

PS4TLA: Privacy Support for the Total Learning Architecture

Specification Document, Volume 3:

Summit Report

**CLEARED
For Open Publication**

Dec 04, 2018

10
Department of Defense
OFFICE OF PREPUBLICATION AND SECURITY REVIEW

Author information

Project lead: Bart P. Knijnenburg, Clemson University

Project team: David Cherry, Yang He, Reza Ghaiumy Anaraky, Moses Namara, Erin Ash

Other contributors to the report: Aigul Kaskina (Fribourg University), Alfred Kobsa (University of California, Irvine), Daricia Wilkinson (Clemson University), Heather Lipford (UNC Charlotte), Heather Patterson (Intel), Henry Sloan (Nyack High School), Hichang Cho (National University of Singapore), Hoda Mehrpouyan (Boise State University), Jen Caltrider (Mozilla Foundation), Kevin Sun (TU Eindhoven), L. Jean Camp (Indiana University), Martijn Willemsen (TU Eindhoven), Nicholas Proferes (University of Kentucky), Pamela Wisniewski (University of Central Florida), Shlomo Berkovsky (CSIRO, Australia), Xinru Page (Bentley University), Yang Wang (Syracuse University), Yaqoub Alsarkal (George Washington University)

Technical point of contact: Andy Johnson, Advanced Distributed Learning (ADL) Initiative

Recommended citation: Knijnenburg, B.P et al. (2018) “Privacy Support for the Total Learning Architecture: Summit Report”. PS4TLA Specification Document, vol. 3, Clemson University, Clemson, SC.

Introduction

How can we reconcile the need for extensive customizability with users’ apparent lack of skills and motivation to manage their own privacy settings? In this report we investigate **User-Tailored Privacy as means to support users’ privacy decision-making**. With User-Tailored Privacy (UTP), a system would first **measure** users’ privacy-related characteristics and behaviors, use this as input to **model** their privacy preferences, and then **adapt** the system’s privacy settings to these preferences (Figure 1). This adaptation could take the form of a default setting or a recommendation, either with or without an accompanying justification.

UTP aims to strike this balance between giving users no control over, or information about, their privacy at all (which will be insufficient in highly sensitive situations and may deter privacy-minded individuals) and giving them full control and information (which makes setting one’s privacy settings unmanageably complex). Arguably, UTP relieves some of the burden of the privacy decision from the user by providing the right privacy-related information and the right amount of privacy control that is useful, but not overwhelming or misleading. This way, it enables them to make privacy-related decisions within the limits of their bounded rationality.




Figure 1: A schematic overview of User-Tailored Privacy (UTP)

With the research on UTP still in its infancy, it is important that the solutions proposed to ADL have broad support from researchers in the privacy and user modeling community. In November 2017, PI Knijnenburg therefore organized a **User-Tailored Privacy Summit** to bring together a group of interested researchers in an effort to standardize existing UTP approaches. The goal of the summit was to garner broader support for user-tailored privacy and to generate this “best practices” report.

1 The Goals of User-Tailored Privacy

Privacy preservation is a serious consideration for users interacting online with various systems and other users. Key in preserving users' privacy is the fact that not all users demand an equal level of privacy in their online interactions. Indeed, while some users would like to minimize their exposure to online systems and peers, others prefer to reap the benefits (e.g. personalization, social networking) that such exposure affords.

As such, privacy experts advocate giving users a considerable amount of control over their privacy, exercised through "privacy settings". The availability of such privacy settings introduces trade-offs users have to make regarding the potential benefits associated with sharing information and potential risks stemming from an unauthorized use of this information [89]. A plethora of research has shown, however, that users are notoriously bad at making risk-benefit trade-offs related to their online privacy. One fundamental problem regarding privacy is thus:

how can we reduce the burden of privacy decision-making?

Alan Westin famously grouped individuals into three camps: "privacy fundamentalists," "privacy pragmatists," and the "privacy unconcerned" ([127] for clearest articulation). However, there have been a number of critiques launched against this categorization [122]. In principle, many of the complaints suggest that this work is overly reductive and ignores that kinds of contexts in which individuals may be comfortable revealing some information to some people under certain circumstances. Helen Nissenbaum's work on privacy as contextual integrity [87] points us to exactly that. To quote Nissenbaum, "Contextual integrity ties adequate protection for privacy to norms of specific contexts, demanding that information gathering and dissemination be appropriate to that context and obey the governing norms of distribution within it" (p. 119).

Research has routinely shown that users want to make varied decisions about the conditions under which their information is shared [79, 80]. However, the reality of many privacy preserving technologies today relies on blanket decision making for opting in or out. As a result, the extant technology is more reminiscent of Westin's model of privacy than the nuanced realities of Nissenbaum's model. A second fundamental problem is thus: **how can we respect individual and contextual differences that call for different privacy-related system behaviors?**

Finally, privacy decisions are not always restricted to individual users. For example, the privacy of photos on online social networks may be determined by the preferences of the people displayed in the photo, and an employee's privacy settings regarding work-related activities may be in part dictated by their employer. Even at the presumably "individual" level, privacy decisions involve a dialectic between the system or person requesting a piece of information, and the user who is asked to provide it. This leads to a third fundamental problem: **how can we reconcile different stakeholder values in making privacy decisions?**

1.1 Tailored Automation

In this paper we introduce automated and user-tailored privacy setting as an adequate solution to these problems. User-Tailored Privacy (UTP) implies the system deciding on behalf of the users what information can be shared, when it can be shared, and who can be granted access to the information. This automation of privacy setting has several notable advantages. First, it substantially reduces the cognitive load associated with privacy-related decision making. Second, it allows default privacy-related system behaviors to be dependent on the context of use rather than uniform across all possible scenarios. Third, it can incorporate individual differences around privacy perceptions and value of information. Fourth, it can encapsulate and reconcile existing legislations, system policies, and values of various stakeholders.

1.2 Tailored Awareness

That said, the risks of automated privacy setting should not be discounted. These can manifest in various ways; namely, openness to manipulations by system designers, increased user vulnerability in case of successful attacks, or potential consequence of inappropriate privacy settings. On top these, another notable risk is the mere over-reliance of users on such an automation, which arguably exacerbates these problems [68, 96]. We argue that the latter may lead to the loss of essential privacy-related decision-making skills, which may turn out valuable in other contexts.

Hence, we advocate here for an implementation of user-tailored privacy that reduces users' decision burden, but at the same time respects the importance of empowering users in their own decision making referring to privacy settings. This requires a delicate balance between automating users' privacy settings, and keeping the user involved in—and informed about—the privacy decision process. We define the former as *tailored automation*, and the latter as *tailored awareness*. Tailored awareness relates to the privacy principle of “transparency”, but rather than aiming to provide complete information about a system's privacy practices (which arguably results in information overload [16]), it selectively makes users aware about those practices that matter to the user in that particular context of use.

1.3 Tailored Guidance

Beyond awareness, empowering users also means enabling them to take manual control over their privacy decisions where desired. However, as privacy settings are often numerous and complex, we introduce the idea of *tailored guidance* as part of UTP. Tailored guidance can involve privacy nudges and suggestions that would highlight the potential benefits or risks of a particular privacy decision and/or point out the relevance of a certain requested piece of information with regard to the stated purpose of its collection. For example, Facebook uses a Privacy Dinosaur as a privacy assistant to help guide users through the available privacy settings. Researchers have also explored the use of nutrition labels [53] and visual cues [101] that serve as warnings of suspicious activity to empower users to make better privacy decisions.

However, users are often inattentive to risk communication and they lack sufficient knowledge to understand potential risks, even when communicated effectively. Furthermore, users' preferences are contextual [18] since they are influenced by individual characteristics such as age [75], culture [26, 74] and personality that would shape how a person views a particular privacy decision as being too risky or not. Therefore, although privacy nudges help to relieve the burden of privacy decision-making, the traditional one-size-fits all nudges are not effective for every user, since what would be perceived as the “right” privacy decision varies greatly among users. The idea of tailored guidance is to acknowledge these unique differences in user privacy preferences and behaviors by tailoring these nudges and suggestions to match the individual preferences and behaviors of each user.

1.4 Tailored Education

While guidance may provide users with information within the context of specific decisions and settings, that guidance may not adequately inform people of the broader privacy-related risks and the variety of ways to mitigate them. Thus, a final step in empowering users is to educate them about privacy, and possible practices for managing their privacy. We acknowledge that users likely have a very limited amount of time and attention they will spend towards learning about privacy and may simply choose to ignore any kind of education or training. Thus, we introduce *tailored education* as part of UTP to emphasize the importance of adapting this privacy education to the user's level of privacy knowledge and their particular values with regards to privacy. Users are more likely to be motivated to engage with educational materials and retain their knowledge if the learning is personalized to their needs and context.

1.5 Selective application

Selective application is a common thread in the various tailored procedures described above. Indeed, at the core of user-tailored privacy is the decision to sometimes automate the privacy settings, and at other times increase users' awareness, guide their decisions, and/or educate them about their privacy. The system can make this decision based on various factors. For one: not all privacy decisions are equally impactful and important—the system should arguably only demand the user's attention when their involvement has a significant impact on their privacy. Moreover, the system may not always be confident about the user's desired privacy setting—in that case handing over control to the user can prevent misspecified settings. Finally, not all users will be equally keen on taking control over their settings. The level of automation versus involvement may thus also be tailored to the user.

2 Contextualizing User-Tailored Privacy

In section 1, we argued that UTP can tailor its support to several contextual factors and the unique needs of individual users. In this chapter, we draw from multiple frameworks and theories that hail from user-centered design and the privacy literature to create an integrated set of considerations and guidelines for contextualizing user-tailored privacy to and beyond end users. Figure 2 summarizes these factors and, in the sections below, we draw from the literature to show how these factors are related, at times overlap, and why they should be considered when tailoring privacy for end users.




Figure 2: Guidelines for Contextual User-Centered Privacy Design

2.1 Social Setting

In Figure 2, the highest level of consideration goes to the social setting in which a technology will be used. For example, social media may be leveraged for both social and functional purposes and cross boundaries for how it is primarily used between home settings and in the workplace [93]. According to the framework of Contextual Integrity [88], social contexts are an organizing principle of social life, structured social settings characterized by “canonical activities, roles, relationships, norms (or rules), and internal values (goals, ends, purposes).” The framework posits that what people desire is not control over their information, but rather an assurance of *appropriate* information sharing, where appropriateness is evaluated against a backdrop of context-specific ends and values that shape what type information, and about whom, may be transmitted by whom and to whom, and under what constraints. The social setting is key to this

analysis, in that it provides an organizing schema for understanding the implications of particular values being threatened. In a healthcare context, for example, a fair allocation of benefits may be risked by wayward information flows, whereas in a workplace context, productivity and freedom from discrimination under the terms of a health-contingent wellness plan may be at issue. Within the home, trust and safety are particularly implicated. Given these considerations, we argue that UTP should tailor its support to the social setting.

It is important to note that a *social setting* is many times associated with a physical location, but it can also be virtual. Furthermore, there is a difference between *space*, or the physical/virtual location, and *place*, the *social setting* or the social meaning created and associated with that space [43]. For example, a given room can be a classroom by day and a social gathering area by evening time. These are two different *places* in the same *space*. The space has not changed but the social meaning ascribed has changed because of the time of day. Similarly, online, the context surrounding the chatroom, newsfeed, virtual world, etc. can change the social meaning and norms associated with how it is used. As such, UTP should not equate social setting with physical location. Instead, it should attempt to integrate various contextual variables to infer the social meaning of the place in which it is employed.

2.2 Technological Domain

The technological domain refers to material or digital manifestation of the technology itself. Even in a given setting (i.e. home), end users may have different privacy preferences based on the unique affordances offered by a given technology [36]. As an example, voice-activated and remotely controlled Internet of Things (IoT) smart home devices are creating widespread privacy concerns [13] as they replace “dumb” appliances that served similar purposes but with fewer capabilities. This may be due to the different information types that can be collected via different technological solutions. For instance, a mobile app for Smart Things or Amazon’s Alexa can both be used to open a homeowner’s garage door, but some users may be less comfortable with Alexa listening to them in order to provide this feature.

This is where it is important to consider the design of a given technology and the importance of the affordances. The word “affordances” represents “the possibilities for goal-oriented actions afforded to specific user groups by technical objects” [78]. In other words, in the offline world a dial is designed to be turnable and a button is designed to be pushable, which is perceived by the user. Similarly, in the digital realm, technology affordances support certain tasks such as allowing users to share content with a broader audience at a lower cost than offline. Privacy researchers have started to investigate how the affordances of digital technologies shape privacy behaviors, attitudes, and expectations [41, 70]. Affordances such as editability and persistence of data impact privacy practices [117, 124]. In this light, UTP should tailor its support to the affordances of the technology to which it is applied. Moreover, as part of tailored guidance (see section 1.3), UTP has the opportunity to selectively emphasize privacy-related affordances that fit users’ privacy practices.

2.3 Actors

The theory of Contextual Integrity [88] also explains that privacy norms depend on the interactions between and among a given situation (i.e., context), the actors (i.e., senders, recipients, and subjects) involved, attributes of the information itself, and transmission principles for how information flows between actors. Actors encompass data subjects, senders, and recipients of information. What often goes awry is the inclusion of new actors as information senders or recipients, which inappropriately extends the reach of these entities into the data subjects' lives. When employers take on the role of wellness program providers, for example, they become behavioral police and cheerleaders in realms that fall outside of the traditional employer-employee relationship, leading to anxiety that one's behavior outside of working hours is now subject to scrutiny. UTP should take existing relationship boundaries into account when providing tailored privacy support [92]. Actors are people or groups of people (i.e., users) who send, receive, or are the subjects of information that flows through technology, but in some cases human actors are now being replaced by autonomous agents. UTP can be considered as part of this class of actors and needs to take the actions of other autonomous agents into account as well.

A number of theories frame privacy as a form of interpersonal boundary regulation, where individuals or groups must negotiate appropriate boundaries with others. Interpersonal boundaries are important because they help users define self, give protection (physically and emotionally), help manage our personal resources, and forge deeper relationships with others [8, 98]. For example, social psychologist Irwin Altman defined privacy as, “an interpersonal boundary process by which a person or group regulates interaction with others,” by altering the degree of openness of the self to others [8]. This process is dialectic in nature, balancing both the restriction and seeking of social interaction with others. The boundary regulation process allows for feedback and readjustment along with a dynamic need for varying levels of separateness and togetherness. According to Altman, boundary mechanisms are behaviors employed in combination and adjusted over time to achieve one's desired level of privacy. Individuals have different mechanisms for erecting boundaries, and they adjust these mechanisms as their needs change. Wisniewski et al. [130] built upon Altman's theory to empirically show how different users have different privacy management profiles on Facebook, which are related to their awareness of the privacy settings and features available to manage one's privacy desires. UTP can provide personalized support for these privacy management strategies.

Building on Altman's conceptualization of privacy, Petronio's Communication Privacy Management Theory (CPM) [98] outlined five suppositions related to disclosure boundaries. First, disclosure privacy deals specifically with the disclosure of private information. Second, a boundary exists between private and public information. Third, individuals have a sense of ownership or control in regard to this private information. Fourth, a rule-based system defines how individuals manage this privacy boundary. The rule management operations associated with this supposition include boundary linkages, co-ownership, and permeability. Boundary linkages are “connections that form boundary alliances” [98]. Co-ownership deals with the privilege to

have joint ownership of one's private information, and permeability deals with "how opened or closed the collective boundaries are once they are formed" [98]. Therefore, disclosure boundaries require a coordination process between co-owners of private information. Fifth, this process is dialectical in nature. Here, Petronio drew from Altman's theory to reiterate that an individual's desire for information privacy may change over time.

CPM delineated between two different interpersonal boundaries: personal and collective. Personal boundaries deal with how one shares private information about one's self, while collective boundaries involve private information shared with others. "A boundary is transformed from a personal to a collective when someone self-discloses to a confidant," explained Petronio [98]. A number of researchers have extended Petronio's CPM theory into the domain of Human-Computer interaction by trying to design interfaces and create models to help users understand and alleviate collective privacy concerns [34, 46]. As mentioned in the section 1, one goal of UTP is to reconcile different stakeholder values. As such it can be used to help the alleviate privacy concerns and expectations of multiple users who are co-managing a collective disclosure boundary. Reconciling the needs and desires of multiple actors is a prominent research problem in decision-support systems, and UTP will have to resolve this problem as well.

2.4 Information Types

Privacy norms also depend on the types of information being exchanged among and between actors. For instance, many users are willing to disclose general information (e.g., their preferences, gender, age, weight) when this information is not personally identifiable, but most are more hesitant sharing sensitive information (e.g., medical and financial data), especially when this data is associated with other information that identifies them as a person [1, 59]. Social media posts are another type of information that users may want to keep private among social connections. Users are also less inclined to allow systems to track their online activities such as their web browsing, email messages and credit card purchases, because they may worry that with this personal information incorrect inferences can be made by the person receiving this data [59]. Moreover, identifiable information can sometimes be inferred from other information disclosed by anonymous users [86] or even their social connections [140]. Recent advances in (wearable) IoT technology can also track users' biometrics, daily activities, and real time behavior—such information is often viewed as too personal to share publicly in raw form [14]. An important task for UTP is to identify what types of information should be collected and how they should be collected, and who this information should be shared with. This depends both on what data is necessary to fulfill the goal of the system or person requesting the information as well as the privacy norms of users and the context of use. This can be a difficult task, because users themselves often do not accurately set their own preferences [77], and their privacy decisions are prone to decision externalities [5, 60]. Finding out how these disclosure tendencies work, and how they correlate with each other can provide a more accurate description of behavior and allow UTP to find what types of information are appropriate to collect while keeping context in mind.

2.5 Norms and Values

Value-sensitive design [38] emphasizes the importance of norms and values in design processes, in general, and highlights how different values (e.g., privacy versus safety) can create tension among different stakeholders or actors. While norms and values are often stable and to a certain extent universal within communities, they differ per culture, and this can have a significant impact on how privacy boundaries are negotiated and perceived [25, 73]. Therefore, UTP must thus not only be cognizant of existing norms and values in the context of use.

A contextual integrity analysis similarly requires that we examine the norms and values that support the operation of any environment as it traditionally exists, so that we may identify how these norms and values may be challenged or strengthened by new technologies. In the context of health care, for example, Americans have historically reveled in the freedom to divulge information about intimate activities to trusted physicians and other practitioners, resting easy with the reasonable assurance that those communications will be confidential. Although shored up by legal protections for patient health information, norms also have root in the importance of the confidentiality principle underlying the Hippocratic oath: “I will respect the privacy of my patients, for their problems are not disclosed to me that the world may know” [115]. How might new technologies change this dynamic? To be effective, many believe that health monitoring must be pervasive, capturing a whole view of a person’s bodily states and behaviors in real time, 24/7. Curtailing a person’s ability to selectively withhold or conceal information about herself, or to contextualize it, may be experienced as an unwanted exposure or scrutiny, leading to guilt or shame. In sum, UTP must be aware of the way the information system it supports may challenge existing norms and values.

2.6 Threats and Benefits

Prior literature suggests that people tend make a tradeoff between the cost and gain of disclosing their personal information, a phenomenon known as privacy calculus [69]. On the “cost” side of this equation, Solove [109] proposed a taxonomy of privacy which includes four categories of “socially recognized privacy violations:”

- **Information collection** refers to the collection of personal or sensitive information about individuals such as their location or financial data.
- **Information processing** refers to the practices of using, storing or manipulating data that has already been collected, for instance, profiling individuals or using the collected data for a different purpose.
- **Information dissemination** refers to propagating or sharing information that has been collected, for instance, selling the collected data to a third party.
- **Invasion** means intrusion into people’s private lives, “disturbing their tranquility or solitude,” for instance, spam emails or stalking.

In practice, people are likely to differ in terms of what kinds of privacy violations that they consider important for themselves. For instance, collecting one's location data might unsettle

some people but not others. Therefore, UTP should take these individual perceptions and valuations of privacy violations into account when helping users with their privacy decisions.

Moreover, UTP should weigh these costs against the benefits of disclosure. We note, though, that the relative weights of these aspects and even the weighting function itself are matters that currently remain unresolved. Moreover, privacy does not always have to be considered a “cost” that is in direct opposition of benefit—it can itself also be considered a benefit. For example, Solove [110] points out that privacy is crucial to freedom of speech and democracy.

Users themselves may not always weigh costs and benefits in what researchers might consider a rational way. Much research and the popular press focus on how privacy concerns present a barrier to technology adoption and use. However, studies that try to predict information disclosure or technology usage have produced mixed results, often showing behavior that does not reflect people’s stated concerns [3, 15, 120]. This mismatch between stated concerns and actual behavior is a widely acknowledged “privacy paradox” [15, 89]. Various streams of research (for recent reviews, see [17, 40, 64]) have tried to uncover additional contextual factors that can help explain the privacy paradox [129, 136], claim that privacy concerns need to be measured differently [71, 123], or point to how human perceptions are subject to heuristics and decision externalities [6, 12, 20, 42, 114].

2.7 Uncertainty and Change

A perennial problem with privacy decisions is that the costs (and sometimes also the benefits) of disclosure are delayed and uncertain [4, 42, 72] which makes it impossible to make purely deterministic cost-benefit calculations. UTP can use stochastic decision models to resolve this issue, although it is difficult to predict the severity of potential unanticipated consequences. Moreover, UTP should also take into account the fact that human decisions are often biased to avoid uncertainty and delayed outcomes [4, 49, 121]—UTP should thus also take the amount of uncertainty and delay itself into account as a negative aspect of the decision context.

Another important phenomenon is the fact that privacy preferences and practices evolve over time [9, 81, 113]. This suggests that the model UTP creates of the user should not be static but instead update continuously to allow for changes in users’ attitudes and behaviors.

2.8 Design and Implementation

Users’ privacy considerations are also strongly influenced by the design and implementation of the system they are using. Many researchers and policymakers acknowledge that designing and implementing systems according to the “privacy-by-design” philosophy can be used to significantly improve users’ privacy [23, 31, 106]. UTP should acknowledge the privacy-related design and implementation characteristics of the system it is implemented in as a baseline for users’ privacy practices and a context and constraint for its own operation. Moreover, UTP can leverage the privacy-by-design philosophy by adding a user-tailored component to it [128].

3 User Privacy Decision-Making

Section 1 denoted UTP as a process that is meant to support users' privacy decision-making practices. Decision processes play a crucial role in decision-making, and privacy decisions are in this sense no different from other complex decision-making processes that involve difficult tradeoffs (e.g., between long and short-term benefits). While existing research has focused on external factors influencing privacy decisions, we argue that studying the mental processes behind privacy decision-making enables us to discern what internal factors drive privacy decision-making and which among these can be influenced or supported by UTP (Figure 3). Such factors may for instance play a role in reconciling between various decision theories and help UTP navigate these different approaches to privacy. For example, some privacy decisions are better handled by habitual, heuristic decision making, which can be fully automated to release the burden of decision-making. Other privacy decisions instead benefit from a more elaborate approach—in these cases UTP can help users make more informed decisions.




Figure 3: Theoretical model of the privacy decision process

3.1 Measuring privacy decision processes

Studying privacy decision processes is challenging because these processes occur in a user's head. Consequently, we can only study outward manifestations of these processes. Self-reported privacy concerns and behavioral intentions do not give accurate insights into users' privacy decision processes—in fact, they have been shown to be unreliable predictors of people's subsequent actual privacy behavior [89]. Researchers should therefore focus on the measurement of actual privacy behavior, ideally *in vivo* as part of people's daily lives and not in a lab or explicit experimental setting since contextual factors influence such behaviors.

Beyond observing behavioral outcomes of the decision process, an attempt can be made to measure lower-level behaviors in order to “trace” these processes. Existing methods for process tracing include eye tracking, mouse tracking, and aspect listing. Eye tracking and mouse tracking assume that individuals’ minds follow their attention, and that their attention can be measured by their gaze or mouse movements. While gaze is usually a more reliable measure of attention than mouse movement, measuring it requires specialized hardware.

Mouse movement can be made more reliable using tools like MouselabWEB [97], which hide on-screen information in boxes that only reveal the information when the cursor is moved over them. This allows researchers to use mouse movement to track what information participants are processing.

The concept of aspect listing is based on the “query theory”, which studies the decision-making process by decomposing a decision into a series of questions [48]. Aspect listing asks participants to list aspects that influenced their decision. By categorizing the listed aspects into the components of the decision and studying the order and number of listed aspects, researchers can track the relative weight and prominence of each of these components in the mental process. For example, in privacy-decisions aspect listing can be used to find out whether users are more focused on the (lack of) benefits or the (lack of) risks in deciding whether to disclose a certain piece of information.

The mentioned techniques for measuring privacy decision processes can be used to conduct research that can inform the user-modeling and adaptation aspects of UTP. Additionally, to the extent that these measurement methods can be implemented by instrumenting existing systems, they can be used to gather direct user input for UTP’s user-modeling efforts.

3.2 Theories related to privacy decision-making

In studying and leveraging privacy decision processes, UTP can rely on several existing theories on decision-making that apply to privacy decision-making as well. The most prominent decision theory is the “privacy calculus”, which argues that users assess the benefits and the costs of disclosure ([69]; see section 2.6). Only when the overall benefits outweigh the costs, users are willing to disclose their personal information [30, 33]. The privacy calculus can be seen as a privacy-specific instance of utility maximization or expectancy-value theory [11, 72, 103, 112]. The expectancy-value theory states that people gather information about various aspects of each choice option, and assign a value to each of these aspects [37]. Utility maximization, in turn, states that people will trade off the different aspects and then choose the option that maximizes their utility [19, 108].

While the privacy calculus assumes that individuals make rational decisions based on a cost/benefit analysis, research has shown that this cost/benefit analysis is affected by heuristic influences such as immediate/instant gratifications [2, 66, 137]. These influences are prominent in privacy decision-making because while the benefits of information disclosure (e.g., social networking, personalization) are perceived to be near and instant, the costs (e.g., the potential

leaking or misuse of the collected data) are often distant, delayed, and uncertain. Hyperbolic discounting [67, 91] explains that people prefer rewards that arrive sooner rather than later as people overrate the present over the future. As such, people are willing to disclose personal information in favor of short-term rewards in spite of perceived long-term privacy risks.

Likewise, Construal Level Theory (CLT, [118, 119]) argues that increased temporal distance makes people focus on abstract thinking, which more confident in dealing with any problems [90]. On the other hand, construing the event in near future involves specificities of the event, which generates uncertainty and increases the perceived risk. Therefore, as people often construe privacy loss as a risk in the distant future, this reduces the perceived likelihood of privacy loss taking place and increases their perceived ability to mitigate the problem.

Other aspects that irrationally influence the risk and benefits in privacy decision-making include “decision framing”, which has been shown to have a strong impact on privacy behavior [2, 66, 137]. Reference point, anchoring, loss aversion and the endowment effect can explain the effect of decision framing. For instance, in Chen et al. [24] subjects were asked to make choices either based on positive or negative framing of risk and found that subjects made more risk averse choices with positive framing in comparison to negative framing. Likewise, studies on opt-in versus opt-out default privacy settings show the relevance of the status quo bias [47]: people overestimate the potential loss and underestimate the potential gain of changing the status quo (i.e., default settings) and thus maintain it instead [12]. Another explanation can be given by the endowment effect, which refers to how the ownership of an object (e.g., an endowed functional benefit) increases its value, thus causing people to perceive a loss if the object is traded (e.g., for increased privacy protection). Both theories rely heavily on prospect theory [50], which explains how individuals are loss-averse with regard to the reference point of a decision.

UTP can use privacy calculus as a prescriptive model for privacy decision-making by treating it as a computational framework for decision support [61]. However, in modeling users *existing* privacy decision-making behaviors, UTP has to take the irrational influences on the privacy calculus into account, lest these irrational influences are attributed to spurious external attributes (see section 4.1).

3.3 External influences

Despite our focus on internal decision-making processes, we need to understand the external factors that influence these processes. These external factors are related to (and influenced by) the contextual factors discussed in section 2, but while these contextual factors play a role at the macro-level, we here cover the micro-level factors influence users’ individual decisions. These external factors include but are not limited to the following:

- **Changes in the Information Boundary**—Users suffer from privacy violations when information is shared beyond its intended boundary [88]. For example, the introduction of Facebook friendship pages aggregated information from different situations onto a single page, causing potential misperceptions and raising users’ concerns [107].

- **Data accumulation and Identifiability**—People have an intuitive understanding that the accumulation of data about them increases the risk and potential severity of negative consequences related to disclosure [58]. Hence, they are more willing to share data anonymously rather than in an identified manner, because identified information tends to accumulate [113]. Note, though, that disclosed data in modern systems often has very high dimensionality, which makes it difficult if not impossible to truly de-identify such data [85]. In other words, accumulation occurs even when data is shared anonymously.
- **Privacy Assurances**—Privacy assurance mechanisms such as privacy and security enhancing technologies [28], certifications and seals of approval [134], and industry self-regulation and government regulations [135] play a key role in increasing individuals' trust and reducing their privacy concerns [132]. This in turn increases their willingness to disclose personal information [105].
- **Group Dynamics and Reciprocity**—Users alter their privacy-related norms to conform and comply with group norms [82] and may adjust their behavior based on interpersonal relationships [55]. Moreover, in social networks, the notion of reciprocity gives users a motivation to disclose information to foster social relationships, despite privacy concerns [65, 99].
- **Interdependent Privacy**— As individuals disclose personal information, they may reveal information that may pertain to more than one individual. This interdependent characteristic makes privacy decision-making a collective issue. Morlok [83] highlights that interdependent privacy concerns affect individuals' disclosure decisions regarding information that is co-owned by others. Conversely, decisions by others may influence the privacy concerns of the co-owners of the information [46].

UTP should take the mentioned external influences into account when modeling users' privacy decisions.

3.4 Feedback

Users' privacy decisions evolve by the virtue of feedback about the consequences of past decisions. Positive feedback highlights desired outcomes, while negative feedback indicates the occurrence of privacy violations. Feedback can be based on real data, but in the case of privacy, where outcomes are delayed and uncertain, expected outcomes can also be used for feedback. In the past, researchers have had success using text, infographics, comics, and video for such risk communication [39, 57, 138, 139]. UTP should give feedback to users about the (expected) consequences of their decisions (or the decisions it made on their behalf) as this will improve their decisions in the long run (cf. education, see section 1.4).

Note though, that negative feedback can result in an overestimation of risks, which can reduce the quality of users' decision-making practices [94, 111]. Hence, UTP should be careful giving negative feedback. For example, rather than simply inform the the user about the (expected) results of a poor decision), it should give constructive feedback by teaching users better ways of making the decision in the future [45].

3.5 Open issues

Unfortunately, research on privacy decision-processes is far from complete. Future work needs to address the following research questions:

- **How do privacy decision processes change over time, and can we best account for these changes?**—Users’ privacy decision processes may change due to a combination of improved knowledge or awareness, evolving norms and values, or changes in existing information boundaries. Not only should UTP evolve its model of the user as a consequence of these changes; as privacy-decisions have delayed consequences, UTP should ideally predict future changes in users’ behaviors and account for them in current adaptation strategies.
- **Can theories help us to decide when to automate and when to empower users in their decision making?**—As mentioned in section 1.5, UTP may tailor the level of automation versus involvement in its approach to privacy support. No existing decision theories provide guidelines regarding the level of decision support to provide.
- **Could privacy negotiation be a good tool for empowering or supporting users?**—Most existing privacy decisions are static, one-sided propositions. *Negotiations* could increase the flexibility of privacy settings and may fit well with the idea of preference construction. While negotiations usually complicate a choice situation, UTP may be able to benefit from their flexibility to provide more fine-grained tailored support.

4 Technical solutions for User-Tailored Privacy

While the ideas behind UTP have been around for a while [63], system and implementation challenges have prevented these ideas from being implemented in real-world systems at a large scale. In this section we address the practical challenges of developing privacy modeling frameworks that are capable of tracking users’ privacy decisions, modeling users’ preferences and behaviors, and providing adaptive decision support in a way that empowers users with a careful balance of automation and personal control. UTP intends to decrease the burden of privacy decision-making by generating user-tailored suggestions such as recommendations, defaults, nudges, adaptations. While designing such systems, designers and developers should carefully choose appropriate mechanisms for user modeling and adaptation. We present some of the considerations regarding these aspects below, guided by the model in Figure 4.




Figure 4: Phases of automate privacy management systems

4.1 Modeling users’ knowledge

Two modeling components are important in provide user-tailored privacy support: knowledge and preferences. Modeling knowledge is important for two reasons: for one, it allows UTP to increase users’ knowledge in a tailored manner (cf. “tailored education”, see section 1.4). Moreover, knowledge plays a role in providing “tailored guidance” (see section 1.3), as such guidance may be more effective if it ties into users’ existing knowledge structures. Finally, a model of the user’s knowledge can help identify misconceptions that result in suboptimal decision-making practices. In other words: knowledge can be the key in disentangling the “rational” component of users’ preferences from irrational influences caused by a lack of knowledge or awareness.

Mental models can be used as a vehicle for knowledge modeling. For example, models such as those proposed by Camp [22] and Wash [125] have been instrumental in the development of warning messages that more closely relate to users’ conceptualizations of privacy [21]. UTP can employ mental models in a similar manner, and we recommend particularly to develop *tailored* models of users’ knowledge. Such tailoring of mental models is important, because mental models tend to differ substantially based on, e.g., users’ level of expertise [10, 27, 35, 54].

4.2 Modeling users' preferences

The other component relates to the modeling of users' privacy behaviors and preferences, usually through observation (see section 3.1), taking into account important contextual variables (see section 2) and influential external factors (see section 3.3). A typical practice is to build "privacy profiles" as digital vector representations of users' privacy behaviors and preferences.

Early versions of user privacy profiles were quantified as *unidimensional* representations [7], often based on the mathematical principles of Item-Response Theory [56, 76]). This mathematical representation reduces users' privacy preferences to a single score, and any profiles developed based on this representation follow one-dimensional pattern representing users' level of concern [84] akin to Westin's foundational classification (see section 1).

A substantial improvement was achieved when researchers started to model users' behavioral tendencies across several dimensions [52, 59, 131]. Clustering users along these dimensions results in profiles that vary not only in extent but also in kind, creating conceptually distinct disclosure profiles or privacy management strategies.

The most common clustering technique used for user privacy modeling assigns each user to the profile that is closest to their preferences or behaviors. Profiles are chosen to best fit the user population, but some users may not follow traditional behavioral patterns and will therefore be misclassified [51, 131]. To account for users who possess inherent characteristics of multiple profiles, user profiling techniques can employ *fuzzy* profiles. Fuzzy profiles avoid discriminative classification, thereby increasing the accuracy of the personalization. It has also been argued that fuzzy methods provide a better means to encapsulate users' cognitive perceptions of complex topics such as privacy [51].

The accuracy of privacy profiles can be further improved by the inclusion of contextual variables (cf. section 2). In fact, in situations where privacy decisions are repeated and plentiful, UTP can implement recommendation algorithms that make personalized and contextualized predictions for each individual user and situation, rather than relying on generic user profiles [44, 73, 95, 102, 104, 116, 126, 133].

4.3 Adaptation through Automation

The implementation of UTP calls for the generation of useful and appropriate adaptive actions or suggestions. This is a crucial component of user-tailored privacy, as all of the many benefits of UTP, from increased satisfaction to greater effective privacy control, rely on the system generating appropriately tailored results. The reason for this is that such adaptations create default (adaptive actions) or framing (adaptive suggestions) effects that not only make it more convenient for users to manage their privacy, but also that can also influence their perceptions or behaviors (see section 3.2). As such, errors made within this process propagate directly to the end user, reducing their satisfaction and potentially creating significant privacy violations.

Even a highly effective adaptive solution has the capacity to create security and ethical concerns entirely due to the nature of the implementation. For one, an over-reliance on automation can turn into a self-fulfilling prophecy and make users vulnerable to manipulations (see section 1.2). Another concern is the privacy concerns that may be generated by UTP's own data collection and automation: machine learning requires huge amounts of data to achieve effectiveness in most contexts, and UTP is no exception, creating an inverse relationship between benefit and privacy. Similarly, unexpected automations of users' privacy decisions may be regarded by the user as intrusive. UTP must take such considerations into account when collecting data and implementing adaptations.

4.4 Adaptation through Design

While in some contexts users will benefit from their privacy decisions being automated, in other cases it is better to give users more control over their privacy decisions. As stated by Camp [22] regarding the lack of adoption of privacy protective technologies, "if naive users are unable to effectively evaluate their own risks and decisions, no amount of technology will be empowering for those users." Indeed, in cases where users are incapable of making a correct cost-benefit estimation, automation can both improve upon users' privacy decision-making practices as well as relieve the burden that comes with privacy decision-making.

However, as discussed in the previous subsection, automation can propagate fallacious decision practices, especially when users' privacy behavior is deviant from their privacy preferences due to heuristic influences. There is also a risk that algorithms could shape users' privacy behaviors in a harmful manner [62]. In such cases "mixed initiative" systems (see section 5.1) can provide a solution: the system automates users' privacy decisions but asks for input or feedback on certain high-impact decisions based on a predefined set of rules.

At such moments, a good human-computer interaction plays a key role, and concepts from "privacy by design" [cite] can be combined with aspects of UTP to create user-friendly and easy-to-digest interfaces that are themselves adapted to the user's privacy management strategies (cf. [128]).

5 Future Challenges for User-Tailored Privacy

In this document we have envisioned UTP as a holistic solution to privacy problems. Consequently, we predict that researchers will continue to investigate and develop new ways to use UTP to support end-users' privacy decision-making practices. In this light, we present the following sociotechnical challenges.

5.1 Empowerment

The ultimate goal of User-Tailored Privacy is to empower users to make better privacy decisions. The philosophy behind this approach is that this cannot be achieved by inundating users with vast swathes of information and labyrinthian controls—neither users' motivation to control their privacy [29] nor their ability to do so effectively behavior [2, 66, 137] is sufficient to make such wholesale application of the “notice and consent” paradigm work [16]. Rather, the philosophy acknowledges that users *are* willing to take control over their privacy in cases where it really matters to them, and able to take control at a conceptual level that makes sense.

UTP does this by conceptualizing privacy as the shared responsibility between the system and the user. While this can likely be achieved in a traditional graphical user interface, an agent-based metaphor is arguably a more appropriate interaction paradigm for such a system. Consider “Petey”, an agent responsible for UTP on a certain platform (e.g. a smartphone or a social network). Petey unobtrusively automates low-impact privacy decisions for which it knows the user's preference with high confidence. But when it detects a decision that actually matters to the user, Petey will hand back control. Moreover, it will do so not via the traditional point-and-click interface displaying numerous intricate privacy settings but by asking the user specific privacy-related questions at an appropriate conceptual level, and not before giving the user an appropriate amount of information to inform this decision. Developing an autonomous privacy support system that seamlessly provides the optimal level of control is the ultimate challenge for UTP.

5.2 Risk management rather than prevention

A UTP system will take proactive preventive measures to prevent a loss of privacy, including warnings and explanations as to why certain user actions are wrong. It is important that a UTP system can make actions that are deemed detrimental difficult to achieve but should never fully prevent a user from taking any action. For example, UTP may explain to the user that they ought to know the person either as a friend or colleague before they should send a friend request, but it should not completely prevent the user from sending the request. Preventing paternalism ensures that users become aware of the consequences of the actions they are about to take and that they remain active participants in the privacy decision-making process.

On a related note, the goal of a UTP system is not to provide 100% privacy but to empower users to make better privacy decisions. In some cases, it is therefore better not to prevent users from

taking detrimental actions but rather figure out how mitigate the problem once the user feels like their privacy is being invaded. For example, if users have an unwanted experience unwanted online that makes them angry, fearful or feel like they lack control [32], a UTP system can champion for mitigation and ensure that the user learns from what went wrong and how to recover from such a breach of privacy. The mitigation strategies can include but are not limited to tightening their privacy settings, reporting the incident, and/or changing future behavior to prevent the same thing from happening again.

5.3 Explore the negative externalities

One of the critical questions for the introduction of any new piece of technology is understanding the scope of its impacts. To cite the work of media-scholar Neil Postman, “Technology giveth and technology taketh away. This means that for every advantage a new technology offers, there is always a corresponding disadvantage. The disadvantage may exceed in importance the advantage, or the advantage may well be worth the cost” [100]. UTP offers the potential to rectify many of the negative externalities associated with other technological inventions. However, an important part of the evaluation of UTP is understanding what the potential negative impacts are.

As part of this evaluation, it will be important for researchers to adopt a socio-technical perspective, because the effects of technological adoption depend on externalities flowing from in-situ adoption and use. For example, it may be important to understand how UTP adoption benefits certain stakeholders more than others. It may be important to understand whether UTP has long-term impacts on user attitudes towards data collection and use (see section 1). And it may be important to consider how UTP itself relies on data and may thus form a potential threat to the user’s privacy (see 4.3). Of course, the evaluation of whether the impacts of a given technology are negative or positive will depend in part on the normative stance adopted by the researcher.

Outcomes of this line of research can be used to help mitigate negative consequences. For example, if it is discovered that the implementation of UTP is negatively impacting the digital literacy of minors, specific policies might be adopted that could help guide appropriate uses.

5.4 How do we show that it works?

Success for privacy-enhancing technology is difficult to define, and even more challenging to measure. We can certainly argue that making the same privacy decisions in less time or with less cognitive burden is beneficial. Yet, if those more efficient decisions and settings are not meeting a user’s needs or matching a user’s preferences, then they are still the wrong decisions. Thus, a key challenge of UTP is to determine how to identify and measure the benefits of tailoring.

Users’ primary goal when interacting with technology is to gain whatever benefits that technology affords, where privacy is one critical value that users often have to weigh against other values. In other words, we cannot simply measure that users end up more private with

UTP, because too few disclosures may prevent technological systems from providing sufficient benefits, and this may actually run counter to users' goals. Privacy itself has multiple definitions, and the definition most appropriate may depend on the type of system and its interactions. Privacy can also be considered a dialectic process, where privacy preferences are not static but change with the interaction and the evolving context. Finally, there are also few standard measurements for privacy and privacy-related values and outcomes. With all those challenges, how should UTP be evaluated? While we believe that further research is needed in privacy evaluation, we have some initial thoughts as to guiding principles.

Privacy decisions can be considered a trade-off between the potential risks of disclosing information weighed against the benefits of interacting and sharing [a privacy calculus reference]. The goal of UTP is thus to help users more easily make this trade-off, empowering them to engage with privacy decision-making when desired, and relieving them of the burden when not desired. Thus, the evaluation of UTP will need to reflect on an individual's balance between caring enough to make an explicit decision, and trusting a system enough to intervene on a user's behalf. Measuring the feelings that relate to this, such as engagement, empowerment, and trust will provide insight into whether a solution is striking the right balance.

Evaluations of UTP will also vary depending on which aspects of UTP are employed—tailored automation, awareness, guidance or education. For example, in cases of tailored guidance or education, the goal may be to help users engage more frequently with appropriate privacy controls, whereas the goal of tailored automation may be to reduce the amount of decision making. All of this means that evaluating UTP is likely to be normative; there are few absolute outcomes that can be measured but instead must be evaluated against an individual's desires and values or against alternatives with tailoring.

Evaluations could also focus on reducing or eliminating the negative outcomes that arise when privacy is violated. For example, for social media applications, UTP could aim to minimize regret of oversharing. Users could also be queried as to their concerns over the risks of information disclosures, or the likelihood of privacy problems. However, we also note that users may not be accurate in their perceptions; their perceptions do not always match real risks. Privacy violations can be rare, and difficult to trace, for many applications.

6 Conclusion

In this report, we have set out to describe the concept of User-Tailored Privacy and elaborated on this concept from the perspective of the sociotechnical context in which it operates, the theories of decision-making that govern its use, and the technical parameters that determine its effective implementation.

UTP offers a model for increasing the precision of privacy decision making, allowing them to dictate the contexts in which they wish to share information—at scale. It offers the opportunity to meet consumer demands without unrealistically increasing their burden. In doing so, UTP offers the potential to increase consumer trust in the ways their data is handled, offers users the ability to make more informed decisions about information disclosure, and offers the possibilities for certain kinds of repeated decision making to be offloaded. UTP is an essential part of the solution to the data privacy problem, that, if complemented by robust technical solutions, can provide a holistic privacy solution.

References

- [1] Ackerman, M.S., Cranor, L.F. and Reagle, J. 1999. Privacy in e-commerce: examining user scenarios and privacy preferences. *Proceedings of the 1st ACM conference on electronic commerce* (Denver, CO, 1999), 1–8.
- [2] Acquisti, A. 2004. Privacy in electronic commerce and the economics of immediate gratification. *Proceedings of the 5th ACM conference on Electronic commerce* (New York, NY, 2004), 21–29.
- [3] Acquisti, A., & Gross, R. 2006. Imagined communities: Awareness, information sharing, and privacy on the Facebook. *Privacy enhancing technologies*. Springer Berlin Heidelberg. 36–58.
- [4] Acquisti, A. and Grossklags, J. 2003. Losses, Gains, and Hyperbolic Discounting: An Experimental Approach to Information Security Attitudes and Behavior. *2nd Annual Workshop on Economics and Information Security* (College Park, MD, 2003).
- [5] Acquisti, A. and Grossklags, J. 2008. What Can Behavioral Economics Teach Us About Privacy? *Digital Privacy: Theory, Technologies, and Practices*. A. Acquisti, S. De Capitani di Vimercati, S. Gritzalis, and C. Lambrinoudakis, eds. Auerbach Publications. 363–377.
- [6] Adjerd, I., Acquisti, A., Brandimarte, L. and Loewenstein, G. 2013. Sleights of Privacy: Framing, Disclosures, and the Limits of Transparency. *Proceedings of the Ninth Symposium on Usable Privacy and Security* (New York, NY, USA, 2013), 9:1–9:11.
- [7] Aïmeur, E., Gambs, S. and Ho, A. 2009. UPP: User Privacy Policy for Social Networking Sites. *2009 Fourth International Conference on Internet and Web Applications and Services* (May 2009), 267–272.
- [8] Altman, I. 1975. *The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding*. Brooks/Cole Publishing.
- [9] Anton, A.I., Earp, J.B. and Young, J.D. 2010. How Internet Users’ Privacy Concerns Have Evolved since 2002. *Security Privacy, IEEE*. 8, 1 (Feb. 2010), 21–27. DOI:<https://doi.org/10.1109/MSP.2010.38>.
- [10] Asgharpour, F., Liu, D. and Camp, L.J. 2007. Mental Models of Security Risks. *Financial Cryptography and Data Security* (2007), 367–377.
- [11] Awad, N.F. and Krishnan, M.S. 2006. The Personalization Privacy Paradox: An Empirical Evaluation of Information Transparency and the Willingness to be Profiled Online for Personalization. *MIS quarterly*. 30, 1 (Mar. 2006), 13–28.
- [12] Baek, Y.M. 2014. Solving the privacy paradox: A counter-argument experimental approach. *Computers in Human Behavior*. 38, (Sep. 2014), 33–42. DOI:<https://doi.org/10.1016/j.chb.2014.05.006>.
- [13] Bannan, C. 2016. The IoT threat to privacy. *TechCrunch*.
- [14] Barcena, M.B., Wueest, C. and Lau, H. 2014. *How safe is your quantified self? Tracking, monitoring, and wearable tech*. Symantech.
- [15] Barnes, S.B. 2006. A privacy paradox: Social networking in the United States. *First Monday*. 11, 9 (2006). DOI:<https://doi.org/10.5210/fm.v11i9.1394>.
- [16] Barocas, S. and Nissenbaum, H. 2009. On notice: The trouble with Notice and Consent. *Proceedings of the Engaging Data Forum: The First International Forum on the Application and Management of Personal Electronic Information* (2009).
- [17] Barth, S. and de Jong, M.D.T. 2017. The privacy paradox – Investigating discrepancies between expressed privacy concerns and actual online behavior – A systematic literature review. *Telematics and Informatics*. 34, 7 (Nov. 2017), 1038–1058. DOI:<https://doi.org/10.1016/j.tele.2017.04.013>.
- [18] Benisch, M., Kelley, P.G., Sadeh, N. and Cranor, L.F. 2011. Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs. *Personal Ubiquitous Computing*. 15, 7 (Oct. 2011), 679–694. DOI:<https://doi.org/10.1007/s00779-010-0346-0>.
- [19] Bettman, J.R., Luce, M.F. and Payne, J.W. 1998. Constructive consumer choice processes. *Journal of consumer research*. 25, 3 (1998), 187–217.
- [20] Brandimarte, L., Acquisti, A. and Loewenstein, G. 2013. Misplaced Confidences: Privacy and the Control Paradox. *Social Psychological and Personality Science*. 4, 3 (2013), 340–347. DOI:<https://doi.org/10.1177/1948550612455931>.
- [21] Bravo-Lillo, C., Cranor, L.F., Downs, J.S. and Komanduri, S. 2011. Bridging the Gap in Computer Security Warnings: A Mental Model Approach. *IEEE Security & Privacy*. 2, 9 (2011), 18–26. DOI:<https://doi.org/10.1109/MSP.2010.198>.

- [22] Camp, L.J. 2009. Mental models of privacy and security. *IEEE Technology and Society Magazine*. 28, 3 (Fall 2009), 37–46. DOI:<https://doi.org/10.1109/MTS.2009.934142>.
- [23] Cavoukian, A. 2009. *Privacy by Design*. Information and Privacy Commissioner of Ontario, Canada.
- [24] Chen, J., Gates, C.S., Li, N. and Proctor, R.W. 2015. Influence of Risk/Safety Information Framing on Android App-Installation Decisions. *Journal of Cognitive Engineering and Decision Making*. 9, 2 (Jun. 2015), 149–168. DOI:<https://doi.org/10.1177/1555343415570055>.
- [25] Cho, H., Knijnenburg, B., Kobsa, A. and Li, Y. 2018. Collective Privacy Management in Social Media: A Cross-Cultural Validation. *ACM Trans. Comput.-Hum. Interact.* 25, 3 (Jun. 2018), 17:1–17:33. DOI:<https://doi.org/10.1145/3193120>.
- [26] Cho, H., Knijnenburg, B.P., Kobsa, A. and Li, Y. 2017. Collaborative Privacy Management in Social Media: A Cross-cultural Validation. *Submitted for journal publication*. (2017).
- [27] Choe, E.K., Jung, J., Lee, B. and Fisher, K. 2013. Nudging people away from privacy-invasive mobile apps through visual framing. *IFIP Conference on Human-Computer Interaction* (2013), 74–91.
- [28] Cimino, J.J., Patel, V.L. and Kushniruk, A.W. 2002. The patient clinical information system (PatCIS): technical solutions for and experience with giving patients access to their electronic medical records. *International Journal of Medical Informatics*. 68, 1 (Dec. 2002), 113–127. DOI:[https://doi.org/10.1016/S1386-5056\(02\)00070-9](https://doi.org/10.1016/S1386-5056(02)00070-9).
- [29] Compañó, R. and Lusoli, W. 2010. The Policy Maker’s Anguish: Regulating Personal Data Behavior Between Paradoxes and Dilemmas. *Economics of Information Security and Privacy*. T. Moore, D. Pym, and C. Ioannidis, eds. Springer US. 169–185.
- [30] Culnan, M.J. and Armstrong, P.K. 1999. Information Privacy Concerns, Procedural Fairness, and Impersonal Trust: An Empirical Investigation. *Organization Science*. 10, 1 (1999), 104–115. DOI:<https://doi.org/10.1287/orsc.10.1.104>.
- [31] Davies, N. and Langheinrich, M. 2013. Privacy By Design [From the Editor in Chief]. *IEEE Pervasive Computing*. 12, 2 (2013), 2–4. DOI:<https://doi.org/10.1109/MPRV.2013.34>.
- [32] Debatin, B., Lovejoy, J.P., Horn, A.-K. and Hughes, B.N. 2009. Facebook and Online Privacy: Attitudes, Behaviors, and Unintended Consequences. *Journal of Computer-Mediated Communication*. 15, 1 (2009), 83–108. DOI:<https://doi.org/10.1111/j.1083-6101.2009.01494.x>.
- [33] Dinev, T. and Hart, P. 2006. An Extended Privacy Calculus Model for E-Commerce Transactions. *Information Systems Research*. 17, 1 (Mar. 2006), 61–80. DOI:<https://doi.org/10.1287/isre.1060.0080>.
- [34] Dourish, P. and Anderson, K. 2006. Collective Information Practice: Exploring Privacy and Security as Social and Cultural Phenomena. *Human-Computer Interaction*. 21, 3 (Sep. 2006), 319–342. DOI:https://doi.org/10.1207/s15327051hci2103_2.
- [35] Egelman, S. and Peer, E. 2015. The Myth of the Average User: Improving Privacy and Security Systems Through Individualization. *Proceedings of the 2015 New Security Paradigms Workshop* (New York, NY, USA, 2015), 16–28.
- [36] Faraj, S. and Azad, B. 2012. The Materiality of Technology: An Affordance Perspective. *Materiality and Organizing: Social Interaction in a Technological World*. P.M. Leonardi, B.A. Nardi, and J. Kallinikos, eds. OUP Oxford.
- [37] Fishbein, M. and Ajzen, I. 1975. *Belief, attitude, intention, and behavior: an introduction to theory and research*. Addison-Wesley Pub. Co.
- [38] Friedman, B., Kahn, P.H. and Borning, A. 2006. Value Sensitive Design and Information Systems. *Human-Computer Interaction and Management Information Systems: Foundations*. M.E. Sharpe (2006), 348–372.
- [39] Garg, V., Camp, L.J., Connelly, K. and Lorenzen-Huber, L. 2012. Risk Communication Design: Video vs. Text. *Privacy Enhancing Technologies* (2012), 279–298.
- [40] Gerber, N., Gerber, P. and Volkamer, M. 2018. Explaining the privacy paradox: A systematic review of literature investigating privacy attitude and behavior. *Computers & Security*. 77, (Aug. 2018), 226–261. DOI:<https://doi.org/10.1016/j.cose.2018.04.002>.
- [41] Gibson, J.J. 2014. *The Ecological Approach to Visual Perception: Classic Edition*. Psychology Press.
- [42] Hallam, C. and Zanella, G. 2017. Online self-disclosure: The privacy paradox explained as a temporally discounted balance between concerns and rewards. *Computers in Human Behavior*. 68, (Mar. 2017), 217–227. DOI:<https://doi.org/10.1016/j.chb.2016.11.033>.

- [43] Harrison, S. and Dourish, P. 1996. Re-place-ing Space: The Roles of Place and Space in Collaborative Systems. *Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work* (New York, NY, USA, 1996), 67–76.
- [44] Ismail, Q., Ahmed, T., Kapadia, A. and Reiter, M.K. 2015. Crowdsourced Exploration of Security Configurations. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2015), 467–476.
- [45] Jessup, R.K., Bishara, A.J. and Busemeyer, J.R. 2008. Feedback Produces Divergence From Prospect Theory in Descriptive Choice. *Psychological Science*. 19, 10 (Oct. 2008), 1015–1022. DOI:<https://doi.org/10.1111/j.1467-9280.2008.02193.x>.
- [46] Jia, H. and Xu, H. 2016. Measuring individuals' concerns over collective privacy on social networking sites. *Cyberpsychology: Journal of Psychosocial Research on Cyberspace*. 10, 1 (May 2016).
- [47] Johnson, E.J., Bellman, S. and Lohse, G.L. 2002. Defaults, Framing and Privacy: Why Opting In ≠ Opting Out. *Marketing Letters*. 13, 1 (2002), 5–15. DOI:<https://doi.org/10.1023/A:1015044207315>.
- [48] Johnson, E.J., Häubl, G. and Keinan, A. 2007. Aspects of endowment: a query theory of value construction. *Journal of experimental psychology: Learning, memory, and cognition*. 33, 3 (2007), 461.
- [49] Kahneman, D., Slovic, P. and Tversky, A. 1982. *Judgment under uncertainty : heuristics and biases*. Cambridge University Press.
- [50] Kahneman, D. and Tversky, A. 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica*. 47, 2 (Mar. 1979), 263–292. DOI:<https://doi.org/10.2307/1914185>.
- [51] Kaskina, A. 2018. Exploring Nuances of User Privacy Preferences on a Platform for Political Participation. *2018 International Conference on eDemocracy eGovernment (ICEDEG)* (Apr. 2018), 89–94.
- [52] Kaskina, A. and Meier, A. 2016. Integrating privacy and trust in voting advice applications. *2016 Third International Conference on eDemocracy & eGovernment (ICEDEG)* (Mar. 2016), 20–25.
- [53] Kelley, P.G., Bresee, J., Cranor, L.F. and Reeder, R.W. 2009. A “nutrition label” for privacy. *Proceedings of the 5th Symposium on Usable Privacy and Security - SOUPS '09* (Mountain View, California, 2009), 1.
- [54] Kelley, T. and Bertenthal, B.I. 2015. Tracking Risky Behavior On The Web: Distinguishing Between What Users ‘Say’ And ‘Do’.”” *Proceedings of the Ninth International Symposium on Human Aspects of Information Security & Assurance* (2015).
- [55] Kimmel, A.J. 2009. *Ethical Issues in Behavioral Research: Basic and Applied Perspectives*. John Wiley & Sons.
- [56] Knijnenburg, B.P. 2015. *A user-tailored approach to privacy decision support*. University of California, Irvine.
- [57] Knijnenburg, B.P. and Cherry, D. 2016. Comics as a Medium for Privacy Notices. *SOUPS 2016 workshop on the Future of Privacy Notices and Indicators* (Denver, CO, Jun. 2016).
- [58] Knijnenburg, B.P. and Kobsa, A. 2013. Making Decisions about Privacy: Information Disclosure in Context-Aware Recommender Systems. *ACM Transactions on Interactive Intelligent Systems*. 3, 3 (2013), 20:1-20:23. DOI:<https://doi.org/10.1145/2499670>.
- [59] Knijnenburg, B.P., Kobsa, A. and Jin, H. 2013. Dimensionality of information disclosure behavior. *International Journal of Human-Computer Studies*. 71, 12 (2013), 1144–1162. DOI:<https://doi.org/10.1016/j.ijhcs.2013.06.003>.
- [60] Knijnenburg, B.P., Kobsa, A. and Jin, H. 2013. Preference-based location sharing: are more privacy options really better? *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France, 2013), 2667–2676.
- [61] Knijnenburg, B.P., Raybourn, E.M., Cherry, D., Wilkinson, D., Sivakumar, S. and Sloan, H. 2017. Death to the Privacy Calculus? *Proceedings of the 2017 Networked Privacy Workshop at CSCW* (Portland, OR, Feb. 2017).
- [62] Knijnenburg, B.P., Sivakumar, S. and Wilkinson, D. 2016. Recommender Systems for Self-Actualization. *Proceedings of the 10th ACM Conference on Recommender Systems* (New York, NY, USA, 2016), 11–14.
- [63] Kobsa, A. 2001. Tailoring Privacy to Users' Needs (Invited Keynote). *User Modeling 2001*. M. Bauer, P.J. Gmytrasiewicz, and J. Vassileva, eds. Springer Verlag. 303–313.
- [64] Kokolakis, S. 2017. Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & Security*. 64, (Jan. 2017), 122–134. DOI:<https://doi.org/10.1016/j.cose.2015.07.002>.
- [65] Krasnova, H., Spiekermann, S., Koroleva, K. and Hildebrand, T. 2010. Online social networks: why we disclose. *Journal of Information Technology*. 25, 2 (2010), 109–125. DOI:<https://doi.org/10.1057/jit.2010.6>.

- [66] Krasnova, H. and Veltri, N.F. 2010. Privacy Calculus on Social Networking Sites: Explorative Evidence from Germany and USA. *2010 43rd Hawaii International Conference on System Sciences (HICSS)* (Jan. 2010), 1–10.
- [67] Laibson, D. 1997. Golden Eggs and Hyperbolic Discounting. *The Quarterly Journal of Economics*. 112, 2 (May 1997), 443–478. DOI:<https://doi.org/10.1162/003355397555253>.
- [68] Lanier, J. 2010. *You Are Not a Gadget: A Manifesto*. Thorndike Press.
- [69] Laufer, R.S. and Wolfe, M. 2010. Privacy as a Concept and a Social Issue: A Multidimensional Developmental Theory. *Journal of Social Issues*. 33, 3 (Apr. 2010), 22–42. DOI:<https://doi.org/10.1111/j.1540-4560.1977.tb01880.x>.
- [70] Leonardi, P.M. 2011. When Flexible Routines Meet Flexible Technologies: Affordance, Constraint, and the Imbrication of Human and Material Agencies. *MIS Quarterly*. 35, 1 (2011), 147–167. DOI:<https://doi.org/10.2307/23043493>.
- [71] Li, H., Luo, X. (Robert), Zhang, J. and Xu, H. 2016. Resolving the privacy paradox: Toward a cognitive appraisal and emotion approach to online privacy behaviors. *Information & Management*. (2016). DOI:<https://doi.org/10.1016/j.im.2017.02.005>.
- [72] Li, Y. 2012. Theories in online information privacy research: A critical review and an integrated framework. *Decision support systems*. 54, 1 (Dec. 2012), 471–481. DOI:<https://doi.org/10.1016/j.dss.2012.06.010>.
- [73] Li, Y., Kobsa, A., Knijnenburg, B.P. and Nguyen, M.C. 2017. Cross-Cultural Privacy Prediction. *Proceedings on Privacy Enhancing Technologies*. 2, (2017), 93–112.
- [74] Li, Y., Kobsa, A., Knijnenburg, B.P. and Nguyen, M.-H.C. 2017. Cross-Cultural Privacy Prediction. *Proceedings on Privacy Enhancing Technologies*. 3, 2 (2017).
- [75] Litt, E. 2012. Knock, Knock. Who’s There? The Imagined Audience. *Journal of Broadcasting & Electronic Media*. 56, 3 (Jul. 2012), 330–345. DOI:<https://doi.org/10.1080/08838151.2012.705195>.
- [76] Liu, K. and Terzi, E. 2010. A Framework for Computing the Privacy Scores of Users in Online Social Networks. *ACM Trans. Knowl. Discov. Data*. 5, 1 (Dec. 2010), 6:1–6:30. DOI:<https://doi.org/10.1145/1870096.1870102>.
- [77] Madejski, M., Johnson, M. and Bellovin, S.M. 2012. A study of privacy settings errors in an online social network. *Fourth International Workshop on Security and Social Networking* (Lugano, Switzerland, 2012), 340–345.
- [78] Markus, M.L. and Silver, M. 2008. A Foundation for the Study of IT Effects: A New Look at DeSanctis and Poole’s Concepts of Structural Features and Spirit. *Journal of the Association for Information Systems*. 9, 10 (Oct. 2008).
- [79] Martin, K. and Shilton, K. 2016. Putting mobile application privacy in context: An empirical study of user privacy expectations for mobile devices. *The Information Society*. 32, 3 (May 2016), 200–216. DOI:<https://doi.org/10.1080/01972243.2016.1153012>.
- [80] Martin, K. and Shilton, K. 2016. Why experience matters to privacy: How context-based experience moderates consumer privacy expectations for mobile applications. *Journal of the Association for Information Science and Technology*. 67, 8 (Aug. 2016), 1871–1882. DOI:<https://doi.org/10.1002/asi.23500>.
- [81] McLaughlin, C. and Vitak, J. 2011. Norm evolution and violation on Facebook. *New Media & Society*. (2011). DOI:<https://doi.org/10.1177/1461444811412712>.
- [82] Min, J. and Kim, B. 2014. How are people enticed to disclose personal information despite privacy concerns in social network sites? The calculus between benefit and cost. *Journal of the Association for Information Science and Technology*. (2014), n/a-n/a. DOI:<https://doi.org/10.1002/asi.23206>.
- [83] Morlok, T. 2016. Sharing is (not) caring—The role of external privacy in users’ information disclosure behaviors on social network sites. *PACIS 2016 Proceedings* (2016), 75.
- [84] Munemasa, T. and Iwaihara, M. 2011. Trend Analysis and Recommendation of Users’ Privacy Settings on Social Networking Services. *Social Informatics* (2011), 184–197.
- [85] Narayanan, A. and Felten, E.W. 2014. *No silver bullet: De-identification still doesn’t work*.
- [86] Narayanan, A. and Shmatikov, V. 2009. De-anonymizing Social Networks. *2009 30th IEEE Symposium on Security and Privacy* (May 2009), 173–187.
- [87] Nissenbaum, H. 2004. Privacy as Contextual Integrity. *Washington Law Review*. 79, (2004), 119–157.
- [88] Nissenbaum, H. 2009. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford Law Books.

- [89] Norberg, P.A., Horne, D.R. and Horne, D.A. 2007. The Privacy Paradox: Personal Information Disclosure Intentions versus Behaviors. *Journal of consumer affairs*. 41, 1 (2007), 100–126. DOI:<https://doi.org/10.1111/j.1745-6606.2006.00070.x>.
- [90] Nussbaum, S., Liberman, N. and Trope, Y. 2006. Predicting the near and distant future. *Journal of Experimental Psychology: General*. 135, 2 (2006), 152–161. DOI:<https://doi.org/10.1037/0096-3445.135.2.152>.
- [91] O’Donoghue, T. and Rabin, M. 2000. The economics of immediate gratification. *Journal of Behavioral Decision Making*. 13, 2 (Apr. 2000), 233–250. DOI:[https://doi.org/10.1002/\(SICI\)1099-0771\(200004/06\)13:2<233::AID-BDM325>3.0.CO;2-U](https://doi.org/10.1002/(SICI)1099-0771(200004/06)13:2<233::AID-BDM325>3.0.CO;2-U).
- [92] Page, X., Kobsa, A. and Knijnenburg, B.P. 2012. Don’t Disturb My Circles! Boundary Preservation Is at the Center of Location-Sharing Concerns. *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media* (Dublin, Ireland, May 2012), 266–273.
- [93] Page, X., Wisniewski, P., Knijnenburg, B.P. and Namara, M. 2018. Social Media’s Have-Nots: An Era of Social Disenfranchisement. *Internet Research*. (Aug. 2018). DOI:<https://doi.org/10.1108/IntR-03-2017-0123>.
- [94] Paich, M. and Serman, J.D. 1993. Boom, Bust, and Failures to Learn in Experimental Markets. *Management Science*. 39, 12 (Dec. 1993), 1439–1458. DOI:<https://doi.org/10.1287/mnsc.39.12.1439>.
- [95] Pallapa, G., Das, S.K., Di Francesco, M. and Aura, T. 2014. Adaptive and context-aware privacy preservation exploiting user interactions in smart environments. *Pervasive and Mobile Computing*. 12, (Jun. 2014), 232–243. DOI:<https://doi.org/10.1016/j.pmcj.2013.12.004>.
- [96] Pariser, E. 2012. *The filter bubble: how the new personalized Web is changing what we read and how we think*. Penguin Books.
- [97] Payne, J.W., Bettman, J.R. and Johnson, E.J. 1993. *The Adaptive Decision Maker*. Cambridge University Press.
- [98] Petronio, S.S. 2002. *Boundaries of Privacy: Dialects of Disclosure*. SUNY Press.
- [99] Posey, C., Lowry, P.B., Roberts, T.L. and Ellis, T.S. 2010. Proposing the online community self-disclosure model: the case of working professionals in France and the U.K. who use online communities. *European Journal of Information Systems*. 19, 2 (Apr. 2010), 181–195. DOI:<https://doi.org/10.1057/ejis.2010.15>.
- [100] Postman, N. 1998. Five Things We Need to Know About Technological Change.
- [101] Rajivan, P. and Camp, J. 2016. Influence of Privacy Attitude and Privacy Cue Framing on Android App Choices. *Symposium on Usable Privacy and Security (SOUPS)* (2016).
- [102] Ravichandran, R., Benisch, M., Kelley, P. and Sadeh, N. 2009. Capturing Social Networking Privacy Preferences: *Privacy Enhancing Technologies*. I. Goldberg and M. Atallah, eds. Springer Berlin / Heidelberg. 1–18.
- [103] Rust, R.T., Kannan, P.K. and Peng, N. 2002. The Customer Economics of Internet Privacy. *Journal of the Academy of Marketing Science*. 30, 4 (Oct. 2002), 455–464. DOI:<https://doi.org/10.1177/009207002236917>.
- [104] Sadeh, N., Hong, J., Cranor, L., Fette, I., Kelley, P., Prabaker, M. and Rao, J. 2009. Understanding and capturing people’s privacy policies in a mobile social networking application. *Personal and Ubiquitous Computing*. 13, 6 (2009), 401–412. DOI:<https://doi.org/10.1007/s00779-008-0214-3>.
- [105] Schoenbachler, D.D. and Gordon, G.L. 2002. Trust and Customer Willingness to Provide Information in Database-Driven Relationship Marketing. *Journal of Interactive Marketing*. 16, 3 (2002), 2–16. DOI:<https://doi.org/10.1002/dir.10033>.
- [106] Shapiro, S.S. 2010. Privacy by Design: Moving from Art to Practice. *Commun. ACM*. 53, 6 (Jun. 2010), 27–29. DOI:<https://doi.org/10.1145/1743546.1743559>.
- [107] Shi, P., Xu, H. and Chen, Y. 2013. Using contextual integrity to examine interpersonal information boundary on social network sites. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2013), 35–38.
- [108] Simon, H.A. 1959. Theories of Decision-Making in Economics and Behavioral Science. *The American Economic Review*. 49, 3 (Jun. 1959), 253–283. DOI:<https://doi.org/10.2307/1809901>.
- [109] Solove, D.J. 2005. *A Taxonomy of Privacy*. Technical Report #ID 667622. Social Science Research Network.
- [110] Solove, D.J. 2007. I’ve Got Nothing to Hide and Other Misunderstandings of Privacy. *San Diego Law Review*. 44, (2007), 745.
- [111] Serman, J.D. 1989. Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment. *Management Science*. 35, 3 (Mar. 1989), 321–339. DOI:<https://doi.org/10.1287/mnsc.35.3.321>.
- [112] Stone, E.F. and Stone, D.L. 1990. Privacy in Organizations: Theoretical Issues, Research Findings, and Protection Mechanisms. *Research in personnel and human resources management*. 8, (1990), 349–411.

- [113] Stutzman, F., Gross, R. and Acquisti, A. 2013. Silent Listeners: The Evolution of Privacy and Disclosure on Facebook. *Journal of privacy and confidentiality*. 4, 2 (Mar. 2013), 7–41.
- [114] Sundar, S.S., Kang, H., Wu, M., Go, E. and Zhang, B. 2013. Unlocking the privacy paradox: do cognitive heuristics hold the key? *CHI '13 Extended Abstracts on Human Factors in Computing Systems* (New York, NY, USA, 2013), 811–816.
- [115] The Hippocratic Oath Today — NOVA | PBS: 2001. <http://www.pbs.org/wgbh/nova/body/hippocratic-oath-today.html>. Accessed: 2018-04-04.
- [116] Toch, E., Cranshaw, J., Drielsma, P.H., Tsai, J.Y., Kelley, P.G., Springfield, J., Cranor, L., Hong, J. and Sadeh, N. 2010. Empirical models of privacy in location sharing. *Proceedings of the 12th ACM international conference on Ubiquitous computing* (Copenhagen, Denmark, 2010), 129–138.
- [117] Treem, J.W. and Leonardi, P.M. 2013. Social Media Use in Organizations: Exploring the Affordances of Visibility, Editability, Persistence, and Association. *Annals of the International Communication Association*. 36, 1 (Jan. 2013), 143–189. DOI:<https://doi.org/10.1080/23808985.2013.11679130>.
- [118] Trope, Y. and Liberman, N. 2010. Construal-Level Theory of Psychological Distance. *Psychological Review*. 117, 2 (Apr. 2010), 440–463. DOI:<https://doi.org/10.1037/a0018963>.
- [119] Trope, Y. and Liberman, N. 2003. Temporal construal. *Psychological review*. 110, 3 (Jul. 2003), 403–421. DOI:<https://doi.org/10.1037/0033-295X.110.3.403>.
- [120] Tufekci, Z. 2008. Can You See Me Now? Audience and Disclosure Regulation in Online Social Network Sites. *Bulletin of Science, Technology & Society*. 28, 1 (Feb. 2008), 20–36.
- [121] Tversky, A. and Kahneman, D. 1974. Judgment under Uncertainty: Heuristics and Biases. *Science*. 185, 4157 (Sep. 1974), 1124–1131. DOI:<https://doi.org/10.1126/science.185.4157.1124>.
- [122] Urban, J. and Hoofnagle, C.J. 2014. The Privacy Pragmatic as Privacy Vulnerable.
- [123] Utz, S. and Kramer, N. 2009. The privacy paradox on social network sites revisited: The role of individual characteristics and group norms. *Cyberpsychology: Journal of Psychosocial Research on Cyberspace*. 3, 2 (2009).
- [124] Vitak, J. and Kim, J. 2014. “You Can’T Block People Offline”: Examining How Facebook’s Affordances Shape the Disclosure Process. *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing* (New York, NY, USA, 2014), 461–474.
- [125] Wash, R. 2010. Folk Models of Home Computer Security. *Proceedings of the Sixth Symposium on Usable Privacy and Security* (New York, NY, USA, 2010), 11:1–11:16.
- [126] Watson, J., Lipford, H.R. and Besmer, A. 2015. Mapping User Preference to Privacy Default Settings. *ACM Transactions on Computer-Human Interaction*. 22, 6 (Nov. 2015), 32:1–32:20. DOI:<https://doi.org/10.1145/2811257>.
- [127] westin, alan 2001. *Privacy on & off the Internet: What consumers want*. Privacy & American Business.
- [128] Wilkinson, D., Sivakumar, S., Cherry, D., Knijnenburg, B.P., Raybourn, E.M., Wisniewski, P. and Sloan, H. 2017. User-Tailored Privacy by Design. *Proceedings of the Usable Security Mini Conference* (San Diego, CA, 2017).
- [129] Wilson, D. and Valacich, J. 2012. Unpacking the Privacy Paradox: Irrational Decision-Making within the Privacy Calculus. *ICIS 2012 Proceedings* (Orlando, FL, Dec. 2012).
- [130] Wisniewski, P.J., Knijnenburg, B.P. and Lipford, H.R. 2017. Making Privacy Personal. *Int. J. Hum.-Comput. Stud.* 98, C (Feb. 2017), 95–108. DOI:<https://doi.org/10.1016/j.ijhcs.2016.09.006>.
- [131] Wisniewski, P.J., Knijnenburg, B.P. and Lipford, H.R. 2017. Making privacy personal: Profiling social network users to inform privacy education and nudging. *International Journal of Human-Computer Studies*. 98, (Feb. 2017), 95–108. DOI:<https://doi.org/10.1016/j.ijhcs.2016.09.006>.
- [132] Wu, K.-W., Huang, S.Y., Yen, D.C. and Popova, I. 2012. The effect of online privacy policy on consumer privacy concern and trust. *Computers in Human Behavior*. 28, 3 (May 2012), 889–897. DOI:<https://doi.org/10.1016/j.chb.2011.12.008>.
- [133] Xie, J., Knijnenburg, B.P. and Jin, H. 2014. Location Sharing Privacy Preference: Analysis and Personalized Recommendation. *Proceedings of the 19th International Conference on Intelligent User Interfaces* (New York, NY, USA, 2014), 189–198.
- [134] Xu, H., Teo, H.-H. and Tan, B.C.Y. 2005. Predicting the Adoption of Location-Based Services: The Role of Trust and Perceived Privacy Risk. *Proceedings of the International Conference on Information Systems* (Las Vegas, NV, Dec. 2005), 861–874.

- [135] Xu, H., Teo, H.-H., Tan, B.C.Y. and Agarwal, R. 2009. The Role of Push-Pull Technology in Privacy Calculus: The Case of Location-Based Services. *Journal of management information systems*. 26, 3 (Dec. 2009), 135–174. DOI:<https://doi.org/10.2753/MIS0742-1222260305>.
- [136] Zafeiropoulou, A.M., Millard, D.E., Webber, C. and O’Hara, K. 2013. Unpicking the privacy paradox: can structuration theory help to explain location-based privacy decisions? *ACM Web Science 2013 (WebSci ’13)* (Paris, France, 2013).
- [137] Zhang, B., Wu, M., Kang, H., Go, E. and Sundar, S.S. 2014. Effects of Security Warnings and Instant Gratification Cues on Attitudes Toward Mobile Websites. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2014), 111–114.
- [138] Zhang-Kennedy, L. and Chiasson, S. 2014. *Using Comics to Teach Users About Mobile Online Privacy*. Technical Report #Technical report TR-14-02. Carleton University.
- [139] Zhang-Kennedy, L., Chiasson, S. and Biddle, R. 2014. Stop Clicking on “Update Later”: Persuading Users They Need Up-to-Date Antivirus Protection. *Persuasive Technology* (2014), 302–322.
- [140] Zheleva, E. and Getoor, L. 2009. To join or not to join: the illusion of privacy in social networks with mixed public and private user profiles. *Proceedings of the 18th international conference on World wide web* (New York, NY, USA, 2009), 531–540.