Simplifying Privacy Decisions: Towards Interactive and Adaptive Solutions

Bart P. Knijnenburg
Donald Bren School of Information and Computer Sciences
University of California, Irvine
bart.k@uci.edu

1. INTRODUCTION

Privacy concerns are an important barrier to the growth of social networks, e-commerce, ubiquitous computing, and location sharing services. The large majority of Internet users takes a pragmatic stance on information disclosure: they trade off the anticipated benefits with the risks of disclosure, a decision process that has been dubbed *privacy calculus* [10,23]. Privacy decisions are inherently difficult though, because they have delayed and uncertain repercussions that are difficult to trade-off with the possible immediate gratification of disclosure [3,5].

How can we help users to balance the benefits and risks of information disclosure in a user-friendly manner, so that they can make good privacy decisions? Existing research has explored two approaches to this problem, but neither provides a satisfying solution. Below I discuss these two approaches, and introduce a new *user-tailored* approach that provides more user-friendly privacy decision support.

2. TRANSPARENCY AND CONTROL

To help users with their privacy calculus, experts recommend giving users comprehensive *control* over what data they wish to share, and more *transparency* about the implications of their decisions [1,22]. However, while users claim to *want* full control over their data, they avoid the hassle of actually *exploiting* this control [8]. Moreover, the privacy controls of systems like Facebook are so complex that users do not even seem to know the implications of their own settings [25]. Similarly, informing users about the rationale behind information requests does not make them more discerning about their privacy decisions, but merely makes them worry about privacy in general. For example, displaying a privacy label on an e-commerce website—a supposed vote of confidence—may *decrease* instead of increase purchases [7].

Evidently, transparency and control do not work well in practice. Due to the complexity of privacy decisions and users' bounded rationality [2,3], an increase in transparency and control often just aggravates the problem by introducing choice overload [12,27] and information overload [11].

3. PRIVACY NUDGES

An alternative approach to support privacy decisions is to introduce subtle yet persuasive *nudges*. Carefully designed nudges make it easier for people to make the right choice, without limiting their ability to choose freely [29]. A *justification*, for example, makes it easier to rationalize decisions and to minimize the regret associated with choosing the wrong option [9]. The effect of justifications in privacy research seems to vary. In my own research I have found that justifications are regarded as helpful, but do not increase users' disclosure or satisfaction but rather *decrease* them [18,19]. Sensible *defaults* are another type of nudge that strongly impact disclosure [4,14,19]. Examples are framing a disclosure decision as either opt-in or opt-in, or changing the order of information requests.

The problem with nudges is that they take a one-size-fits-all approach to privacy: They assume that the "true cost" [13] of disclosure is roughly the same for every user, piece of information, and situation. But privacy decisions are highly user- and context-dependent: The fact that one person has no problems disclosing a certain item in a particular context does not mean that disclosure is equally likely for a different person, a different item, or in a different context [16,24]. Likewise, what is a convincing justification to disclose a certain item in a particular context for a certain person, may be a completely irrelevant reason for a different person, a different item, or a different context [6,21]. What we need, then, is *personalized* privacy decision support.

4. EXPLICATING PRIVACY

The first step towards personalized privacy decision support is to *explicate* the privacy calculus: to move beyond a mere description towards deeper understanding of people's cognitive decision-making process. What kind of benefits and threats do users consider when making disclosure decisions? What is the relative weight of each of these aspects? Can the weights be influenced by a justification or a default, and if so, in what context(s)? More research is needed to answer these questions.

For example, I showed in [19] that the effect of justifications on information disclosure decisions is mediated by users' perceptions of help, trust and self-anticipated satisfaction with the system. In [17], I demonstrated that the effect of decision context (i.e. the available options) in a location-sharing service depends on users' perception of the privacy and benefits of the available options. Finally, in [20] we show that perceived risk and perceived relevance mediate users' evaluation of the purpose-specificity of information disclosure requests.

5. CONTEXTUALIZING PRIVACY

The second step towards a personalized privacy decision support is to *contextualize* the privacy calculus: to determine how stable information disclosure is across people, items and situations, and, importantly, where it is context-dependent.

For example, my research shows that although justifications generally do not increase disclosure or satisfaction, tailoring justifications to the user can reduce this negative effect [21]. Such tailored justifications are *personalized privacy nudges*: they intelligently choose the correct justification for the respective user, or decide to not show any justification at all.

Similarly, *personalized defaults* can be set up in a way that anticipates people's disclosure behavior, thereby making the disclosure decisions easier and more convenient. My work and that of others shows that even though privacy preferences vary considerably across users, distinct subgroups of users with similar privacy preferences can be identified in many domains [16,26]. Moreover, these subgroups can be mapped to demographics (e.g. age) and other behaviors (e.g. mobile Internet usage). My recent work shows that these personalized defaults may also be tailored to the website requesting the information: in [20] we show that people are more likely to disclose information that matches the *purpose* of the website requesting the information.

Finally, in [17] I show that privacy decisions are influenced by the *available options* to choose from ("context effects", cf. [15,28,30]). In that study, users of a location-sharing service decided whether to share their location with friends, colleagues and third party applications, with the following options: no sharing, city, city block, or exact location. We manipulated the availability of the "city" and "exact location" options, and showed that their absence or presence had a strong impact on how many users would choose each of the other available options.

6. PRIVACY ADAPTATION PROCEDURE

The ultimate purpose of this contextualized and explicated understanding of users' privacy calculus is to develop a Privacy Adaptation Procedure to support people's privacy decisions. Using recommender system algorithms, the procedure predicts users' privacy preferences based on their known characteristics. It then provides automatic "smart default" settings in line with users' disclosure profile. Smart defaults reduce the burden of control, but at the same time respect users' inherent privacy preferences. Similarly, it provides tailored disclosure justifications, but only to users who can be expected to react rationally to them, so that they will not cause privacy scares in the other users.

This Privacy Adaptation Procedure relieves some of the burden of the privacy decision from the user by providing the right amount of information and control that is useful but not overwhelming or misleading. It thus enables users to make privacy-related decisions within the limits of their bounded rationality.

7. REFERENCES

1. Acquisti, A. and Gross, R. Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook. In *Privacy Enhancing Technologies*. 2006, 36–58.
2. Acquisti, A. and Grossklags, J. Privacy and Rationality in Individual Decision Making. *IEEE Security & Privacy 3*, 1 (2005), 26–33.
3. Acquisti, A. and Grossklags, J. What Can Behavioral Economics Teach Us About Privacy? In A. Acquisti et al., eds., *Digital Privacy*. Taylor & Francis, 2008, 363–377.
4. Acquisti, A., John, L.K., and Loewenstein, G. The Impact of Relative Standards on the Propensity to Disclose. *Journal of Marketing Research 49*, 2 (2012), 160–174.
5. Acquisti, A. Privacy in Electronic Commerce and the Economics of Immediate Gratification. *Proceedings of the ACM Conference on Electronic Commerce*, (2004), 21–29.
6. Besmer, A., Watson, J., and Lipford, H.R. The impact of social navigation on privacy policy configuration. *Proceedings of SOUPS*, (2010).
7. Bustos, L. Best Practice Gone Bad: 4 Shocking A/B Tests. *GetElastic*, 2012. http://www.getelastic.com/best-practice-gone-bad-4-shocking-ab-tests/.
8. Compañó, R. and Lusoli, W. The Policy Maker's Anguish: Regulating Personal Data Behavior Between Paradoxes and Dilemmas. In T. Moore, D. Pym and C. Ioannidis, eds., *Economics of Information Security and Privacy*. Springer US, New York, NY, 2010, 169–185.
9. Connolly, T. and Zeelenberg, M. Regret in decision making. *Current directions in psych. science 11*, 6 (2002), 212–216.
10. Culnan, M.J. "How Did They Get My Name?": An Exploratory Investigation of Consumer Attitudes toward Secondary Information Use. *MISQ 17*, 3 (1993), 341–363.
11. Eppler, M.J. and Mengis, J. The Concept of Information Overload: A Review of Literature from Organization Science, Accounting, Marketing, MIS, and Related Disciplines. *The Information Society 20*, 5 (2004), 325–344.
12. Iyengar, S.S. and Lepper, M.R. When choice is demotivating: Can one desire too much of a good thing? *J. of Personality and Social Psychology 79*, 6 (2000), 995–1006.
13. John, L.K., Acquisti, A., and Loewenstein, G. Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information. *J of Consumer Research 37*, 5 (2011), 858–873.
14. Johnson, E.J., Bellman, S., and Lohse, G.L. Defaults, Framing and Privacy: Why Opting In ≠ Opting Out. *Marketing Letters 13*, 1 (2002), 5–15.
15. Kahneman, D. and Tversky, A. Prospect Theory: An Analysis of Decision under Risk. *Econometrica 47*, 2 (1979), 263–292.
16. Knijnenburg, B.P., Kobsa, A., and Jin, H. Dimensionality of information disclosure behavior. *International Journal of Human-Computer Studies*, (2013).
17. Knijnenburg, B.P., Kobsa, A., and Jin, H. Preference-based location sharing: are more privacy options really better? *Proceedings of CHI*, (2013), 2667–2676.
18. Knijnenburg, B.P., Kobsa, A., and Saldamli, G. Privacy in Mobile Personalized Systems: The Effect of Disclosure Justifications. *Proceedings of U-PriSM*, (2012).
19. Knijnenburg, B.P. and Kobsa, A. Making Decisions about Privacy: Information Disclosure in Context-Aware Recommender Systems. *ACM Transactions on Interactive Intelligent Systems 3*, 3 (2013).
20. Knijnenburg, B.P. and Kobsa, A. Counteracting the negative effect of form auto-completion on the privacy calculus. *Proceedings of ICIS*, (2013).
21. Knijnenburg, B.P. and Kobsa, A. Helping users with information disclosure decisions: potential for adaptation. *Proceedings of IUI*, (2013), 407–416.
22. Kolter, J. and Pernul, G. Generating User-Understandable Privacy Preferences. *2009 International Conference on Availability, Reliability and Security*, IEEE Computer Society (2009), 299–306.
23. Laufer, R.S. and Wolfe, M. Privacy as a Concept and a Social Issue: A Multidimensional Developmental Theory. *Journal of Social Issues 33*, 3 (1977), 22–42.
24. Li, H., Sarathy, R., and Xu, H. Understanding situational online information disclosure as a privacy calculus. *Journal of Computer Information Systems 51*, 1 (2010), 62–71.
25. Liu, Y., Gummadi, K.P., Krishnamurthy, B., and Mislove, A. Analyzing facebook privacy settings: user expectations vs. reality. *Proc. SIGCOMM 2011*, ACM (2011), 61–70.
26. Phelps, J., Nowak, G., and Ferrell, E. Privacy Concerns and Consumer Willingness to Provide Personal Information. *Journal of Public Policy & Marketing 19*, 1 (2000), 27–41.
27. Scheibehenne, B., Greifeneder, R., and Todd, P.M. Can There Ever Be Too Many Options? A Meta-Analytic Review of Choice Overload. *Journal of Consumer Research 37*, 3 (2010), 409–425.
28. Simonson, I. Choice Based on Reasons: The Case of Attraction and Compromise Effects. *Journal of Consumer Research 16*, 2 (1989), 158–174.
29. Thaler, R.H. and Sunstein, C. *Nudge : improving decisions about health, wealth, and happiness*. Yale University Press, New Haven, NJ & London, U.K., 2008.
30. Tversky, A. Elimination by aspects: A theory of choice. *Psychological Review 79*, 4 (1972), 281–299.