Chapter 5

The Development of Privacy Norms

Nicholas Proferes

Abstract This chapter addresses how we develop, revisit, and negotiate norms around privacy when confronted with new technologies. The chapter first examines Nissenbaum's (Washington Law Review 79(1):119–157, 2004) theory of privacy as contextual integrity, a framework that helps unpack how context-relevant norms for appropriateness and transmission can be challenged by new technologies. It then reviews how social norms develop as we build mental models of how a technology works during its diffusion process. The chapter concludes with suggestions for designers about approaches for thinking through implications when a design may challenge a preexisting social norm, or where there is no socially agreed upon norm. This includes careful reflection on who challenges to the current social norms may benefit and who they may hurt.

5.1 Introduction

Privacy norms shape what we consider appropriate information flow in any given situation. For example, while we might voluntarily share sensitive health information about ourselves with a friend, it would likely violate our sense of privacy for that friend to turn around and tell our boss the same information. As new technologies such as cellphone cameras, social media platforms, and persistently listening technologies such as digital personal assistants have been introduced into society, prior privacy norms have been thrown into question. This chapter addresses how we develop, revisit, and negotiate norms around privacy when faced with new technologies.

Privacy norms have evolved over time. Historically, we defined privacy as relating to a narrow set of situations and circumstances. However, novel inventions challenged these prior privacy regimes, throwing both the norms and the concept

N. Proferes (✉)

School of Social and Behavioral Sciences, Arizona State University, Tempe, AZ, USA

e-mail: nicholas.proferes@asu.edu

© The Author(s) 2022

B. P. Knijnenburg et al. (eds.), *Modern Socio-Technical Perspectives on Privacy*,

https://doi.org/10.1007/978-3-030-82786-1_5

of privacy itself into disarray. This chapter introduces Nissenbaum's [1] model of privacy as contextual integrity as a way to help us make sense of these challenges. Nissenbaum posits that, in a given setting, contextual integrity is maintained when the norms of information appropriateness and norms of information transmission are respected; when it's not, our sense of privacy is violated. The contextual integrity framework helps unpack how context-relevant norms for appropriateness and transmission can be challenged by new technologies.

Next, the chapter details how we develop and evaluate social norms for information appropriateness and transmission in relation to new technologies. First, we build an internal mental model of how a technology works, the kinds of information it collects and transmits, and the kinds of actions and outcomes it can afford us. We then take that picture and paste it to our understandings of particular social situations. In each of these social situations, we consider the social roles associated with a given context, our own expectations and those of others, and the possible actions and practices of others, as informed by history, culture, law, and social convention. Based on this, we engage in a kind of calculus, weighing the benefits of particular technological ends afforded to us versus any challenges to established norms. Thus, at the individual level, social privacy norms are continuously revised in relation to the perceived benefits of particular uses of a technology. As we will see, this is not always a rational process and is often filled with risk and uncertainty. This process of social negotiation scales as a technology diffuses and begins to involve not just users but also social leaders, policy-makers, and designers.

The chapter concludes with suggestions for designers about approaches for thinking through implications when a design may challenge a preexisting social norm, or where there is no socially agreed upon norm. This includes careful reflection on who challenges to the current social norms may benefit and who they may hurt.

Key Takeaways

- Privacy norms shape our expectations for what's appropriate in a given situation.
- Privacy norms are socially constructed and evolve overtime. This means they may vary culturally and may change, even within a culture.
- A new technology can create new social contexts, which challenge preexisting norms. Our beliefs about how a given technology works are a key component of how we build and adapt our privacy norms.

5.2 Privacy and Challenges in Relation to Technology

In the United States, the regulatory genesis for privacy rights can be found in rules meant to protect people from government agents (such as rights against unreasonable search and seizure), restricting access to specific kinds of information about one's self (such as rights against being forced to testify against one's self), and the establishment of particular physical locations that could be considered

private (such as the privacy of one's home) [1]. These frameworks have historically helped create certain baseline expectations for privacy. For more on how privacy laws and frameworks vary internationally, see Chap. 2. However, the relevance and applicability of these frameworks have also been challenged by the development of new technologies.

Smith [2] traces the conceptualization of privacy as it has existed from the founding of the United States to contemporary outlooks, noting, "each time when there was renewed interest in protecting privacy it was in reaction to new technology" (p. 6). For example, when the handheld snap camera became widely available in the late 1800s, it became more readily possible to invade the privacy of others from a distance and to use someone else's likeness without their permission, challenging both our previous beliefs about what privacy should protect and the norms of what is appropriate. New practices emerged, such as newspaper photographers "feeding an 'unseemly gossip' industry by taking and publishing candid shots of people without their consent" [3]. In response to these practices, Brandeis and Warren [4] argued "solitude and privacy have become more essential to the individual; but modern enterprise and invention have, through invasions upon his privacy, subjected him to mental pain and distress, far greater than could be inflicted by mere bodily injury" (p. 263). As part of their seminal *Harvard Law Review* article, "The Right to Privacy," the pair argued for a new vision of privacy as the right to be "let alone."

Smith [2] observes that as computer databases meant for commercial usages became more ubiquitous in the 1960s, public concern about privacy grew rapidly. Digital technologies were increasing the ability for actors to collect, store, aggregate, and transmit information in ways that extended beyond the interruption of another's seclusion. Americans became particularly concerned with "informational privacy" (p. 6). New policies emerged in response to demand from a worried public. For example, the 1973 issuance of the Fair Information Practice Principles (FIPPs) and a US Secretary Advisory Committee report entitled *Records, Computers, and the Rights of Citizens* and the 1974 Privacy Act were enacted to try to put some controls on commercial practices [5]. At the same time, researchers such as Alan Westin [6, 7] began to push for what are some of the basic underpinnings of online privacy protections used today, such as informed consent for transmission of personal information.

Today, privacy scholars such as Solove [8] have argued that privacy is a concept in disarray. Older models have failed to keep pace with the actual practices enabled by contemporary tech [9]. Technologies such as cellphone cameras, social media platforms, and persistently listening digital personal assistants have complicated our earlier notions of what privacy should protect and what is normatively appropriate. For example, cellphone cameras have raised questions about whether or not the practice of taking pictures of strangers in public places and then circulating them online for entertainment—a practice known as posting "strangershots" [10]—is a violation of one's privacy. Social media platforms such as Facebook have raised questions about the kinds of information resharing with third parties that are socially permissible (for more on this, read [11] on the Cambridge Analytica scandal). And Internet of Things assistive technologies such as Alexa rely on human workers that

often listen to voice recordings captured in Echo owners' homes and offices, "as part of an effort to eliminate gaps in Alexa's understanding of human speech and help it better respond to commands" [12], a practice not always known to users.

We have collectively struggled with questions about whether or not these technological practices violate our privacy and what appropriate social norms are for activities involving these technologies. Each new tool creates a new social context which complicates our reliance on earlier privacy practices. Nissenbaum's [1] conceptualization of privacy as contextual integrity offers an alternative framework that helps us understand not just why new technologies constantly cause us to revisit our privacy norms but additionally offers insights into where the norms for privacy come from.

Key Takeaways

- Our understanding of what privacy is and what it should protect has historically evolved in response to the introduction of new technologies.
- Narrow definitions of privacy, such as only considering it "the right to be let alone," are being challenged by new technologies, such as cellphone cameras, social media platforms, and persistently listening digital personal assistants.
- Nissenbaum's theory of privacy as "contextual integrity" can help unpack why new technology uses challenge privacy.

5.3 Privacy as Contextual Integrity

Nissenbaum's [1] contextual integrity framework allows us to unpack how context-relevant norms for information gathering and information flow form a basis for prescriptive evaluations of privacy-related situations. One of the main premises of contextual integrity is that, rather than only having specific areas of life where privacy is a concern, there are "no arenas of life not governed by norms of information flow, no information or spheres of life for which 'anything goes.'" Almost everything—things that we do, events that occur, transactions that take place—happens in a context not only of place but of politics, convention, and cultural expectation" (p. 119). Importantly for our conversation, social norms around information flow are always in play because we always exist in some kind of contextual position. We are never outside of the social. Each life situation we find ourselves in contains its own distinct norms, which are dictated by our social role, expectations, actions, and practices. For example, we may find ourselves in the social role of "patient seeing their doctor." The norms of information for these social roles are defined by history, culture, law, and convention. Information flow in this situation is dictated by norms of how the medical field has historically treated patient information, standardized medical practices, and various laws that govern how medical information is collected.

Nissenbaum argues that there are two primary types of informational privacy norms: norms of informational appropriateness and norms of information flow.

Contextual integrity, and our underlying sense of privacy, is upheld when **both** types of norms are upheld. Norms of information appropriateness generally govern the match between the type of information being requested and the context of the request. For example, it would be perfectly reasonable for a doctor to ask a patient about their health condition. However, it might be unreasonable to go up to strangers in a park and ask them about the same question. Again, roles, expectations, actions, and practices as informed by history, culture, law, and social convention will tell us what is an appropriate request versus an inappropriate request. New technologies can challenge preexisting norms of information appropriateness. For example, a social media profile generator may make requests for information about things such as location, political affiliation, birthdays, tastes, etc. As the social media platform constitutes a new social context, users may struggle in determining what is normatively appropriate.

Norms of distribution govern “movement, or transfer of information from one party to another or others” (p. 122). It might be perfectly reasonable for our doctors to gather sensitive health information from us, but if they exchange it with our bosses without our consent, this would likely violate our normative expectations. Cellphone cameras have made it possible to record and rebroadcast the activities of others in public, challenging our earlier notions of “privacy via obscurity” [13]. However, in addition to these new information flows, what makes norms of distribution particularly tricky is that many of the information flows enabled by contemporary tools are not transparent to users. As a result, users may be broadly unaware that particular flows exist, and thus, when they are revealed, can cause considerable consternation.

Nissenbaum’s framework helps us identify reasons why particular practices may violate our expectations for privacy. For example, new technologies might demand kinds of information that we are uncomfortable sharing, thus violating information appropriateness. A technology could also be used to transmit information in ways that violate our expectations for information transmission and our imagined audiences, such as when social media posts are shown to our bosses [14, 15]. However, new technologies can also create the potential for privacy harm when we have incomplete understandings of how they work and when norms of information appropriateness and transmission are still being socially negotiated. The next section talks about how we build expectations for technology that we then use to negotiate appropriate social norms.

Key Takeaways

- Contextual integrity argues that every moment of our lives is governed by situationally informed information norms (i.e., contexts).
- In contextual integrity, there are two types of informational norms that impact our privacy evaluations: norms of information appropriateness and norms of information flow. Both must be met for us not to feel our privacy has been violated.
- The norms we use to evaluate information appropriateness and information flow are influenced by our social roles, expectations, actions, and practices, which

are themselves shaped by history, culture, law, and convention. As a result, individuals may have different evaluations of whether or not the same practice is a privacy violation, depending on these factors.

5.4 Building Expectations

Contextually relevant social norms for information appropriateness and transmission do not fall from the sky. They have their genesis in individual expectations about how technologies work and what they can afford in terms of actions and outcomes, applied to social contexts (which, again, are made up social roles, expectations, actions, and practices, as shaped by history, culture, law, and convention). We develop our understandings of what a technology does and how it might enable certain information flows through three sources: our direct interactions with a technology, watching others use a technology, and consuming discourse about a technology (e.g., reading a newspaper article about a new technology) [16]. Once we have this “mental model” in place, we situate our understandings of these information flows against broader preexisting social norms. Social norms then develop out of scaled expectations about appropriate behavior that occurs in the context [17].

For technology designers, what is particularly important in this process is how individuals develop their mental models about what a technology affords in terms of information flows. The term “affordance” originally comes from the perceptual psychologist J. J. Gibson [18], who argues that the meanings of objects in an environment can be directly perceived and that these perceptions can then be mentally linked to the possible actions that can be taken in an environment. For example, in perceiving a large leafy tree, the individual may observe that this object creates shade on a sunny day. After perceiving this affordance within the environment, the individual may take the action of sitting down under the tree to cool off (realizing this affordance in action).

Norman [19] and Gaver [20] are the two authors who are generally credited for taking Gibson’s concept from psychology and importing it into the study of technological artifacts and technological design. Gaver [20] observes that any given technology provides a set of affordances that exist in relationship with that technology’s users. These affordances “are properties of the world that are compatible with and relevant for people’s interactions” (p. 79). This is to say technologies can afford us certain interactions and outcomes within the world. For example, social media sites commonly afford various degrees of visibility to users [21]. However, for the individual to realize an affordance in action, the affordance must first be *perceptible*. It is only when technological affordances are perceptible to the individual that there can be a direct link between perception and action [20]. When the affordances of technology are not perceivable (such as when they are hidden) or are perceived incorrectly by an individual, this can lead to mistakes. Poor design choices can hinder the perceptibility of a technology’s affordances and,

hence, why badly designed technology is more likely to lead to user failures and frustration. Perceptibility of information collection and information transmission is critical for the social negotiation of appropriate norms.

Once an individual has perceived a technology, but before action, they often build a conceptual model for a technology [22]. These conceptual models are used to “test” how a technological object should work. When the individual adds in the context of the environment, themselves, and other objects in relationship with the technology to this internal picture, the individual arrives at what Norman [22] calls a mental model. Mental models are internal representations of the world that people use to model and predict the world around them. These models provide “predictive and explanatory power for understanding the interaction” ([22], p. 7).

Individuals’ mental models facilitate prediction and realizing affordances in different scenarios. Yet, most of us operate without a fully developed mental model of every technology we use. In fact, an individual’s mental models need not be fully accurate with respect to how a technology works to be functional. For example, an individual may not know the full details of how Google’s PageRank algorithm works, but that individual can likely still use Google search bar to look for websites. But this is also where the potential for violating expectations around information appropriateness and information flow can crop up. If individuals do not actively perceive information flows made possible by certain technologies and have them incorporated into their mental models, they may feel as though their privacy is violated when those flows are later revealed. For example, with persistently listening digital personal assistants, individuals may believe that a device is constantly listening for trigger words that cause the device to “wake” for requests. They may not have it built into their mental models, however, that recordings can be made by these digital assistants outside of the “trigger word” scenarios and that these recordings could be listened to by humans. This can result in their privacy feeling violated as expectations for information transmission are thrown into conflict.

There can be numerous reasons why users do not perceive particular information flows. Users have a tendency to develop understandings of how information flows from feedback mechanisms within a design interface that they directly experience. For example, Proferes [23] shows that Twitter users have more accurate understandings of how features such as hashtags, retweeting, following, and direct messaging work than Twitter’s APIs or Twitter’s data-gathering techniques that rely on the use of tracking cookies. In the absence of clear feedback mechanisms about how particular flows work, users will sometimes try to fill in gap, inferring, correctly or not. For example, Eslami et al. [24] found Facebook users “wrongly attributing the composition of their feeds to the habits or intent of their friends and family” (p. 161) rather than interventions made by the News Feed algorithm.

Outside of misunderstandings stemming from opaque design elements, users can also develop misunderstandings from the discourse that they consume. How a company talks to users about a product working is a critical part of the way individuals build mental models. For example, the messages about a company’s products communicated by its founders, CEOs, or other representatives are often picked up in the media and rebroadcast. These become important framing mechanisms for people

looking to make sense of a new tool [25]. If these speakers leave out key details (e.g., how they collect, share, or sell user data), users may develop incomplete mental models.

Once we have our mental model of the informational requests a technology makes and the kinds of information flows a technology enables, we situate these models against existing norms and engage in a kind of prescriptive evaluation. We may choose to not use a technology, use it only in certain ways, or use it wholesale and expect others to do the same. These social norms are developed out of expectations about appropriate behavior that occurs in social contexts and are built on our own beliefs about roles within the social context, our own expectations and the expectations of others, and practices, as informed by history, culture, law, and social convention. The whole-sale adoption of these norms depends on a longer process of social negotiation that takes place as the technology diffuses throughout society.

Key Takeaways

- People carry internal pictures of how they think technologies work, called “mental models.”
- People develop their mental models of technology three ways: direct interaction with a technology, watching others use a technology, or consuming messages about a technology.
- When users’ mental models don’t match actual practices of information collection and transmission, this can violate contextual integrity and their sense of privacy.

5.5 Negotiating Norms and Negotiating Technology

Once individuals have a sense of what a technology is and what it affords, they may evaluate and interpret the use of a particular technology in light of extant social norms. Depending on what they perceive the particular benefits of the technology to be in the social context, they may choose to forgo, revise, or stick to earlier normative behaviors. For example, sharing certain kinds of information publicly on social media may have been seen as socially unfathomable 50 years ago but is seen as an acceptable social practice today. Many users find that the benefits they derive from such sharing outweigh what might have earlier been seen as violating norms of appropriateness. On the flipside, many individuals have chosen to forgo the use of digital personal assistants because of fears about the kinds of data they collect [26].

However, users’ decisions about the adherence to privacy norms are not always rational. Early privacy analysis from behavioral economics often tried on the use of rational choice theory to explain why users make certain privacy decisions. In reality, users’ decisions are often made for less than rational reasons [27]. For more on this, see Chap. 4. Complicating matters, there is often a high degree of

information asymmetry at play, and users cannot see into the future to look at the actual end consequences of their use of particular technologies. Instead, they must assess risk and uncertainty and face ambiguity in deciding whether or not to adhere to particular norms. This is also where beliefs about roles within the social context, our own expectations and the expectations of others, and practices, as informed by history, culture, law, and social convention, come into play. For example, if we have diminished social power in our contextual roles, we may be more likely to rely on already established social norms rather than relying on our own normative evaluations for information disclosure. Or, if we are apprehensive about sharing data through a new tool because we aren't entirely aware of what the third-party data flows look like, being in a social context where regulatory frameworks punish companies who misuse user data can help give us that added bit of trust.

As technologies diffuse, social practices involving those technologies are negotiated, stabilized, and become more obdurate in nature. Individual choices about adoption and use scale within social groups and the meaning and rituals associated with a given technology either fall by the wayside or become adopted [28]. Once the meanings and uses of a given technology are socially agreed to through adoption practices, the meaning of the technological artifact becomes stabilized within the social setting, and normatively acceptable practices emerge.

While the arrow of time has seemed to have resulted in more relaxed privacy norms, so much so that some have called it the “death of privacy” (see, e.g., [29–31]), users can and will push back against technologies that fall too far outside of negotiated norms. For example, introduced in late 2007, Facebook's Beacon program tracked the purchases of Facebook users on third-party websites and subsequently broadcasted messages about those purchases to those users' friends. However, many users were unhappy with this development, finding it “an intrusive form of advertising that took online surveillance and targeted marketing too far” ([32], p. 12). Soon, users created a petition protesting the new feature, and soon after, amid public outcry, Facebook pulled the plug on the program.

When normative violations of information collection and information flows occur, different social actors with the social context will attempt to respond. Individuals will make choices about whether or not to use a particular technology or otherwise augment their behavior, social groups may express unease in the discursive field, policy-makers may consider passing new laws or regulations, and designers will consider whether or not to change the technology to more closely align with the perceived norm.

Key Takeaways

- Individuals weigh the value of benefits of challenges to existing privacy norms against the perceived benefits, but this process is often less than fully rational, and individuals are often having to make decisions with incomplete mental models.
- Our social position, expectations and the expectations of others, and practices as informed by history, culture, law, and social convention, all play a role in whether or not we adhere to existing norms.

- If a technology is too far out of joint with social norms around privacy, different social actors will “push back” in a myriad of ways, including non-adoption, augmented use, complaint, or regulation.

5.6 Conclusion

Privacy is an ever-evolving, contextually shaped phenomenon. Privacy norms are constantly evaluated and reevaluated in light of the new situations enabled by novel technology. Privacy norms are also not universal. This chapter has grounded much of its narrative in talking about privacy norms as they have existed and changed in the United States, which is a limitation. The stories and histories of evolving privacy norms vary both internationally, but also at the more microlevel.

It is important for designers to understand the multitude of situational, social dynamics at play, and not to consider their own experience with local privacy norms as universal. Hubs of innovation in Silicon Valley, for example, often considerably lack diversity [33], which can create a challenge for appreciating how different groups prioritize privacy and the need to protect themselves from certain kinds of information gathering and information flows. Privacy harms are not evenly distributed across the population. For example, the use of real-name policies on social media platforms can create the potential for privacy harms for transgender and gender-variant users, drag queens, Native Americans, abuse survivors, and others [34]. Designers must carefully reflect on who challenges to the current social norms may benefit and who they may hurt. Careful technology design must consider the ways in which challenging existing privacy norms carries with it ethical implications. For more on the ethical implications of privacy work, see Chap. 17.

Working with privacy advocates can help developers think through the challenges that a novel technology may raise. Collaborations between users, academics, nonprofits, and industry can further the responsible development of tech in ways that maximize benefit while minimizing potential harms. Doing so can help designers avoid “creepy” technology [35] and actions out of joint with the social context for which they are creating tools.

Key Takeaways

- Designers must be careful not to take their own privacy norms as universal and to consider the different social contexts in which a technology will be deployed.
- The impacts of changes to privacy norms are not evenly distributed and frequently present outsized risks to disenfranchised groups. Thus, ethical evaluation should be considered in the design process.
- Working with privacy advocates can benefit technology development.

References

1. Nissenbaum, H. 2004. Privacy as contextual integrity. *Washington Law Review* 79 (1): 119–157.
2. Smith, R.E. 2004. Ben Franklin's web site: Privacy and curiosity from plymouth rock to the internet. *Privacy Journal*.
3. Singer, N. 2010, December 11. Online privacy races against technology. *The New York Times*. Retrieved from <https://www.nytimes.com/2010/12/12/business/12stream.html>
4. Brandeis, L., and S. Warren. 1890. The right to privacy. *Harvard Law Review* 4 (5): 193–220.
5. Gellman, R. 2014. Fair information practices: A basic history. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.2415020>
6. Westin, A.F. 1970. *Privacy and freedom*. New York: Atheneum.
7. Westin, A.F., and M.A. Baker. 1972. *Databanks in a Free Society: Computers, Record Keeping and Privacy*. Times Books.
8. Solove, D.J. 2010. *Understanding Privacy*. Harvard University Press.
9. Cate, F.H. (2006). The failure of fair information practice principles. *Consumer Protection in the Age of the Information Economy*.
10. Cagle, L.E. 2019. Surveilling strangers: The disciplinary biopower of digital genre assemblages. *Computers and Composition* 52: 67–78. <https://doi.org/10.1016/j.compcom.2019.01.006>.
11. Isaak, J., and M.J. Hanna. 2018. User data privacy: Facebook, Cambridge analytica, and privacy protection. *Computer* 51 (8): 56–59.
12. Day, M., G. Turner, and N. Drozdiak. 2019. *Is Anyone Listening to You on Alexa? A Global Team Reviews Audio*. Retrieved May 10, 2019, from Bloomberg website: <https://www.bloomberg.com/news/articles/2019-04-10/is-anyone-listening-to-you-on-alexa-a-global-team-reviews-audio>
13. Zimmer, M. 2008. The externalities of Search 2.0: The emerging privacy threats when the drive for the perfect search engine meets Web 2.0. *First Monday* 13: 3. Retrieved from <http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/2136/1944>
14. Litt, E., and E. Hargittai. 2016. The imagined audience on social network sites. *Social Media + Society* 2 (1): 2056305116633482.
15. Marwick, A.E., and D. Boyd. 2010. I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society* 13 (1): 114–133. <https://doi.org/10.1177/1461444810365313>.
16. Rogers, E.M. 2003. *Diffusion of Innovations*. 5th ed. New York, NY: Free Press.
17. McDonald, R.I., and C.S. Crandall. 2015. Social norms and social influence. *Current Opinion in Behavioral Sciences* 3: 147–151.
18. Gibson, J.J. 1977. The concept of affordances. *Perceiving, Acting, and Knowing*: 67–82.
19. Norman, D.A. 1988. *The Psychology of Everyday Things*. New York: Basic Books.
20. Gaver, W.W. 1991. Technology affordances. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 79–84.
21. Boyd, D. 2010. Social network sites as networked publics: Affordances, dynamics, and implications. In *A Networked Self*, pp. 47–66. Routledge.
22. Norman, D.A. 1983. Some observations on mental models. *Mental Models* 7 (112): 7–14.
23. Proferes, N. 2017. Information flow solipsism in an exploratory study of beliefs about Twitter. *Social Media + Society*, 3, 1. <https://doi.org/10.1177/2056305117698493>
24. Eslami, M., A. Rickman, K. Vaccaro, A. Aleyasen, A. Vuong, K. Karahalios, . . . C. Sandvig. 2015. “I always assumed that I wasn't really that close to [her]”: Reasoning about invisible algorithms in the news feed. In *Proceedings of the 33rd Annual SIGCHI Conference on Human Factors in Computing Systems*, 153–162.
25. Hoffmann, A.L., N. Proferes, and M. Zimmer. 2018. “Making the world more open and connected”: Mark Zuckerberg and the discursive construction of Facebook and its users. *New Media & Society* 20 (1): 199–218.

26. Shulevitz, J. 2018, November. Alexa, should we trust you? *The Atlantic*. Retrieved from <https://www.theatlantic.com/magazine/archive/2018/11/alexa-how-will-you-change-us/570844/>.
27. Acquisti, A., and J. Grossklags. 2007. What can behavioral economics teach us about privacy. *Digital Privacy: Theory, Technologies and Practices* 18: 363–377.
28. Bijker, W.E., Hughes, T., & Pinch, T. (1987). *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology*.
29. Andrews, L. 2012. *I Know Who You Are and I Saw What You Did: Social Networks and the Death of Privacy*. Simon and Schuster.
30. Froomkin, A.M. 2000. The death of privacy. *Stanford Law Review* 52 (5): 1461–1543.
31. Garfinkel, S. 2000. *Database Nation: The Death of Privacy in the 21st Century*. Sebastopol, CA: O'Reilly.
32. Cohen, N.S. 2008. The valorization of surveillance: Towards a political economy of Facebook. *Democratic Communiqué* 22 (1): 5–22.
33. Wong, J.C. 2017. Segregated Valley: The ugly truth about Google and diversity in tech. In *The Guardian*. <https://www.theguardian.com/technology/2017/aug/07/silicon-valley-google-diversity-black-women-workers>.
34. Haimson, O.L., and A.L. Hoffmann. 2016. Constructing and enforcing “authentic” identity online: Facebook, real names, and non-normative identities. *First Monday* 21 (6).
35. Tene, O., and J. Polonetsky. 2013. A theory of creepy: Technology, privacy and shifting social norms. *Yale JL & Technology*. 16: 59.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.