Taking Control of Household IoT Device Privacy

A White Paper for the Sociotechnical Cybersecurity Workshop

Bart P. Knijnenburg, Clemson University — bartk@clemson.edu
Alfred Kobsa, University of California, Irvine — kobsa@uci.edu
Martijn C. Willemsen, Eindhoven University of Technology — m.c.willemsen@tue.nl

Household IoT devices are intended to collect information in the home, and interact with each other, to create powerful new applications that support our day-to-day activities. For example, a smart fridge detects when certain groceries are running out, and triggers a shopping console to order fresh supplies. Technical solutions can ascertain that the data used for such functionality are minimized [13, 20, 24], but arguably, any functionality requires at least some amount of personal data. Therefore, users will have to trade off privacy and functionality: a solution that is fully privacy preserving will be limited in functionality, while a fully functional IoT solution may require far-reaching data collection and communication. Each user may have a different personal preference on this privacy-functionality tradeoff continuum; the challenge is to have them pick a setting that is pareto-optimal, i.e., one that provides the highest level of privacy for a given level of functionality (and vice versa). This is a different data minimization challenge: not on the system side, but on the *user side*.

However, one can question if users will be able to make such optimal privacy decisions by themselves. Each device will have its own fine-grained privacy settings, and there exist many interdependencies between devices—both in privacy and functionality. Therefore, there are many possibilities for users to make inconsistent privacy decisions that limit functionality (e.g. the shopping console cannot order food if it is not allowed to communicate with the smart fridge) or that do not protect privacy in the end (e.g. if the fridge is not allowed to send open/close events but the oven is, it still allows the server to know about the users' whereabouts).

Moreover, while privacy researchers argue that users employ a *privacy calculus* [7, 16]—i.e. that they make disclosure decisions by trading off the anticipated benefits with the risks of disclosure [23]—others have demonstrated that findings from *behavioral economics* apply to privacy decision-making as well [2, 3]—i.e. that users are often prone to take mental shortcuts (i.e. *heuristics*) that to not consider the privacy-benefit ratio, but are instead prone to numerous decision biases [4, 10, 11, 15]. Beyond demonstrating the existence of these decision biases, little effort has been put into understanding the underlying mechanisms and decision processes that cause them, which is essential to understand how to better support privacy decision making.

Research suggests that users' privacy decisions in a complex IoT usage scenario are very likely to be suboptimal. Our grand challenge is thus to improve household IoT users' privacy decisions. We note that prior privacy research has not focused on the underlying decision processes to better understand *how*, *why* and *when* these decisions are suboptimal. We thus call for evidence-based socio-technical research to a) uncover users' decision processes, and b) use this knowledge to design adequate decision support mechanisms.

Why is this a grand challenge?

To our best knowledge, this challenge requires a major advance in user-centric privacy research, because very little work in this area has considered improving the privacy decision process itself. Such research is required, though, because traditional approaches to user-centric privacy research are likely inadequate to solve household IoT privacy problems.

For example, the traditional approach of "transparency and control" [9, 14, 18, 21, 27] requires a rich user interface that enables users to consumer privacy-related information and control disclosure-related settings. But IoT devices are often controlled by voice command or mobile device, and even then the interaction is purposefully reduced due to the pervasive and ubiquitous nature of the technology.

Similarly, privacy nudges [1, 5, 6, 25] take a “one-size-fits-all” approach to privacy [22]: They assume that the “true cost” [10] of disclosure is roughly the same for every user, for every piece of information, in every situation. Household IoT scenarios consist of unique combinations of devices though, which precludes making “one-size-fits-all” inferences about the best privacy settings. But even disregarding such unique configurations, different users are likely to have different preferences, and the same users may indeed have different preferences depending on the specific situation.

We propose a more thorough investigation of the variability and context-dependency of privacy decisions (i.e., on what dimensions do people differ in their information disclosure behavior, and which contextual variables influence this decision?) as well as the (sub-)optimality of such decisions (i.e., in what situations do users' actual decisions deviate from their longer-term preferences?). This will allow IoT developers to offer a more succinct set of privacy controls that can be supported within the limited interaction bandwidth users tend to have with their household IoT devices, without oversimplifying the available control to the extent that it leaves users vulnerable to privacy threats.

This challenge also requires a major advance in decision-making research, because unlike e.g. product decisions, privacy decision outcomes are vague, uncertain, and emotionally laden. This means that such decisions can often not be captured in conventional decision matrices.

A socio-technical approach

This challenge requires a significant academic research effort to carefully study household IoT users' privacy decisions in both controlled and real-world environments. Methods such as process tracing [26] can improve our theoretical understanding of the cognitive processes that are underlying classic phenomena such as loss aversion and context effects, while part-worth-utility mapping [12] can enable the selection of a set of IoT privacy management profiles that efficiently spans the risk-benefit spectrum, thereby reducing choice overload.

It also requires policy-makers to determine the boundaries of admissible data collection. IoT devices are expected to operate in many different regulatory environments, and server-side device-to-device communications likely involves unprecedented volumes of data traversing several legal jurisdictions [8, 17, 19]. From a regulatory perspective, this makes privacy management a global concern. The coordination of admissible data collection and sharing practices between legal jurisdictions is a formidable task.

Finally, it requires industry efforts to incorporate the control templates proposed by academic researchers and policy-makers into existing IoT management platforms such as Microsoft's Azure IoT Suite, Apple's HomeKit, and Samsung's new IoT real-time operating system. The challenge of this implementation lies in automatically adjusting these control templates to the set of devices that has been connected to the platform, and in working within the boundaries of the available settings of these connected devices.

Conclusion

Privacy research that focuses directly on decision processes can have a transformatory effect on the difficulty of privacy decisions and the usability of privacy interfaces. Arguably, the current focus on transparency and control merely makes privacy decisions easier for expert users, as interpreting the vast amount of information and harnessing the provided control requires skills that many users do not possess. Like nudges, our suggested approach will make privacy decisions easier for all users, not just for the ones who happen to have the requisite digital skills. Unlike nudges though, this approach avoids a one-sided paternalistic approach in favor of empowering users to make better privacy decisions themselves.

References

[1] Acquisti, A. 2009. Nudging Privacy: The Behavioral Economics of Personal Information. *IEEE Security and Privacy*. 7, (Nov. 2009), 82–85.

[2] Acquisti, A. and Grossklags, J. 2005. Privacy and Rationality in Individual Decision Making. *IEEE Security & Privacy*. 3, 1 (2005), 26–33.

[3] 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.

[4] 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 (2012), 160–174.

[5] Almuhimedi, H., Schaub, F., Sadeh, N., Adjerid, I., Acquisti, A., Gluck, J., Cranor, L.F. and Agarwal, Y. 2015. Your Location Has Been Shared 5,398 Times!: A Field Study on Mobile App Privacy Nudging. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2015), 787–796.

[6] Balebako, R. and Cranor, L. 2014. Improving App Privacy: Nudging App Developers to Protect User Privacy. *IEEE Security & Privacy*.

[7] Culnan, M.J. 1993. "How Did They Get My Name?": An Exploratory Investigation of Consumer Attitudes toward Secondary Information Use. *MIS Quarterly*. 17, 3 (1993), 341–363.

[8] Greenleaf, G. 2012. The influence of European data privacy standards outside Europe: implications for globalization of Convention 108. *International Data Privacy Law*. 2, 2 (May 2012), 68–92.

[9] 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 (Mar. 2007), 19–33.

[10] 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 (Feb. 2011), 858–873.

[11] 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.

[12] 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.

[13] Kobsa, A. 2007. Privacy-Enhanced Personalization. *Communications of the ACM*. 50, 8 (2007), 24–33.

[14] Kolter, J. and Pernul, G. 2009. Generating User-Understandable Privacy Preferences. *Conf. on Availability, Reliability and Security* (Fukuoka, Japan, 2009), 299–306.

[15] Lai, Y.-L. and Hui, K.-L. 2006. Internet Opt-In and Opt-Out: Investigating the Roles of Frames, Defaults and Privacy Concerns. *Proceedings of the 2006 ACM SIGMIS Conference on Computer Personnel Research* (Claremont, CA, 2006), 253–263.

[16] Laufer, R.S. and Wolfe, M. 1977. Privacy as a Concept and a Social Issue: A Multidimensional Developmental Theory. *Journal of Social Issues*. 33, 3 (1977), 22–42.

[17] LeSieur, F. 2012. Regulating cross-border data flows and privacy in the networked digital environment and global knowledge economy. *International Data Privacy Law*. 2, 2 (May 2012), 93–104.

[18] Metzger, M.J. 2006. Effects of Site, Vendor, and Consumer Characteristics on Web Site Trust and Disclosure. *Communication Research*. 33, 3 (2006), 155–179.

[19] Moerel, L. 2011. Back to basics: when does EU data protection law apply? *International Data Privacy Law*. 1, 2 (May 2011), 92–110.

[20] Pfitzmann, A. and Köhntopp, M. 2001. Anonymity, Unobservability, and Pseudonymity: A Proposal for Terminology. *Anonymity 2000*. H. Federrath, ed. Springer-Verlag. 1–9.

[21] 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 (2005), 339–360.

[22] Spiekermann, S., Grossklags, J. and Berendt, B. 2001. E-privacy in 2nd Generation E-Commerce: Privacy Preferences versus Actual Behavior. *Proceedings of the 3rd ACM conference on Electronic Commerce* (Tampa, FL, 2001), 38–47.

[23] Taylor, D., Davis, D. and Jillapalli, R. 2009. Privacy concern and online personalization: The moderating effects of information control and compensation. *Electronic Commerce Research*. 9, 3 (Sep. 2009), 203–223.

[24] Verykios, V.S., Bertino, E., Fovino, I.N., Provenza, L.P., Saygin, Y. and Theodoridis, Y. 2004. State-of-the-art in Privacy Preserving Data Mining. *SIGMOD Rec*. 33, 1 (Mar. 2004), 50–57.

[25] Wang, Y., Leon, P.G., Acquisti, A., Cranor, L.F., Forget, A. and Sadeh, N. 2014. A Field Trial of Privacy Nudges for Facebook. *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems* (Toronto, Canada, 2014), 2367–2376.

[26] Willemsen, M.C., Böckenholt, U. and Johnson, E.J. 2011. Choice by value encoding and value construction: Processes of loss aversion. *Journal of Experimental Psychology: General*. 140, 3 (2011), 303–324.

[27] Xu, H. 2007. The effects of self-construal and perceived control on privacy concerns. *ICIS 2007 Proceedings* (2007), paper 125.