Death To The Privacy Calculus?

Bart P. Knijnenburg
Clemson University
Clemson, SC, USA
bartk@clemson.edu

Elaine Raybourn*
Sandia National Laboratories
Albuquerque, NM, USA
emraybo@sandia.gov

David Cherry
Daricia Wilkinson
Saadhika Sivakumar
Clemson University
Clemson, SC, USA
dcherry@clemson.edu
dariciw@clemson.edu
ssivaku@clemson.edu

Henry Sloan
Nyack High School
Nyack, NY, USA
henryksloan@gmail.com

This work has been supported by DoD research award W911QY-16-C-0105.

*Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. The views expressed are those of the author and do not necessarily represent the views or policies of the Advanced Distributed Learning (ADL) Initiative or the Office of the Deputy Assistant Secretary of Defense.

Abstract

The “privacy calculus” has been used extensively to describe how people make privacy-related decisions. At the same time, many researchers have found that such decisions are often anything but calculated. More recently, the privacy calculus has been used in service of machine learning approaches to privacy. This position paper discusses the practical and ethical questions that arise from this use of the privacy calculus.

Author Keywords

Privacy calculus; user-tailored privacy; ethics.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

Introduction

Laufer and Wolfe [28,29] coined the term “calculus of behavior” to refer to the cognitive process that underlies people’s disclosure decisions. Many researchers have since used the term “privacy calculus” to describe privacy-related decision behaviors [10,11,13,30,33,52], and it has become a well-established concept in privacy research [31,37,42]. Other researchers, however, have demonstrated that people rarely take a truly calculative approach to privacy decision making, and are often prone to take mental shortcuts instead [2,48].

We discuss these departures from rationality, how they come about, and the impact they have on the presumed normative justifications for existing privacy solutions. This will lead us to a relatively new type of privacy solution, *user-tailored privacy*, which addresses some of the ethical questions raised by existing solutions. User-tailored privacy uses the privacy calculus *prescriptively*, with the risk/benefit tradeoff serving as an objective function for machine learning algorithms [7,14,20]. We will argue that this use of the privacy calculus raises its own set of practical and ethical questions that may cause ethical dilemmas. In outlining these questions, we hope to spark a discussion of the ethical concerns regarding user-tailored privacy.

Privacy Calculus as a Descriptive Theory?

The privacy calculus is commonly operationalized as a tradeoff between *risk* and *benefit*. The psychological process behind this tradeoff is often seen as a conscious and rational decision process. For example, Li [31] argues that the privacy calculus can be seen as a privacy-specific instance of utility maximization or expectancy-value theory [5,40,46]. These specific decision theories have been criticized for making unrealistic assumptions about the rationality of decision-makers [12,41], and a similar criticism can be leveled against the privacy calculus itself [18,19].

Rather than being rational, people's privacy decisions are influenced by various heuristics, such as information on others' privacy decisions (i.e. "social proof" [3]), the order of sensitivity in which decisions are being made ("foot in the door" and "door in the face" [3]), the overall professionalism of the privacy-setting user interface ("affect heuristic" [17]), the available options to choose from ("context non-invariance" [24]), and the default setting and phrasing of privacy-related requests ("default" and "framing" effects [22,27]).

Given these well-documented departures from rationality, it is surprising that the privacy calculus is such a prominent theory of privacy decision making. This may be because most research on privacy decision making asks users to evaluate risk and benefit using a *retrospective* and *holistic* approach rather than looking at the level of individual decisions [9,13,15,16,39,51,52]. Using this approach, it is hard to invalidate the privacy calculus, because these retrospective evaluations are just as likely to be post hoc rationalizations as they are to be the true motivations behind users' behaviors.

Indeed, users' privacy decisions are much more akin to "plans" in Activity Theory [6]: both risk and benefit are *anticipated* (in that users will usually not know the consequences of their decision up front and can thus only base their judgments on past outcomes) and *contextualized* (in that they have to regard the consequences of taking a *specific* action with regard to a *specific* recipient in a *specific* context) [10,32,39,43]. This contextualized anticipatory nature of privacy decisions is also at the core of Altman's *privacy regulation theory* [4], Nissenbaum's *contextual integrity* [34], and Petronio's *communication privacy management* [38]. In other words, privacy decisions are much more complex than the privacy calculus presumes them to be. This has consequences for the two main privacy paradigms in place today: notice and choice, and privacy nudging.

Consequences for Notice and Choice

Notice and choice are *prerequisites* of the privacy calculus: notice enables us to assess risks and benefits, and choice is needed to make meaningful tradeoffs. However, the contextualized nature of privacy behaviors means that users need to make separate choices for each context, resulting in complex privacy-setting in-

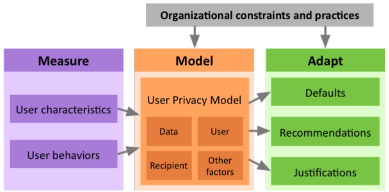

Figure 1: A schematic representation of user-tailored privacy:

The system first measures users' characteristics and privacy-related behaviors.

It uses these measurements to create a personalized model of the users' willingness to disclose different types of data to different types of recipients, in the context of other factors that may influence their decision.

Finally, it adapts the user interface to the predicted privacy decision, by changing the default privacy setting, giving an explicit recommendation, and/or providing a context-based justification for the predicted behavior.

terfaces. Similarly, the anticipated nature of privacy means that even with extensive notice, users have imperfect knowledge about the consequences of their actions. Complexity and incomplete information often result in heuristic decision-making [8]. Notice and choice may thus seem like an ethical way of providing privacy protection from a privacy calculus perspective, but if you see privacy behaviors as contextualized anticipatory reflections, then notice and choice are not enough to protect users' privacy.

Consequences for Privacy Nudging
Privacy nudging attempts to make it easier to take privacy-preserving actions by creating a *choice architecture* that promotes benefit and avoids risk [1,47]. A privacy nudge would promote safe features (e.g. highlighting or enabling them by default) and dissuade users from using risky features (e.g. hiding or disabling them by default). However, because privacy behaviors are contextualized, users' actions are based on complex identities that include their culture, world view, life experience, personality, intent, and so on, and they may thus perceive different features as "risky" and "safe" [25,50]. Moreover, any given user's preferences may change if the context changes. Nudging may seem like an ethically justifiable practice from a privacy calculus perspective, but if you see privacy behaviors as contextualized anticipatory reflections, then it becomes clear that nudges are rarely good for everyone, and may thus threaten consumer autonomy [44,45].

Privacy Calculus as a Prescriptive Theory?

How can we move beyond the "one-size-fits-all" approach to privacy embodied in both nudges and notice and choice? A more recent paradigm is that of "user-tailored privacy" (see Figure 1), which provides personalized decision support by first predicting users' privacy preferences and behaviors and then providing *adaptive nudges* (e.g. automatic initial default settings). The most prominent examples of user-tailored privacy use the privacy calculus in a *prescriptive* manner, with the risk/benefit tradeoff serving as an objective function for machine learning algorithms [7,14,20]. In this prescriptive approach, the user is no longer responsible for determining the risks and benefits, and making the tradeoff; instead, an algorithm will automatically make this tradeoff, taking the context, the user's known characteristics, their decision history, and the decision history of like-minded other users into account.

The reliance on machine learning means that the system will alleviate the decision burden via a nudge that presumably has no normative "valence" but is instead based on each users' actual preferences within the decision context [20]. This approach raises its own set of practical and ethical questions though. These questions and their normative consequences are discussed below.

What contextual variables should be included?
Earlier we suggested contextual variables that influence users' privacy decision behavior: the user, the information, and the recipient. Research shows that even when these parameters are equal, each user still shows variable behavior from one instance to the next [36]. It is thus possible that there are other contextual variables that should be included in the model as well. However, measuring too many contextual variables will turn the procedure itself into a threat to user privacy.

How should risk and benefit be determined?
One way to determine the risk of a privacy-related behavior is to measure its prominence among users [20].

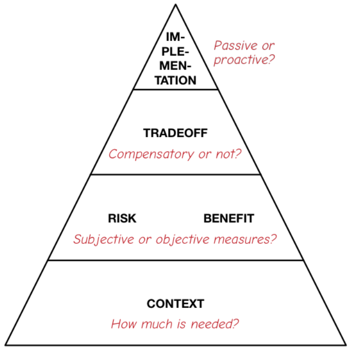

Figure 2: The moral dilemmas regarding user-tailored privacy:

How much information is needed to accurately model risk and benefit in context?

Should risk and benefit be measured in a subjective or objective manner?

Should the risk/benefit tradeoff be modeled as a compensatory or a non-compensatory decision?

Should the user-tailored adaptation take a passive or active form?

Behavior may confound risk with other factors, which will need to be disentangled [14]. But even when measured carefully, behavior is still open to external influences (as discussed earlier), creating an imbalance between attitudes and behaviors (i.e., the "privacy paradox" [35,48]). One could also measure risk *perceptions*. These may differ per user, though, which may result in a computationally intractable definition of risk. Finally, one could opt for expert opinions of risk, but getting contextualized expert risk estimates is challenging, given the vast range of possible contexts.

How should benefit be determined?
If information is collected for personalization purposes, then it may be possible to specify an objective benefits calculation, driven by the predicted utility of the information for the system [20]. Adaptive systems can often capitalize on unanticipated correlations between personal information and preferences, so this "objective" benefit may sometimes be quite different from users' perception of benefit. Adequate explanations or justifications can reduce the conflict that this may generate. Systems in which disclosure has a less well-defined benefit must rely on perceived benefit regardless.

How should the tradeoff be modeled?
One possible implementation of a risk/benefit tradeoff is a linear function of the two [7]. In this function the relative weight of risk versus relevance can be dynamically estimated for each user, or there may be different user-tailored weights for various types of information, since privacy behaviors are multidimensional [25,50]. A linear function of risk and benefit models a compensatory decision strategy (i.e. high levels of benefit can compensate high levels of risk). Alternatively, a non-compensatory threshold model puts a user-tailored upper bound on the maximum tolerable level of risk. Recent work shows this to be a preferable solution due to its predictably bounded behavior [20].

How should the adaptation be presented?
The outcome of the risk/benefit tradeoff can be used to compare possible privacy-related behaviors and determine which behavior is most beneficial to the user. Subsequently, the system has several opportunities to act upon this knowledge. The most passive action it can take is to provide the user suggestions, or to highlight the most beneficial options [21,23]. A more proactive approach would be to prioritize information requests, or to set default settings in line with this knowledge [23]. Care needs to be taken to give users a certain amount of autonomy, without overburdening them.

Discussion and Conclusion

These questions give rise to a normative discussion about the true purpose—the objective function—of user-tailored privacy (see Figure 2). For example, using behavioral or perceptual measurements of risk and benefit makes the normative assumption that the system should tailor to the user's *current* privacy practices or attitudes. While this avoids nudging users into using features they do not want to use, one could question whether some users' attitudes and behaviors are simply a product of their lack of awareness [49]. Alternatively, one could make a normative case for a version of user-tailored privacy that promotes features that the user is currently *not* using, in an effort make them more aware of these features. Such "self-actualizing" [26] privacy recommendations would arguably need to be paired with a presentation method that is less proactive, lest we inadvertently nudge users into privacy behaviors that are antithetical to their core values.

References

1. Alessandro Acquisti. 2009. Nudging Privacy: The Behavioral Economics of Personal Information. *IEEE Security and Privacy* 7: 82–85. https://doi.org/10.1109/MSP.2009.163
2. Alessandro Acquisti and Jens Grossklags. 2005. Privacy and Rationality in Individual Decision Making. *IEEE Security & Privacy* 3, 1: 26–33. https://doi.org/10.1109/MSP.2005.22
3. Alessandro Acquisti, Leslie K John, and George Loewenstein. 2012. The Impact of Relative Standards on the Propensity to Disclose. *Journal of Marketing Research* 49, 2: 160–174. https://doi.org/10.1509/jmr.09.0215
4. Irwin Altman. 1975. *The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding.* Brooks/Cole Publishing Company, Monterey, CA.
5. Naveen Farag Awad and M. S. Krishnan. 2006. The Personalization Privacy Paradox: An Empirical Evaluation of Information Transparency and the Willingness to be Profiled Online for Personalization. *MIS Quarterly* 30, 1: 13–28.
6. Jakob E. Bardram. 1997. Plans as Situated Action: An Activity Theory Approach to Workflow Systems. *Proceedings of the Fifth European Conference on Computer Supported Cooperative Work*: 17–32. https://doi.org/10.1007/978-94-015-7372-6_2
7. Michael Benisch, Patrick Gage Kelley, Norman Sadeh, and Lorrie Faith Cranor. 2011. Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs. *Personal Ubiquitous Computing* 15, 7: 679–694. https://doi.org/10.1007/s00779-010-0346-0
8. James R. Bettman, Mary Frances Luce, and John W. Payne. 1998. Constructive Consumer Choice Processes. *Journal of Consumer Research* 25, 3: 187–217. https://doi.org/10.1086/209535
9. R. K Chellappa and R. G Sin. 2005. Personalization versus privacy: An empirical examination of the online consumer’s dilemma. *Information Technology and Management* 6, 2: 181–202. https://doi.org/10.1007/s10799-005-5879-y
10. Mary J Culnan and Robert J Bies. 2003. Consumer Privacy: Balancing Economic and Justice Considerations. *Journal of Social Issues* 59, 2: 323–342. https://doi.org/10.1111/1540-4560.00067
11. Tamara Dinev and Paul Hart. 2006. An Extended Privacy Calculus Model for E-Commerce Transactions. *Information Systems Research* 17, 1: 61–80. https://doi.org/10.1287/isre.1060.0080
12. Gerd Gigerenzer and Daniel G. Goldstein. 1996. Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review* 103, 4: 650–669. https://doi.org/10.1037/0033-295X.103.4.650
13. Il-Horn Hann, Kai-Lung Hui, Sang-Yong Lee, and Ivan Png. 2007. Overcoming Online Information Privacy Concerns: An Information-Processing Theory Approach. *Journal of Management Information Systems* 24, 2: 13–42. https://doi.org/10.2753/MIS0742-1222240202
14. Ron Hirschprung, Eran Toch, Frank Bolton, and Oded Maimon. 2016. A methodology for estimating the value of privacy in information disclosure systems. *Computers in Human Behavior* 61: 443–453. https://doi.org/10.1016/j.chb.2016.03.033
15. Shuk Ying Ho and Kar Tam. 2006. Understanding the Impact of Web Personalization on User Information Processing and Decision Outcomes. *MIS Quarterly* 30, 4: 865–890.

16. Kai-Lung Hui, Bernard C. Y. Tan, and Chyan-Yee Goh. 2006. Online information disclosure: Motivators and measurements. *ACM Transactions on Internet Technology* 6, 4: 415–441. https://doi.org/10.1145/1183463.1183467
17. Leslie K. John, Alessandro Acquisti, and George Loewenstein. 2011. Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information. *Journal of consumer research* 37, 5: 858–873. https://doi.org/10.1086/656423
18. Flavius Kehr, Tobias Kowatsch, Daniel Wentzel, and Elgar Fleisch. 2015. Thinking Styles and Privacy Decisions: Need for Cognition, Faith into Intuition, and the Privacy Calculus. In *12th International Conference on Wirtschaftsinformatik*.
19. Flavius Kehr, Daniel Wentzel, and Peter Mayer. 2013. Rethinking the Privacy Calculus: On the Role of Dispositional Factors and Affect. In *ICIS 2013 Proceedings*.
20. Bart P. Knijnenburg. 2015. A user-tailored approach to privacy decision support. University of California, Irvine, Irvine, CA. Retrieved from http://search.proquest.com/docview/1725139739/abstract
21. Bart P Knijnenburg and Hongxia Jin. 2013. The Persuasive Effect of Privacy Recommendations. In *Twelfth Annual Workshop on HCI Research in MIS*. Retrieved from http://aisel.aisnet.org/sighci2013/16
22. Bart P. Knijnenburg and A. Kobsa. 2014. Increasing Sharing Tendency Without Reducing Satisfaction: Finding the Best Privacy-Settings User Interface for Social Networks. In *ICIS 2014 Proceedings*. Retrieved from http://aisel.aisnet.org/icis2014/proceedings/ISSecurity/4
23. Bart P. Knijnenburg and Alfred Kobsa. 2013. Helping users with information disclosure decisions: potential for adaptation. In *Proceedings of the 2013 ACM international conference on Intelligent User Interfaces*, 407–416. https://doi.org/10.1145/2449396.2449448
24. Bart P. Knijnenburg, Alfred Kobsa, and Hongxia Jin. 2013. Preference-based location sharing: are more privacy options really better? In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2667–2676. https://doi.org/10.1145/2470654.2481369
25. Bart P. Knijnenburg, Alfred Kobsa, and Hongxia Jin. 2013. Dimensionality of information disclosure behavior. *International Journal of Human-Computer Studies* 71, 12: 1144–1162. https://doi.org/10.1016/j.ijhcs.2013.06.003
26. Bart P. Knijnenburg, Saadhika Sivakumar, and Daricia Wilkinson. 2016. Recommender Systems for Self-Actualization. In *Proceedings of the 10th ACM Conference on Recommender Systems*, 11–14. https://doi.org/10.1145/2959100.2959189
27. Yee-Lin Lai and Kai-Lung Hui. 2006. Internet Opt-In and Opt-Out: Investigating the Roles of Frames, Defaults and Privacy Concerns. In *Proceedings of the 2006 ACM SIGMIS Conference on Computer Personnel Research*, 253–263. https://doi.org/10.1145/1125170.1125230
28. Robert S Laufer, Harold M Proshansky, and Maxine Wolfe. 1973. Some Analytic Dimensions of Privacy. In *Proceedings of the Lund Conference on Architectural Psychology*.
29. Robert S Laufer and Maxine Wolfe. 1977. Privacy as a Concept and a Social Issue: A Multidimensional Developmental Theory. *Journal of Social Issues*

33, 3: 22–42. https://doi.org/10.1111/j.1540-4560.1977.tb01880.x

30. H. Li, R. Sarathy, and H. Xu. 2010. Understanding situational online information disclosure as a privacy calculus. *Journal of Computer Information Systems* 51, 1: 62–71.
31. Yuan Li. 2012. Theories in online information privacy research: A critical review and an integrated framework. *Decision Support Systems* 54, 1: 471–481. https://doi.org/10.1016/j.dss.2012.06.010
32. Naresh K Malhotra, Sung S Kim, and James Agarwal. 2004. Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Nomological Framework. *Information Systems Research* 15, 4: 336–355. https://doi.org/10.1287/isre.1040.0032
33. George R. Milne and Mary Ellen Gordon. 1993. Direct Mail Privacy-Efficiency Trade-offs within an Implied Social Contract Framework. *Journal of Public Policy & Marketing* 12, 2: 206–215. https://doi.org/10.2307/30000091
34. Helen Nissenbaum. 2004. Privacy as Contextual Integrity. *Washington Law Review* 79: 119–157.
35. Patricia A. Norberg, Daniel R. Horne, and David A. Horne. 2007. The Privacy Paradox: Personal Information Disclosure Intentions versus Behaviors. *Journal of Consumer Affairs* 41, 1: 100–126. https://doi.org/10.1111/j.1745-6606.2006.00070.x
36. Sameer Patil, Roman Schlegel, Apu Kapadia, and Adam J. Lee. 2014. Reflection or Action?: How Feedback and Control Affect Location Sharing Decisions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 101–110. https://doi.org/10.1145/2556288.2557121
37. Paul A Pavlou. 2011. State of the Information Privacy Literature: Where Are We Now and Where Should We Go. *MIS Quarterly* 35, 4: 977–988.
38. Sandra Petronio. 1991. Communication Boundary Management: A Theoretical Model of Managing Disclosure of Private Information Between Marital Couples. *Communication Theory* 1, 4: 311–335. https://doi.org/10.1111/j.1468-2885.1991.tb00023.x
39. Joseph Phelps, Glen Nowak, and Elizabeth Ferrell. 2000. Privacy Concerns and Consumer Willingness to Provide Personal Information. *Journal of Public Policy & Marketing* 19, 1: 27–41. https://doi.org/10.1509/jppm.19.1.27.16941
40. Roland T. Rust, P. K. Kannan, and Na Peng. 2002. The Customer Economics of Internet Privacy. *Journal of the Academy of Marketing Science* 30, 4: 455–464. https://doi.org/10.1177/009207002236917
41. Herbert A. Simon. 1982. *Models of Bounded Rationality: Empirically grounded economic reason*. MIT Press.
42. H. Jeff Smith, Tamara Dinev, and Heng Xu. 2011. Information Privacy Research: An Interdisciplinary Review. *MIS Quarterly* 35, 4: 989–1016.
43. H. Jeff Smith, Sandra J Milberg, and Sandra J Burke. 1996. Information Privacy: Measuring Individuals' Concerns about Organizational Practices. *MIS Quarterly* 20, 2: 167–196. https://doi.org/10.2307/249477
44. N. Craig Smith, Daniel G. Goldstein, and Eric J. Johnson. 2013. Choice Without Awareness: Ethical and Policy Implications of Defaults. *Journal of Public Policy & Marketing* 32, 2: 159–172. https://doi.org/10.1509/jppm.10.114

45. Daniel J. Solove. 2013. Privacy Self-Management and the Consent Dilemma. *Harvard Law Review* 126: 1880–1903.

46. Eugene F Stone and Dianna L Stone. 1990. Privacy in Organizations: Theoretical Issues, Research Findings, and Protection Mechanisms. *Research in Personnel and Human Resources Management* 8: 349–411.

47. Richard H Thaler and Cass Sunstein. 2008. *Nudge : improving decisions about health, wealth, and happiness*. Yale University Press, New Haven, NJ & London, U.K.

48. Dave Wilson and Joseph Valacich. 2012. Unpacking the Privacy Paradox: Irrational Decision-Making within the Privacy Calculus. In *ICIS 2012 Proceedings*. Retrieved from http://aisel.aisnet.org/icis2012/proceedings/ResearchInProgress/101

49. Pamela J. Wisniewski, Bart P. Knijnenburg, and Heather Richter Lipford. 2016. Making Privacy Personal: Profiling Social Network Users to Inform Privacy Education and Nudging. *International Journal of Human-Computer Studies*. https://doi.org/10.1016/j.ijhcs.2016.09.006

50. Pamela J. Wisniewski, Bart P. Knijnenburg, and Heather Richter Lipford. 2017. Making privacy personal: Profiling social network users to inform privacy education and nudging. *International Journal of Human-Computer Studies* 98: 95–108. https://doi.org/10.1016/j.ijhcs.2016.09.006

51. Heng Xu, Xin (Robert) Luo, John M Carroll, and Mary Beth Rosson. 2011. The personalization privacy paradox: An exploratory study of decision making process for location-aware marketing. *Decision Support Systems* 51, 1: 42–52. https://doi.org/10.1016/j.dss.2010.11.017

52. Heng Xu, Hock-Hai Teo, Bernard C. Y. Tan, and Ritu Agarwal. 2009. The Role of Push-Pull Technology in Privacy Calculus: The Case of Location-Based Services. *Journal of Management Information Systems* 26, 3: 135–174. https://doi.org/10.2753/MIS0742-1222260305