Obstacles to Transparency in Privacy Engineering

Kiel Brennan-Marquez
Information Law Institute
New York University School of Law
Information Society Project
Yale Law School

Daniel Susser
Information Law Institute
New York University School of Law
Philosophy Department
San Jose State University

Abstract—Transparency is widely recognized as indispensable to privacy protection. However, producing transparency for end-users is often antithetical to a variety of other technical, business, and regulatory interests. These conflicts create obstacles which stand in the way of developing tools which provide meaningful privacy protections or from having such tools adopted in widespread fashion. In this paper, we develop a "map" of these common obstacles to transparency, in order to assist privacy engineers in successfully navigating them. Furthermore, we argue that some of these obstacles can be successfully avoided by distinguishing between two different conceptions of transparency and considering which is at stake in a given case—transparency as providing users with insight into what information about them is collected and how it is processed (what we call transparency as a "view under-the-hood") and transparency as providing users with facility in navigating the risks and benefits of using particular technologies.

Keywords—privacy, transparency, privacy engineering, privacy management, information ethics

I. Introduction

One key component of privacy engineering is transparency for end users. To exercise meaningful choice over privacy-related decisions, users must be aware of how privacy protections are (or are not) operating in the background. Put otherwise, if users are not privy to the makeup of privacy protections, privacy engineering will be an elite, rather than democratic, enterprise—responsive not to user preferences and interests, but exclusively to those of governments and corporations.

Yet the need for transparency collides with various prerogatives common to for-profit firms, and this collision gives rise to various obstacles for privacy engineers. Because much work on privacy engineering occurs in academic or otherwise free and/or open-source contexts, it is tempting (and sometimes appropriate) to brush such worries aside. If, however, privacy engineers and technical privacy managers hope to have any widespread impact in the private sector they must take stock of common obstacles to transparency.

With that background in mind, this paper makes two analytic contributions, both of which seek to facilitate transparency-focused privacy engineering in practice. *First*, we offer a "map" of different obstacle-types that privacy engineers are likely to face, grouped according to the institutional dynamics that create them. The goal of this "map" is to help engineers within firms (as well as academic commentators) understand *why* obstacles to transparency arise, and to equip engineers to better navigate the obstacles in practice.

Second, we suggest a strategy for circumventing obstacles to transparency. Stated formally, the strategy is to focus on engineering solutions that maximize the *functional goal* of transparency—greater user facility regarding the collection and use of information—while minimizing "under-the-hood exposure" for the firm. In some cases, the two go hand in hand; the former depends on the latter. But in other cases, user facility can be improved without significantly expanding a firm's "under-the-hood" exposure, and in those cases, we think that privacy engineers and privacy managers have a distinctive—and crucial—role to play in advancing transparency.

II. Privacy and Transparency

From the beginning, transparency has been front and center in discussions about information privacy. The original (1973) U.S. Department of Health, Education, and Welfare report on *Records, Computers, and the Rights of Citizens* made two transparency-related recommendations: (1) "There must be no personal data record-keeping systems whose very existence is secret," and (2) "There must be a way for an individual to find out what information about him is in a record and how it is used." Those recommendations were then translated into the (1980) Organization for Economic Cooperation and Development (OECD) Guidelines on the Protection of Privacy and Transborder Flows of Personal Data, becoming the "Openness Principle": "There should be a general policy of openness about developments, practices and policies with respect to personal data. Means should be readily available of establishing the existence and nature of personal data, and the main purposes of their use, as well as the identity and usual residence of the data controller."

The U.S. Federal Trade Commission's Fair Information Privacy Practices require transparency or openness through their "Notice/Awareness" principle: "Web sites would be required to provide consumers clear and conspicuous notice of their information practices, including what information they collect, how they collect it (e.g., directly or through non-obvious means such as cookies), how they use it, how they provide Choice, Access, and Security to consumers, whether they disclose the information collected to other entities, and whether other entities are collecting information through the site." Finally in 2012 the Obama Administration released a

DOI 10.1109/SPW.2016.18

"Consumer Privacy Bill of Rights," which contains an explicit "Transparency Principle": "Consumers have a right to easily understandable and accessible information about privacy and security practices. At times and in places that are most useful to enabling consumers to gain a meaningful understanding of privacy risks and the ability to exercise Individual Control, companies should provide clear descriptions of what personal data they collect, why they need the data, how they will use it, when they will delete the data or de-identify it from consumers, and whether and for what purposes they may share personal data with third parties."

These policy discussions have suffused the theory and practice of privacy engineering. Janic et al., for instance, catalog sixteen privacy enhancing technologies (PETs) specifically designed to provide end-user transparency, which they refer to as "transparency enhancing tools" (TETs) (2013). Each of these tools, they argue, either provides users with information about "how service providers claim to handle user's personal information" or with information about "how service providers actually handle user's personal information," and do so in "an accurate and for an average Internet user comprehensible way" (Janic et al. 2013). As Christian Zimmermann puts it, the goal of TETs is "providing users with insight into a data controller's intended and actual data handling behavior" (2015).

III. Obstacles to Transparency

The impulse to provide users with information about what is happening "behind the scenes" conflicts with a number of other imperatives that businesses face. Given the institutional environment in which for-profit firms operate, it is only natural that whether or not new technologies are privacy-protective will depend, in practice, on the extent to which firms feel able to deliver transparency without (substantially) compromising other goals. And this, in turn, will depend partly on the ability of engineers and privacy advocates to *convince* firms of that possibility. To that end, we propose a map of obstacles to transparency that privacy engineers are likely to face in practice, with the hope that greater clarity about the source of different kinds of obstacles will aid in their navigation.

For analytic purposes, obstacles to transparency can be divided into three types. First, technical or design considerations might cause transparency to hamper the purpose of either the overall service or the privacy protections themselves. Second, institutional or regulatory considerations might render transparency, under some circumstances, unworkable; disclosing certain information about internal privacy protocols may invite liability—in the form of either legal sanctions or consumer ire—down the line. Third, strategic or economic considerations could make businesses wary about granting public access to (at least some portion of) their internal workings. We consider each in turn.

A. Technical and Design Considerations

There are a number of different ways that demands for transparency can conflict with technical or design considerations. One is that giving users a window into how a system works can, at least in principle, render the system more vulnerable to intrusion or attack. Second, forthrightness about what information is being collected (or about how it is being processed) can undermine functionality. One example of where this problem is likely to arise is in online controlled experiments, often referred to as A/B testing. In such tests, two or more different versions of a website or user interface are presented to different users (divided into "control" and "treatment" groups). Information is then collected about the types of user interaction and levels of user engagement elicited from each variation, and that information is used to refine the design, determine the return on investment (ROI) of new features, and so on. Nearly all major web-based companies, including Microsoft, Amazon, eBay, Etsy, Facebook, Google, and LinkedIn, use A/B testing extensively (Kahavi et al. 2013). Yet this kind of experiment requires an absence of transparency. If users knew how the interface they were engaging with was being manipulated it would no doubt skew the experimental results.

A related, though deeper, example brings the issue into even starker relief. There is a burgeoning field of websites and apps designed to facilitate experiments in psychology research. For instance, in 2014 cognitive scientists at Oxford University and the University of Birmingham developed a free app containing four classic psychology experiments in the form of short games. In the first month after the app was made publicly available 20,800 people used it to participate in their studies (Brown et al. 2014). When designing and conducting such experiments it is often necessary to mislead participants, since tendencies such as the "Hawthorne effect"—where research subjects change their behavior to meet perceived expectations—can bias results. Although some psychologists argue that misleading experiment subjects is wrong for both moral and methodological reasons, it remains a widespread practice (Bortolotti and Mameli 2006).

Whether for the sake of improving interface usability or for the sake of scientific research, experiments conducted online through websites or apps layer information privacy concerns on top of existing ethical questions regarding informed consent and participant autonomy. How is transparency to be provided for end-users when user awareness undermines functionality?

B. Legal, Regulatory, and Reputational Obstacles

Another obstacle to transparency is that firms face potential liability for representations they make about information practices. Such liability comes in formal and informal varieties. On the formal side, consumer protection statutes and regulations leave firms vulnerable to lawsuits alleging "misrepresentation" of business practices, including information practices. (To pick just one example by way of illustration, LinkedIn recently faced a lawsuit from subscribers to its "premium" service, alleging, among other things, that LinkedIn's Privacy Policy misrepresented the sturdiness of its data security infrastructure.) On the informal side, firms face potential criticism from watchdog groups and media outlets, as well as scrutiny from government agencies, such as the Federal Trade Commission. Further compounding the problem is the multi-jurisdictional nature of liabilities connected to disclosure. Someday, a federal statutory scheme may comprehensively regulate information privacy. But until then,

firms are left to contend with consumer protection laws from fifty-one discrete jurisdictions—all fifty states, plus the federal system. And this is not even to mention transnational legal issues.

Perversely, the complex regulatory environment occupied by information firms often makes them either (1) less inclined to engage in discretionary transparency, or (2) prone to flooding users with information—for example, in the form of byzantine "terms of use" agreements or privacy policies—that satisfy formal disclosure obligations but fail to enrich user understanding. The latter route is particularly concerning if, as we argue below, user understanding is the crux of transparency. But it is also unsurprising, given the realities of privacy law. After all, courts "almost invariably look to the terms of use agreement or to the privacy policy," as opposed, for example, to a user's "specif[ied] privacy preferences," when assessing the scope of privacy expectations (Harzog 2011).[1] In light of this, it only stands to reason that firms would dedicate an outsized amount of attention to formal instruments of disclosure.

C. Strategic and Economic Obstacles

A final obstacle to transparency—in some ways, the most intractable one—is that firms often have a generalized economic interest in keeping their practices opaque. Sometimes, this is true in a formal sense, for example under trade secret laws that make "secrecy" a factual prerequisite of legal protection. More commonly, it is true in an informal sense, when firms derive an economic benefit from keeping their information practices—including privacy practices—outside the public view. We pass no judgment on whether information firms are *justified* in ascribing strategic or economic interests to opacity. In some cases, they likely are; in other cases, they likely are not, but firms (perhaps justifiably) err on the side of claiming broad business interest. In either case, it seems clear that a reflexive aversion to transparency, on economic grounds, operates as something of an industry norm.

What is more, reflexive aversion to transparency seems especially rampant among information companies. Apple, for example, has grown infamous for its obsession with secrecy, even going so far as to shield large swaths of information from its own employees (Fox 2013). Of course, secrecy about things like product design does not necessarily correspond to secrecy about information privacy practices. But once secrecy becomes a norm, it naturally permeates a firm's culture. Similarly, Silicon Valley companies, particularly in their early stages, are notorious for limiting public information about their internal operations—for example, by forcing employees to sign aggressive non-disclosure and non-compete agreements (Lobel 2013). Again, the point is not that these measures necessarily lead to secrecy about information privacy practices, but rather that they reflect a distinctive culture among information firms today.

[1] It bears noting that some scholars, including Professor Harzog, advocate expanding consumer deception laws to encompass privacy "disclosures" effectuated through interface design. In this paper, our analysis is confined to the law as it currently exists.

IV. Circumventing Obstacles to Transparency

The three obstacle-types identified in the last section are endemic to privacy design in the sense that they flow from technical, regulatory, and business climates in which information firms operate. Although not every attempt to increase transparency will meet with all three obstacles—and some attempts to increase transparency may meet with none of them—the obstacles are here to stay. But the obstacles can often be circumvented in practice, and the actors in the best position to appreciate that fact and to guide the circumvention effort, will be privacy engineers.

A. Strategies Based on Obstacle-Type

To begin with, we briefly note some obvious strategies for circumventing each of the obstacles identified in the last part. Our goal here is simply to demonstrate our map's utility; we assume that the most effective strategies will come from privacy practitioners themselves.

Consider first technical or design obstacles. If informing end-users about how information is being collected and processed threatens to undermine the functionality of the program or system itself, one might decide from the start to reevaluate one's approach. There is always more than one way to solve a technical problem; if the strategy a team has adopted is incompatible with providing end-user transparency the team would do well to rethink its strategy at the outset, rather than trying to bandage its effects piecemeal later on.

On the legal or regulatory side, the obstacle stems from making explicit claims about how a technology or system works. As such, one strategy to avoid the obstacle in question is to indicate information about the system to users implicitly rather than explicitly, so as to avoid legal liability and compliance requirements. Examples of such strategies include "skeumorphic" interface design and the reintroduction of familiar audio indicators to technologies where such indicators have been rendered obsolete, such as the "click" sound accompanying a smartphone camera's shutter (Calo 2012).

Finally, faced with a strategic or economic obstacle, where transparency conflicts with a firm's desire to keep its practices secret, privacy engineers might develop tools for indicating to users *what* a system does with information about them without revealing *how* the system does it. They could, in other words, provide partial information rather than a full accounting of a firm's data practices. In contrast to design obstacles (where even partial information could undermine functionality) and legal obstacles (where it could invite liability), strategic obstacles could potentially be avoided simply by circumscribing the amount of information provided to users.

B. An Overall Strategy: Retooling "Transparency"

Beyond these obstacle-specific strategies, there is also an overall strategy that would help, in our view, to enable greater transparency across the board. Namely, firms should attend more carefully to *the type of transparency that is stake*. Here, privacy engineers have an important role to play; they can help focus a firm's attention on why transparency matters.

On this front, we would distinguish between two types of transparency. The first is "user facility," which refers to the extent of user understanding (and user autonomy) regarding the general uses to which information—especially sensitive information—is being put, and, as a corollary, the risks associated with sharing information with a given firm. The second type of transparency could be called "under-the-hood exposure," referring to the publication of information about a firm's specific practices.

In some cases, the two notions of transparency more or less fully overlap. If, for example, users demand information about a data breach—what was the nature of the breach, what kind of data was hacked, and so forth—it will likely be necessary for the firm to "lift the hood," and provide users (or the general public) insight into its information practices. In other words, there may be no way of fostering "user facility"—of communicating to users, for example, what risks they face—without significant under-the-hood exposure.

In many cases, however, the two notions of transparency do not fully overlap—and what ultimately matters is not under-the-hood exposure, but user facility. For example, when it comes to what information is transmitted to third-party advertisers, and how that information is packaged (e.g., in anonymized form or not), the user likely does not care about the intricacies of a firm's practices. Presumably, what she cares about is being able to anticipate likely "information flows" (Nissenbaum 2010), and being able to decide, in light of anticipated information flows, whether or not she is comfortable sharing her information (and if so, which information). In fact, the point can be sharpened further. In many circumstances, apprising a user about the intricacies of a firm's practices may end up *impeding* her ultimate "facility," because they will distract from, or even interfere with, a lay understanding of how information is being managed. If the goal of transparency is to provide average users with "accurate" and "comprehensible" heuristics (Janic et al. 2013) for engaging in what Daniel Solove has helpfully termed "privacy self-management" (Solove 2013), technical details are as likely to occlude as they are to illuminate. Indeed, even if average users can make sense of technical details—which is unlikely—privacy decisions are especially susceptible to cognitive bias; for example, even sophisticated users often fail to account for aggregation effects when contemplating the significance of information-sharing (Solove 2013; compiling cognitive science research to this effect). In light of this, under-the-hood exposure is as prone to lead users astray as it is to facilitate their understanding.

By conceiving of transparency in terms of user facility (rather than under-the-hood exposure), not only will privacy engineers better vindicate the normative goals of transparency; in many contexts, they will *also* circumvent the obstacles identified in the last section. Many obstacles that fall under the "legal and regulatory" category, as well as the "strategic and economic" category, become significantly milder as attention shifts to user facility. On the first front, interface design typically falls beyond the scope of what courts consider relevant to liability (Harzog 2011), improvements to interface design can foster transparency with few, if any, legal or regulatory hiccups. Likewise, on the second front, business concerns often come down to the (perceived) risks of "opening the hood," not user understanding of what information practices support a given tool. If anything, increased user facility stands to yield business *benefits*, as user concern about information privacy grows.

V. Conclusion

When all is said and done, the purpose of transparency, from the average user's perspective, can be stated very simply. Transparency helps avoid the "black-box" character of much consumer technology today (Pasquale 2015). And in so doing, it helps attenuate the pronounced lack of user autonomy (and corresponding sense of corporate paternalism) that plagues digital participation.

To create the kind of information privacy environment we want, increasing user facility is important for both intrinsic and instrumental reasons, and privacy engineers have a crucial role to play in this effort. But we must be clear about what the effort entails. In many cases, the answer is not to impart technical knowledge. It is to clarify, in an accessible way, what potential risks and what potential benefits come along with using particular tools. Do that—and the rest will follow.

Acknowledgments

The authors are grateful to Seda Gürses, Karen Levy, Ira Rubinstein, and Joris van Hoboken for their comments.

References

[1] L. Bortolotti and M. Mameli, "Deception in psychology: moral costs and benefits of unsought self-knowledge." Accountability in Research: Policies and Quality Assurance 13(3), 2006.

[2] H. Brown, P. Zeidman, P. Smittenaar, R. Adams R, F. McNab, R. Rutledge, R. Dolan. "Crowdsourcing for cognitive science–the utility of smartphones." PLoS ONE 9(7):e100662, 2014.

[3] R. Calo, "Against notice skepticism in privacy (and elsewhere)," Notre Dame Law Review 87, pp. 1027-1072, 2012.

[4] J. Fox, "Why Apple has to become more open," Harvard Business Review (Blog), 2013.

[5] R. Gellman, "Fair information practices: a basic history," unpublished. Available http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2415020

[6] W. Harzog, "Website design as contract," American University Law Review 60, pp. 1635-71, 2011.

[7] M. Janic, J. Wijbenga, and T. Veugen, "Transparency enhancing tools (TETs): an overview," Third Workshop on Socio-Technical Aspects in Security and Trust (STAST), pp. 18-25, 2013.

[8] R. Kohavi, A. Deng, B. Frasca, T. Walker, Ya Xu, and N. Pohlmann, "Online controlled experiments at large scale," Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1168-1176, 2013.

[9] O. Lobel, Talent Wants To Be Free: Why We Should Learn to Love Leaks, Raids, and Free-Riding. New Haven: Yale University Press, 2013.

[10] H. Nissenbaum, Privacy in Context: Technology, Policy, and the Integrity of Social Life. Palo Alto: Stanford University Press, 2010.

[11] F. Pasquale, Black Box Society: The Secret Algorithms That Control Money and Information. Cambridge: Harvard University Press, 2013.

[12] D. Solove, "Privacy self-management and the consent dilemma," Harvard Law Review 126, pp. 1880-1903, 2013.

[13] C. Zimmermann, "A categorization of transparency-enhancing technologies," unpublished, 2015. Available http://arxiv.org/abs/1507.04914