

Nudges for Privacy and Security: Understanding and Assisting Users' Choices Online

ALESSANDRO ACQUISTI, Carnegie Mellon University
IDRIS ADJERID, University of Notre Dame
REBECCA BALEBAKO, RAND Corporation
LAURA BRANDIMARTE, University of Arizona
LORRIE FAITH CRANOR, Carnegie Mellon University
SARANGA KOMANDURI, Civis Analytics
PEDRO GIOVANNI LEON, Banco de Mexico
NORMAN SADEH, Carnegie Mellon University
FLORIAN SCHAUB, University of Michigan
MANYA SLEEPER, Google
YANG WANG, Syracuse University
SHOMIR WILSON, University of Cincinnati

Advancements in information technology often task users with complex and consequential privacy and security decisions. A growing body of research has investigated individuals' choices in the presence of privacy and information security trade-offs, the decision making hurdles affecting those choices, and ways to mitigate such hurdles. This article provides a multi-disciplinary assessment of the literature pertaining to privacy and security decision making. It focuses on research on assisting individuals' privacy and security choices with soft paternalistic interventions that nudge users towards more beneficial choices. The article discusses potential benefits of those interventions, highlights their shortcomings, and identifies key ethical, design, and research challenges.

CCS Concepts: •**Security and privacy** → **Human and societal aspects of security and privacy**;
•**Human-centered computing** → *Human computer interaction (HCI)*; *Interaction design*;

Additional Key Words and Phrases: Privacy, Security, Nudge, Soft Paternalism, Behavioral Economics.

ACM Reference Format:

Acquisti, A., Adjerid, I., Balebako, R., Brandimarte, L., Cranor, L. F., Komanduri, S., Leon, P. G., Sadeh, N., Schaub, F., Sleeper, M., Wang, Y., Wilson, S. 2016. Nudges for Privacy and Security: Understanding and Assisting Users' Choices Online. *ACM Comput. Surv.* 1, 1, Article 1 (September 2016), 41 pages.
DOI: 0000001.0000001

1. INTRODUCTION

As they go about their online activities, individuals are faced with an increasing number of privacy and security decisions. Those decisions range from configuring visibility on social networking sites to deciding whether to download a mobile app based on the access to sensitive data it requests; from determining whether to trust a website to clicking on—or ignoring—a link in an email. Such decisions arise in both personal and

This research has been supported by the National Science Foundation under grant CNS-1012763 (Nudging Users Towards Privacy), as well as grants CNS-0627513 and CNS-0905562, and by a Google Focused Research Award. This research has also been supported by CMU CyLab under grants DAAD19-02-1-0389 and W911NF-09-1-0273 from the Army Research Office, the IWT SBO SPION Project, Nokia, France Telecom, and the CMU/Portugal Information and Communication Technologies Institute.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2016 ACM. 0360-0300/2016/09-ART1 \$15.00

DOI: 0000001.0000001

work contexts: workers who used to rely on workplace system administrators to manage enterprise security often find that they have to configure some security settings on their own, be it in the context of “Bring Your Own Device” workplaces or while interacting with an ever more diverse set of services where system administrators are not available to help.

Many factors render privacy and security decisions especially complex online. First, information technologies and online threats constantly evolve, leaving users in what economists would describe as a condition of incomplete and asymmetric information [Akerlof 1970]. It is often difficult for users to determine how much of their data may be collected and how it might be used; and it is nearly impossible for them to fully assess what security and privacy vulnerabilities they might expose themselves to, if they decide to interact with a given system. Second, the trade-offs associated with privacy decisions (e.g., sharing or hiding data) or security decisions (e.g., choosing the “right” type of protection for a 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 Herb Simon referred to as “bounded rationality” [Simon 1957]. Both of these problems are exacerbated by the inherent uncertainty, and sometimes ambiguity, associated with the trade-offs involved in privacy and security choices, where laxer settings may be key to unlocking additional functionalities and to deriving (seemingly) more value from a given system. Third, decisions involving the disclosure of information or protection of information systems are 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 not just based on rational calculus, but also based on heuristics [Simon 1982], mental shortcuts, feelings, and emotions. Heuristics are not necessarily bad: they can prove quite effective most of the time—but can also lead to regrettable decisions, ranging from over-sharing to increased exposure to cyber-attacks. In fact, interfaces of popular online services and systems sometimes exploit these heuristics and biases in order to nudge users to act in ways that are not always aligned with the users’ own intentions [Stutzman et al. 2013].

In recent years, various streams of research have tried to understand and assist people’s privacy and security decisions and behaviors. Usability research has attempted to help users by making security and privacy interfaces more usable [Sasse et al. 2001; Cranor and Garfinkel 2005; Garfinkel and Lipford 2014]. Behavioral economics and decision research have analyzed the hurdles individuals face when making privacy or security decisions online [Acquisti 2004; Acquisti et al. 2015]. In addition, a growing body of work has started examining how interfaces and services can be designed to counter biases responsible for disadvantageous security and privacy decisions [Chiasson et al. 2008; Acquisti 2009; Wang et al. 2014; Almuhiemedi et al. 2015]. This line of research takes up the concept of “weak,” “libertarian,” or “soft” paternalism [Thaler and Sunstein 2008]. Soft paternalism applies lessons from behavioral research to design policies, systems, and architectures of choice intended to nudge users towards better decisions without restricting their options. This article draws from results across this array of disciplines to summarize the research aimed at understanding hurdles faced by individuals in online privacy and security decision making, and developing ways of effectively assisting individuals in their privacy and security decisions.

First, we review research in relevant fields (such as behavioral decision research, behavioral economics, and experimental psychology) to gain insights into the impact of cognitive and behavioral biases on online security and privacy decision making (Section 2). Then, we review interventions developed in various fields (such as decision research, human-computer interaction, and persuasive technologies) aimed at helping

users make “better” online security and privacy decisions—that is, decisions that minimize adverse outcomes or are less likely to be regretted. We show how this work shares similarities with mechanisms developed to nudge people in a variety of other domains, such as health and retirement planning. We broadly refer to these efforts as “nudging research,” regardless of the originating field of study. We posit that all these efforts can be largely viewed as implementations of soft paternalistic concepts, whereby interventions are intended to gently guide users towards safer practices rather than imposing particular decisions. We suggest that prior work on the design of user interface technologies for security and privacy can be examined from a nudging perspective: every design decision potentially nudges users in one direction or another. Furthermore, we point to examples of existing interfaces that nudge individuals either towards more protective behaviors (Section 3) or, sometimes, towards riskier ones (Section 4). We further discuss practical and ethical questions associated with nudging for security and privacy, along with a discussion of design and research challenges in this area (Section 5). Finally, we conclude with a summary of insights identified in this review (Section 6).

2. HURDLES IN (PRIVACY AND SECURITY) DECISION MAKING

The theory of rational choice models individuals as agents with stable preferences that guide their decision making: decision making is assumed to be rational, in that individuals aim to maximize their utility given their preferences and other constraints they may be subject to [Becker 1976]—for instance, the available budget. Since privacy and security decisions also give rise to complex trade-offs, such as the costs and benefits associated with disclosure or protection of sensitive information, it seems reasonable to use the rational model of decision making in these domains as well. Indeed, attempts to apply economic approaches to privacy date back to the late 1970s and early 1980s, when Chicago School scholars started arguing that privacy protection may actually represent an obstacle to the achievement of efficient market equilibria [Posner 1978; 1981; Stigler 1980], and continued in the information systems literature through the concept of “privacy calculus” [Laufer and Wolfe 1977; Smith et al. 2011]. Similar efforts to link economic research with security research began in the late 1990s [Anderson 2001].

As stylized above, the traditional model of rational choice glosses over a number of factors that affect actual privacy and security choices of real consumers and decision makers. First, in the late 1990s research into human factors and usability started uncovering the role of interfaces in affecting—and sometimes impairing—users’ awareness of privacy and security features or threats, as well as their ability to make sound choices [Whitten and Tygar 1999; Adams and Sasse 1999; Cranor and Garfinkel 2005; Garfinkel and Lipford 2014]. Second, in the early 2000s behavioral economics and behavioral decision research started highlighting the role of biases and heuristics in affecting privacy and security behaviors [Acquisti 2004; Acquisti and Grossklags 2005], and the role of emotions and cognition in disclosure decisions [Li et al. 2008; Li et al. 2011]. Biases do not imply that individuals’ behavior is erratic, irrational (in the sense of absurd), or wrong: biases are simply systematic, and therefore predictable, deviations from rational choice theory [Camerer et al. 2003]. Heuristics are shortcuts in decision making which may lead to success, but can also lead to errors. Standard assumptions of economic rationality do not account for the observation that individuals have limited resources to allocate to decision processes, and that when information, experience, or time is limited, they have to rely on heuristics, especially when the actual value of potential choices is uncertain [Simon 1982].

In this section, we draw from findings in human factors research, usability research, behavioral economics, decision research, and experimental psychology, to provide an

overview of hurdles users face when making privacy and security decisions. We focus on how decision making is affected by incomplete and asymmetric information, heuristics and bounded rationality, as well as cognitive and behavioral biases. For each type of hurdle, we start with a general description of the phenomenon and then discuss its relevance to security and privacy decision making. To classify those hurdles, we draw on the distinction made by Tversky and Kahneman [1975] between heuristics and biases, and add to those notions the economic concepts of incomplete and asymmetric information. However, rather than aiming for an exhaustive description of decision making hurdles, we highlight those that are most representative and consequential in privacy and security decision making, and have been the focus of research at the intersection of privacy, security, and psychology.¹

2.1. Incomplete and Asymmetric Information

The term “incomplete information” can be used to refer to situations where economic agents lack information pertaining to a transaction. The term “asymmetric information” can be used to refer to situations where agents who interact in a transaction have differential access to information relevant to that transaction: one party may have more or better information than other parties. The party who is better informed is often in an advantageous position when completing the transaction; however, asymmetric information can cause widespread adverse effects and inefficiencies. A classic example is the “market for lemons” [Akerlof 1970]. Consider the market for used cars. The seller of a used car knows the car’s actual history and quality, while its potential buyers have limited information. The uncertainty regarding the car’s quality leads potential buyers to discount the value of used cars sold on the market: because the used car could be a “lemon,” buyers’ willingness to pay for the *average* car is lower than what it would have been under symmetric, complete information. As a result, owners of *good* cars may stop selling them, when their reservation price is higher than buyers’ average evaluation of used cars. Thus, asymmetric information may crowd the best cars out of the market, leaving only lower quality cars for sale. Anticipating this result, potential buyers further reduce their willingness to pay for the average car, giving rise to a vicious circle that may result in the market’s failure.

Situations of asymmetric or incomplete information are the norm in the domains of privacy and information security. For instance, the defender may not know which vector the adversary will use for the attack. Or, the data holder (that is, the agent receiving personal data about another agent, the data subject) has more information than the data subject regarding the purposes and conditions of future use of that data. For example, when subscribing to a mailing list, most times we cannot know whether the mailing list might be sold to a third party that could send us unsolicited messages [Varian 1996]. In this situation we may end up receiving junk mail, spam emails, and phone calls that represent a cost to us in terms of the time and energy required for dealing with them.

2.2. Heuristics and Bounded Rationality

Many situations require individuals to make judgments about uncertain events over time and under information constraints. Similarly, security or privacy contexts often require users to evaluate the probability of adverse security or privacy events that are highly uncertain and often far in the future. Such probability judgments may not be

¹More exhaustive lists of decision heuristics and biases have been presented by Rabin [1998] and Camerer et al. [2011]. For a shorter review of the role of behavioral economics in privacy research, see Acquisti and Grossklags [2007]. For a review of psychology and security, see Schneier [2007]. For a recent review of the behavioral literature on privacy, see Acquisti et al. [2015].

trivial and can require considerable cognitive effort and information, leading individuals to lean on heuristics.

Heuristics, or rules of thumb, are shortcuts in decision making. Individuals use heuristics when bounded rationality prevents the exploration of all possible outcomes. The rational choice model used in traditional economics views economic agents as rational decision makers who maximize their utilities when making a decision, regardless of its complexity. In contrast to this view of humans as precise calculators, who have unlimited computational resources at their disposal, Simon [1957] acknowledged human limitations and “bounded rationality,” and described the decision process as a series of simplification steps. People first simplify the choices available by using heuristics that may have very little to do with rational decision making in the economic sense, and only then apply their rationality to select the best option among the remaining ones. This process may or may not lead to the same choice that the classic rational economic agent would have made.

Numerous examples of bounded rationality exist in everyday life. Consider financial savings: perfectly rational agents would evaluate savings decisions via a trade-off between the value they place on current consumption versus the discounted value of future consumption and their predictions about future earnings potential. However, this calculation is difficult and involves many uncertain factors. Thus, individuals may rely on decision making shortcuts or heuristics to simplify this decision, such as choosing the savings plan that has been selected as a default by their employer with the (often incorrect) impression that the default was set based on normative judgment [Choi et al. 2004; Madrian and Shea 2001]. Consider a privacy-related example: when we need to buy a certain product, we may not necessarily consider all possible alternatives available (such as purchasing from a different store, buying online, using cash, using a gift card, using a credit card, etc.) and their respective privacy implications. Instead, we may buy the desired product using whatever form of payment is most convenient at the time of purchase. Similarly, it is unlikely that we will invest the effort to understand the privacy policies or terms of service of the online store selling that product: most users accept the policies without reading them [Good et al. 2006; McDonald and Cranor 2008]. We may even ignore relevant information: clear and succinct notices about spyware in computer programs may not prevent users from installing potentially harmful software, arguably acting against their own best interests [Good et al. 2007].

The availability heuristic is one way in which individuals simplify probability judgments. It refers to the misestimation of probabilities depending on the salience (i.e., availability) of relevant points of comparison. For example, a younger individual may underestimate the likelihood of being diagnosed with prostate cancer due to a lack of accessible examples. Conversely, an older individual may know someone who has been diagnosed with prostate cancer and may, therefore, overestimate his own chance of being diagnosed with prostate cancer. When facing privacy decisions, individuals are often uncertain about the risks associated with disclosure, and may lean on heuristics to judge the probability of adverse events. The availability heuristic may come into play when users are heavily influenced by salient cues that may or may not be effective signals of the probability of adverse events. For instance, they may attempt to estimate the risk of disclosure by evaluating the probability of others disclosing personal information in the same or similar contexts.

Relatedly, the representativeness heuristic refers to overestimating the probability of events that are similar to their population. For instance, in a sequence of random coin flips between heads (H) and tails (T), individuals may perceive the sequence THTHT as less likely than a seemingly random draw from the same distribution (e.g., HHTHT) because the latter is more in line with the characteristic of the population.

In the context of privacy decisions, the representativeness heuristic may lead users to perceive privacy intrusions as low-probability events, because they are often not directly observable, and therefore do not conform to the perceived distribution of events online.

Behavioral privacy research suggests that these heuristics are important factors for probability judgments on adverse privacy and security events. Acquisti et al. [2012] found that disclosure is sensitive to herding effects, where disclosure of personal information can be influenced by information about others' disclosures. This is particularly problematic in contexts where disclosures by others are extremely salient, while consequences of inappropriate disclosures may be opaque and therefore less encountered by individuals, resulting in individuals underestimating the potential risks of over-disclosure. Spiekermann et al. [2001] explained that observed inconsistencies between stated privacy concerns and actual disclosure behaviors in an e-commerce experiment could have happened due to the availability of positive memories with shop assistants. The impact of salient contextual cues may be particularly pronounced when individuals are unable to gauge the appropriate level of privacy concern in a given environment. Ur et al. [2012b] showed that people make privacy judgments regarding online tracking based on a company's name, but not necessarily based on its privacy practices or the practice of tracking itself. Related to that, Metzger [2006] found that while vendor reputation has an important effect on trust, privacy policies do not affect trust or disclosure. Hoofnagle and King [2008] found that the presence of privacy policies on websites, regardless of their content, can lead people to think that websites protect their personal information. Thus, they may rely on contextual factors, such as the look-and-feel of a website or the manner in which the questions about personal information are solicited, to gauge the appropriate level of privacy concern and their subsequent level of disclosure. In the area of security, it has been shown that users often ignore security warnings, even in the context of sensitive websites, such as banking or health [Schechter et al. 2007; Sharek et al. 2008]. If users have not been (or are not aware of someone who has been) hurt due to lax security behaviors, they are unlikely to heed security advice.

2.3. Cognitive and Behavioral Biases

Whereas heuristics are shortcuts in decision making, biases affect rational decision making in general, regardless of a decision's complexity. Cognitive and behavioral biases are systematic errors in judgments and behaviors. As the term "systematic" suggests, biases do not necessarily imply odd or "wrong" behavior. Rather, they represent deviations from the stylized economically rational behavior predicted by rational choice theory. These biases and their application to economic decision making have been studied since the late 1950s [Tversky and Kahneman 1981; Simon 1957; 1982; Kahneman and Tversky 1979]. Since then, research in behavioral economics, behavioral decision making, and cognitive psychology has uncovered cognitive and behavioral biases affecting decision making in many domains. Here we focus on biases that have been shown to affect online privacy and security decisions.

2.3.1. Anchoring. When making a decision or assessment, we often consider information that may or may not be relevant to the situation at hand. That information creates a reference point from which we make small adjustments to make a decision for the specific circumstance faced. This cognitive bias is known as anchoring. Ariely et al. [2003] illustrated anchoring with a behavioral experiment in which participants were first asked to provide the last two digits of their social security number before estimating the price of a bottle of wine or luxury chocolates. Results showed that participants

provided price estimates that were close to the two digits they had entered. The two digits anchored their responses, although they had no relation to the products.

Anchoring may significantly affect privacy decision making. For instance, when deciding what to post on a social media service, one may be vastly affected by what others post, and set that as an anchor regardless of the actual comfort one may initially feel in revealing personal information, or regardless of the consequences that revealing personal information may entail. Deciding what to reveal and what to keep private is not trivial, as it involves assessments related to how others may perceive, interpret and use one's revelations, whether these revelations may have future negative consequences for oneself or others, and so on. Given the complexity of the decision process, people may tend to take the example of their trusted peers as a reference point for what is appropriate to post, and emulate them.

Acquisti et al. [2012] showed anchoring effects in the willingness to disclose sensitive information through a survey. Participants who were told that other participants had made sensitive disclosures were more likely to disclose sensitive information than those who were not told anything about other participants' disclosures. Participants were further influenced by order effects as they tended to disclose more sensitive information when the survey started with privacy-intrusive questions and gradually reduced in sensitivity, compared to the survey starting with less privacy-intrusive questions and gradually increasing in sensitivity. Chang et al. [2016] showed that exposing participants to more provocative selfie images changed their perception of what information is appropriate to share and significantly increased the probability of them disclosing personal information. Another example is the common practice among U.S. telecommunication and Internet service providers of requesting social security numbers from new customers, regardless of whether they need that information. This may result in people anchoring their expectations. Thus if another company asks for an SSN, based on prior experiences the person may be willing to reveal this sensitive information without thinking about its necessity or potential consequences.

2.3.2. Loss aversion. According to classical economic theory, one should derive as much utility from a gain as one derives disutility from a tantamount loss [Friedman and Savage 1948]. However, this property does not seem to hold true in experiments, as people tend to dislike losses more than they like equivalent gains. In other words, people tend to be loss averse. Prospect theory, developed as an alternative to classical expected utility theory, offers an explanation: the value function (corresponding to the classical utility function) is assumed concave for gains but convex for losses [Kahneman and Tversky 1979]. This feature of the value function implies loss aversion. In a typical experiment, a group of people are endowed with an amount of money and are asked to state the highest amount they would be willing to pay for a certain object (e.g., a mug), while another group of people are endowed with the object and asked to state the least amount of money they would sell it for. On average, people endowed with the object (and thus in a position where they would lose it in exchange for money) ask for more money than those endowed with money are willing to pay for it [Kahneman et al. 1990].

Loss aversion also applies to privacy and security issues. While people may not be willing to pay for services that would help protect their privacy and security online [Shostack 2003], they report high privacy concerns about companies gathering their personal information [Hoffman and Novak 1997; Taylor 2003; Pew Research Internet Project 2013]. In behavioral experiments, Grossklags and Acquisti [2007] and Acquisti et al. [2013] showed that people are willing to accept more money in exchange for disclosing personal information than they are willing to pay to regain control over the same information. In other words, when people feel in possession of their personal

information they value it more (e.g., they resist losing it), but when they feel they have already lost it, they value it less.

2.3.3. Framing effects. Prospect theory incorporates and explains other decision making biases, such as framing effects, preference reversals, and inconsistencies caused by the wording of the options available. For instance, whether an option is presented as a loss or as a gain significantly affects people's choice. Tversky and Kahneman [1981] showed that treatment A, which will save 200 out of 600 people infected with a deadly disease, is preferred (on average) over treatment B, which is expected to result in 400 people dying, despite the same number of people being saved with each treatment. Positive framing emphasizes the gain and thus moderates loss aversion, whereas negative framing makes losses salient.

Framing effects affect privacy decision making in a similar fashion. Framing a certain privacy notice as more or less protective compared to a reference point (e.g., a competitor's privacy notice or the status quo) can make the policy itself more or less desirable to users, and can ultimately affect their willingness to share personal information [Adjerid et al. 2013]. Research has shown that different ways of framing control over disclosure of personal information affect users' disclosure decisions, even when the objective privacy risks remain the same. Specifically, in the context of an online social network, participants who were offered stronger privacy controls disclosed more personal information compared to those who were offered weaker control [Brandimarte et al. 2013]. Greater control has also been found to reduce privacy concerns regarding online personalization [Taylor et al. 2009].

2.3.4. Hyperbolic discounting. Inter-temporal choices, such as refraining from a snack now in order to maintain your appetite for the next meal, often deviate from rational models of decision making. Traditional economic models of inter-temporal decision making rely on discount utility models that employ a constant discount rate to allow individuals to weigh the trade-offs of consumption now versus delayed consumption. More generally, discount rates are used by people to discount future expected pay-offs (e.g., cost and benefits) and make life decisions in the present, such as decisions concerning savings, work effort, exercise, health habits, etc. Such a constant discount rate has implications for individual behavior. For example, it implies that individuals choosing between consumption now or in one week will use the same discount rate as individuals choosing between consumption in one year or in one year and one week. However, behavior that violates this implication of discount utility theory has been observed in multiple experiments [Frederick et al. 2002]. For example, Thaler [1981] observed a 345% discount rate in a one month horizon, a 120% discount rate over a one year horizon, and a 19% discount rate over a five year horizon. Moreover, rational accounts of inter-temporal decision making do not account for a need for immediate gratification which may lead to overvaluation of present consumption.

Hyperbolic time discounting, an alternative model to discount utility theory, accounts for observed patterns by considering a discount rate that is decreasing over time. Such a deviation from traditional discount utility theory may affect individuals' financial, physical, and social well-being. For example, hyperbolic time discounting could help explain some individuals' choice to overeat or forego saving at the expense of future health and financial risks that may reduce their life expectancy and quality [Ainslie and Haslam 1992; Laibson 1997; Loewenstein and Prelec 1992; Shapiro 2005].

This deviation from rational choice also has significant implications for security and privacy decision making, due to the inter-temporal trade-offs involved [Acquisti 2004]. For instance, privacy decision making may involve disclosure, which brings individuals some immediate gratification, such as social interaction or access to desired services, while simultaneously subjecting them to privacy costs that may be incurred months

or even years later. For example, Jentzsch et al. [2012] found that given a choice of two movie ticket providers requiring different amounts of personal information and charging different prices, people preferred monetary savings (an immediate gain) over privacy (a longer term risk). However, when the tickets were offered at the same price, the privacy-friendly provider gained the larger market share, suggesting that people do care about privacy, but they often heavily discount the risk of disclosing personal information. A positive present monetary gain often trumps privacy concerns. In an e-commerce experiment, immediate benefits of disclosure (e.g., product recommendation and discounts) were suggested as the probable cause that led privacy-concerned participants to provide the same amount of personal details as less concerned participants [Spiekermann et al. 2001]. Individuals who are susceptible to hyperbolic time discounting may discount the future heavily relative to the value they can instantly receive from disclosing information. This *may* lead to repercussions for individuals later in life. For instance, sharing information about recent purchases may generate enthusiastic comments from friends, but could also result in future price discrimination; posting a salacious picture or a revealing tweet might cause a moment of hilarity among friends, but could lead to regret [Wang et al. 2011]; publicly disclosing innocuous information, such as religion or family status, on social media sites could result in bonding among friends, but also denied job applications in the future [Acquisti and Fong 2014]. In the case of information security, neglecting to install security updates or to update antivirus software may save time and effort now, but increases the vulnerability to future security breaches.

2.3.5. Optimism bias and overconfidence. Optimism bias and overconfidence are instances of incorrect estimates of subjective probabilities. Optimism bias is an underestimation of the chances that one might be subject to a negative event. For example, traders may show unrealistic optimism in estimating the probability of a loss based on market conditions. Similarly, people may underestimate the probability of becoming victims of cybercrime or identity theft, which may lower their alertness online. Overconfidence is an overestimation of the accuracy of one's judgments, resulting in excessive confidence in them. Lichtenstein et al. [1982] showed how people who were certain about their answers to general knowledge questions were right only 80% of the time, revealing a 20% overestimation of their accuracy.

People may be overconfident in their assessment of privacy or security risks. For example, they may be convinced that their antivirus software is fully effective against all possible threats, while its efficacy may be substantially lower. As a consequence, they engage in insecure online behaviors that lead to compromised devices and data loss. Internet users often accept "unjustified" risks that enable successful attacks [Xia and Brustoloni 2005]. While insecure behavior may not always be the result of overconfidence and can also follow well-reasoned decisions [Herley 2009; Vila et al. 2003], cybersecurity and privacy attacks do affect users who underestimate the risks.

2.3.6. Post-completion errors. Tasks that involve a series of steps, often to be repeated several times, consist of primary and secondary goals. For instance, photocopying a document entails taking the document to the copy machine, making the copy, and finally removing the original from the machine. People are prone to post-completion errors, where they succeed in achieving the primary goal (copying the document) but fail in the secondary goal (removing the original from the copier) [Li et al. 2005].

These errors are likely to occur when a person is interacting with a machine, and when this interaction involves a series of procedures one has to perform consistently every time [Curzon and Blandford 2004]. Forgetting to log off of a shared computer at the end of one's primary task, or omitting to close a browsing session and cleaning one's

browsing history may have significant privacy and security consequences, especially if the primary goal involved sensitive or personal material.

2.3.7. Status quo bias. The status quo bias refers to individuals' affinity for default choices. For example, an opt-in approach to organ donation results in significantly fewer organ donors relative to an opt-out approach, in which the default is to be an organ donor [Johnson and Goldstein 2004]. Dhingra et al. [2012] identify a "default pull, in that not only do people have a propensity to keep the default, but the option identified as the default also shapes which alternatives they choose, should they deviate from the default. Explanations that have been proposed for the impact of defaults on behavior include inattention, transaction cost of choosing an alternative, loss aversion, defaults reflecting normative judgments, and defaults having a role in the construction of uncertain preferences [Anderson 2003; Kahneman et al. 1991; Kahneman and Miller 1986].

Privacy decisions have several characteristics that may make default options especially impactful on behavior. First, information about privacy risks can often be difficult to find and understand, and thus attention may be a limiting factor in decision making, allowing defaults to go unchanged for many individuals. Moreover, preference uncertainty may be pronounced in the context of privacy decision making [John et al. 2011], and the trade-offs associated with disclosure are often difficult to quantify and evaluate. Thus, normative judgments of defaults or their propensity to shape uncertain preferences may have a significant impact on behavior. Furthermore, users of privacy and security tools may assume that the default configurations of those tools protect them, without reviewing the settings [Leon et al. 2012].

3. ASSISTING PRIVACY AND SECURITY CHOICES

Incomplete information, heuristics, and cognitive or behavioral biases can lead users to make decisions that result in undesirable outcomes. If the decisions involve privacy or security, those outcomes may include security breaches, privacy invasions, and regrettable disclosures. For example, Calo [2010] discussed privacy harms that can arise out of individuals' lack of awareness regarding how their data is collected or used; Wang et al. [2011] studied user regrets caused by unintended disclosures on Facebook; Sleeper et al. [2013] studied regrets related to Twitter posts; and Anderson et al. [2013] estimated information security costs associated with cybercrime, partially facilitated by poor security decisions.

In this section, we provide an overview of privacy and security interventions—that is, strategies and approaches aimed at assisting privacy and security decision making. In particular, we focus on a broad space of interventions aimed at countering the decision making hurdles discussed in the previous section. Such interventions range from education and transparency (which may address problems associated with incomplete information), to usable privacy and security tools (which may address problems associated with bounded rationality), to nudging interventions (which may address problems associated with cognitive and behavioral biases).

3.1. Soft Paternalism and Nudging

The interventions we consider are drawn from different research fields, including usability, persuasive technologies, and behavioral decision research. They have often evolved separately from each other, and at different times. For instance, the role of educational interventions and usability research in privacy and security decision making has been explored for over 15 years in the computer science literature (see Bishop [2000], Langheinrich [2002], and Cranor and Garfinkel [2005], respectively); in comparison, the exploration of soft paternalism or nudging to support online security and

privacy decision making is more recent [Acquisti 2009]. However, these interventions are best regarded as complements rather than substitutes. For instance, greater transparency and awareness interventions can help overcome information hurdles. Yet, more information is not always the answer to privacy and security problems [Adjerid et al. 2013], and may even have a negative effect [Ben-Shahar and Schneider 2010]: too much information and too many settings can be overwhelming. In turn, usable security and privacy research aims to overcome decision complexity through the design of interfaces that offer users manageable and easy-to-understand options. Yet, usability does not guarantee better decision making and may not account for other decision making hurdles, such as heuristics and cognitive or behavioral biases, that may trump any benefits afforded by more usable interfaces. An example would be default settings that lead users to select configurations that are not well aligned with their objectives [Stutzman et al. 2013]. This is where behavioral research can help in reframing options made available to users and thereby ameliorate their decision making.

3.1.1. Paternalism, Libertarianism, and Soft Paternalistic Interventions. One way to connect these different interventions from various streams of research is to focus on their role in informing and guiding users' decisions towards safer, better choices, without imposing a particular decision. Such "soft paternalistic" approaches stand between two opposite solutions to policy problems: strong paternalism and strictly libertarian approaches.

Paternalistic approaches impose decisions on users that are believed to be beneficial for them [Camerer et al. 2003; Thaler and Sunstein 2008]. A common example of paternalism is government regulation. While regulating behavior (e.g., forcing individuals or organizations to act in a specific way) is a form of coercion with potentially negative effects (such as limiting autonomy, leading to suboptimal outcomes for specific groups, or stifling innovation), policy makers have traditionally resorted to this approach in life-threatening situations (for example, banning cigarettes and alcohol for underage individuals or the mandatory use of seat belts). Paternalistic regulations impose disciplinary, financial, or criminal penalties for failing to comply with the desired behavior, but do not address systemic decision making problems directly. On the opposite end of the spectrum, strictly libertarian approaches rely on self-regulatory solutions, such as market dynamics and economic competition, in which organizations police themselves to prevent failures, and individuals are expected to make choices in their own best interest. A libertarian system design takes a neutral view and just provides the user with a set of available options, without regard for whether some options are potentially detrimental to the user's interests [Thaler and Sunstein 2008]. Neither approach (regulation or self-regulation) is guaranteed to achieve the stated objectives. Regulation might fail (or cause unintended consequences), and self-regulatory approaches may not adequately address problems—such as smoking among minors, or privacy protection [Acquisti et al. 2015].

Soft paternalistic interventions, instead, attempt to influence decision making in order to improve individual well-being, without actually limiting individual choices, but in fact preserving freedom of choice. Soft paternalistic approaches achieve this by reframing the choices available to users to increase the likelihood that users will make decisions beneficial to them. Soft paternalistic interventions are often referred to as "nudges." Thaler and Sunstein [2008] define a nudge as "any aspect of the choice architecture that alters people's behavior in a predictable way without forbidding any options or significantly changing their economic incentives" While this definition may still leave room for confusion as to "what exactly is to be regarded as a nudge, [Hansen 2016] by and large nudges attempt to account for behavioral and decision making bi-

ases that may result in choices with potentially adverse outcomes. Nudging acknowledges that users are often unaware of these biases and can be affected by subtle differences in system design, e.g., defaults, saliency of features, or feedback. Nudging users through improved system designs and architectures aims at mitigating, or sometimes leveraging, human biases to lead to more beneficial outcomes for users. This approach can be leveraged to nudge users to make decisions that are better aligned with their security and privacy objectives.

3.1.2. Dimensions of Nudges. We present various soft-paternalistic interventions that support privacy and security decision making in information systems. Using the term broadly, we refer to these interventions collectively as “privacy and security nudges.” We acknowledge, however, that the discussed interventions may be at different positions along the spectrum between strict paternalistic and libertarian approaches—with some, like transparency, being designed to provide limited influence on an individual’s choice, and others, like defaults, more forcefully and predictably affecting users’ decisions. In order to structure different interventions, we start from Thaler and Sunstein’s NUDGES acronym [Thaler and Sunstein 2008]: iNcentives, Understand mappings, Defaults, Give feedback, Expect errors, Saliency/Structuring complex choices. We extend this categorization by considering design recommendations from other fields, such as usability research, debiasing, and persuasive technologies.² Debiasing refers to attempting to assist human decision making by reducing bounded rational behavior. Debiasing includes devising interventions that will reduce the effect on an identified judgmental bias, or creating situations where the bias does not emerge at all [Fischhoff 1981]. Persuasive technology research is concerned with “a computing system, device, or application intentionally designed to change a person’s attitude or behavior in a predetermined way” [Fogg 2002]. Persuasive technologies have been applied in many domains and for many purposes: to encourage healthier behaviors through exercise [Segerstahl and Oinas-Kukkonen 2007] and physical activity [Consolvo et al. 2006], to reduce electricity consumption among college students [Petersen et al. 2007], and to help people quit smoking [Graham et al. 2006].³

As a result, we group the various nudging interventions around the following six categories: Information, Presentation, Defaults, Incentives, Reversibility and error resiliency, as well as Timing. We note that these categories are interrelated and not mutually exclusive. For instance, feedback can be given at various points in time, and the particular timing can affect decision making. Our grouping of nudging dimensions provides a narrower and more application-oriented categorization of techniques than other more general approaches to nudging proposed previously. These dimensions, shown in Table I, allow us to directly map various interventions and nudging techniques with design recommendations.

Below, we first introduce respective interventions in general computer systems or non-computing-related societal systems, before focusing more specifically on privacy and security-related nudging approaches found in both academic literature and commercial applications. While we focus on how nudges can be used to overcome decision making hurdles, we also discuss, at the end of this section, how cognitive and behav-

²Others, e.g., Dolan et al. [2012], have categorized nudges with a focus on assisting public policy decisions.

³We do not attempt to define differences between nudging and these other disciplines. In fact, we argue that some implementation of persuasive technology and usability can be construed as nudges, and vice versa. Similarly, we do not adopt a strict definition of nudges as interventions that do not change economic incentives, but rather a broad and flexible one that allows for different degrees of forcefulness and slight changes in incentives. Nevertheless, one difference from persuasive technology and usability is the theoretical foundation underpinning nudging. That is, nudging has its origins in behavioral decision research and focuses on countering biases and cognitive limitations (see Section 2).

Dimensions	Subdimensions	Targeted Hurdles
Information. <i>Reduces information asymmetries and provides a realistic perspective of risks.</i>	Education Feedback	Asymmetric and incomplete information, availability heuristic Asymmetric and incomplete information, bounded rationality, availability heuristic, optimism bias and overconfidence
Presentation. <i>Provides necessary contextual cues in the user interface to reduce cognitive load and convey the appropriate level of risk.</i>	Framing	Loss aversion, optimism bias and overconfidence, representativeness heuristic
	Ordering	Post-completion errors, anchoring
	Saliency	Availability heuristic, optimism bias and overconfidence
	Structure	Bounded rationality, availability heuristic, representativeness heuristic
Defaults. <i>Reduce user effort by configuring the system according to users' expectations.</i>	—	Status quo bias, diversification heuristic
Incentives. <i>Motivate users to behave according to their stated preferences.</i>	Increasing cost	Loss aversion
	Rewards / Punishment	Hyperbolic discounting, loss aversion
Reversibility (error resiliency). <i>Limits the impact of mistakes.</i>	—	None in particular. The goal is to allow users to recover from suboptimal decisions potentially caused by behavioral biases.
Timing. <i>Defines the right moment to nudge</i>	—	Each nudging technique may be needed at different points in time.

Table I: Overview of nudging dimensions and the relevant hurdles that they mitigate or exploit. Dimensions are not necessarily mutually exclusive.

ioral biases are exploited to lead users away from making optimal privacy and security decisions.

3.2. Nudging with Information

The provision of information aims at mitigating negative effects of asymmetric information and at overcoming availability and overconfidence biases that may lead to suboptimal decisions. When providing information to create awareness about privacy and security risks, system designers need to consider that users are subjected to different cognitive biases, as discussed in Section 2.3. These biases can either be leveraged to encourage beneficial behavior or need to be mitigated to prevent unintended outcomes. Effective communication uses clear, short, and relevant messages to support users' decision making.

3.2.1. Education and feedback. We distinguish two general approaches for providing information: education and feedback. Education provides information before the user engages with the system or a specific feature. Thus, education primarily supports future decisions. Feedback is commonly provided alongside system usage.

Nutrition labels are examples of educational interventions. Nutrition labels, such as the one mandated by the U.S. Nutrition Label and Education Act (NLEA), inform about and help to compare the nutritional content of food products. Studies about the impact of nutrition labels, however, have concluded that the ones who benefit the most from these labels are people who are already motivated to investigate nutritional content [Drichoutis et al. 2006].

Feedback can make users aware of the consequences of their actions and assist them in making better decisions over time. Feedback can also inform about expected and actual outcomes before or immediately after making a decision. In the latter case, well-designed nudges allow users to revise their decision if the outcome does not meet their expectations, as we discuss in more detail in Section 3.6. A system could also purposefully delay a decision from taking effect to facilitate decision making.

Users' actions are typically intended to change a system's current state. Feedback allows users to either verify that the state was changed as expected or notice unexpected results. It is further useful to remind users about a system's current state. Feedback improves decision making by mitigating the effects of asymmetric information and bounded rationality. Feedback could include suggestions for beneficial actions, along with an explanation why those actions were suggested. Feedback can further prevent suboptimal decisions that arise when users employ simplifying strategies based on availability heuristics that ignore complex, but in some cases strictly better, options.

Providing feedback when users are engaging in risky behavior can serve as a deterrent. For example, positioning displays showing the current speed of cars close to highway speed limit signs helps drivers to realize when they are speeding, and also allows them to compare their speed with that of other drivers, reducing the false sense that everyone is speeding up [Hansen 2012].

3.2.2. Education in privacy and security. In the context of computing systems, users can be educated about the benefits and risks associated with a system before they use it. Users are often required to make privacy and security decisions as they use systems or applications. Those decisions can be supported by relevant information about the consequences of the available options. For example, Facebook users must decide whether to accept friend and application requests, what to share, and with whom to share it. Similarly, Internet users have to decide what websites to trust when shopping online, whether an online banking website is legitimate, or whether opening an email attachment is safe. We argue that users can benefit if they are given relevant information to make those decisions, including the available options and the expected results.

Grossklags et al. [2010] developed an economic model for security decisions based on the amount of information available to users who were deciding whether to invest in individual security mechanisms, such as firewalls or anti-virus, and self-insurance mechanisms, such as backing up their data. Users were described as naive or expert based on their understanding of externalities and the impact of the entire network and others' decisions on their risk. The amount of information available to experts or naive users impacted their decisions to protect or self-insure.

Notices and education are not the same. Morgan and Henrion [1992] discuss how warnings about risk should be tailored to match users mental models, in order to improve their decision making. Raja et al. [2011] applied this to warnings for firewalls. By leveraging "known metaphors," such as locked doors and bandits, in warnings, users were better informed about risks and made more security-protective decisions.

Privacy policy documents, such as those provided by websites, are examples of notices that might overwhelm the reader. Education and information can be achieved through better design [Schaub et al. 2015]. For example, policies can be formatted in a readable and concise manner, such as the "nutrition labels format suggested by Kelley et al. [2010], who found that standardized privacy labels can have a significant impact on users' understanding of privacy policies. Participants were more accurate and faster in reading the standardized notices, and could better compare different policies. Tsai et al. [2011] found that purchasing decisions were impacted when search engines included information about merchants' privacy practices in their results. Similarly,

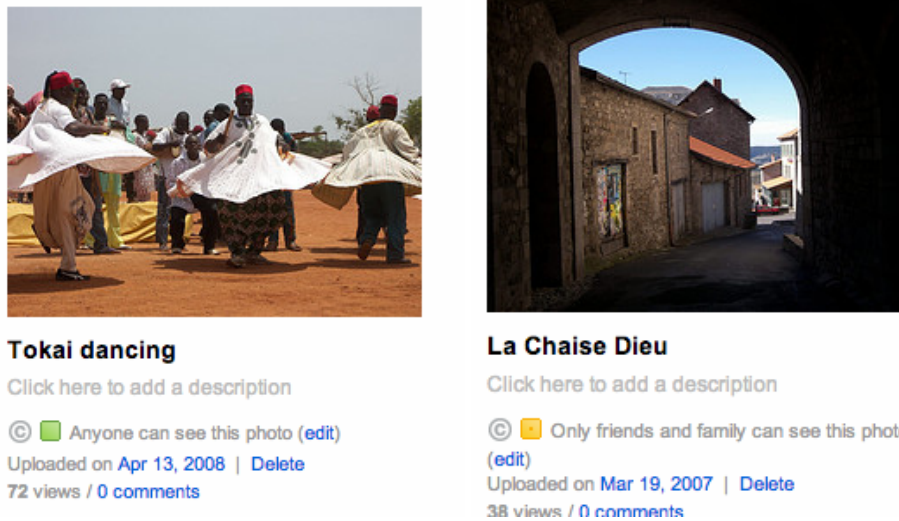


Fig. 1: Flickr concisely shows privacy settings for images below the image with the text 'Anyone can see this photo' (left) or 'Only friends and family can see this photo' (right).

Kelley et al. [2013] found that privacy information displayed in a simulated mobile app store influenced the selection of apps that study participants would download.

Concise information can inform users about privacy settings. Figure 1 shows an example from flickr.com. In this case, the two pictures have different privacy settings, indicated by an icon and text shown below the picture. Options to edit settings are readily available. Bugnosis [Alsaid and Martin 2003] is an educational tool that raises information about “web bugs, i.e., images or scripts placed on web pages to track users. It signals both audibly and visually when a web bug is detected. The goal of the tool was to inform journalists and policy makers about web tracking, nudging them to act. More recent web tracking privacy tools, such as Ghostery (www.ghostery.com/), PrivacyBadger (<https://www.eff.org/privacybadger>), and Lightbeam (<http://www.mozilla.org/lightbeam/>), provide notifications about tracking and some of them can block specific trackers.

Tan et al. [2014] studied the effects of developer-specified explanations in apps' permission requests on iOS. They found that permission requests that included an explanation were more likely to be approved. However, the content of the provided explanation had no observable impact in their study. In a mobile field study, Shih et al. [2015] found that participants shared the most when requests contained no information about the requesting mobile app or the purpose of access. When a purpose explanation was given, study participants disclosed less for purposes that were vague or only benefitted the developer, but disclosed more when the purpose was specific and benefitted them.

Notifications can take forms other than simple text or icons. They can be visceral notices that allow a user to experience information [Calo 2012]. For example, digital cameras make a clicking sound similar to analog cameras. These sounds alert subjects that a photo is being taken and may help avoid surreptitious photos. Groom and Calo [2011] investigated how visceral notices impacted privacy concerns and disclosure. The study included several notices designed to increase privacy concern (anthropomorphic agents; personalized information based on IP or history), and one that was designed

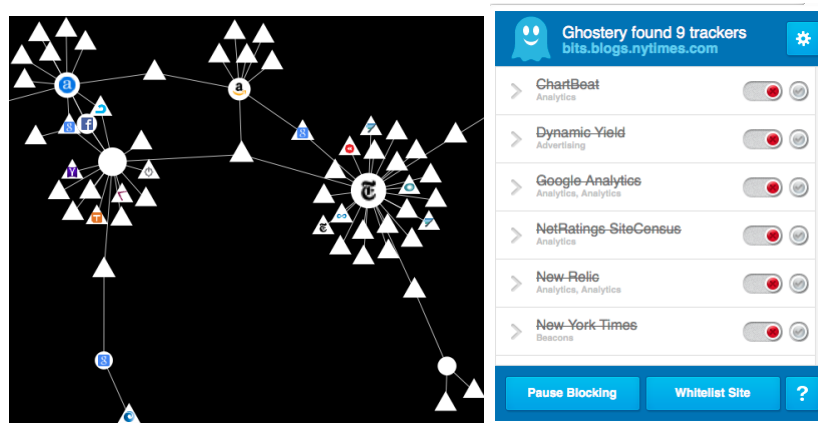


Fig. 2: Browser plugins providing feedback on online tracking. Lightbeam (*left*) visualizes connections that trackers can make across websites. Ghostery (*right*) shows the trackers present on the visited webpage.

to weaken concern, informality. They found that a notice that appealed to users' social responses could elicit more privacy-protective behavior than traditional notices.

The Privicons project [König and Schallaböck 2011] proposed the inclusion of privacy preferences in email messages. The recipient's email client interprets the encoded preference information and displays icons that convey how the sender expects the email to be handled, e.g., keeping it secret, or deleting it after reading. Könings et al. [2011] propose a system that enables users to share information over social networking services without the service provider learning the information. They employ Privicons as privacy nudges to make recipients aware of how users expect the shared information to be used. While those icons are non-binding, they provide the recipient with an indication of the sender's expectations. PriPref Broadcaster is a mobile app that employs a similar approach to enable mobile users to make others in their physical proximity aware of their privacy needs and preferences [Könings et al. 2014]. For instance, users can anonymously broadcast their privacy preferences for appearing in photos or on video to other nearby users of the app.

While educating users and providing clear concise information is a desirable goal, information and education are no guarantees that users can achieve their stated preferences and maximize their welfare in privacy-sensitive scenarios. Even useful and salient information may be ignored when the hurdles and biases discussed above enter into the decision making process.

3.2.3. Feedback for privacy and security. Leon et al. [2012] demonstrated the importance of feedback with a usability study of tools for limiting online behavioral advertising. They found that a lack of feedback can lead to suboptimal tool configuration. Furthermore, participants sometimes thought they had turned on settings that were not turned on—an issue that Clark et al. [2011] also identified in the context of encrypting sensitive radio communication.

Tools have been built to increase awareness about online tracking by providing feedback about how users' online activities are being tracked by third parties. Figure 2 shows the different visual representations of online tracking provided by the Lightbeam and Ghostery browser extensions, respectively. Such tools enable users to understand and reason about online tracking and its potential consequences.

Usable and in-time privacy notices make websites data practices more salient and can nudge Internet users to use websites with more privacy-respectful data practices. Privacy Bird was built with the idea of warning users whenever a website engages in practices that do not align with users' expectations [Cranor et al. 2006]. Relatedly, personalized examples can improve the effectiveness of risk communication for security and privacy decisions [Harbach et al. 2014].

Password meters, shown as the user creates the password, can be effective nudges in helping users create more secure passwords by providing immediate feedback on password strength. Ur et al. [2012a] found that users who were shown password meters created longer passwords, and with some meters also created passwords that were more difficult to guess. However, users were also likely to consider the stricter meters as "annoying," suggesting that nudges that try to push users' existing mental models might not be as effective. Komanduri et al. [2014] built an in-time feedback tool that informs users about the predictability of their passwords as they create them and effectively encourages the creation of stronger passwords.

Several studies have investigated feedback in the context of location sharing applications [Sadeh et al. 2009; Lederer et al. 2004; Tsai et al. 2009; Iachello et al. 2005; Raento and Oulasvirta 2005; Jedrzejczyk et al. 2010]. This includes interfaces that allow users to see who requested their information and possibly where they were at the time. Sadeh et al. [2009] and Tsai et al. [2009] both report significant changes in privacy settings by users who are given such feedback. In both studies, users are found to refine their settings and selectively open up, thereby deriving more value from the location sharing app while having a better sense of control over their privacy. Almuhimedi et al. [2015] have shown in a field study how daily nudges, informing Android users about the frequency with which their mobile apps access sensitive data, can motivate users to review and modify permission settings. These experiments have shown that simple nudges can have a powerful effect in helping users converge towards settings that are better aligned with their privacy objectives.

Social navigation has also been used to provide users with feedback about the suitability of their privacy settings. Social navigation leverages the user's community to provide guidance for the user's decisions. In the context of privacy, control interfaces can be enriched with information about friends' privacy settings. Patil et al. [2011] enhanced the privacy settings dialog of an instant messenger to show the dominant disclosure settings of the user's contacts. Besmer et al. [2010] extended the interface for adding applications on Facebook. For each information item to be shared with a new application, their interface shows how many other friends have given this application access. Goecks and Mynatt [2005] visualize how many other community members have visited a website and if they blocked its cookies. They note that social navigation facilitates learning from the behavior of others but also carries the risk of just following what others do. They suggest engaging experts to help guide behavior.

3.3. Nudging with Presentation

Many of the hurdles discussed in Section 2 have to do with how information and choices are presented. In this section we discuss a number of related yet nuanced presentation concepts that have been used effectively in nudge design. Framing of provided information and choices can either mitigate or increase behavioral biases such as loss aversion, unrealistic optimism and overconfidence, and representativeness heuristics. The order in which alternatives are presented can affect users' choices, due to anchoring and adjustment bias. If one of the options presented is more salient than the others, availability bias might result in incorrectly weighting this option. This bias can be put to advantage, however, when salient examples are employed in user education to reduce overconfidence. Finally, users' cognitive limitations will cause failures when

dealing with a large or overly complex set of alternatives. When such choices are required, presenting alternatives in a structured way can mitigate unforeseen effects of simplification strategies and representativeness bias.

3.3.1. Framing. Framing is concerned with how benefits and risks of choices are communicated. Framing effects have been demonstrated consistently across several studies [Levin et al. 1998]. Information presented to users might be framed so that users can easily assess the risks involved in a particular transaction. In situations where users underestimate risks due to overconfidence bias, system designers could compensate for this by framing the risks in a way that exaggerates them. Relying on loss aversion bias, this could be achieved by framing the risks as a clear loss for the user. A classic example of overconfidence is speeding while driving; drivers underestimate the risks of speeding and overestimate the associated time savings [Pe'er 2011]. One design feature that can be used to emphasize the risks taken when driving fast is the use of visual illusions [Hansen 2012; Thaler and Sunstein 2008]. For example, virtual speed bumps or potholes painted on dangerous roads can increase the perceived risk and cause drivers to slow down.

To mitigate the effect of representativeness heuristics and to help users assess risks more accurately, information should be framed in a way that supports users in distinguishing between the representativeness of an event and its probability of occurrence.

3.3.2. Ordering and saliency. The order in which options are shown or required to be followed impacts users' behavior. Research on anchoring and adjustment bias suggests that the first piece of information offered affects a user's perception of the remaining options. For example, the order in which dishes are listed on a restaurant's menu likely impacts what guests order and how much they will be paying [Wisdom et al. 2010].

Another presentation component related to both framing and ordering is saliency, which can either mitigate or leverage availability and overconfidence biases by making people focus on and think about certain elements, rather than others, when making a decision. For example, saliency can reduce overconfidence and assist teenagers in assessing the risks of speeding while driving. The Pennsylvania driver's manual includes photographs of accidents that were caused by speeding, in combination with the message, "The most frequent crash type for 16-year-old drivers in Pennsylvania is a single-vehicle, run-off-the-road crash. From 2007 to 2009, there were 2,969 crashes and 28 people killed" [Pennsylvania Department of Transportation 2013]. The saliency and vivid nature of this picture and message can help teenagers reduce overconfidence about their driving skills. Other government measures have purposely used saliency to effectively raise awareness of risk. The Tennessee Department of Transportation displays weekly statistics of fatalities on digital road signs along interstates and highways with the intention of increasing awareness about highway deaths and of saving lives [Tennessee Department of Transportation 2014].

The availability bias can also be leveraged by system designers for the benefit of users. As a simple example, a wet floor warning including a human silhouette slipping backwards can help users to visualize and internalize the risks of stepping on a wet floor.

Saliency does not always succeed in overcoming availability biases. After the 9/11 terrorist attacks, many airlines were struggling since people were wary of flying. Although the probability of airplane-based terrorist attacks is very low, travelers were affected by an availability bias caused by the attack that was leading them to assess the risk of flying to be much higher than it really was [Huddy et al. 2003; Floyd et al. 2004]. People perceive higher risks for situations that are spectacular, rare, and beyond their control. For example, more people die from car accidents or food poisoning every year than from terrorist attacks, but people have higher perceptions of risk for

the latter [Schneier 2007]. We can think of additional airport security measures by the US Government following the 9/11 attacks (which were very salient for every user) as attempts to mitigate the perceived risk of flying. However, research has shown that people tended to fly less as a result of such security measures [Blalock et al. 2007], arguably because of the added inconvenience.

3.3.3. Structure. Whereas framing, ordering, and saliency can be used to address decision making hurdles affecting a specific choice or behavior, providing structure is an essential strategy for addressing hurdles in complex decisions requiring the evaluation of multiple options and alternatives. How information and options are structured can mitigate the negative effects of bounded rationality biases, simplifying strategies, and representativeness heuristics. For example, an individual in the market for a new credit card may end up choosing the card with the lowest Annual Percentage Rate (APR)—using a simplifying strategy driven by an availability heuristic—and ignore many other fees and hidden costs. Similarly, users might select suboptimal cell phone plans. Decision making in complex situations with multiple alternatives can be supported by facilitating structured evaluation and comparison of those alternatives, or providing better search options. Examples of structured presentation are tools that help users understand the total cost of available alternatives given their specific usage patterns [Financial Consumer Agency of Canada 2014] or tailored search, recommendation, or filtering options. The representativeness heuristic can also be leveraged to assist users in making decisions. For example, an online music store that categorizes music into genres could provide an iconic song as a sample for each genre, thereby helping users to create a mapping of song attributes to their musical interests as a guide to find what they are looking for.

3.3.4. Presentation nudges for privacy and security. Braunstein et al. [2011] found that wording a survey question to remind users that they are revealing sensitive information impacts how much they are willing to reveal. In particular, a reminder that data is sensitive will result in less disclosure. This indicates the importance of framing disclosure decisions in light of associated risks, as well as the effectiveness of making sensitivity and risks salient to individuals. In addition to wording, other parts of a system or an interface may frame privacy choices in a way that nudges the user. For example, Brandimarte et al. [2013] found a control paradox in the responses of users deciding to whether to share personal data. When the presentation introduced sources of uncertainty about whether their profile would be published online, the users revealed less sensitive information than users who were certain their profile would be published. When users perceived that they had more control over their data, they underestimated the sources of risk that they didn't have control over. This further indicates the need for structured presentation in sharing decisions, especially in complex sharing environments with multiple recipients, such as social media platforms.

The way that choices are framed and structured also impacts how people balance costs and privacy concerns. Egelman et al. [2013] found that people were willing to pay more for Android apps that requested fewer permissions when they had several options for price and permissions. However, when only given one choice, participants were not as willing to pay for privacy. Therefore, applications that only give users the option of installation with a fixed set of permissions may be nudging users away from privacy, indicating the need for more structured choices, such as systems that provide users with selective permission controls [Almuhimedi et al. 2015].

Bravo-Lillo et al. [2013] found that by making the name of the software publisher more salient in warning dialogs they could nudge people not to download suspicious software. While visual nudges were somewhat effective, the most effective nudges required users to interact with the salient publisher field. In this example, saliency is

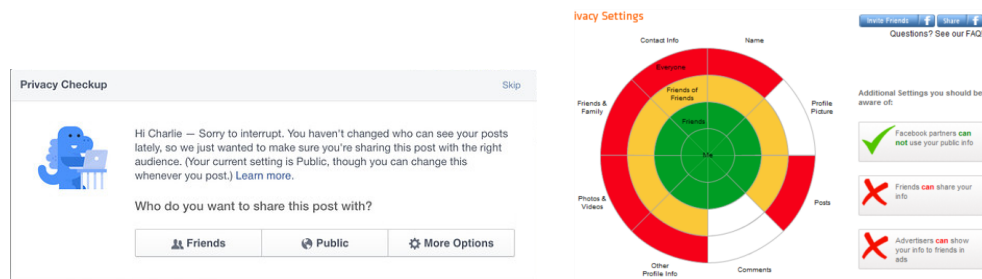


Fig. 3: Facebook-related privacy nudges. Facebook’s privacy dinosaur (*left*) reminds users that they are about to post publicly and nudges them to check their privacy settings. PrivacyDefender (*right*) visualizes the audience of information on Facebook.

used without framing: the user’s attention is drawn to the publisher name without emphasizing risks or benefits of specific publishers. Wang et al. [2014] evaluated a modified Facebook user-interface that combined saliency and ordering effects. Their interface provided feedback about a post’s audience and allowed users to modify or cancel posts within 10 seconds. They found that users perceived the new interface as helpful in preventing unintended disclosures. In 2014, Facebook deployed a related privacy tool, shown on the left-hand side in Figure 3, which warns users before posting publicly when they haven’t posted in a while, but their current privacy setting is to share publicly [Albergotti 2014]. The associated risk is made salient, the ordering of the warning before the post is published reduces subsequent regret, and the available options are framed and structured in a way that forces an explicit selection of the post’s audience instead of just dismissing the warning. Park et al. [2014] propose a game-theoretic approach to assess oversharing risks which could be used to inform when such nudges should be shown to users. Gao et al. [2012] built FreeBu, a tool that assists Facebook users to configure friends groups based on attributes that were found to be commonly used among Facebook users. Similarly, Mazzia et al. [2012] built a tool that helps Facebook users understand the visibility of their profiles among different groups. Previously, several third parties had introduced software tools that ease the management of Facebook privacy settings. The example shown on the right-hand side in Figure 3 helps people visualize who can see their Facebook activity and personal information [Wondrasek 2010]. All these tools aim to help users understand and properly manage their privacy settings by presenting the information in a more salient, timely, and structured manner.

3.4. Nudging with Defaults

As discussed in Section 2, the status quo bias suggests that individuals often stick with default choices. As a result, system designers have the opportunity to influence people’s choices by carefully setting defaults to best serve individuals needs and align with their privacy expectations. Moreover, in complex systems, it is unreasonable to assume that individuals will make an optimal decision for every available option or even make a choice for all of them. The nudging approach suggests that system designers need to consider users’ limitations as well as their expectations, to decide what options to provide and whether a particular default is preferable over a required choice approach in which the user must make an explicit decision before using the system.

3.4.1. Defaults in general. System defaults have been shown to have a significant impact on how individuals make decisions in many areas, including retirement savings, organ donation, food consumption, and health care [Madrian and Shea 2001; Choi et al. 2004;

Johnson and Goldstein 2004; Downs et al. 2009; Halpern et al. 2007]. For instance, many employees believe that they should be saving more for retirement, yet many of them fail to enroll in 401(k) contribution plans [Beshears et al. 2009]. Research found that participation rates in one contribution plan under the opt-in approach were only 65% after 36 months of employment, but when automatic enrollment was the default choice enrollment of new employees increased to 90% [Madrian and Shea 2001; Choi et al. 2004]. This example demonstrates how the status quo bias can be leveraged to assist users in achieving the outcomes they desire by using appropriate defaults.

In the context of computing and information systems, users can be overwhelmed with an excessive number of configuration or customization options and may lack the required knowledge to make the correct decision for each option. Mackay [1991] investigated how 51 users customized software applications over a period of four months, and found that lack of time and knowledge were the two principal barriers to customizing software. Specifically, participants decided to spend most of the time “using” the software to perform their tasks as opposed to customizing it. Carefully considering the defaults during the design of information systems can have a significant influence on people’s interaction with them, especially with respect to privacy and security.

3.4.2. Defaults in privacy and security. Defaults for privacy and security settings are of particular importance because they mandate how the application protects and shares users’ information. System designers should set defaults that protect the information of conventional users while allowing expert users to customize the application to their specific needs. Ontario’s former privacy commissioner Ann Cavoukian has been a strong advocate of privacy by design, suggesting that privacy settings should be enabled by default to guarantee that users’ privacy is protected, even if they don’t explicitly act to protect it [Cavoukian 2009]. Acquisti and Gross [2006] found that although the Facebook default settings at the time of their study allowed profile information to be publicly searchable, the majority of the surveyed Facebook users had not changed these settings, providing more evidence that users hardly change default settings—with potentially detrimental effects to privacy.

Lai and Hui [2006] studied the difference between opt-in and opt-out for receiving websites’ newsletters, as well as the role played by default settings. For example, in the opt-in option, users were prompted with either the text “Please send me newsletters,” which is unchecked by default, or the text “Please do not send me newsletters” which is checked by default. They found that participants were more likely to choose the default option. However, they found a moderation effect of existing privacy concerns, in that users with high privacy concerns were less likely to be swayed by the default setting. The Microsoft privacy guidelines for software development recognize the importance of privacy protective defaults and recommend a “no collection” option to be pre-selected when collecting user opt-in consent for data collection [Microsoft 2008].

In 2012, Microsoft announced that it would turn on the “Do Not Track” option by default in the release of its browser Internet Explorer 10. All users who did not change the default setting would send a flag to online advertisers that they did not want their browsing history to be tracked. The ensuing debate showed how crucial the default setting is, as advertisers balked and stated that such a default setting could not be honored [Angwin 2012], as it may not reflect the user’s actual preference.

Shilton et al. [2008] discuss how to protect privacy in urban sensing studies with mobile phones. They suggest to arrive at appropriate default settings for a project by facilitating a dialog between users and designers through participatory design.

Goldstein et al. [2008] distinguish between “mass defaults, in which everyone gets the same default, and personalized defaults tuned to the user’s needs. However, personalized defaults require some information about the user. They should therefore be

created in a way that respects the user's privacy. For example, Fang and LeFevre [2010] suggest a privacy wizard for social networking that asks users to confirm privacy settings for individual friends. Responses are used to build a classifier to categorize the remaining friends. This is an example of creating a personal default after a minimum of configuration has been completed. Wilson et al. [2013] investigated the effect of privacy profiles that help users specify their location sharing preferences and found that these profiles can considerably affect users' sharing behavior. Liu et al. [2016] leverage similar profiles to provide users with personalized recommendations for mobile app permission settings.

3.5. Nudging with Incentives

Loss aversion and hyperbolic discounting can be ameliorated with incentives, in the form of either punishments or rewards. Providing the right incentives can help mitigate hyperbolic discounting bias that often prevents users from appreciating negative long-term consequences. For example, people tend to underestimate the future consequences of CO₂ emissions and too often make inefficient use of motorized vehicles. High Occupancy Vehicle (HOV) lanes are intended to reward those drivers who are willing to carpool by enabling a faster commute, and in the process help reduce CO₂ emissions and the number of cars on the road.

Loss aversion can be leveraged to shape the perception and effectiveness of rewards and punishments. Namely, framing a punishment as a loss might have a stronger effect on users' behavior than framing it as a simple penalty. Similarly, giving users an a priori reward that can be lost if the expected behavior is not met is potentially more effective than simply offering an a posteriori reward. Brennan et al. [2009] studied the effect of incentives in a virtual search task, in which participants were asked to identify weapons using luggage scanners. They found that punishments (losing points) led to better performance (accurate detection of weapons inside luggage) than rewards (gaining points).

3.5.1. Rewards and punishments. Non-financial incentives and rewards involving social systems and peer pressure can have strong effects on users' behaviors [Lindbeck 1997]. The use of virtual badges can be a strong reward for users who exhibit a desired behavior. This idea is similar to product reviews in which good products receive ratings that make them look better against competitors. Research has shown that social incentives act as an important motivator for using location sharing applications [Lindqvist et al. 2011]. For example, Foursquare, a location sharing application, grants "mayor" badges to those users who check in the most at a given place. Foursquare further incentivizes users to "check in" by providing coupons and giving access to custom reviews of the place where a user has just checked in.

Increasing the cost or difficulty of choosing specific configurations, even at the simple level of requiring multiple confirmations, is a "punishment" that could prevent non-experts or inexperienced users from selecting risky settings. On the other hand, secure or less risky configurations can be made easier to select, encouraging users to choose them. Punishments and rewards can further be applied in response to users' choices.

Online users or users of mobile devices have currently no way to measure the costs of their decisions about privacy or security. Security costs may be borne by the company or enterprise, as opposed to the individual. As discussed by Herley [2009], ignoring security risks may actually be a rational decision for a user who is unlikely to experience consequences. However, individual insecure behaviors cause negative externalities that can affect other Internet users. Here, a nudge may include information about costs and negative externalities of insecure behaviors in a salient yet non-interruptive manner.

3.5.2. Costs of privacy and security. A particular aspect of security and privacy with regard to incentives is potential inconvenience (or cost) associated with certain privacy or security measures. Egelman et al. [2010] studied users' willingness to tolerate increased costs (in this case, time) for security benefits. They found that users who were exposed to a delay and who did not have an explanation of the security reasons for the delay were likely to cheat or drop out of the study. However, this suggests that when costs are necessary to improve security, they can be ameliorated with detailed explanations. Grossklags and Acquisti [2007] found that users will set different prices on privacy depending on whether they are asked to pay for protection or are offered payment for the same information. They found that the "willingness-to-accept" price, at which individuals are willing to sell information, is substantially higher than the average "willingness-to-protect" price, which individuals would be willing to pay to protect their information. Thus, the method used to reward users for privacy or security may act as a nudge that influences their decision.

3.5.3. Incentives for security. Incentives can be used to improve security as part of corporate strategy. Users who handle email attachments in a careless manner can be punished by restricting their email accounts to official communications. A user who is more cautious with email attachments can be rewarded by allowing use of the corporate account for personal emails. Brustoloni and Villamarín-Salomón [2007] developed audited security dialogs, in which users were held accountable for their decisions to open email attachments. Those who took unjustified risks could be "subject to a variety of sanctions, such as being unable to use the application for increasing periods of time, having to pay increasing fines, or having to pass remedial training." Their user study found that these dialogs resulted in significantly fewer unjustified risks.

Since 2000, the economics of information security has been a popular field of research. This field started with the observation that liability for security failures is often diffused and that this creates a moral hazard [Varian 2000; Anderson 2001]. Since then, the field has expanded to study many security and privacy issues from an economic perspective [Moore and Anderson 2011]. In one study, Christin et al. [2012] paid participants a small amount to download an unknown executable onto their computer. They increased the financial incentive each week, in order to determine how that impacted participants' security decisions. At the maximum payment of \$1.00, 43% of participants were willing to download and execute the file.

In the corporate environment, Herath and Rao [2009] studied employees' security behaviors and found that subjective norms, peers' secure behaviors, and perceived effectiveness of their actions were incentives that motivated secure behavior.

3.6. Facilitating Reversibility and Error Resiliency

Human beings are prone to error and can benefit from designs that ease error correction. Error-resilient systems aim to protect people from inadvertent errors that can be caused by decision making hurdles. For example, framing effects or hyperbolic time discounting can lead individuals to make purchases that they do not really want. As an effort to protect consumers who might purchase products in a "hot" state, the Federal Trade Commission's "cooling-off" rule gives customers in the United States three days to cancel purchases of \$25 or more when they purchase an item in their home or at a location that is not the seller's permanent business address [Federal Register 2004].

3.6.1. Forced action or automated completion. Nudging systems can mitigate post-completion errors originating from tasks that have both primary and secondary goals by either forcing people to take an action or automatically completing secondary tasks. These nudging systems are present in many aspects of our lives, such as driving. Drivers' primary task is to drive to their destination. However, there are many as-

sociated secondary tasks that some drivers tend to forget, such as switching on their headlights. Light sensors that automatically turn on the headlights in low light conditions provide automated completion of this secondary task. Systems also alert users to fasten seat belts or to turn off the headlights when the engine has stopped. The latter is a good example of a situation prone to post-completion errors, as the driver has achieved the primary goal.

Similarly, systems have been designed for automatically turning off the lights in office spaces when nobody is around, switching off electronic equipment when risky situations arise, and to alert users when further action is needed to successfully complete a secondary task. Many bank ATMs now require users to remove their bank card before money is dispensed or even before any request can be performed, in order to help users avoid forgetting their cards in the ATM after receiving the requested cash.

3.6.2. Reversibility for privacy and security. In the field of computing systems, system recovery features offered in some operating systems, such as Microsoft's System Restore [Microsoft 2014] or Apple's Time Machine [Apple Inc. 2014], allow users to revert to a stable configuration in cases where software has been corrupted, the system has been affected by a computer virus, or data loss has occurred. Reversibility can provide fault protection against security or privacy errors. This may include an option to "undo" an action, in which a user moves from the current, undesired state, back to a previous state [Loregian and Locatelli 2008]. For example, Facebook, Twitter, and Gmail allow users to delete or "unsent" regrettable posts, tweets, or emails to help them manage their privacy after accidentally revealing too much information. A user study on email retraction found that retracting an email was preferred in cases where it had not already arrived in the recipient's inbox. In other cases, including when the recipient had already seen the email, an apology was preferred [Cabitza and Loregian 2008]. Besmer and Lipford [2010] studied how people respond to Facebook photo tags. They found that individuals may not want their photos shared on the social network by others. A photo tagging tool built by the researchers allowed users to request that their friends remove tags from photos. They found that users actually preferred an undo tag over negotiating with their friends to remove tags.

3.7. Timing of Nudges

As the examples provided in the previous sections suggest, nudges can be applied at different points in time. For example, feedback can be provided before and during system usage, and feedback provided at the appropriate time may be more likely to influence behavior. Similarly, incentives can be provided at different points in time. Depending on the circumstances, even the system structure might be required to change over time to better assist users. Timing is a dimension that should be considered in combination with the other nudge dimensions discussed above. We point out a couple of scenarios in which timing can be a particularly effective aspect of the nudge.

Some tools provide nudges at specific times, when the user may be in a "hot" state or under the influence of alcohol. Examples include the Social Media Sobriety Test [Webroot 2010] and Mail Goggles for Gmail [Perlow 2008]. Both allow the user to select specific hours of the week in which posting on social network sites or sending emails is blocked until the user completes a dexterity or cognitive test, such as solving a small math problem. Just-in-time notices providing relevant, clear, and contextual information can help mitigate information asymmetries [Patrick and Kenny 2003; Schaub et al. 2015]. Balebako et al. [2011] discuss further nudging techniques applied at different points in time that also enable users to bypass them if needed.

Other nudges, such as privacy indicators, may have a greater impact based on when they are seen. Egelman et al. [2009] investigated whether participants in a lab study




Fig. 4: Twitter nudges users to check their application access settings, right after they change their password—making it more likely for users to act on the nudge.

were more likely to pay a premium for websites with good privacy practices. They found that the timing of the privacy notice was important; viewing privacy indicators before visiting the website had a greater impact than seeing the indicators once the users already arrived at the website. Balebako et al. [2015] found that individuals pay more attention to privacy notices shown in the context of a mobile app compared to in the app store.

In the context of security, Figure 4 shows a Twitter nudge that encourages users to review the applications that have access to their accounts. The nudge is shown at a time when users are likely to be receptive to security advice—when they have just changed their password. The blue button encourages the user to proceed and the presented list makes relevant information about each application salient.

4. NUDGING AWAY FROM PRIVACY AND SECURITY

So far, we have discussed how interventions and nudges can ameliorate the negative effects of decision making hurdles. Some of the presented nudging approaches even exploit cognitive and behavioral biases to promote more beneficial security and privacy behavior. However, these biases and cognitive hurdles can also be exploited to achieve the opposite effect. Marketers have learned to exploit consumers' representativeness bias to better sell their products. For example, consider a commercial featuring a top model wearing a purple bikini and advertising a weight-loss supplement. Although it might be unlikely that this woman used the advertised product, many consumers, subject to representativeness bias, will likely associate the supplement with a slim and beautiful body, leading them to buy it.

In the context of privacy and security, some interventions exploit decision hurdles to nudge users to disclose more information or behave detrimentally to their interests [Bösch et al. 2016]. Below, we discuss examples of exploitative nudges employed by industry or cyber-criminals.

As we have discussed already, most users never change default settings. While privacy-friendly defaults can ameliorate this effect, defaults have also been used by companies to encourage users to make unintended purchases or disclose personal information for the company's benefit [Dinner et al. 2011; Goldstein et al. 2008]. An example of such nudging with defaults is the installation wizard for Java, which is designed such that the default option installs the Ask.com toolbar in the user's browser and even sets Ask.com as the default search provider, even though this toolbar plays no part in the smooth functioning of the Java package.

Framing and presentation can further promote disclosure by creating choice architectures that influence users away from privacy. The ease or attractiveness of one option may nudge users towards choosing it—even if the available options are otherwise



Fig. 5: Pandora’s choice architecture highlights the option to keep the profile information public, using wording that implies defaults as well as more prominent button colors and positioning.

equal. Many existing choices are designed so that the most obvious, brightest, or easiest option may discourage information privacy. For example, in promotional emails, the option to unsubscribe is placed at the bottom, in small text and bland colors.

An example is Pandora’s sign-up dialog from 2010, shown in Figure 5. This dialog asks users if they would like to “keep their profile page public, or “make it private. The text implies that keeping the profile public is the default option. The “keep public button is also in a brighter color, making it more attractive, and on the right side of the dialog box, a position that is typically used for buttons that imply forward movement. Another example of design pushing users away from privacy is Facebook’s app permission request dialog, which was changed to have only one button stating “Play Game” instead of options to allow or deny the app access to personal information [Charkham 2012]. Similarly, using data from a 6-year longitudinal study of 5,076 Facebook users, Stutzman et al. [2013] document changes in Facebook interface between 2009 and 2010 that may have nudged users towards more disclosure. When updating to Windows 8.1, users were made to believe that creating a Microsoft account is required as part of the installation. The account creation requires providing detailed personal information to Microsoft. The first page provides no option to use an existing email account from a different provider. The only way to avoid creating a new Microsoft account is by clicking “Create Account”; the following page then offers the choice to use an existing (non-Microsoft) account. In addition, the option to use an existing account is displayed in a color that blends with the background, thereby making it inconspicuous.

Other websites, such as LinkedIn, display progress bars on a user’s profile indicating “profile completion” in order to nudge users to disclose more information. Some free mobile apps lead users away from privacy by making the benefits of sharing information clear, while hiding potential risks. The Facebook News Feed may create the impression that others are constantly sharing information, which could incentivize the viewer to post more to comply with group norms.

Some privacy nudges may also be ambiguous. For example, Flickr’s privacy indicators shown in Figure 1 consist of two inconsistent nudging aspects. The text clearly conveys who can see the photo. Yet, the icon color could potentially act as a nudge to share photos more openly. Making an image public is indicated with a green icon, restrictive sharing is indicated by an orange icon. Users may perceive green as the preferable choice, if they are conditioned by the traffic light metaphor. However, it is not clear if this ambiguity is intentional or due to an inconsistency in perspective on that metaphor: the designer assigning colors according to the flow of information (“green” = free flowing) versus users mapping colors to benefit (“green” = good for me).

While companies, especially those providing seemingly free services, exploit decision making hurdles to nudge their users to disclose more information, cyber-criminals employ similar strategies for phishing attacks, scams, identity theft, and fraud. Stajano and Wilson [2011] analyzed hundreds of fraud cases and identified a number of cognitive and behavioral biases as likely explanations for why users become fraud victims.

The widespread availability of industry nudges towards more disclosure, or cyber-criminals' nudges towards increased security risks, underlines that interventions targeting users' decision hurdles can not only nudge towards beneficial behavior but can also influence users to behave against their interests. The Dark Patterns website (www.darkpatterns.org) maintains a list of deceptive user interfaces, some of which we have discussed, with the aim of eliminating such practices. The dichotomous potential of nudging raises ethical implications, which we discuss in the next section.

5. TO NUDGE OR NOT TO NUDGE

Nudging can be most useful in contexts where users are confronted with uncertain, subjective trade-offs and scenarios that involve weighing possibly conflicting considerations. Many security and privacy decisions fall into that category, as users often need to weigh privacy and security considerations against other priorities or even against each other. In the workplace, these priorities often have to do with an employee's ability to perform (e.g., accessing a public WiFi network to increase one's productivity, using a USB device to share a presentation, or clicking on a link in a customer email). In private contexts, conflicting priorities arise from considerations such as the desire to use a new mobile app, desire for social interaction, or peer pressure to share information on a social media site.

However, because in many of these situations there is no obviously right or best choice, it is fair to ask how far nudges should go in attempting to influence a user's security or privacy decisions, and whether it is appropriate to nudge users at all. In this section we discuss some of the design and research challenges associated with security and privacy nudges: Should we nudge at all (Section 5.1)? If yes, towards what outcomes (Section 5.2)? And who should implement those nudges (Section 5.3)? We argue for approaches in which users are nudged towards decisions that are consistent with their own preferences or objectively improve their welfare. In addition, we highlight additional research challenges in nudging. While not exhaustive, these discussion points can serve as a meaningful starting point for designers, developers, and researchers to evaluate the acceptability of nudging for a given purpose, and the appropriateness of nudge designs.

5.1. Every Design Choice Is a Nudge

A first design (as well as research) challenge is whether we should nudge users at all. Sunstein [2012] discusses various arguments against nudging under any circumstance, and concludes, however, that the most powerful of those arguments requires unrealistic assumptions about human behavior. Similarly, Selinger and Whyte [2011] argue that "social anxieties associated with nudges" are overstated, although ethical considerations remain. To address those considerations, we will first observe that, whenever a designer creates a system, every design choice inescapably influences the user in some way. As Thaler and Sunstein [2008] describe, it is necessary to choose, at some point, between available design options. They provide the example of a cafeteria, in which the layout of the food influences user choice. The food must be set out. So, a nudge of some type is inevitable. Accordingly, Thaler and Sunstein [2008] refer to those involved in the design of nudges as "choice architects," and to the process and end-product of their design as "choice architecture." Any choice architecture, whether

it was intentionally designed to affect users' behavior or not, will impact how users interact with a system.

Just as physical design impacts choices, digital design decisions can affect how online users behave. Thus, we argue that most User Interface design decisions can be viewed as nudges of some kind. Small adjustments to website design can cause dramatic behavioral changes. For example, minor changes to the sign-up button and family photo on Barack Obama's 2008 presidential campaign website resulted in a 40% increase in sign-ups [Christian 2012]. Similarly, the choice of defaults on a webpage can cause drastic differences in user behavior. When a European rail website started including seat reservations as part of the ticket by default, instead of as an additional option, the number of tickets that included seat reservations rose from 9% to 47% [Goldstein et al. 2008].

Accordingly, a key question for designers and researchers may not just be to what extent it is ethical to nudge users, but whether a system's inevitable nudges are in fact ethical. Some design decisions may benefit the system owner more than the user, and other design decisions may be intentionally malicious, as discussed in Section 4. Are the nudges leading users to make decisions or adopt behaviors that are not in their best interests? To what extent should we expect system designers to exercise restraint? To what extent should they be held accountable for the way in which their interfaces may nudge users in one direction or another? And conversely, to what extent can system designers presume that they know what is best for users and attempt to nudge users accordingly?

Beyond obviously malicious uses of nudging, all nudges have the inherent quality of encouraging an individual to make a particular choice. Thus, nudge design assumes that a particular entity, for example a policy maker or system designer, is in a better position to decide what choice the individual should make [Glaeser 2005; Beaulier and Caplan 2007]. The ethical implications of nudging are complicated by a number of factors [Ly et al. 2013]. It is difficult to optimize the benefit of nudges for all users. Nudges may be optimized for individual or societal benefit, or their impact may vary for different types of users [Johnson et al. 2012; Buckley 2005]. This is further complicated by nudge designers having to decide how transparent to make the nudge to users, and by the fact that the impact of nudges can be difficult to evaluate [Bovens 2009; Johnson et al. 2012].

Absent nudging efforts, the design process involves complex, often interrelated choices [Norman 2013]. Thus, nudges for security and privacy must also be considered in the context of other design elements and concerns in order to create a user-friendly, cost-effective system overall.

5.2. Evaluating Nudges and Desired Outcomes

A second challenge relates to the direction of the nudge: towards what goal, or outcome, should the nudge be aimed? Proponents of nudging argue that policy makers should implement nudges when the net benefit for "irrational" individuals exceeds the aggregate costs (e.g., changing a default setting) both to rational individuals and to any other affected entities such as businesses or taxpayers [Camerer et al. 2003]. However, arguments centered on nudging when the net benefit to individuals is positive have a significant shortcoming: What criteria should be used to measure benefit? Loewenstein and Haisley [2007] suggest multiple potential criteria, such as experienced utility or self-reported happiness. However, the latter criterion is limited by policy makers' ability to measure individual happiness and the influence of habituation on self-reported happiness. In contrast, contexts in which strictly dominant alternatives exist may be good opportunities for nudging. For example, nudging employees to invest in a 401(k) retirement plan with company matching is reasonable, assuming that all else being

equal, individuals prefer more money over less. Another approach may be to simply allow individuals to opt-in to nudge interventions that target sub-optimal decision making [Ly et al. 2013]. The limitation in this approach is that, in some cases, those in need of such remedies may be unlikely to recognize or acknowledge the need for decision support.

Moreover, the context of privacy and security may also present unique challenges with respect to nudging. While other contexts, such as saving money or eating healthier, may have relatively clear desired outcomes to nudge towards [Downs et al. 2009], privacy and security decision making may not necessarily share this feature, as the costs of intrusions are often intangible or difficult to measure. An overly simplistic approach may be to minimize disclosure by individuals or encourage users to always implement stringent security controls. This approach, however, is flawed, as disclosures by individuals, while increasing privacy risks, may also lead to some economic, social, or personal gain. Similarly, overly stringent security mechanisms may lead to limited functionality and hassle for users. Alternatively, we propose two potential desired outcomes for guiding the implementation of privacy and security nudges: minimizing regret and aligning individuals' behavior to their stated preferences.

Minimizing regret. We may consider instances in which disclosures by individuals are likely to be regretted in the future as nudge-worthy. For example, disclosures made under the influence of alcohol or drugs, using offensive language, or discussing contentious political issues could be appropriate instances for privacy nudges as they have a higher likelihood of being regretted in the future [Wang et al. 2011].

Aligning behavior with stated preferences. Another potential desired outcome for privacy nudges is to better align individuals' behavior with their stated preferences. For example, we may find that a majority of individuals are concerned about disclosing their political or religious affiliation to potential employers. Thus, nudges that influence users to limit such disclosures to employers, or nudges encouraging users to limit access to such disclosures, may be justified. Along these lines, we may also consider contexts ripe for nudging in which individuals incur a steep privacy or security cost in return for limited functionality.

5.3. Who Should Implement Privacy and Security Nudges?

Even if we are able to identify opportunities for nudge interventions, a third challenge arises: it is rarely clear who, in fact, should design and implement such nudges. For instance, it is not clear that market forces would drive private companies to implement nudges that aim to align individuals' behavior with their privacy preferences. On the one hand, companies seeking to retain their customers may have incentive to limit future regret stemming from the use of their service. This, however, assumes that companies can actually be held accountable for adverse effects stemming from their collection and use of individual personal information and that these adverse effects are salient to individuals. Moreover, for companies with consumer personal information at the core of their business model, privacy considerations can often be at odds with business incentives, and some companies that profit from personal data are not consumer-facing (e.g., data brokers).

Alternate approaches may include end-users using self-commitment strategies to nudge their own future behaviors (for instance, installing apps on their computers that limit the amount of time they can spend on a given day browsing social media sites); and regulators requiring companies that collect sensitive personal information to implement nudges aimed at improved consumer privacy decision making. While regulatory approaches have downsides, including the drawback of generally being broad-brushed [Johnson et al. 2012; Mullane and Sheffrin 2012], this approach circumvents the lack of sufficient market forces to encourage companies to implement

privacy nudges and also can be more desirable to companies when compared to more heavy-handed regulatory approaches that, for example, simply prohibit certain kinds of collection and use of personal information.

In situations where neither market forces nor regulation can easily attain a desired privacy or security outcome, privacy and security nudges may still be provided by third parties, such as non-profits, privacy activists, or businesses specializing in security and privacy protection. For instance, browser extensions or dedicated mobile apps can assist users with nudges in making beneficial privacy decisions.

5.4. Guidelines for Ethical Nudge Design

Based on the above discussion, we highlight some of the many facets to consider when deciding whether or not to deploy privacy or security nudges:

Are the normative judgments expressed through the nudge appropriate for the privacy and security concerns of the user population? This is perhaps the central question and the most difficult to answer, since it is possible that a specific nudge is appropriate for the majority of users and yet inappropriate for a different subset of users. For instance, within the context of online social networks, subpopulations may vary in their expectations of formality, emotional expressiveness, and disclosure of personal information. When no harm is imminent, nudging against a user's norms may alienate them from the intent behind future nudges or cause them to try to bypass the nudge.

How likely are unintended adverse consequences? System designers are already responsible for the consequences of design decisions, but nudging is especially applicable to features with ethical repercussions. Because of this, increased scrutiny for possible unintended consequences is appropriate. A nudge that encourages a user to avoid businesses with suboptimal privacy policies, for example, may discourage a person from a transaction that would contribute to their well-being if no optimal alternatives exist. Arguments have been made that soft paternalism, which underlies nudging, reduces an individual's freedom to occasionally err, and that such freedom is important, both for personal liberty (which, they posit, transcends mere "choice preservation") and for welfare [Wright and Ginsburg 2012]. Glaeser [2005] argues that soft paternalism may stigmatize certain choices, leading to negative social consequences for some individuals. The applicability of these arguments to nudging for privacy and security may depend on the severity of a mistake in a particular domain, as well as the relative value of learning as a result of a mistake.

Does nudging transfer responsibility in an inappropriate way? System designers may wish to avoid creating nudges that reduce the user's responsibility for important decisions. Bovens [2009] warns of the risks of "infantilization," where individuals become unable to make their own decisions to fit their preferences, as they come to rely on nudges for guidance. Such an outcome would reduce an individual's awareness of the value of privacy and security, potentially leading them to underestimate new threats in contexts not covered by the nudges they are accustomed to. Because of this, it is possible that a transfer of culpability is a greater ethical risk than a suboptimal decision.

Could the nudge appear to present a conflict of interest, and does the nudging party have an appropriate responsibility to nudge? Glaeser [2005] warns that the soft paternalism behind nudging, as a form of persuasion, is open to abuse. If an entity providing software is perceived as neutral in a multiparty transaction, then nudging might not be appropriate. For example, consider a loan comparison tool provided by a privacy advocacy group. Nudging users to a specific bank based on unstated sources of information, even for benign reasons, would not be appropriate. Moreover, the appropriateness of nudging requires an alignment between the purpose of the nudge and the role of the nudging entity. Nudging from a government entity to secure identity information (e.g., one's social security number), for instance, may be seen as appropri-

ate, while nudging from the same entity to secure less vital information (such as social networking data on a specific site) introduces concerns of misplaced or overreaching responsibility. Equally problematic are nudges that have been made obligatory in spite of the implementer holding a conflict of interest. Willis [2014] argues that “a push can easily overwhelm a nudge”: requiring companies to allow consumers to opt out from behavioral advertising, for example, would likely fail to obtain its purported benefits to consumer privacy if companies motivate consumers to accept the opt-in default.

Once a decision has been made to nudge for privacy or security reasons, the system designer must then choose from a variety of nudging techniques that may vary in (soft) paternalistic nature, each with its own ethical considerations. Although the choice of technique is subject to the specific application, some guidelines seem appropriate:

The direction, salience, and firmness of a nudge should be proportional to the user's benefit from the suggested course of action. Dire warnings for items of small risk, such as sharing one's email address, are at best inappropriate, and at worst can be deceptive. Similarly, a mild nudge is a poor design choice when the user may take actions with disastrous consequences, such as sharing one's credit card number with a suspected phishing website.

The user's choices should be respected. In spite of the best intentions of a nudge (or the worst consequences of ignoring it), it must remain auxiliary to the user's actions. If the user chooses to disregard a nudge, it is appropriate to assume that she has a reason; if no reason can possibly exist, then providing the choice is inappropriate in the first place. For example, it may be appropriate to warn social media users that their posts may be inflammatory or offensive to others, but making it difficult or impossible to post would be inappropriate. System designers may consider allowing users to customize the nudge, for instance by letting them configure under what circumstances the nudge will be triggered.

Nudging techniques should respect ethical norms applicable to persuasion. Nudges that utilize default effects, for example, should nudge towards choices that are explicitly endorsed [Dinner et al. 2011], to remove ambiguity over the presence of persuasion.

Users' expectations for truthful information should be respected. Nudges should not distort or present incomplete information, unless there is a strong, defensible rationale to do so. An example of such an intervention is the sending of fake phishing emails to users. This approach has been used to nudge users to learn how to better protect themselves from phishing attacks [Kumaraguru et al. 2007]. Given the number of people who fall for such attacks and the economic impact, this particular use of deception can be deemed admissible, with many organizations now reporting that they use simulated phishing attacks to train employees [Gartner Group 2014]. Use of deception to nudge users should however be carefully weighed, as it may negatively impact user trust.

5.5. Additional Research Challenges

We conclude this section with an overview of some of the more interesting challenges as well as opportunities for research in nudging for privacy and security.

A first challenge is that current approaches to nudging seem to rely on a one-size-fits-all model, where the same intervention is used across diverse members of a population. As individuals within a population differ in terms of preferences, awareness, knowledge, but also personality traits or susceptibility to biases and heuristics, individually tailored nudges may result in more effective interventions [Egelman and Peer 2015].

A second challenge relates to expanding the research on decision making hurdles that affect privacy and security behaviors. In particular, are there biases and heuristics

that are specific, or even unique, to the privacy and security domain, and that have not yet received sufficient attention in behavioral research?

A third challenge relates to how designers can implement and integrate nudges within their systems in unobtrusive ways, avoiding user habituation or fatigue; conversely, how can users stay alert about, or become aware of, subtle attempts by system designers to nudge and influence their behaviors? As system providers accumulate increasing amounts of user data, their ability to subtly and invisibly influence behavior may grow. Such ability may—as discussed above in this section—be used both to better serve and to take advantage of the end user.

6. CONCLUSION

Making the “right” online privacy or security decision—that is, balancing revelation and protection of data in manners that maximize a person’s or an organization’s welfare and minimize future regrets or losses—is increasingly difficult, due to the growing complexity of information technologies and the variety of activities we conduct online. In recent years, much research attention from scholars operating across diverse scientific disciplines has been devoted to understanding individuals’ privacy and security behaviors. The traditional economic view of individuals engaging in a privacy (or security) calculus, rationally trading off costs and benefits of their actions, has been complemented by a more behaviorally grounded analysis of the heuristics and biases that may also influence behavior, and not always for the better. In parallel, researchers have been devoting efforts to developing more usable privacy and security tools, as well as designing choice architectures that take into consideration those decision making hurdles in order to ameliorate people’s choices. In this article, we have reviewed both streams of research: the one investigating hurdles in online privacy and security decision making, and the one investigating ways of assisting that decision making. We have focused on the emerging field of asymmetric or “soft” paternalism. The idea behind soft paternalism is that lessons learnt about the psychological processes underlying behavior can be used to actually aid that behavior, by designing tools and policies that enhance choice, without restricting it. The ultimate goal of such interventions is to increase individual, organizational, and societal welfare by either mitigating or exploiting people’s behavioral biases in a way that leads people to make more informed and more desirable autonomous choices. In doing so, we have shown how existing concurrent research (such as usability research and persuasive computing) can be seen through the lenses of nudging and soft paternalism.

Our analysis of both the behavioral literature in general and the privacy and security literature in particular has highlighted a vast array of factors that may affect and impair end-users’ privacy and security choices. As shown in Section 2, the effects of various biases and heuristics on privacy and security choices have already started to be analyzed in a growing body of empirical research. The effect of other biases and heuristics—well known to behavioral researchers, less so to privacy and security specialists—is currently only conjectural, but the hypotheses we present in that section may help drive future research efforts. Furthermore, examples of soft-paternalistic interventions in the field of privacy and security have started to arise both in research and in actual commercial products. As highlighted in Section 3, we can find growing evidence both of tools aimed at making people reflect on their disclosure or security actions before they take them, and of tools and interface designs that nudge individuals towards more (or more open) disclosures.

In this survey article, our goal has been to document ongoing efforts in this area, discuss some of their limitations, and highlight their potential. We view this new emerging area as one that could lead to the design of a variety of tools and systems that

effectively assist humans with online security and privacy decisions without imposing overly prescriptive models of what the “right” decisions might be.

The studies we covered in this review have attempted to highlight the human processes that drive privacy and security behaviors, and how those processes can be (and are being) influenced by tools, interfaces, and choice architectures—even when they remain agnostic regarding the appropriateness of such interventions. As noted in Section 5, judging appropriateness is outside the scope of our review—it is, instead, the domain of society’s and individuals’ autonomous valuations. In stating that, however, we have also argued that far from seeing nudging interventions as an invasion on individuals’ otherwise pristine and untouched autonomy, we should realize that every design decision behind the construction of every online (e.g., software, online social networks, online blogs, mobile devices and applications, etc.) or offline (e.g., conference rooms, vehicles, food menus, etc.) system or tool we use has the potential to influence users’ behaviors, regardless of whether the designer, or the user, is fully aware of those influences and their consequences. In simple terms, there is no such thing as a neutral design in privacy, security, or anywhere else. Therefore, we argue for conscious and cautious design of choice architectures and nudges that are inherent to any system, as well as the use of nudging to help users overcome cognitive and behavioral hurdles that may impact their privacy and security decisions.

REFERENCES

- Alessandro Acquisti. 2004. Privacy in electronic commerce and the economics of immediate gratification. In *Proc. 5th ACM Conference on Electronic Commerce*. ACM, New York, NY, USA, 21–29.
- Alessandro Acquisti. 2009. Nudging privacy: The behavioral economics of personal information. *Security & Privacy, IEEE* 7, 6 (2009), 82–85.
- Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347, 6221 (2015), 509–514.
- Alessandro Acquisti and Christina M Fong. 2014. An experiment in hiring discrimination via online social networks. *Available at SSRN 2031979* (2014), 1–81.
- Alessandro Acquisti and Ralph Gross. 2006. Imagined communities: Awareness, information sharing, and privacy on the Facebook. In *Proc. 6th Int. Workshop Privacy Enhancing Tech. (PET '06)*. Springer, 36–58.
- Alessandro Acquisti and Jens Grossklags. 2005. Privacy and rationality in individual decision making. *IEEE Security & Privacy* 2 (2005), 24–30.
- Alessandro Acquisti and Jens Grossklags. 2007. What can behavioral economics teach us about privacy. In *Digital Privacy: Theory, Technologies and Practices*, Sabrina De Capitani di Vimercati, Stefanos Gritzalis, Costas Lambrinoudakis and Alessandro Acquisti (Eds.). Auerbach Publications, 363–377.
- Alessandro Acquisti, Leslie K. John, and George Loewenstein. 2012. The impact of relative standards on the propensity to disclose. *Journal of Marketing Research* 49, 2 (2012), 160–174.
- Alessandro Acquisti, Leslie K. John, and George Loewenstein. 2013. What is privacy worth? *The Journal of Legal Studies* 42, 2 (2013), 249–274.
- Alessandro Acquisti, Curtis R. Taylor, and Liad Wagman. 2016. The economics of privacy. *Journal of Economic Literature* 52, 2 (2016).
- Anne Adams and Martina Angela Sasse. 1999. Users are not the enemy. *Commun. ACM* 42, 12 (1999), 40–46.
- Idris Adjerid, Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2013. Sleights of privacy: framing, disclosures, and the limits of transparency. In *Proc. Symp. Usable Privacy and Security (SOUPS '13)*. ACM, 1–11.
- George Ainslie and Nick Haslam. 1992. Hyperbolic discounting. In *Choice over time*, G. Loewenstein and J. Elster (Eds.). Russell Sage Foundation, New York, NY, USA, 57–92.
- George A. Akerlof. 1970. The market for “lemons”: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics* 84, 3 (1970), 488–500.
- Reed Albergotti. 2014. Facebook’s Blue Dino Wants You to Mind Your Posting. WSJ Blog. (April 2014). <http://blogs.wsj.com/digits/2014/04/01/facebooks-blue-dino-wants-you-to-mind-your-posting>

- Hazim Almuhammedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Faith Cranor, and Yuvraj Agarwal. 2015. Your location has been shared 5,398 times! A field study on mobile app privacy nudging. In *Proc. Conf. Human Factors in Comp. Sys. (CHI '15)*. ACM, 1–10.
- Adil Alsaïd and David Martin. 2003. Detecting web bugs with Bugnosis: Privacy advocacy through education. In *Proc. 2nd international conference on Privacy enhancing technologies*. Springer, 27–31.
- Christopher J Anderson. 2003. The psychology of doing nothing: Forms of decision avoidance result from reason and emotion. *Psychological bulletin* 129, 1 (2003), 139–167.
- Ross Anderson. 2001. Why information security is hard: An economic perspective. In *Proc. 17th Annual Computer Security Applications Conference (ACSAC '01)*. IEEE, New York, NY, USA, 358–365.
- Ross Anderson, Chris Barton, Rainer Böhme, Richard Clayton, Michel van Eeten, Michael Levi, Tyler Moore, and Stefan Savage. 2013. Measuring the cost of cybercrime. In *The Economics of Information Security and Privacy*. Springer-Verlag, Berlin, Heidelberg, 265–300.
- Julia Angwin. 2012. Microsoft's "Do Not Track" Move Angers Advertising Industry. Wall Street Journal Blog. (2012). <http://blogs.wsj.com/digits/2012/05/31/microsofts-do-not-track-move-angers-advertising-industry/>
- Apple Inc. 2014. Mac Basics: Time Machine backs up your Mac. (2014). <http://support.apple.com/kb/ht1427>.
- Dan Ariely, George Loewenstein, and Drazen Prelec. 2003. "Coherent arbitrariness": Stable demand curves without stable preferences. *Quarterly Journal of Economics* 118, 1 (2003), 73–106.
- Rebecca Balebako, Pedro G. Leon, Hazim Almuhammedi, Patrick Gage Kelley, Jonathan Mugan, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. 2011. Nudging users towards privacy on mobile devices. In *Proc. CHI 2011 Workshop on Persuasion, Nudge, Influence and Coercion*. ACM, 1–4.
- Rebecca Balebako, Florian Schaub, Idris Adjerid, Alessandro Acquisti, and Lorrie Cranor. 2015. The impact of timing on the salience of smartphone app privacy notices. In *Proc. CCS Workshop on Security and Privacy in Smartphones and Mobile Devices (SPSM '15)*. ACM, New York, NY, USA, 63–74.
- Scott Beaulier and Bryan Caplan. 2007. Behavioral economics and perverse effects of the welfare state. *Kyklos* 60, 4 (2007), 485–507.
- Gary S. Becker. 1976. *The economic approach to human behavior*. University of Chicago Press, Chicago, IL.
- Omri Ben-Shahar and Carl E. Schneider. 2010. The failure of mandated discourse. *U. Pa. L. Rev.* 159, 3 (2010), 647.
- John Beshears, James J. Choi, David Laibson, and Brigitte C. Madrian. 2009. The importance of default options for retirement saving outcomes: Evidence from the United States. In *Social security policy in a changing environment*. University of Chicago Press, Chicago, IL, USA, 167–195.
- Andrew Besmer and Heather Richter Lipford. 2010. Moving beyond untagging: Photo privacy in a tagged world. In *Proc. Conference on Human Factors in Computing Systems (CHI '10)*. ACM, 1563–1572.
- Andrew Besmer, Jason Watson, and Heather Richter Lipford. 2010. The impact of social navigation on privacy policy configuration. In *Proc. Symp. Usable Privacy and Security (SOUPS '10)*. ACM, 1–10.
- Matt Bishop. 2000. Education in information security. *Concurrency, IEEE* 8, 4 (2000), 4–8.
- Garrick Blalock, Vrinda Kadiyali, and Daniel H. Simon. 2007. The impact of post-9/11 airport security measures on the demand for air travel. *Journal of Law and Economics* 50, 4 (2007), 731–755.
- Christoph Bösch, Benjamin Erb, Frank Kargl, Henning Kopp, and Stefan Pfattheicher. 2016. Tales from the dark side: Privacy dark strategies and privacy dark patterns. *Proc. Privacy Enhancing Technologies* 2016, 4 (2016), 237–254.
- Luc Bovens. 2009. The ethics of nudge. In *Preference change: Approaches from philosophy, economics and psychology*, Till Grne-Yanoff and S.O. Hansson, (Eds.). Springer, Berlin, 207–220.
- Laura Brandimarte, Alessandro Acquisti, and George Loewenstein. 2013. Misplaced confidences: Privacy and the control paradox. *Social Psychological and Personality Science* 4, 3 (2013), 340–347.
- Alex Braunstein, Laura Granka, and Jessica Staddon. 2011. Indirect content privacy surveys: Measuring privacy without asking about it. In *Proc. Symp. Usable Privacy and Security (SOUPS '11)*. ACM, 1–14.
- Cristian Bravo-Lillo, Saranga Komanduri, Lorrie Faith Cranor, Robert W. Reeder, Manya Sleeper, Julie Downs, and Stuart Schechter. 2013. Your attention please: Designing security-decision UIs to make genuine risks harder to ignore. In *Proc. Symp. Usable Privacy and Security (SOUPS '13)*. ACM, 1–12.
- Patricia C. Brennan, Poornima Madhavan, Cleotilde Gonzalez, and Frank C. Lacson. 2009. The impact of performance incentives during training on transfer of learning. *Proc. Human Factors and Ergonomics Society Annual Meeting* 53, 26 (2009), 1979–1983.
- José Carlos Brustoloni and Ricardo Villamarín-Salomón. 2007. Improving security decisions with polymorphic and audited dialogs. In *Proc. Symp. Usable Privacy and Security (SOUPS '07)*. ACM, 76–85.
- Frank H. Buckley. 2005. Perfectionism. *Supreme Court Economic Review* 13 (2005), 133–163.

- Federico Cabitza and Marco Loregian. 2008. Much undo about nothing?: Investigating why email retraction is less popular than apologizing. In *Proc. Nordic Conf. HCI (NordiCHI '08)*. ACM, 431–434.
- Ryan Calo. 2010. The boundaries of privacy harm. *Indiana Law Journal* 86, 3 (2010), 1–31.
- Ryan Calo. 2012. Against notice skepticism in privacy (and elsewhere). *Notre Dame Law Review* 87, 3 (2012), 1027–1072.
- Colin Camerer, Samuel Issacharoff, George Loewenstein, Ted O'Donoghue, and Matthew Rabin. 2003. Regulation for conservatives: Behavioral economics and the case for 'asymmetric paternalism'. *U Penn Law Rev* 151, 3 (2003), 1211–1254.
- Colin F. Camerer, George Loewenstein, and Matthew Rabin. 2011. *Advances in behavioral economics*. Princeton University Press, Princeton, NJ, USA.
- Ann Cavoukian. 2009. *Privacy by design: Take the challenge*. Information and Privacy Commissioner of Ontario, Canada, Toronto, ON, Canada. <http://privacybydesign.ca>
- Daphne Chang, Erin L. Krupka, Eytan Adar, and Alessandro Acquisti. 2016. Engineering information disclosure: Norm shaping designs. In *Proc. Conf. Human Factors in Comp. Sys. (CHI '16)*. ACM, 587–597.
- Avi Charkham. 2012. 5 design tricks Facebook uses to affect your privacy decisions. TechCrunch. (Aug. 2012). <https://techcrunch.com/2012/08/25/5-design-tricks-facebook-uses-to-affect-your-privacy-decisions/>
- Sonia Chiasson, Alain Forget, Robert Biddle, and Paul C. van Oorschot. 2008. Influencing users towards better passwords: Persuasive cued click-points. In *Proc. 22nd British HCI Group Annual Conference on People and Computers*. British Computer Society, Swinton, UK, 121–130.
- James J. Choi, David Laibson, Brigitte Madrian, and Andrew Metrick. 2004. For better or for worse: Default effects and 401(K) savings behavior. In *Perspectives on the economics of aging*, David Wise (Ed.). University of Chicago Press, Chicago, IL, USA, 81–121.
- Brian Christian. 2012. The A/B test: Inside the technology that's changing the rules of business. Wired. (April 2012). <http://www.wired.com/business/2012/04/ff.abtesting/>
- Nicolas Christin, Serge Egelman, Timothy Vidas, and Jens Grossklags. 2012. It's all about the Benjamins: An empirical study on incentivizing users to ignore security advice. In *Proc. Conference on Financial Cryptography and Data Security (FC '12)*. Springer-Verlag, Berlin, Heidelberg, 16–30.
- Sandy Clark, Travis Goodspeed, Perry Metzger, Zachary Wasserman, Kevin Xu, and Matt Blaze. 2011. Why (special agent) Johnny (still) can't encrypt: A security analysis of the APCO Project 25 two-way radio system. In *USENIX Security Symposium*. USENIX Association, Berkeley, CA, USA.
- Sunny Consolvo, Katherine Everitt, Ian Smith, and James A Landay. 2006. Design requirements for technologies that encourage physical activity. In *Proc. Conf. Hum. Factors in Comp. Sys. (CHI '06)*. ACM, 457–466.
- Lorrie Faith Cranor and Simson Garfinkel. 2005. *Security and usability: Designing secure systems that people can use*. O'Reilly Media, Inc., Sebastopol, CA.
- Lorrie Faith Cranor, Praveen Guduru, and Manjula Arjula. 2006. User interfaces for privacy agents. *ACM Transactions on Computer-Human Interaction (TOCHI)* 13, 2 (2006), 135–178.
- Paul Curzon and Ann Blandford. 2004. Formally justifying user-centred design rules: A case study on post-completion errors. In *Proc. 4th Int. Conf. Integrated Formal Methods*. Springer, 461–480.
- Nikhil Dhingra, Zach Gorn, Andrew Kener, and Jason Dana. 2012. The default pull: An experimental demonstration of subtle default effects on preferences. *Judgm. Decis. Mak.* 7, 1 (2012), 69–76.
- Isaac Dinner, Eric J. Johnson, Daniel G. Goldstein, and Kaiya Liu. 2011. Partitioning default effects: Why people choose not to choose. *Journal of Experimental Psychology: Applied* 17, 4 (2011), 332–341.
- Paul Dolan, Michael Hallsworth, David Halpern, D. King, R. Metcalfe, and Ivo Vlaev. 2012. Influencing behaviour: The mindspace way. *Journal of Economic Psychology* 33, 1 (2012), 264–277.
- Julie S. Downs, George Loewenstein, and Jessica Wisdom. 2009. Strategies for promoting healthier food choices. *The American Economic Review* 99, 2 (2009), 159–164.
- Andreas C. Drichoutis, Panagiotis Lazaridis, and Rodolfo M. Nayga. 2006. Consumers' use of nutritional labels: A review of research studies and issues. *Academy of Marketing Science Rev.* 10, 9 (2006), 1–25.
- Serge Egelman, David Molnar, Nicolas Christin, Alessandro Acquisti, Cormac Herley, and Shriram Krishnamurthi. 2010. Please continue to hold: An empirical study on user tolerance of security delays. In *Proc. Workshop on the Economics of Information Security (WEIS)*.
- Serge Egelman, Adrienne Porter Felt, and David Wagner. 2013. Choice architecture and smartphone privacy: There's a price for that. In *The Economics of Information Security*. Springer, 211–236.
- Serge Egelman and Eyal Peer. 2015. The myth of the average user: Improving privacy and security systems through individualization. In *Proceedings of the 2015 New Security Paradigms Workshop*. ACM, 16–28.

- Serge Egelman, Janice Tsai, Lorrie Faith Cranor, and Alessandro Acquisti. 2009. Timing is everything?: The effects of timing and placement of online privacy indicators. In *Proc. 27th International Conference on Human Factors in Computing Systems (CHI '09)*. ACM, New York, NY, USA, 319–328.
- Lujun Fang and Kristen LeFevre. 2010. Privacy wizards for social networking sites. In *Proc. 19th International Conference on World Wide Web (WWW '10)*. ACM, New York, NY, USA, 351–360.
- Federal Register. 2004. 16 Code of Federal Regulation Part 429. <http://www.archives.gov/federal-register/cfr/subject-title-16.html>. (2004).
- Financial Consumer Agency of Canada. 2014. Credit Card Selector Tool. <http://itools-ioutils.fcac-acfc.gc.ca/stcv-osvc/ccst-oscc-eng.aspx>. (April 2014).
- Baruch Fischhoff. 1981. *Debiasing*. Technical Report. DTIC Document.
- Myron F. Floyd, Heather Gibson, Lori Pennington-Gray, and Brijesh Thapa. 2004. The effect of risk perceptions on intentions to travel in the aftermath of September 11, 2001. *Journal of Travel & Tourism Marketing* 15, 2-3 (2004), 19–38.
- B. J. Fogg. 2002. *Persuasive technology* (1 ed.). Morgan Kaufmann, Burlington, MA, USA.
- Shane Frederick, George Loewenstein, and Ted O'Donoghue. 2002. Time discounting and time preference: A critical review. *Journal of Economic Literature* 40, 2 (2002), 351–401.
- Milton Friedman and Leonard J. Savage. 1948. The utility analysis of choices involving risk. *The Journal of Political Economy* 56, 4 (1948), 279–304.
- Bo Gao, Bettina Berendt, Dave Clarke, Ralf De Wolf, Thomas Peetz, Jo Pierson, and Rula Sayaf. 2012. Interactive grouping of friends in OSN: Towards online context management. In *Proc. Int. Conf. Data Mining Workshops*. IEEE, 555–562.
- Simson Garfinkel and Heather Richter Lipford. 2014. Usable security: History, themes, and challenges. *Synthesis Lectures on Information Security, Privacy, and Trust* 5, 2 (2014), 1–124.
- Gartner Group. 2014. *Magic quadrant for security awareness computer-based training vendors*. Technical Report. Gartner Group.
- Edward L. Glaeser. 2005. *Paternalism and psychology*. NBER Working Paper No. 11789. National Bureau of Economic Research.
- Jeremy Goecks and Elizabeth Mynatt. 2005. Social approaches to end-user privacy management. In *Security and usability: Designing secure systems that people can use*, L. F. Cranor and S. Garfinkel (Eds.). O'Reilly, Sebastopol, CA, 523–547.
- Daniel G. Goldstein, Eric J. Johnson, Andreas Herrmann, and Mark Heitmann. 2008. Nudge your customers toward better choices. *Harvard Business Review* 86, 12 (2008), 99–105.
- Nathaniel S. Good, Jens Grossklags, Deirdre K. Mulligan, and Joseph A. Konstan. 2007. Noticing notice: A large-scale experiment on the timing of software license agreements. In *Proc. SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, New York, NY, USA, 607–616.
- Nathaniel S. Good, Jens Grossklags, David Thaw, Aaron Perzanowski, Deirdre K. Mulligan, and Joseph A. Konstan. 2006. User choices and regret: Understanding users' decision process about consensually acquired spyware. *I/S: A Journal of Law and Policy* 2, 2 (2006), 283–344.
- Connor Graham, Peter Benda, Steve Howard, James Balmford, Nicole Bishop, and Ron Borland. 2006. Heh—keeps me off the smokes...: Probing technology support for personal change. In *Proc. 18th Australia Conference on Computer-Human Interaction (OzCHI '06)*. ACM, New York, NY, USA, 221–228.
- Victoria Groom and Ryan Calo. 2011. Reversing the privacy paradox: An experimental study. *Available at SSRN 1993125* (2011).
- Jens Grossklags and Alessandro Acquisti. 2007. When 25 cents is too much: An experiment on willingness-to-sell and willingness-to-protect personal information. In *Proc. Workshop on the Economics of Information Security (WEIS)*. 1–22.
- Jens Grossklags, Benjamin Johnson, and Nicolas Christin. 2010. When information improves information security. In *Proc. Int. Conf. Financial Cryptography and Data Security (FC '10)*. Springer, 416–423.
- Scott D. Halpern, Peter A. Ubel, and David A. Asch. 2007. Harnessing the power of default options to improve health care. *New England Journal Of Medicine* 357 (2007), 1340–1344.
- Pelle Guldborg Hansen. 2012. Nudging traffic safety by visual illusions. iNudgeYou. (2012). <http://inudgeyou.com/en/archives/504>
- Pelle Guldborg Hansen. 2016. The definition of nudge and libertarian paternalism: does the hand fit the glove?. *European Journal of Risk Regulation* 7, 1 (2016), 155–174.
- Marian Harbach, Markus Hettig, Susanne Weber, and Matthew Smith. 2014. Using personal examples to improve risk communication for security & privacy decisions. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 2647–2656.

- Tejaswini Herath and H. Raghav Rao. 2009. Encouraging information security behaviors in organizations: Role of penalties, pressures and perceived effectiveness. *Decision Support Syst.* 47, 2 (2009), 154–165.
- Cormac Herley. 2009. So long, and no thanks for the externalities: the rational rejection of security advice by users. In *Proc. Workshop on New Security Paradigms (NSPW '09)*. ACM, New York, NY, USA, 133–144.
- Donna L. Hoffman and Thomas P. Novak. 1997. A new marketing paradigm for electronic commerce. *The Information Society: An International Journal* 13, 1 (1997), 43–54.
- Chris Hoofnagle and Jennifer King. 2008. What Californians understand about privacy online. Available at SSRN 1262130 (2008), 1–33.
- Leonie Huddy, Stanley Feldman, Gallya Lahav, and Charles Taber. 2003. Fear and terrorism: Psychological reactions to 9/11. In *Framing terrorism: The news media, the government and the public*, Pippa Norris, Montague Kern, and Marion Just (Eds.). Routledge, New York, NY, USA, 255–278.
- Giovanni Iachello, Ian Smith, Sunny Consolvo, Mike Chen, and Gregory D. Abowd. 2005. Developing privacy guidelines for social location disclosure applications and services. In *Proc. Symposium on Usable Privacy and Security (SOUPS '05)*. ACM, New York, NY, USA, 65–76.
- Lukasz Jedrzejczyk, Blaine A. Price, Arosha K. Bandara, and Bashar Nuseibeh. 2010. On the impact of real-time feedback on users' behaviour in mobile location-sharing applications. In *Proc. Sixth Symposium on Usable Privacy and Security (SOUPS '10)*. ACM, New York, NY, USA, 1–12.
- Nicola Jentzsch, Sören Preibusch, and Andreas Harasser. 2012. *Study on monetising privacy: An economic model for pricing personal information*. European Union Agency for Network and Inf. Sec. (ENISA).
- Leslie K. John, Alessandro Acquisti, and George Loewenstein. 2011. Strangers on a plane: Context-dependent willingness to divulge sensitive information. *J. Consum. Res.* 37, 5 (2011), 858–873.
- Eric J. Johnson and Daniel G. Goldstein. 2004. Defaults and donation decisions. *Transplantation* 78, 12 (2004), 1713–1716.
- Eric J. Johnson, Suzanne B. Shu, Benedict G. C. Dellaert, Craig Fox, Daniel G. Goldstein, Gerald Häubl, Richard P. Larrick, John W. Payne, Ellen Peters, David Schkade, Brian Wansink, and Elke U. Weber. 2012. Beyond nudges: Tools of a choice architecture. *Marketing Letters* 23, 2 (2012), 487–504.
- Daniel Kahneman, Jack L. Knetsch, and Richard H. Thaler. 1990. Experimental tests of the endowment effect and the Coase theorem. *Journal of Political Economy* 98, 6 (1990), 1325–1348.
- Daniel Kahneman, Jack L. Knetsch, and Richard H. Thaler. 1991. Anomalies: The endowment effect, loss aversion, and status quo bias. *The Journal of Economic Perspectives* 5, 1 (1991), 193–206.
- Daniel Kahneman and Dale T. Miller. 1986. Norm theory: Comparing reality to its alternatives. *Psychological Review* 93, 2 (1986), 136–153.
- Daniel Kahneman and Amos Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47, 2 (1979), 263–291.
- Patrick Gage Kelley, Lucian Cesca, Joanna Bresee, and Lorrie Faith Cranor. 2010. Standardizing privacy notices: An online study of the nutrition label approach. In *Proc. Conf. Human Factors in Computing Systems (CHI '10)*. ACM, 1573–1582.
- Patrick Gage Kelley, Lorrie Faith Cranor, and Norman Sadeh. 2013. Privacy as part of the app decision-making process. In *Proc. Conf. Human Factors in Computing Systems (CHI '13)*. ACM, 3393–3402.
- Saranga Komanduri, Richard Shay, Lorrie Faith Cranor, Cormac Herley, and Stuart Schechter. 2014. Telepathwords: Preventing weak passwords by reading Users minds. In *Proc. 23rd USENIX Security Symposium*. USENIX Association, Berkeley, CA, USA, 591–606.
- Ulrich König and Jan Schallaböck. 2011. Privacy preferences for e-mail messages. IETF draft. (2011). <http://tools.ietf.org/html/draft-koenig-privicons-01>
- Bastian Könings, David Piendl, Florian Schaub, and Michael Weber. 2011. PrivacyJudge: Effective privacy controls for online published information. In *Proc. Conf. Privacy, Security, Risk and Trust (PASSAT '11)*. IEEE, 935–941.
- Bastian Könings, Sebastian Thoma, Florian Schaub, and Michael Weber. 2014. PriPref broadcaster: Enabling users to broadcast privacy preferences in their physical proximity. In *Proc. 13th International Conference on Mobile and Ubiquitous Multimedia (MUM '14)*. ACM, New York, NY, USA, 133–142.
- Ponnurangam Kumaraguru, Yong Rhee, Alessandro Acquisti, Lorrie Faith Cranor, Jason Hong, and Elizabeth Nunge. 2007. Protecting people from phishing: The design and evaluation of an embedded training email system. In *Proc. Conf. Human Factors in Computing Systems (CHI '07)*. ACM, 905–914.
- Yee-Lin Lai and Kai-Lung Hui. 2006. Internet opt-in and opt-out: Investigating the roles of frames, defaults and privacy concerns. In *Proc. Conf. Computer Personnel Research (CPR '06)*. ACM, 253–263.
- David Laibson. 1997. Golden eggs and hyperbolic discounting. *Quart. J. Economics* 112, 2 (1997), 443–478.
- Marc Langheinrich. 2002. A privacy awareness system for ubiquitous computing environments. In *Proc. Conference on Ubiquitous Computing (UbiComp '02)*. Springer-Verlag, London, 237–245.

- Robert S. Laufer and Maxine Wolfe. 1977. Privacy as a concept and a social issue: A multidimensional developmental theory. *Journal of social Issues* 33, 3 (1977), 22–42.
- Scott Lederer, Jason I. Hong, Anind K. Dey, and James A. Landay. 2004. Personal privacy through understanding and action: Five pitfalls for designers. *Personal and Ubiquitous Comp.* 8, 6 (2004), 440–454.
- Pedro Giovanni Leon, Blase Ur, Rebecca Balebako, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Why Johnny can't opt out: A usability evaluation of tools to limit online behavioral advertising. In *Proc. Conf. Human Factors in Computing Systems (CHI '12)*. ACM, 589–598.
- Irwin P. Levin, Sandra L. Schneider, and Gary J. Gaeth. 1998. All frames are not created equal: A typology and critical analysis of framing effects. *Organ. Behav. Hum. Dec.* 76, 2 (1998), 149–188.
- Han Li, Rathindra Sarathy, and Heng Xu. 2011. The role of affect and cognition on online consumers' decision to disclose personal information to unfamiliar online vendors. *Decis. Support Syst.* 51, 3 (2011), 434–445.
- Han Li, Rathindra Sarathy, and Jie Zhang. 2008. The role of emotions in shaping consumers' privacy beliefs about unfamiliar online vendors. *Journal of Information Privacy & Security* 4, 3 (2008), 36–62.
- Simon Y. W. Li, Ann Blandford, Paul Cairns, and Richard M. Young. 2005. Post-completion errors in problem solving. In *Proc. 27th Annual Conference of the Cognitive Science Society*. 1–6.
- Sarah Lichtenstein, Baruch Fischhoff, and Lawrence D. Phillips. 1982. Calibration of probabilities: The state of the art to 1980. In *Judgment under uncertainty: Heuristics and biases*, Daniel Kahneman, Paul Slovic, and Amos Tversky, (Eds.). Cambridge Univ. Press, UK, 306–334.
- Assar Lindbeck. 1997. Incentives and Social norms in household behavior. *Am. Econ. Rev.* 87, 2 (1997), 370–377.
- Janne Lindqvist, Justin Cranshaw, Jason Wiese, Jason Hong, and John Zimmerman. 2011. I'm the mayor of my house: Examining why people use foursquare—a social-driven location sharing application. In *Proc. Conf. Human Factors in Computing Systems (CHI '11)*. ACM, 2409–2418.
- Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhammedi, Shikun Zhang, Norman Sadeh, Alessandro Acquisti, and Yuvraj Agarwal. 2016. Follow my recommendations: A personalized privacy assistant for mobile app permissions. In *Symp. Usable Privacy and Security (SOUPS '16)*. USENIX Association, Denver, CO, 27–41.
- George Loewenstein and Emily Celia Haisley. 2007. The economist as therapist: Methodological ramifications of 'light' paternalism. *Available at SSRN 962472* (2007), 1–50.
- George Loewenstein and Drazen Prelec. 1992. Anomalies in intertemporal choice: Evidence and an interpretation. *The Quarterly Journal of Economics* 107, 2 (1992), 573–597.
- Marco Loregian and Marco P. Locatelli. 2008. An experimental analysis of undo in ubiquitous computing environments. In *Proc. Conf. Ubiquitous Intelligence and Computing*. Springer, 505–519.
- Kim Ly, Nina Mazār, Min Zhao, and Dilip Soman. 2013. A practitioner's guide to nudging. Rotman School of Management. University of Toronto. (March 2013).
- Wendy E. Mackay. 1991. Triggers and barriers to customizing software. In *Proc. SIGCHI Conference on Human Factors in Computing Systems (CHI '91)*. ACM, New York, NY, USA, 153–160.
- Brigitte C. Madrian and Dennis F. Shea. 2001. The power of suggestion: Inertia in 401(K) participation and savings behavior. *Quarterly Journal of Economics* 116, 4 (2001), 1149–1187.
- Alessandra Mazzia, Kristen LeFevre, and Eytan Adar. 2012. The PViz comprehension tool for social network privacy settings. In *Proc. Symp. Usable Privacy and Security (SOUPS '12)*. ACM, 1–12.
- Aleecia M. McDonald and Lorrie Faith Cranor. 2008. The cost of reading privacy policies. *I/S: A Journal of Law and Policy* 4, 3 (2008), 540–565.
- Miriam J. Metzger. 2006. Effects of site, vendor, and consumer characteristics on web site trust and disclosure. *Communication Research* 33, 3 (2006), 155–179.
- Microsoft. 2008. Privacy guidelines for developing software products and services. (2008).
- Microsoft. 2014. What are the system recovery options in Windows? (2014). <http://windows.microsoft.com/en-us/windows/what-are-system-recovery-options>
- Tyler Moore and Ross Anderson. 2011. *Economics and Internet security: A survey of recent analytical, empirical and behavioral research*. Tech. Rep. TR-03-11. Dept. Computer Science, Harvard Univ.
- M. Granger Morgan and Max Henrion. 1992. *Uncertainty: A guide to dealing with uncertainty in quantitative risk and policy analysis*. Cambridge University Press, New York, NY, USA.
- Maggie Mullane and Steven Sheffrin. 2012. *Regulatory nudges in practice*. White paper. Department of Economics and Murphy Institute, Tulane University.
- Donald A. Norman. 2013. *the design of everyday things: Revised and expanded*. Basic Books, New York.

- Joon S. Park, Kevin A. Kwiat, Charles A. Kamhoua, Jonathan White, and Sookyoung Kim. 2014. Trusted Online Social Network (OSN) services with optimal data management. *Computers & Security* 42 (2014), 116–136.
- Sameer Patil, Xinru Page, and Alfred Kobsa. 2011. With a little help from my friends: Can social navigation inform interpersonal privacy preferences?. In *Proc. Conf. Computer Supported Cooperative Work (CSCW '11)*. ACM, 391–394.
- Andrew S. Patrick and Steve Kenny. 2003. From privacy legislation to interface design: Implementing information privacy in human-computer interactions. In *Proc. Workshop Privacy Enhancing Techn. (PET '03)*. Springer, 107–124.
- Eyal Pe'er. 2011. The time-saving bias, speed choices and driving behavior. *Transportation Research Part F: Traffic Psychology and Behaviour* 14, 6 (2011), 543–554.
- Pennsylvania Department of Transportation. 2013. PA Driver's manual. Chapter 3 – Learning to drive. (March 2013).
- Jon Perlow. 2008. New in Labs: Stop sending mail you later regret. Official Gmail Blog. (October 2008). <http://gmailblog.blogspot.com/2008/10/new-in-labs-stop-sending-mail-you-later.html>.
- John E. Petersen, Vladislav Shunturov, Kathryn Janda, Gavin Platt, and Kate Weinberger. 2007. Dormitory residents reduce electricity consumption when exposed to real-time visual feedback and incentives. *International Journal of Sustainability in Higher Education* 8, 1 (2007), 16–33.
- Pew Research Internet Project. 2013. Anonymity, Privacy, and Security Online. <http://www.pewinternet.org/2013/09/05/anonymity-privacy-and-security-online/>. (September 2013).
- Richard A. Posner. 1978. The right of privacy. *Georgia Law Review* 12, 3 (1978), 393–422.
- Richard A. Posner. 1981. The economics of privacy. *The American Economic Review* 71, 2 (1981), 405–509.
- Matthew Rabin. 1998. Psychology and economics. *Journal of Economic Literature* 36, 1 (1998), 11–46.
- Mika Raento and Antti Oulasvirta. 2005. Privacy management for social awareness applications. In *Proc. Workshop on Context Awareness for Proactive Systems*. 105–114.
- Fahimeh Raja, Kirstie Hawkey, Steven Hsu, Kai-Le Clement Wang, and Konstantin Beznosov. 2011. A brick wall, a locked door, and a bandit: A physical security metaphor for firewall warnings. In *Proc. Seventh Symposium on Usable Privacy and Security (SOUPS '11)*. ACM, New York, NY, USA, 1–20.
- Norman Sadeh, Jason Hong, Lorrie Faith Cranor, Ian Fette, Patrick Gage Kelley, Madhu Prabaker, and Jinghai Rao. 2009. Understanding and capturing people's privacy policies in a mobile social networking application. *Personal and Ubiquitous Computing* 13, 6 (2009), 401–412.
- M. Angela Sasse, Sacha Brostoff, and Dirk Weirich. 2001. Transforming the 'weakest link': A human/computer interaction approach to usable and effective security. *BT Technol. J.* 19, 3 (2001), 122–131.
- Florian Schaub, Rebecca Balebako, Adam L. Durity, and Lorrie Faith Cranor. 2015. A design space for effective privacy notices. In *Proc. Symp. Usable Privacy and Security (SOUPS '15)*. USENIX Association.
- Stuart E Schechter, Rachna Dhamija, Andy Ozment, and Ian Fischer. 2007. The emperor's new security indicators. In *IEEE Symposium on Security and Privacy*. IEEE, New York, NY, USA, 51–65.
- Bruce Schneier. 2007. The psychology of security. *Commun. ACM* 50, 5 (2007), 128.
- Katarina Segerstahl and Harri Oinas-Kukkonen. 2007. Distributed user experience in persuasive technology environments. In *Proc. 2nd International Conference on Persuasive Technology*. Springer, 80–91.
- Evan Selinger and Kyle Whyte. 2011. Is there a right way to nudge? The practice and ethics of choice architecture. *Sociology Compass* 5, 10 (2011), 923–935.
- Jesse M. Shapiro. 2005. Is there a daily discount rate? Evidence from the food stamp nutrition cycle. *Journal of Public Economics* 89, 2 (2005), 303–325.
- David Sharek, Cameron Swofford, and Michael Wogalter. 2008. Failure to recognize fake internet pop up warning messages. *Proc. Human Factors and Erg. Society Ann. Meeting* 52, 6 (2008), 557–560.
- Fuming Shih, Ilaria Liccardi, and Daniel J. Weitzner. 2015. Privacy tipping points in smartphones privacy preferences. In *Proc. Conf. Human Factors in Computing Systems (CHI '15)*. ACM, 807–816.
- Katie Shilton, Jeffrey A. Burke, Deborah Estrin, Mark Hansen, and Mani Srivastava. 2008. *Participatory privacy in urban sensing*. Technical Report. Center for Embedded Network Sensing.
- Adam Shostack. 2003. Paying for privacy: Consumers and infrastructures. In *Proc. 2nd Annual Workshop on Economics and Information Security*.
- Herbert A. Simon. 1957. *Models of man, social and rational: Mathematical essays on rational human behavior in a social setting*. Wiley, New York, NY, USA.
- Herbert A. Simon. 1982. *Models of bounded rationality: Empirically grounded economic reason*. MIT Press.

- Manya Sleeper, Justin Cranshaw, Patrick Gage Kelley, Blase Ur, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. 2013. "I read my Twitter the next morning and was astonished": A conversational perspective on Twitter regrets. In *Proc. Conf. Human Factors in Comp. Sys. (CHI '13)*. ACM, 3277–3286.
- H. Jeff Smith, Tamara Dinev, and Heng Xu. 2011. Information privacy research: An interdisciplinary review. *MIS Quarterly* 35, 4 (2011), 989–1015.
- Sarah Spiekermann, Jens Grossklags, and Bettina Berendt. 2001. E-privacy in 2nd generation e-commerce: Privacy preferences versus actual behavior. In *Proc. Conf. Electronic Commerce (EC '01)*. ACM, 38–47.
- Frank Stajano and Paul Wilson. 2011. Understanding scam victims: Seven principles for systems security. *Commun. ACM* 54, 3 (March 2011), 70–75.
- George J. Stigler. 1980. An introduction to privacy in economics and politics. *J. Legal Stud.* 9, 4 (1980), 623–644.
- Fred Stutzman, Ralph Gross, and Alessandro Acquisti. 2013. Silent listeners: The evolution of privacy and disclosure on Facebook. *Journal of Privacy and Confidentiality* 4, 2 (2013), 7–41.
- Cass R. Sunstein. 2012. The Storrs Lectures: Behavioral economics and paternalism. *Yale Law J.* 122, 7 (2012), 1826.
- Joshua Tan, Khanh Nguyen, Michael Theodorides, Heidi Negrón-Arroyo, Christopher Thompson, Serge Egelman, and David Wagner. 2014. The effect of developer-specified explanations for permission requests on smartphone user behavior. In *Proc. Conf. Hum. Fact. Comp. Sys. (CHI '14)*. ACM, 91–100.
- David G. Taylor, Donna F. Davis, and Ravi Jillapalli. 2009. Privacy concern and online personalization: The moderating effects of information control and compensation. *Electr. Commer. Res.* 9, 3 (2009), 203–223.
- Humphrey Taylor. 2003. *Most people are "privacy pragmatists" who, while concerned about privacy, will sometimes trade it off for other benefits*. Technical Report. Harris Interactive.
- Tennessee Department of Transportation. 2014. Tennessee Highway Fatalities. <http://www.tdot.state.tn.us/ghso/thf.htm>. (June 2014).
- Richard H. Thaler. 1981. Some empirical evidence on dynamic inconsistency. *Econ. Lett.* 8, 3 (1981), 201–207.
- Richard H. Thaler and Cass R. Sunstein. 2008. *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press, New Haven, CT, USA.
- Janice Y. Tsai, Serge Egelman, Lorrie Faith Cranor, and Alessandro Acquisti. 2011. The effect of online privacy information on purchasing behavior: An experimental study. *Information Systems Research* 22, 2 (2011), 254–268.
- Janice Y. Tsai, Patrick Gage Kelley, Paul Drielsma, Lorrie Faith Cranor, Jason Hong, and Norman Sadeh. 2009. Who's viewed you?: The impact of feedback in a mobile location-sharing application. In *Proc. Conf. Human factors in computing systems (CHI '09)*. ACM, 2003–2012.
- Amos Tversky and Daniel Kahneman. 1975. Judgment under uncertainty: Heuristics and biases. In *Utility, probability, and human decision making*, Dirk Wendt and Charles Vlek (Eds.). Springer, 141–162.
- Amos Tversky and Daniel Kahneman. 1981. The framing of decisions and the psychology of choice. *Science* 211, 4481 (1981), 453–458.
- Blase Ur, Patrick Gage Kelley, Saranga Komanduri, Joel Lee, Michael Maass, Michelle Mazurek, Timothy Passaro, Richard Shay, Timothy Vidas, Lujo Bauer, Nicholas Christin, and Lorrie Faith Cranor. 2012a. How does your password measure up? The effect of strength meters on password creation. In *Proc. USENIX Security Symposium*. USENIX Association, Berkeley, CA, USA, 1–16.
- Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012b. Smart, useful, scary, creepy: Perceptions of online behavioral advertising. In *Proc. Eighth Symposium on Usable Privacy and Security (SOUPS '12)*. ACM, New York, NY, USA, 1–15.
- Hal R. Varian. 1996. Economic aspects of personal privacy. In *Privacy and self-regulation in the information age*. U.S. Department of Commerce.
- Hal R. Varian. 2000. Economic scene: Managing online security risks. *New York Times* (June 2000).
- Tony Vila, Rachel Greenstadt, and David Molnar. 2003. Why we can't be bothered to read privacy policies: Models of privacy economics as a lemons market. In *Proc. Conf. Electr. Commer. (EC '03)*. ACM, 403–407.
- Yang Wang, Gregory Norcie, Saranga Komanduri, Alessandro Acquisti, Pedro Giovanni Leon, and Lorrie Faith Cranor. 2011. I regretted the minute I pressed share: A qualitative study of regrets on Facebook. In *Proc. Symp. Usable Privacy and Security (SOUPS '11)*. ACM, 1–16.
- Yang Wang, Pedro Giovanni Leon, Alessandro Acquisti, Lorrie Faith Cranor, Alain Forget, and Norman Sadeh. 2014. A field trial of privacy nudges for Facebook. In *Proc. 32nd Annual ACM Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 2367–2376.
- Webroot. 2010. Social media sobriety test. http://www.webroot.com/En_US/sites/sobrietytest/. (2010).

- Alma Whitten and J. Doug Tygar. 1999. Why Johnny can't encrypt: A usability evaluation of PGP 5.0. In *Proc. USENIX Security Symposium*. USENIX Association, Berkeley, CA, USA, 1–16.
- Lauren E. Willis. 2014. Why not privacy by default? *Berkeley Technology Law Journal* 29, 1 (2014), 1–57.
- Shomir Wilson, Justin Cranshaw, Norman Sadeh, Alessandro Acquisti, Lorrie Faith Cranor, Jay Springfield, Sae Young Jeong, and Arun Balasubramanian. 2013. Privacy manipulation and acclimation in a location sharing application. In *Proc. Conf. Pervasive and ubiquitous comp. (Ubicomp '13)*. ACM, 549–558.
- Jessica Wisdom, Julie S. Downs, and George Loewenstein. 2010. Promoting healthy choices: Information versus convenience. *American Economic Journal: Applied Economics* 2, 2 (2010), 164–178.
- Evan Wondrasek. 2010. Take control of your Facebook privacy with PrivacyDefender. (June 2010). <http://www.makeuseof.com/tag/control-facebook-privacy-privacydefender/>
- Joshua D. Wright and Douglas H. Ginsburg. 2012. Behavioral law and economics: Its origins, fatal flaws, and implications for liberty. *Northwestern University Law Review* 106, 3 (2012), 12–63.
- Haidong Xia and José Carlos Brustoloni. 2005. Hardening Web browsers against man-in-the-middle and eavesdropping attacks. In *Proc. Conf. World Wide Web (WWW '05)*. ACM, 489–498.

Received TBD; revised TBD; accepted TBD