarwal (SCS) http://www.synergylabs.org/yuvraj/
L. Bauer (CIT) http://www.ece.cmu.edu/~lbauer/bio.html
A. Blum (SCS) http://www.cs.cmu.edu/~avrim/
T. Breaux (SCS) http://www.cs.cmu.edu/~breaux/biography.html
L. Cranor (CIT/SCS) http://lorrie.cranor.org/bio.html
A. Datta (CIT/SCS): http://www.andrew.cmu.edu/user/danupam/bio.txt
L. Jia (CIT) https://www.andrew.cmu.edu/user/liminjia/
J. Peha (CIT) http://users.ece.cmu.edu/~peha/bio.html
N. Sadeh (SCS) http://www.normsadeh.org/short-narrative/
D. Sicker (CIT) https://www.epp.cmu.edu/people/bios/sicker.html
T. Sandholm (SCS) http://www.cs.cmu.edu/~sandholm/