Think before you click: how reflective patterns contribute to privacy

Arnout Terpstra

Tilburg Institute of Law, Technology & Society (TILT), Tilburg University, a.c.terpstra@tilburguniversity.edu

SURF, arnout.terpstra@surf.nl

Paul Graßl

iHub, Radboud University, p.grassl@bsi.ru.nl.

Hanna Schraffenberger

iHub, Radboud University, h.schraffenberger@ru.nl

The digital economy thrives on personal data. This drives companies, which often are legally obliged to ask for users' consent before collecting their data, to employ manipulative design patterns that unconsciously steer users to provide more data: dark patterns. As an opposing force, scholars have explored the idea of *bright* patterns (nudging users towards more privacy-friendly choices). In this paper, however, we argue that nudging users through design towards any privacy choice (whether dark or bright) is morally problematic. We make a case for *reflective* design patterns, which stimulate reflective thought processes through designing-in friction. Through reflection, individuals examine how digital products might impact their privacy and how it fits their world views. We discuss how users can or cannot be expected to deliberate about their privacy choices and what role design plays in this context. We end with a case study of the identity management system IRMA to illustrate how reflective design patterns can be implemented in practice. This paper draws from and integrates previous work by the authors [Graßl et al. 2021, Terpstra et al. 2019, Jacobs and Schraffenberger 2020].

CCS CONCEPTS • Human Computer Interaction (HCI) • Interaction design • Human and societal aspects of security and privacy

Additional Keywords and Phrases: privacy, dark patterns, bright patterns, reflective patterns, reflection, friction

1 INTRODUCTION

In today's highly digitized society, citizens face enormous challenges to protect their privacy. Digital products are often financed by personalised advertising, provoking incentives for their creators to collect as much personal data from users as possible. In many cases, users have to agree to the collection of their personal data beforehand, leading to a vast amount of consent requests an average user has to face. Against this background, design strategies emerged which go by the name *dark patterns,* steering users towards a certain choice through persuasive interface design [Brignull n.d., Gray et al. 2018]. Within the context of privacy, dark patterns can be understood as evil design nudges that push users towards the least privacy-friendly option. Nudging means guiding people towards decisions in line with their best interest by making minor changes in the choice environment without compromising freedom of choice [Thaler and Sunstein 2009]. A famous example of a nudge is the fly painted on urinals in public men's toilets to prevent urine spillage. There are many examples of dark patterns within the context of privacy, as outlined by [Graßl et al. 2021] or covered in the media [Pardes 2020]. A visual example of a dark pattern is presented below [Figure 1].

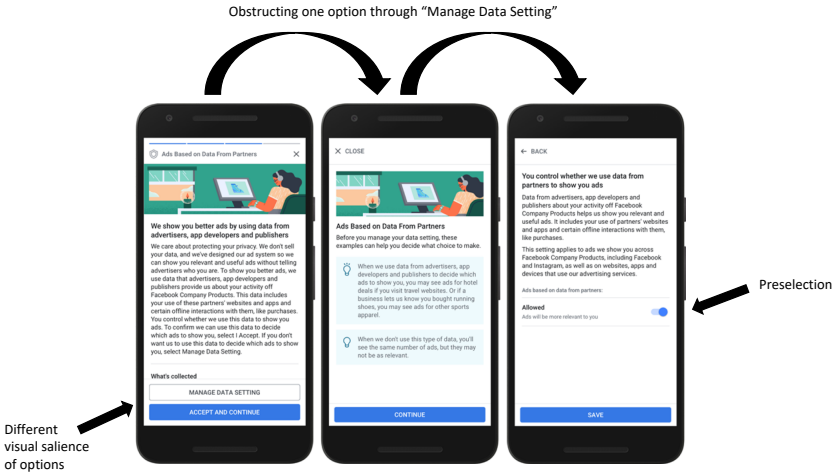

Figure 1: Example of dark patterns used by Facebook

One could plot design strategies that steer privacy choices along a line which begins in the dark (at 0% brightness) and ends in the light (100% brightness). The more a choice architecture is built with the best interest of the individual in mind, the lighter it becomes. One could then easily conclude that in order to turn around the slow erosion of online privacy [Koops and Leenes 2006], all one needs to do is ensure design strategies used in interaction design become brighter and brighter. However, our main thesis for this workshop is that by focusing on the *brightness* of design patterns, a crucial and fundamental characteristic of privacy (regulation) is overlooked: individual choice is reduced to one of compliancy [Terpstra et al. 2019]. This affects how individuals deal with and learn about privacy and privacy issues: by depreciating the moral value of (privacy) choice, individuals will feel resigned [Turow et al. 2015, Hargittai and Marwick 2016] and cannot become morally responsible agents [Brownsword 2005, Mezirow 2003, 1997]. Thus, designers should develop and apply patterns that appeal to a user's reflective ability: *reflective patterns*. Through reflective thinking, users can (deliberately) reason about their privacy preferences and learn from past decisions how to proceed with current privacy choices. In this paper, we briefly discuss how reflective patterns contribute to privacy decision-making and finish with an example of an application built specifically with reflective patterns in mind: IRMA (I Reveal My Attributes).

2 PRIVACY DECISIONS AND REFLECTIVE PATTERNS

Online privacy is increasingly regulated, for example through the General Data Protection Regulation in Europe [GDPR] or the California Consumer Privacy Act in (a part of) the U.S. [CCPA]. These regulations assume a ‘knowledge gap’ between an organisation offering a service and an individual using that service [Calo 2014, Hoofnagle and Urban 2014]. Through principles such as transparency (or notice), organisations are legally obliged to be transparent about what data they collect and for what purposes, such that this knowledge gap can be reduced. By reading privacy policies, prospective users are expected to understand in what ways their privacy might be affected. By expanding individuals’ rights and furthering their control (or consent) over the use of data and their privacy, individuals are expected to make well-informed, rational (privacy) decisions.

However, many problems have been identified with transparency and control. In sum: they do not accomplish their goal of effectively informing users nor giving users meaningful control over their data [Terpstra et al. 2019]. In order to learn and reason about how using certain digital products might affect one's privacy (i.e., to make 'good' privacy decisions) individuals need to think reflectively [Terpstra et al. 2019]. Reflective thinking, or (critical) reflection, is the examination of previous experiences and assumptions as input for future actions, assumptions and decisions [White et al. 2006, ten Dam and Volman 2004, Ennis 1991]. It starts with an awareness that existing assumptions may need to be (re-)examined, initiated by a *disorienting dilemma* [Mezirow 2006, 2000]. A disorienting dilemma, or breakdown [Baumer 2015], refers to experiences, beliefs or feelings that are incompatible with an individual's current perspectives. In a second step, a phase of inquiry [Baumer 2015] takes place during which one actively examines what caused the disorienting dilemma and how it differs from one's existing way of making sense of the world. Lastly, the reflective thinking process ends with transformation [Baumer 2015]: after re-evaluating one's assumptions, changes are made in beliefs, attitudes and behaviour.

The notion that mindful usage of and continuous reflection on (digital) products are beneficial towards individuals and society is acknowledged by several design theories, such as 'reflective design' [Sengers et al. 2005], 'adverserial design' [DiSalvo 2012], 'critical design' [Dunne and Raby 2001] and 'slow design' [Grosse-Hering et al. 2013]. Individuals should thus be encouraged – through the design of the products and services they use – to make individual choices and maintain a moral position by using their reflective capabilities to reflect on how they think and feel about privacy before, during, and after interacting with digital technology. This requires at least three components: (1) a way to challenge one's habitual behaviour and thoughts, e.g., through friction [Terpstra et al. 2019]; (2) a phase of reflective thinking about one's own privacy behaviour and its consequences, supported by the digital product, e.g., by giving the user information on privacy issues; (3) meaningful controls: the possibility to put newly learned thoughts and ideas into action.

Within the context of privacy, reflection should not only occur at decision time. Provoking reflection should *fuse* with the interaction(s) of digital products, such that users are invited to reflect on their interactions and privacy at different moments, whenever it suits them. To accomplish this, we suggest, for example, deliberately slowing down interactions, breaking up interactions into separate smaller parts, or deliberately making interactions more complex (similar strategies are proposed by Distler et al. [2020] to reduce risk-taking behaviour). Other strategies include deliberately designing-in ambiguous content, different and contrasting opinions or perspectives [Vasalou, et al. 2015], or asking specific questions that lead individuals to consider other perspectives [Broockman and Kalla 2016]. Doing so will (likely) cause friction, allowing users to escape habits and become conscious of their interactions, as well as simply give users more time to think (contrasting with 'optimising' the interface to reduce the number of clicks and increase the speed at which users accomplish their goal).

To be clear, we acknowledge it is unrealistic to assume that people should only make deliberate privacy decisions all the time. Repetitive privacy decisions of lower importance (e.g., cookie settings) could, for instance, be addressed by regulation with a global browser setting that websites have to respect. However, in those situations where the importance is high or the potential (negative) consequences can be significant (e.g., allowing Facebook to use face recognition to identify you in photos) and an informed choice is necessary, 'reflective' patterns should be applied.

3 CASE STUDY: IRMA

IRMA is a mobile app for identity management (developed by the Privacy by Design Foundation [PbD]), which takes the form of a virtual wallet. In this wallet, users collect cards that contain personal information and which are issued

by trusted sources. Users can selectively disclose the information from the cards online, e.g., to log in to a website or sign a document. In the following, we discuss the IRMA app as an example of a design that uses reflective patterns. The section revisits and extends ideas presented by [Jacobs and Schraffenberger 2020] in an earlier popular science article. On the technological level, the IRMA app has been designed with a *Privacy by Design* approach. This entails that the user's data is stored exclusively on the user's phone and that data revealed via IRMA can be limited to what is absolutely necessary, thereby supporting *data minimization*. To access, for instance, an age-restricted game with IRMA, one only reveals that one is older than 18 instead of one's birthdate. On the interaction design level, the app follows a *Privacy by Deliberation* approach that encourages people to make careful decisions and take control over the release of their personal data [Jacobs and Schraffenberger 2020]. The design includes three strategies that contribute to deliberate data-sharing decisions:

1. **Real-world metaphors:** IRMA uses the metaphor of a wallet that contains cards with personal information (similar to ID cards). This encourages users to treat IRMA and the information in it with the same carefulness that they apply to their physical wallet and ID card. (The municipality of Amsterdam contributed this metaphor in 2020 as part of a bigger redesign.)
2. **Friction:** Revealing personal data with IRMA requires the user to carry out a series of steps, which stimulate people to consider whether sharing their data is really necessary. Users need to unlock the IRMA app with their PIN, start the IRMA-disclosure session (e.g., by scanning a QR or clicking a link on the website), and confirm the disclosure of their data. By requiring these steps and interrupting their primary flow, people – to some degree – may also be prevented from sharing data in the 'heat of the moment' or by accident (e.g., 'clicking faster than thinking').
3. **What You See Is What You Share (WYSIWYS):** Confronting people with their actual data (e.g., *their* name, *their* date of birth) helps to make the disclosure-decision as concrete as possible. Thus, when a party requests information from the user, the app informs the user about *who* is requesting information from them (the URL or name of the requesting party) and presents a visual overview of the exact information that is requested. We call this principle: "What you see is what you share". The principle also entails that for requests that do not fit on the screen, users need to scroll through the entire list before agreeing to share the information.

These three principles contribute to deliberate data sharing practices. However, we realize that users still might share sensitive information without giving this much thought. Reasons for this might be that the disclosure screen highlights the option to share data with a big blue button. Also, users have already put in quite some effort when they face the disclosure screen. This might motivate them to accept the request, so their previous effort was not 'for nothing'. Furthermore, asking users to go through the same series of steps with every single data-request likely contributes to habituation, which might harm the effectiveness of friction [Distler et al. 2020]. While IRMA provides a sound basis, we thus see room for improvement. Based on our personal experience and informal observations from various user tests, we envision improvements of IRMA as well as of data-sharing requests in general by **considering:**

1. **...in which cases users actually should make a decision themselves**. E.g., in the case of IRMA, it might be better if *overasking* requests (requests for data that are not necessary for a certain service) would never reach the users. An idea to prevent such overasking requests when they concern sensitive data has been presented by [Jacobs and Schraffenberger 2020]: Requesting parties could be required to obtain a certificate before they can request information like a citizen number from IRMA users. To obtain such a certificate, the requesting

party would have to make explicit why they need this data. This way, only justified requests for sensitive data would reach the end user.

2. **...how much user reflection should go into a disclosure decision**. As a guideline, we suggest that the reflection should be proportional to the risks associated with sharing the data. In the case of IRMA, sharing a citizen service number poses more risks than sharing a non-attribute like "above 18". A first idea to implement a form of *attribute-dependent friction*, is to color-code requests according to the sensitivity of the requested data – to stimulate reflection when it matters most [Jacobs and Schraffenberger 2020].
3. **...the level on which a decision has to be made**. If decisions likely are made on a *general* level (e.g., "I never want cookies"), one should not promote deliberation on a case-by-case basis (e.g., asking users about this on every single website). Similarly, if IRMA users likely want to agree to *all* requests of a specific *type* by a *specific party* (e.g., all >18 requests by a gaming site), we should not bother users to reflect on each of those requests individually [Jacobs and Schraffenberger 2020].

As these considerations show, the IRMA app has provided us with a rich use case and live laboratory for exploring data disclosure requests. Our insights about how to design for deliberate decisions likely translate to other data sharing requests. The core idea is to combine the principle of data minimization with helping users to make their own deliberate decisions – by promoting reflection through design. For the latter, the most important realization is that designing for deliberation requires design patterns that – unlike many existing patterns – do not strive for efficiency but instead stimulate the user to reflect [Jacobs and Schraffenberger 2020].

4 CONCLUSION

This paper has argued that the field of HCI needs to come up with reflective design patterns that help people make deliberate privacy decisions. We invite participants of this workshop to develop proposals for such reflective patterns. In the context of reflective patterns, it is essential to revisit and reconsider widely accepted and often-desired design goals such as consistency, efficiency and effectiveness critically. Consistency likely causes people to 'automate' processes, which can be harmful if we want people to stop and think. Likewise, while traditional interaction design often aims at efficiency, design for deliberation can mean we need to slow people down intentionally and give them the time they need to make a decision. Furthermore, we have to be careful when measuring the effectiveness of reflective patterns. While it is tempting to measure whether people can successfully accept and decline a request, a reflective pattern is only truly effective when users also make the choices that they do not later regret. This does not mean we need to throw overboard everything we (interaction designers) know. Existing principles like visibility might still apply. For instance, making the actual data one is asked to share visible can make a request for data much more concrete.

ACKNOWLEDGEMENTS

We thank prof. dr. Bart Jacobs for his valuable comments on earlier versions of this paper.

REFERENCES

Graßl, P., Schraffenberger, H., Zuiderveen Borgesius, F., & Buijzen, M. (2021). Dark and Bright Patterns in Cookie Consent Requests. Journal of Digital Social Research, 3(1), 1-38. https://doi.org/10.33621/jdsr.v3i1.54

Terpstra, A., Schouten, A. P., de Rooij, A., & Leenes, R. E. (2019). Improving privacy choice through design: How designing for reflection could support privacy self-management. First Monday, 24(7). https://doi.org/10.5210/fm.v24i7.9358

Jacobs, B., & Schraffenberger, H. (2020). Friction for privacy. why privacy by design needs user experience design. European Cyber Security

Perspectives 2020, 12–14.

Brignull, H. (n.d.). Dark patterns. Retrieved from https://darkpatterns.org/

Gray, C. M., Kou, Y., Battles, B., Hoggatt, J., & Toombs, A. L. (2018). The dark (patterns) side of UX design. In R. Mandryk, M. Hancock, M. Perry, & A. Cox (Eds.), Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18 (pp. 1–14). New York, New York, USA: ACM Press. https://doi.org/10.1145/3173574.3174108

Thaler, R. H., & Sunstein, C. R. (2009). Nudge: Improving decisions about health, wealth, and happiness (Rev. and expanded ed). New York: Penguin Books.

Pardes, A. (2020, August 12). How Facebook and Other Sites Manipulate Your Privacy Choices. Wired. https://www.wired.com/story/facebook-social-media-privacy-dark-patterns/

Koops, B.-J., Leenes, R., 2006. "'Code' and the slow erosion of privacy," Michigan Telecommunications and Technology Law Review, volume 12, number 1, pp. 115–188, and at http://repository.law.umich.edu/mttlr/vol12/iss1/3

Turow, J., Hennessy, M., Draper, N., 2015. "The tradeoff fallacy: How marketers are misrepresenting American consumers and opening them up to exploitation," Annenberg School of Communication, University of Pennsylvania, at https://www.asc.upenn.edu/sites/default/files/TradeoffFallacy_1.pdf

Hargittai, E., Marwick, A., 2016. "'What can I really do?' Explaining the privacy paradox with online apathy," International Journal of Communication, volume 10, at https://ijoc.org/index.php/ijoc/article/view/4655

Brownsword, R., 2005. "Code, control, and choice: Why east is east and west is west," Legal Studies, volume 25, number 1, pp. 1–21. doi: https://doi.org/10.1111/j.1748-121X.2005.tb00268.x

Mezirow, J., 2003. "Transformative learning as discourse," Journal of Transformative Education, volume 1, number 1, pp. 58–63. doi: https://doi.org/10.1177/1541344603252172

GDPR. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). Official Journal L, 119, 1–88. https://eur-lex.europa.eu/eli/reg/2016/679/oj

CCPA. (2018). California Consumer Privacy Act of 2018. https://leginfo.legislature.ca.gov/faces/codes_displayText.xhtml?division=3.&part=4.&lawCode=CIV&title=1.81.5

Mezirow, J., 1997. "Transformative learning: Theory to practice," New Directions for Adult & Continuing Education, volume 1997, number 74, pp. 5–12. doi: https://doi.org/10.1002/ace.7401

Calo, R., 2014. "Code, nudge, or notice?" Iowa Law Review, volume 99, number 2, pp. 773–802, and at https://ilr.law.uiowa.edu/assets/Uploads/ILR-99-2-Calo.pdf

Hoofnagle, C.J., Urban, J.M., 2014. "Alan Westin's privacy Homo Economicus," Wake Forest Law Review, volume 49, number 2, pp. 261–317, and at https://scholarship.law.berkeley.edu/facpubs/2395/

White, S., Fook, J., Gardner, F., 2006. "Critical reflection: A review of contemporary literature and understandings," In: S. White, J. Fook, and F. Gardner (editors). Critical reflection in health and social care. Maidenhead: Open University Press, pp. 3–20.

ten Dam, G., Volman, M., 2004. "Critical thinking as a citizenship competence: Teaching strategies," Learning and Instruction, volume 14, number 4, pp. 359–379. doi: https://doi.org/10.1016/j.learninstruc.2004.01.005

Ennis, C.D., 1991. "Discrete thinking skills in two teachers' physical education classes," Elementary School Journal, volume 91, number 5, pp. 473–487. doi: https://doi.org/10.1086/461670

Mezirow, J., 2006. "An overview on transformative learning," In: P. Sutherland and J. Crowther (editors). Lifelong learning: Concepts and contexts. London: Routledge, pp. 90–105.

Mezirow, J., 2000. "Learning to think like an adult: Core concepts of transformation theory," In: J. Mezirow and Associates (editors). Learning as transformation: Critical perspectives on a theory in progress. San Francisco: Jossey-Bass, pp. 3–33.

Baumer, E.P.S., 2015. "Reflective informatics: Conceptual dimensions for designing technologies of reflection," CHI '15: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, pp. 585–594. doi: https://doi.org/10.1145/2702123.2702234

Sengers, P., Boehner, K., David, S., Kaye, J., 2005. "Reflective design," CC '05: Proceedings of the Fourth Decennial Conference on Critical Computing: Between Sense and Sensibility, pp. 49–58. doi: https://doi.org/10.1145/1094562.1094569

DiSalvo, C., 2012. Adversarial design. Cambridge, Mass.: MIT Press.

Dunne, A., Raby, F., 2001.Design Noir: The Secret Life of Electronic Objects. Birkhäuser.

Grosse-Hering, B., Mason, J., Aliakseyeu, D., Bakker, C., & Desmet, P., 2013. Slow design for meaningful interactions. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 3431-3440).

Distler, V., Lenzini, G., Lallemand, C., & Koenig, V. (2020, October). The Framework of Security-Enhancing Friction: How UX Can Help Users Behave More Securely. In New Security Paradigms Workshop 2020 (pp. 45-58).

Vasalou, A., Oostveen, A.-M., Bowers, C., Beale, R., 2015. "Understanding engagement with the privacy domain through design research," Journal of the Association for Information Science and Technology, volume 66, number 6, pp. 1,263–1,273. doi: https://doi.org/10.1002/asi.23260

Broockman, D., Kalla, J., 2016. "Durably reducing transphobia: A field experiment on door-to-door canvassing," Science, volume 352, number 6282 (8 April), pp. 220–224. doi: https://doi.org/10.1126/science.aad9713

PbD, Privacy by Design Foundation, https://privacybydesign.foundation/