Privacy support for personalized learning/training systems

Bart P. Knijnenburg
Clemson University School of Computing
Clemson, SC
bartk@clemson.edu

ABSTRACT

This paper discusses the effective management of privacy in the development of personalized learning/training systems. While such systems employ ubiquitous data collection and user to enable highly personalized and pervasive training recommendations, the data collection and user modeling may cause privacy threats that act as a barrier to their adoption. Two effective privacy management strategies are discussed. First of all, in line with the *privacy by design* philosophy, the impact of several operational characteristics of personalized learning/training systems is analyzed. Beyond this, the paper discusses the idea of *user-tailored privacy* (UTP) to support learners' privacy decisions, as well as the factors that can be used as input for UTP and the adaptation mechanisms that UTP can employ.

ABOUT THE AUTHOR

Bart P. Knijnenburg is an Assistant Professor in Human-Centered Computing at the Clemson University School of Computing. He holds a B.S. in Innovation Sciences and an M.S. in Human-Technology Interaction from the Eindhoven University of Technology, The Netherlands, an M.A. in Human-Computer Interaction from Carnegie Mellon University, and a PhD in Information and Computer Sciences from UC Irvine. Dr. Knijnenburg works on privacy decision-making and user-centric evaluation of adaptive systems, and has led trailblazing efforts in the development of user-tailored privacy. In 2013, he was the recipient of the first Google PhD Fellowship in Privacy.

Privacy support for personalized learning/training systems

Bart P. Knijnenburg
Clemson University School of Computing
Clemson, SC
bartk@clemson.edu

INTRODUCTION

Personalized learning/training systems employ ubiquitous data collection (e.g. by interfacing with social media activity and smartphone sensors) and user modeling (e.g. by collecting highly detailed user-training interaction behavior) to enable highly personalized and pervasive ("anytime, anywhere") training recommendations (Raybourn and Regan 2011). Moreover, these systems often contain a "social network" component that allows learning materials, activities, and outcomes to be shared across learners (enabling peer interactions) and other learning systems (allowing for an extensible learning environment).

Privacy threats have shown to be an important barrier to the adoption of personalized systems (Chellappa and Sin 2005; van de Garde-Perik et al. 2008; Kobsa 2007; Phelps et al. 2000; Sutanto et al. 2013; Treiblmaier and Pollach 2007), and it is therefore of utmost importance that such threats are minimized in such systems. From a security perspective, strict security requirements regarding personal and sometimes mission-critical information are at odds with the philosophy that learners should have ownership over their own learner data. From a privacy perspective, the social capital-based advantages of sharing learner profiles are at odds with the fact that these learner profiles are also used for sensitive employment decisions regarding placement, selection and promotion. On top of this, the international deployment of a personalized learning/training system introduces prominent cultural variation in privacy concerns and social etiquette (Cockcroft and Rekker 2015; Cremonini and Valeri 2003).

As a result of all this, users of personalized learning/training systems have to carefully navigate a multi-dimensional array of privacy concerns, carefully balancing the benefits and risks of disclosing or allowing access to their personal information. However, users of complex information systems have been consistently incapable of effectively managing their own privacy (Acquisti et al. 2012; John et al. 2011; Johnson et al. 2002; Knijnenburg et al. 2013a; Lai and Hui 2004; Liu et al. 2011; Madejski et al. 2012), leaving them vulnerable to perceived and real privacy threats.

This paper employs the philosophy of *privacy by design* (Cavoukian 2010; Langheinrich 2001; Shapiro 2009; Spiekermann 2012) to investigate the impact of various operational characteristics of personalized learning/training systems on users' privacy, which allows developers and researchers of such systems to select the characteristics that best alleviate users' concerns. Moreover, given the inherent focus of personalized learning/training systems on user modeling, there exists another opportunity to implement *user-tailored privacy* (UTP) (Knijnenburg 2015; Knijnenburg and Jin 2013; Knijnenburg and Kobsa 2013a, 2014; Watson 2015), i.e., to model users' privacy concerns and provide them with adaptive privacy decision support.

PRIVACY BY DESIGN

Privacy by design is a design philosophy in which privacy aspects are addressed early in the system design and development process, rather than after the system has been developed ("post hoc privacy") (Cavoukian 2010; Langheinrich 2001; Shapiro 2009; Spiekermann 2012). While post hoc privacy solutions typically try to mitigate privacy problems that exist within a system, privacy by design tries to avoid problems from occurring at all.

To explore the opportunity of privacy by design, we investigate the impact of the operational characteristics of the personalized learning/training systems on users' privacy. Developers and researchers can use this analysis to make informed decisions about the operational characteristics of their personalized learning/training systems.

Ownership Model

To provide an extensible learning environment, a personalized learning/training system may provide the ability to share learning outcomes across multiple independent learning applications. Learning outcomes can be considered to be personal information, and the privacy dynamics of the system will be different depending on the ownership model for this information: beyond a traditional centralized approach, this could range from a partially decentralized architecture in which a central control application operates as a portal to the learning applications (in this case, the privacy dynamic resembles those of *app stores* (Jeon et al. 2012; Wang et al. 2011a)), to a a completely distributed architecture with portable learner models (in this case, the privacy consequences of *client-side personalization* are relevant (Kobsa et al. 2016)).

In the fully centralized approach, users would make a single set of privacy decisions for the entire system, and not be able to control how the system further distributes the personal information among external learning applications. The portal approach may give users considerably more control over their privacy settings, but research has shown that users are notoriously bad at deciding what information to disclose to which application (Jeon et al. 2012; Kelley et al. 2012; Wang et al. 2011a). A good privacy design would include easily understandable privacy notices describing the implications of disclosure to each individual application (Kelley et al. 2010).

Finally, the portable learner model approach offloads the privacy management to the user themselves. This precludes the need for a central control application, but is not without further problems: users of portable user models worry about what happens when the device that contains their portable user model gets lost or stolen; not only would this potentially cause a data breach, but it would also mean a loss of their user model (Kobsa et al. 2016). This problem is of course exacerbated if learner models contain mission-critical information. A good privacy design would allow users to remotely wipe their user model, and to restore old models from encrypted backups (this is similar to how Apple users can mitigate the problems of a lost or stolen iPhone) (Kobsa et al. 2016).

Extent of Mining and Tracking

Most personalized learning/training systems employ a highly adaptive learner model that proactively mines and tracks a variety of information sources to provide personalized learning experiences. In the military and other government organizations, the goal of these systems is to train employees on the job, adapting presented training modules to personal capabilities, mission requirements, and available time and other resources.

The extent of mining and tracking involved in this highly adaptive behavior will have profound effects on users' privacy perceptions. Particularly, it has been shown that users are more concerned about personal information that is collected automatically compared to manually provided information (Knijnenburg et al. 2013b; Knijnenburg and Kobsa 2013b; Kobsa et al. 2016). Similarly, the tracking of real-world activities is arguably more privacy-invasive (and difficult to control) than the tracking of in-system behaviors (Niu et al. 2010). The mining and tracking activities may also be regulated by government privacy regulations (Donley 2007).

The effect of the extent of mining and tracking on users' privacy perceptions is a classical example of the *personalization-privacy paradox* (Awad and Krishnan 2006; Chellappa and Sin 2005): it is impossible to provide high-quality personalized training recommendations without extensive tracking by some sort of personalization server, especially when centralized goals are expected to be taken into account in the personalization process. A good design compromise would be a two-tier personalization approach: On the first tier, resources, mission goals, and users' previous learning outcomes are used to decide *what* training modules to recommend to the user. At the second tier, behavioral tracking can be used to decide *when* is the optimal time for the user to engage in each of the recommended modules. The personalization at this second tier can occur client-side, preventing the user's behavioral patterns from being processed and stored by a centralized server.

Learning Recommendation Mechanism

Part of the highly adaptive learner model vision is a strong focus on *proactive* recommendation. Note, though, that privacy also means "the right to be left alone". This means that the learning recommendation mechanism can have consequences for user privacy perceptions. Particularly, the stronger a system's suggestion, the higher the chance that users may feel that their personal space is invaded by the system (Cosley et al. 2003; McKenzie et al. 2006). In this sense, a system that *highlights* recommended learning options is less obtrusive than a system that uses *push notifications* to suggest certain training exercises. A system that *automatically* takes action (or that *sorts* or *filters* available options) without notifying the user is arguably even less obtrusive, but research has shown that users rarely change the default setting (Johnson and Goldstein 2003). Such a system would thus have to deal with the ethical implications of making decisions for their users without explicit consent (Smith et al. 2013), which is a privacy concern in the sense that users lose control over their actions. These different mechanisms also differ in the extent with which they can persuade the user (McKenzie et al. 2006), and it is important to find the right balance between persuasion and different types of privacy.

A good privacy design takes into account the *confidence* the system has in its recommendations. Recommendations that have a very high likelihood of getting accepted (i.e. they adhere to strongly engrained user behavior patterns) can be effected automatically, while more controversial recommendations (e.g. recommendations based on mission requirements rather than user preference) should be presented via notifications.

Sharing Capabilities

Many personalized learning/training systems emphasize the importance of social connections in learning. Particularly, with Open Social Learner Models, users are able to compare their learning progress (performance) with other users. Beyond that, users are of course expected to share their learning portfolios with their superiors to support placement and selection (Raybourn and Regan 2011). These sharing capabilities have important consequences for users' privacy.

To the extent that users can share content and performance data with their peers, the privacy perceptions of sharing on social networks apply (Wang et al. 2011b; Wisniewski et al. 2014). Recent research has shown that users employ a multitude of behaviors to protect their privacy, withholding information being just one of them. As different users employ different privacy management strategies (Wisniewski et al. 2014), it serves to implement many of these mechanisms. Note that more sharing is not always better: while extensive peer sharing supports users in building social capital (Ellison et al. 2007), giving users the amount of privacy they want (rather than pushing them to be more open) actually increases their overall social connectedness (Wisniewski et al. 2015). As for sharing within an organization, a recent overview of the privacy field has acknowledged that there is a severe lack of scientific knowledge regarding this activity (Smith et al. 2011).

Social Dynamics

Finally, the social dynamics of a personalized learning/training system may have an impact on the privacy perceptions and behavior its users. Most notably, organizations may want to promote *collaboration* and support in learning and training, but they can also introduce game-based aspects, which may instill a social dynamic of *competition*. Users' sharing behaviors within these different dynamics have to date not been researched extensively.

BEYOND PRIVACY BY DESIGN

The privacy by design methodology is not without criticism. One critique is that the methodology is lacking broader integration with other considerations that need to be addressed in the software development cycle (Spiekermann 2012). Another critique is that the current main approaches to privacy by design—"transparency and control" and "privacy nudging"—are lacking.

The Shortcomings of Transparency and Control

Proponents of transparency and control argue that users should be given comprehensive control over what data they wish to share, while at the same time providing them with more information about the implications of their decisions (Hui et al. 2007; Kolter and Pernul 2009; Rifon et al. 2005; Xu 2007). This makes intuitive sense. At least some minimum level of control over one's disclosure is necessary to engage in a privacy-related decision-making: without control, the user does not have any influence on the risk/benefit tradeoff. Moreover, people can only make an informed tradeoff between benefits and risks if they are given adequate information. Information enables them to make an accurate assessment of the possible risks and benefits of disclosure. Based on this reasoning, advocates of transparency and control argue that it *empowers* users to regulate their privacy at the desired level (Bulgurcu 2012; Cavusoglu et al. 2013; Sadeh et al. 2009).

Unfortunately, though, there is ample evidence that users often have difficulties navigating privacy settings. Specifically, users' privacy decisions turn out to be more heuristic than rational (Acquisti and Grossklags 2005, 2008), and they fall prey to a wide array of decision fallacies (Acquisti et al. 2012; John et al. 2011; Johnson et al. 2002; Knijnenburg et al. 2013a; Lai and Hui 2004). Moreover, as managing privacy is not the core task of a learner, there is likely to be a lack of motivation. Indeed, recent work shows that while users claim to want control over their privacy, they often do not devote their full attention to the provided mechanisms of transparency and control (Compañó and Lusoli 2010; Knijnenburg et al. 2013c). Consequently, it no surprise that most users of Facebook—which has one of the most extensive privacy control mechanisms in the industry—are incapable of correctly identifying their own privacy settings (Liu et al. 2011; Madejski et al. 2012). If personalized learning/training systems were to employ similar "transparency and control" mechanisms, they would likely undergo a similar faith.

The Shortcomings of Privacy Nudging

Proponents of privacy nudging, on the other hand, argue that privacy by design solutions should *relieve* some of the decision-making burden by making it easier for users to process and execute the information disclosure decisions, without taking away users' control altogether. Nudges are subtle yet persuasive cues that make people more likely to decide in one direction or the other (Thaler and Sunstein 2008). Carefully designed nudges make it easier for people to make the right choice, without limiting their ability to choose freely. Nudges ostensibly turn people's decision fallacies into mechanisms that help them (Acquisti 2009): they exploit these fallacies to create a *choice architecture* that encourages wanted behavior and inhibits unwanted behavior (Thaler and Sunstein 2008).

Effectively nudging users' privacy design is difficult, though, especially in situations or systems where there are both strong benefits and important risks to information disclosure. Personalized learning/training systems are a good example of such systems. First of all, there are strong benefits in tracking users' behaviors (i.e., learning personalization; anytime, anywhere training) and having them share resources and accomplishments (i.e., social learning; creating organizational knowledge). Uniformly nudging users towards more privacy may make it more difficult for less privacy-minded users to enjoy these benefits. At the same time, the potential risks are also larger. Personalized learning/training systems may deal with classified military information and data that can affect employment and promotion decisions. Uniformly nudging people towards more disclosure may increase the risk that these data fall into the wrong hands. Consequently, a simple, one-sided nudge is not a feasible solution.

Towards a solution

There are two solutions to the problems with transparency, control, and nudges. First of all, their shortcomings can be mitigated by *integrating* the two approaches, applying privacy nudging wherever users' preferences are clear and uniform (or where there exist organizational constraints), and providing transparency and control wherever a plurality of user preferences exist. Integrating transparency, control, and nudges in this way requires a careful analysis of the user base of the learning/training system. Specifically, surveys, interviews, or user tests should be performed to distinguish between the privacy aspects for which users have low motivation and/or a uniform opinion, and those that are considered most critical and/or highly contested. At ha higher level, controlled experiments can be employed to study this attention differential in across various systems. As an example, previous work has used the *elaboration likelihood model* as a means to explain why users are not always highly motivated to make privacy-decisions (Angst and Agarwal 2009; Kobsa et al. 2016; Lowry et al. 2012).

The other solution is to go beyond privacy by design; i.e., to employ *user-tailored privacy* as a privacy solution. User-tailored privacy combines the convenience of privacy nudges with the respect for users' privacy preferences (which are likely to range on a wide spectrum (Harris et al. 2003; Knijnenburg et al. 2013b)) that is inherent in transparency and control. Unlike transparency and control, user-tailored privacy does not assume that users are rational and highly motivated decision-makers when it comes to privacy. Unlike nudges, user-tailored privacy avoids a one-sided paternalistic approach in favor of learning from the user's past decisions to best respect their inherent preferences. Combined, transparency and control, privacy nudging, and user-tailored privacy provide a comprehensive toolbox for designing privacy support for personalized learning/training systems.

USER-TAILORED PRIVACY

While privacy by design can prevent many privacy problems, some design questions regarding personalized learning/training systems will not have a universal solution. This means that most personalized learning/training systems will need a plethora of settings that allow users to *customize* their desired level of privacy. Unfortunately, though, there is ample evidence that users often have difficulties navigating privacy settings. User-tailored privacy (UTP) is an approach to privacy that models users' privacy concerns and provides them with adaptive privacy decision support. By providing user-tailored support, it reconciles the need for extensive customizability with users' lack of skills and motivation to manage their own privacy settings (Knijnenburg 2015; Knijnenburg and Jin 2013; Knijnenburg and Kobsa 2013a, 2014; Watson 2015). With user-tailored privacy, 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.

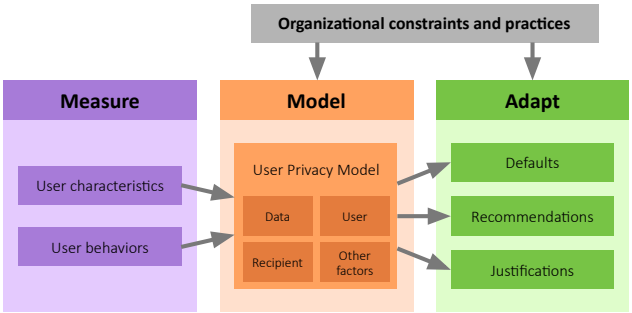

Figure 1: A schematic overview of User-Tailored Privacy (UTP)

The bulk of the existing work on UTP is algorithmic, and focuses on how to predict users' privacy preferences and behaviors from user characteristics and behaviors (Fang and LeFevre 2010; Pallapa et al. 2014; Ravichandran et al. 2009). This work shows that privacy preferences can indeed be modeled with off-the-shelf machine learning components, but is otherwise purely theoretical in nature; research on the benefits and drawbacks of actual adaptive decision-support strategies is less common (Knijnenburg 2015). The remainder of this section discusses how user characteristics, behaviors, and contextual factors influence users' privacy decisions in personalized

learning/training systems. It will also analyze potential adaptation strategies that can be used to support users' privacy decisions, and address potential organizational constraints and practices that can be taken into account in the modeling process.

User Characteristics, Behaviors, and Contextual Factors

Existing work has shown that user-tailored privacy critically depends on the evaluation of user- and context-related factors that influence users' privacy concerns and behaviors. This work has shown that data, user, and recipient are important, but also that for many applications there are additional system-specific factors (Dong et al. 2015; Lusoli et al. 2012; Olson et al. 2005; Patil and Lai 2005; Xie et al. 2014) that influence users' decisions. Due to this context-specific nature of users' privacy decisions, it stands to reason to also make the User Privacy Model underlying user-tailored privacy context-specific. This means taking into account contextual variables that have been shown to influence users' privacy concerns and behavior.

One such variables is the data itself. Several researchers have found that people's privacy concerns are multi-dimensional, meaning that they have different preferences for different types of information (Lusoli et al. 2012; Olson et al. 2005; Spiekermann et al. 2001). Furthermore, research shows that these preferences can be summarized into distinct *profiles* (Knijnenburg et al. 2013b; Olson et al. 2005; Wisniewski et al. 2014). The recipient of the information seems to play an important role as well, both in "commercial" and "social" privacy settings (Knijnenburg et al. 2013c; Knijnenburg and Kobsa 2014; Patil and Lai 2005), and recipients can also be grouped to simplify the privacy decision problem. For example, on social networks the optimal grouping seems to be Friends, Family, Classmates, Colleagues, and Acquaintances (Knijnenburg et al. 2014), but this clustering might be different for recipients in personalized learning/training systems. Furthermore, in certain types of systems, privacy preferences may depend on other contextual factors. For example, researchers have found that time (weekday or weekend, daytime or evening) is an important determinant of users' willingness to disclose their location (Dong et al. 2015; Xie et al. 2014). Finally, note that user-tailored privacy can operate within organizational constraints, taking into account existing rules, as well as common practices. This way, user-tailored privacy helps users to select settings that are not only in line with their own preferences, but that also take into account the value for the organization and existing rules.

Adaptation strategies

A previous section noted that there are different ways to present learning recommendations, and the same is true for privacy recommendations. Specifically, a user-tailored privacy module may *highlight* recommended privacy decisions (or hide the ones that are less likely to be chosen) (Knijnenburg and Jin 2013), make *justifications* for certain privacy-related actions (Knijnenburg and Kobsa 2013b), *automatically* take action (making use of the fact that users rarely change the default setting (Johnson et al. 2002; Lai and Hui 2004)), or *sort* information requests in an order that balances their sensitivity with their usefulness for the system (Knijnenburg 2015). As mentioned previously, these different mechanisms differ in the extent with which they interrupt and persuade the user (Knijnenburg and Jin 2013), so decisions regarding the adaptation strategy should be considered carefully.

Examples and Discussion

A number of military/government-related examples may help illustrate the concept of user-tailored privacy (UTP):

Example 1—A certain personalized learning/training system normally tracks users' location (Data) in order to give context-relevant vigilance training exercises (Organizational practice). However, the user-tailored privacy procedure of the system has learned that like many young mothers (User characteristic), Mary (User) does not want her location (Data) tracked outside work hours (Other factor). It therefore turns the location tracker off by default when Mary is not on the clock (Default).

Example 2—David is a professional translator who needs to decide how to share his recent milestones—two certificates in the Arabic and Farsi language he has recently earned (Data)—within his organization (Recipient). Due to the rules of his employer (Organizational constraint), UTP requires him to share these milestones with his direct supervisor (Recipient). Moreover, from his previous interactions (User behaviors), UTP knows that David keeps close ties to several other military divisions. UTP therefore suggests (Recommendation) that he should share his new certifications with the heads of these divisions (Recipient) as well, arguing they are likely to be interested in exploiting his newly gained skills in an upcoming mission (Justification).

User-tailored privacy aims to strike this balance between giving users no control over, or information about, their privacy at all, and giving them full control and information. Arguably, user-tailored privacy 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 (Knijnenburg 2015).

CONCLUSION AND FUTURE WORK

This paper discussed the effective management of privacy in the development of personalized learning/training systems. In line with the *privacy by design* philosophy, it analyzed the impact of the ownership model, extent of mining/tracking, recommendation mechanism, and sharing capabilities on users' privacy concerns, and suggested design solutions for potential problems. Beyond this, the paper discussed the idea of *user-tailored privacy* (UTP), arguing that adapting highlighted features, default settings, justifications, or request orders to the users' characteristics, behaviors, and contextual factors can significantly decrease the burden on users' decision strategy without doing away with the notion that privacy preferences vary across the user base.

The suggestions provided in this paper are taken from general privacy literature, and assumed to generalize to personalized learning/training systems. Actual privacy studies of such systems would reveal the extent to which such generalization do and do not hold. Such work is scarce, though, so we recommend that the developers of personalized learning/training systems conduct their own evaluations to supplement the knowledge provided in this paper. Even if these systems particularly focus on government/military systems, the work may still be very useful for the privacy community, especially in light of Smith et al.'s (2011) remark very little work has focused on that organization-level privacy dynamics. The military and other governmental organizations are a useful test bed for these approaches.

REFERENCES

Acquisti, A. 2009. "Nudging Privacy: The Behavioral Economics of Personal Information," *IEEE Security and Privacy* (7), pp. 82–85.
Acquisti, A., and Grossklags, J. 2005. "Privacy and Rationality in Individual Decision Making," *IEEE Security & Privacy* (3:1), pp. 26–33.
Acquisti, A., and Grossklags, J. 2008. "What Can Behavioral Economics Teach Us About Privacy?," in *Digital Privacy: Theory, Technologies, and Practices*, A. Acquisti, S. De Capitani di Vimercati, S. Gritzalis, and C. Lambrinoudakis (eds.), New York/London: Auerbach Publications, pp. 363–377.
Acquisti, A., John, L. K., and Loewenstein, G. 2012. "The Impact of Relative Standards on the Propensity to Disclose," *Journal of Marketing Research* (49:2), pp. 160–174.
Angst, C. M., and Agarwal, R. 2009. "Adoption of electronic health records in the presence of privacy concerns: the elaboration likelihood model and individual persuasion," *MIS Quarterly* (33:2), pp. 339–370.
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), pp. 13–28.
Bulgurcu, B. 2012. "Understanding the information privacy-related perceptions and behaviors of an online social network user," Ph.D. Thesis, Vancouver, BC: University of British Columbia.
Cavoukian, A. 2010. *Privacy by Design*, Information and Privacy Commissioner of Ontario, Canada.
Cavusoglu, H., Phan, T., and Cavusoglu, H. 2013. "Privacy Controls and Content Sharing Patterns of Online Social Network Users: A Natural Experiment," in *ICIS 2013 Proceedings* Milan, Italy.

Chellappa, R. K., and Sin, R. G. 2005. "Personalization versus privacy: An empirical examination of the online consumer's dilemma," *Information Technology and Management* (6:2), pp. 181–202.
Cockcroft, S., and Rekker, S. 2015. "The relationship between culture and information privacy policy," *Electronic Markets*, pp. 1–18.
Compañó, R., and Lusoli, W. 2010. "The Policy Maker's Anguish: Regulating Personal Data Behavior Between Paradoxes and Dilemmas," in *Economics of Information Security and Privacy*, T. Moore, D. Pym, and C. Ioannidis (eds.), New York, NY: Springer US, pp. 169–185.
Cosley, D., Lam, S. K., Albert, I., Konstan, J. A., and Riedl, J. 2003. "Is Seeing Believing?: How Recommender System Interfaces Affect Users' Opinions," in *SIGCHI Conference on Human Factors in Computing Systems*, Ft. Lauderdale, FL: ACM, pp. 585–592.
Cremonini, L., and Valeri, L. 2003. "Benchmarking Security and Trust in Europe and the US," No. IST 2000-26746.
Dong, C., Jin, H., and Knijnenburg, B. P. 2015. "Predicting Privacy Behavior on Online Social Networks," in *Ninth International AAAI Conference on Web and Social Media* AAAI Publications, April 21, pp. 91–100.
Donley, M. B. 2007. "Department of Defense Privacy Program," No. DoD 5400.11-R, Washington, D.C.: Department of Defense.
Ellison, N. B., Steinfield, C., and Lampe, C. 2007. "The Benefits of Facebook 'Friends:' Social Capital and College Students' Use of Online Social Network Sites," *J. of Computer-Mediated Communication* (12:4), pp. 1143–1168.
Fang, L., and LeFevre, K. 2010. "Privacy Wizards for Social Networking Sites," in *19th International Conference on World Wide Web*, Raleigh, NC: ACM, pp. 351–360.
van de Garde-Perik, E., Markopoulos, P., de Ruyter, B., Eggen, B., and Ijsselsteijn, W. 2008. "Investigating Privacy Attitudes and Behavior in Relation to Personalization," *Social Science Computer Review* (26:1), pp. 20–43.
Harris, M. M., Hoye, G. V., and Lievens, F. 2003. "Privacy and Attitudes Towards Internet-Based Selection Systems: A Cross-Cultural Comparison," *International Journal of Selection and Assessment* (11:2-3), pp. 230–236.
Hui, K.-L., Teo, H. H., and Lee, S.-Y. T. 2007. "The Value of Privacy Assurance: An Exploratory Field Experiment," *MIS Quarterly* (31:1), pp. 19–33.
Jeon, J., Micinski, K. K., Vaughan, J. A., Fogel, A., Reddy, N., Foster, J. S., and Millstein, T. 2012. "Dr. android and mr. hide: fine-grained permissions in android applications," in *second ACM workshop on Security and privacy in smartphones and mobile devices* ACM, pp. 3–14.
John, L. K., Acquisti, A., and Loewenstein, G. 2011. "Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information," *Journal of consumer research* (37:5), pp. 858–873.
Johnson, E. J., Bellman, S., and Lohse, G. L. 2002. "Defaults, Framing and Privacy: Why Opting In ≠ Opting Out," *Marketing Letters* (13:1), pp. 5–15.
Johnson, E. J., and Goldstein, D. 2003. "Do Defaults Save Lives?," *Science* (302:5649), pp. 1338–1339.
Kelley, P. G., Cesca, L., Bresee, J., and Cranor, L. F. 2010. "Standardizing privacy notices: an online study of the nutrition label approach," in *28th Conf. on Human Factors in Computing Systems,* Atlanta, GA: pp. 1573–1582.
Kelley, P. G., Consolvo, S., Cranor, L. F., Jung, J., Sadeh, N., and Wetherall, D. 2012. "A Conundrum of Permissions: Installing Applications on an Android Smartphone," in *Financial Cryptography and Data Security*, Lecture Notes in Computer Science, J. Blyth, S. Dietrich, and L. J. Camp (eds.), Springer Berlin, pp. 68–79.
Knijnenburg, B. P. 2015. "A user-tailored approach to privacy decision support," Ph.D. Thesis, Irvine, CA: University of California, Irvine.
Knijnenburg, B. P., and Jin, H. 2013. "The Persuasive Effect of Privacy Recommendations," in *Twelfth Annual Workshop on HCI Research in MIS* Milan, Italy, pp. 16:1–16:5.
Knijnenburg, B. P., and Kobsa, A. 2013a. "Helping users with information disclosure decisions: potential for adaptation," in *ACM international conference on Intelligent User Interfaces* Santa Monica, CA: pp. 407–416.
Knijnenburg, B. P., and Kobsa, A. 2013b. "Making Decisions about Privacy: Information Disclosure in Context-Aware Recommender Systems," *ACM Transactions on Interactive Intelligent Systems* (3:3), pp. 20:1–20:23.
Knijnenburg, B. P., and Kobsa, A. 2014. "Increasing Sharing Tendency Without Reducing Satisfaction: Finding the Best Privacy-Settings User Interface for Social Networks," in *ICIS 2014 Proceedings* Auckland, New Zealand.

Knijnenburg, B. P., Kobsa, A., and Jin, H. 2013a. "Preference-based location sharing: are more privacy options really better?," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* Paris, France: ACM, pp. 2667–2676.
Knijnenburg, B. P., Kobsa, A., and Jin, H. 2013b. "Dimensionality of information disclosure behavior," *International Journal of Human-Computer Studies* (71:12), pp. 1144–1162.
Knijnenburg, B. P., Kobsa, A., and Jin, H. 2013c. "Counteracting the Negative Effect of Form Auto-completion on the Privacy Calculus," in *ICIS 2013 Proceedings* Milan, Italy.
Knijnenburg, B. P., Kobsa, A., and Jin, H. 2014. *Segmenting the Recipients of Personal Information*, Submitted to the SOUPS2014 workshop on Privacy Personas and Segmentation.
Kobsa, A. 2007. "Privacy-Enhanced Personalization," *Communications of the ACM* (50:8), pp. 24–33.
Kobsa, A., Cho, H., and Knijnenburg, B. P. 2016. "The Effect of Personalization Provider Characteristics on Privacy Attitudes and Behaviors: An Elaboration Likelihood Model Approach," *Journal of the Association for Information Science and Technology*.
Kolter, J., and Pernul, G. 2009. "Generating User-Understandable Privacy Preferences," in *Conf. on Availability, Reliability and Security*Fukuoka, Japan: IEEE Computer Society, pp. 299–306.
Lai, Y.-L., and Hui, K.-L. 2004. "Opting-in or opting-out on the Internet: Does it Really Matter?," in *ICIS 2004: Twenty-Fifth International Conference on Information Systems* Washington, D.C., pp. 781–792.
Langheinrich, M. 2001. "Privacy by Design: Principles of Privacy-Aware Ubiquitous Systems," in *Ubicomp 2001*, G. D. Abowd, B. Brumitt, and S. A. N. Shafer (eds.), (Vol. LNCS 2201) Springer-Verlag, pp. 273–291.
Liu, Y., Gummadi, K. P., Krishnamurthy, B., and Mislove, A. 2011. "Analyzing facebook privacy settings: user expectations vs. reality," in *ACM conference on Internet measurement conference*, pp. 61–70.
Lowry, P. B., Moody, G., Vance, A., Jensen, M., Jenkins, J., and Wells, T. 2012. "Using an elaboration likelihood approach to better understand the persuasiveness of website privacy assurance cues for online consumers," *Journal of the American Society for Information Science and Technology* (63:4), pp. 755–776.
Lusoli, W., Bacigalupo, M., Lupiáñez-Villanueva, F., Andrade, N., Monteleone, S., and Maghiros, I. 2012. "Pan-European Survey of Practices, Attitudes and Policy Preferences as Regards Personal Identity Data Management," SSRN Scholarly Paper No. ID 2086579, Rochester, NY: Social Science Research Network.
Madejski, M., Johnson, M., and Bellovin, S. M. 2012. "A study of privacy settings errors in an online social net-work," in *Fourth International Workshop on SECurity and SOCial Networking* Lugano, Switzerland, pp. 340–345.
McKenzie, C. R. M., Liersch, M. J., and Finkelstein, S. R. 2006. "Recommendations Implicit in Policy Defaults," *Psychological Science* (17:5), pp. 414–420.
Niu, Y., Shi, E., Chow, R., Golle, P., and Jakobsson, M. 2010. "One Experience Collecting Sensitive Mobile Data," in *SOUPS 2010 Usable Security Experiment Reports (USER) Workshop*.
Olson, J. S., Grudin, J., and Horvitz, E. 2005. "A study of preferences for sharing and privacy," in *CHI '05 Extended Abstracts* Portland, OR: ACM, pp. 1985–1988.
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), pp. 232–243.
Patil, S., and Lai, J. 2005. "Who Gets to Know What when: Configuring Privacy Permissions in an Awareness Application," in *SIGCHI Conference on Human Factors in Computing Systems*Portland, OR: ACM, pp. 101–110.
Phelps, J., Nowak, G., and Ferrell, E. 2000. "Privacy Concerns and Consumer Willingness to Provide Personal Information," *Journal of Public Policy & Marketing* (19:1), pp. 27–41.
Ravichandran, R., Benisch, M., Kelley, P., and Sadeh, N. 2009. "Capturing Social Networking Privacy Preferences:," in *Privacy Enhancing Technologies*, Lecture Notes in Computer Science, I. Goldberg and M. Atallah (eds.), (Vol. 5672) Springer Berlin / Heidelberg, pp. 1–18.
Raybourn, E., and Regan, D. 2011. "Exploring e-portfolios and Independent Open Learner Models: Toward Army Learning Concept 2015," in *I/ITSEC Proceedings, Florida USA*.
Rifon, N. J., LaRose, R., and Choi, S. M. 2005. "Your Privacy Is Sealed: Effects of Web Privacy Seals on Trust and Personal Disclosures," *Journal of Consumer Affairs* (39:2), pp. 339–360.
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), pp. 401–412.
Shapiro, S. S. 2009. "Privacy by design: moving from art to practice," *Commun. ACM* (53), pp. 27–29.
Smith, H. J., Dinev, T., and Xu, H. 2011. "Information Privacy Research: An Interdisciplinary Review," *MIS Quarterly* (35:4), pp. 989–1016.

Smith, N. C., Goldstein, D. G., and Johnson, E. J. 2013. "Choice Without Awareness: Ethical and Policy Implications of Defaults," *Journal of Public Policy & Marketing* (32:2), pp. 159–172.
Spiekermann, S. 2012. "The Challenges of Privacy by Design," *Commun. ACM* (55:7), pp. 38–40.
Spiekermann, S., Grossklags, J., and Berendt, B. 2001. "E-privacy in 2nd Generation E-Commerce: Privacy Preferences versus Actual Behavior," in *3rd ACM conference on Electronic Commerce*Tampa, FL, pp. 38–47.
Sutanto, J., Palme, E., Tan, C.-H., and Phang, C. W. 2013. "Addressing the Personalization-Privacy Paradox: An Empirical Assessment from a Field Experiment on Smartphone Users," *MIS Quarterly* (37:4), pp. 1141–1164.
Thaler, R. H., and Sunstein, C. 2008. *Nudge : improving decisions about health, wealth, and happiness*, New Haven, NJ & London, U.K.: Yale University Press.
Treiblmaier, H., and Pollach, I. 2007. "Users' Perceptions of Benefits and Costs of Personalization," in *ICIS 2007 Proceedings*.
Wang, N., Xu, H., and Grossklags, J. 2011a. "Third-party apps on Facebook: privacy and the illusion of control," in *5th ACM Symposium on Computer Human Interaction for Management of Information Technology*.
Wang, Y., Norcie, G., Komanduri, S., Acquisti, A., Leon, P. G., and Cranor, L. F. 2011b. "'I regretted the minute I pressed share': a qualitative study of regrets on Facebook," in *Proceedings of the Seventh Symposium on Usable Privacy and Security*Pittsburgh, PA: ACM, pp. 10:1–10:16.
Watson, J. 2015. "Predicting privacy settings with a user-centered approach," in *2015 International Conference on Collaboration Technologies and Systems (CTS)*, June, pp. 499–500.
Wisniewski, P., Islam, A. K. M. N., Knijnenburg, B. P., and Patil, S. 2015. "Give Social Network Users the Privacy They Want," in *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, CSCW '15, Vancouver, Canada: ACM, pp. 1427–1441.
Wisniewski, P., Knijnenburg, B. P., and Richter Lipford, H. 2014. "Profiling Facebook Users' Privacy Behaviors," in *SOUPS2014 Workshop on Privacy Personas and Segmentation* Menlo Park, CA.
Xie, J., Knijnenburg, B. P., and Jin, H. 2014. "Location Sharing Privacy Preference: Analysis and Personalized Recommendation," in *19th International Conference on Intelligent User Interfaces*, pp. 189–198.
Xu, H. 2007. "The effects of self-construal and perceived control on privacy concerns," in *ICIS 2007 Proceedings*.