Stony Brook University

Hermeneutic Privacy: On Identity, Agency, and Information

A Dissertation Presented

by

Daniel Susser

to

The Graduate School

in Partial Fulfillment of the

Requirements

for the Degree of

Doctor of Philosophy

in

Philosophy

Stony Brook University

May 2015

Stony Brook University

The Graduate School

Daniel Susser

We, the dissertation committee for the above candidate for the

Doctor of Philosophy degree, hereby recommend

acceptance of this dissertation.

Eduardo Mendieta – Dissertation Advisor
Professor, Department of Philosophy

Don Ihde – Chairperson of Defense
Distinguished Professor Emeritus, Department of Philosophy

Robert Crease
Professor, Department of Philosophy

Serene Khader
Associate Professor, Department of Philosophy, Brooklyn College

Victoria Hesford
Associate Professor, Department of Cultural Analysis and Theory

This dissertation is accepted by the Graduate School.

Charles Taber
Dean of the Graduate School

Abstract of the Dissertation

Hermeneutic Privacy: On Identity, Agency, and Information

by

Daniel Susser

Doctor of Philosophy

in

Philosophy

Stony Brook University

2015

The dominant approach in privacy theory defines information privacy as individual control over personal information. Against this view, I argue that the idea of controlling personal information is both incoherent and impracticable. That is because personal information is indistinguishable from non-personal information, and information (of any kind) is nearly impossible to control. Instead of understanding information privacy exclusively in terms of information control, I argue that we ought to think more broadly about the ways people use information to shape how others perceive and understand who they are—what I call social self-authorship. In addition to trying to control which particular pieces of information about us other people have, we work to contextualize and guide the interpretation of that information. I argue that our capacity to do that is central to our ability to draw interpersonal boundaries, and that our ability to draw such boundaries is a necessary condition for social and political agency. In order to protect information privacy in the Information Age, we therefore have to respect what I call norms of hermeneutic privacy. I articulate those norms, and I discuss how they might be realized in technology design, technology education, and technology law.

For Auntie Brenda,

and for Tim.

Contents

Introduction 1

1. Out of Control 15

    1.1. Control Theories of Privacy 18
    1.2. On the Very Idea of Personal Information 20
    1.3. Our Global Information Society 26
    1.4. Privacy in the First Place 32
    1.5. Information Privacy Without Control 40

2. Between You and Me 45

    2.1. Social Selves (and How We Author Them) 49
    2.2. On Facebook-Work 61

3. Acting with Others 73

    3.1. Social and Political Agency 78
    3.2. The Stakes of Self-Authorship 85
    3.3. Identity and Algorithms 95

4. Hermeneutic Privacy 106

    4.1. From Consent to Due Representation 113
    4.2. Hermeneutic Privacy by Design 124
    4.3. Privacy and Technology Literacy 132
    4.4. Information Privacy Law 142

Conclusion 153

References 162

Introduction

Forget Forgetting

> "They say that removing something from the internet is about as easy as removing urine from a swimming pool, and that's pretty much the story."
>
> - Alex Kozinski[1], "The Dead Past"

California passed a law in 2013 that came to be called the "Eraser Law." It requires, amongst other things, that websites and online service providers remove content posted online by California minors, should they request it.[2] The main sponsor of the law, State Senator Darrell Steinberg, argued that future college admissions officers and potential employers will be able to dig things up about applicants that they did and publicized when they were too young to know better. We don't usually hold people responsible in adulthood for mistakes they made when they were children, Steinberg argues, and we ought not to let the internet undermine that norm (Martino 2013; Southwell 2013). "The thinking, say supporters of the new 'eraser' law, is that boys will be boys (and girls, well, girls) and that the indiscretions of youth shouldn't haunt them down the road" (Alexander and Ho 2013).

The basic idea behind the California law is not new. In 1995 the European Union enshrined in EU law a set of principles governing the collection, storage, and use of personal

---

[1] Judge, United States Court of Appeals for the Ninth Circuit.

[2] California Senate Bill No. 568, "An act to add Chapter 22.1 (commencing with Section 22580) to Division 8 of the Business and Professions Code, relating to the Internet." http://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=201320140SB568

information, which were first articulated two decades prior.[3] And in 2014, the European Court of Justice ruled that those principles guaranteed each citizen the "right to be forgotten" (Streitfeld 2014). In today's world, in what has come to be called the Information Age, to say that we have a right to have certain things about us "forgotten" is to say that we have a right to have the digital records of those things erased or deleted.[4] Much like the California law, the right to be forgotten thus obligates websites, internet service providers (ISPs), and even search engines to remove information and references to information which there is "no legitimate reason for keeping," whenever the person or persons identified by the information requests it.[5]

Each of these laws—the California law and the European Union law—represents an effort on the part of legislators to protect people's privacy. Specifically, they aim to protect what philosophers and legal scholars call *information privacy,* which is the privacy we expect around information about us. Like other kinds of privacy, such as the privacy of our homes, bodily privacy, and privacy around personal decisions, information privacy is believed by many to be a basic, necessary feature of democratic society, and it has been recognized as such in legal discussions for more than a hundred years.[6] Today, however, with so much information about us being generated, collected, and stored, and with little transparency about who has access to that information or how they use it, information privacy is considered to be both extremely

---

[3] They were articulated first in 1973, by the US Department of Health, Education, and Welfare (HEW) in a report titled, *Records, Computers, and the Rights of Citizens*, and were subsequently adopted by the Organization for Economic Cooperation and Development (OECD) in its 1980 report, *Guidelines on the Protection of Privacy and Transborder Flows of Personal Data,* which became the basis for international governmental and inter-governmental data privacy regulations. See chapter 4 for a more detailed discussion of these principles and their history.

[4] See Mayer-Schönberger (2011).

[5] See the European Commission "Factsheet on the 'Right to be Forgotten' Ruling (C-131/12)," http://ec.europa.eu/justice/data-protection/files/factsheets/factsheet_data_protection_en.pdf

[6] See chapter 1.

vulnerable and more important than ever. A great deal of attention is therefore starting to turn toward privacy theory, toward figuring out what we need to do exactly to protect information privacy in the Information Age.

In order to protect information privacy we first have to spell out what it means, what it demands, and what practices and policies help to achieve it. This, of course, is a matter of considerable debate. "Eraser" laws and the right to be forgotten typify the dominant position, which says that information privacy means having *control* over information about ourselves—being able to determine who has access to which pieces of information about us and what they are allowed to do with it. On this view, respecting our right to information privacy means obtaining our consent before collecting and using information about us. And if we give that consent initially, and then we decide down the road that we want to withdraw it, "eraser" laws and the right to be forgotten demand that we be allowed to exert control over information about ourselves by compelling those who have it to delete it.

These laws will almost certainly fail to deliver on their promises. The California law is so riddled with exceptions it isn't even clear how it is meant to succeed in the first place. It only covers content posted by the individual petitioner, so if someone posts incriminating photos of their friends on Facebook, the friends—the people who are identified in the photos—have no right of erasure. There is no protection for content which depicts illegal activity. And it only covers the platform on which the content was originally posted. If the information is copied or archived elsewhere online, the law won't touch it (Ferenstein 2013). By contrast, the EU ruling is far more comprehensive. It includes content which was originally posted by the petitioner and then reposted elsewhere by a third party, and it allows not only for the removal of content, but

also for the removal of *references* to the content, such as in search engine results (Rosen 2012). Despite its broad reach, however, the EU directive isn't going to work either. That is because, as the epigraph above so colorfully illustrates, destroying information is extremely difficult.

Consider the case of Caitlin Seida, who describes in a 2013 *Salon* essay how a Halloween picture of her dressed as Lara Croft: Tomb Raider, which she posted on Facebook, was taken, without her knowledge, captioned "Lara Croft: Fridge Raider" and turned into a viral internet meme. Seida suffers from polycystic ovarian syndrome and a failing thyroid gland, and is therefore, as she puts it "larger than someone my height should be." The article is mostly about Seida's confrontation with internet fat shaming, but the way she describes what happened to her personal information—the picture—is instructive.

By the time she discovered what had been done with it, the picture was already all over the internet, having "metastasized through reposts on Twitter, Tumblr, Reddit, 9Gag, FailBlog." Nevertheless, Seida set about issuing copyright violation notices to each website and service where it appeared. "I would have to issue hundreds of them," she writes, "My work as a paralegal had given me some training in this regard, but it was tedious, like pulling weeds out of the planet's largest garden. I had to seek out each instance of the image and sift around until I could find contact information." While Seida succeeded in having many instances of the picture taken offline, she quickly realized that the task was ultimately hopeless. "I got a fair number of them taken down, but once something like this spreads, it's out there forever. I still go through the less tasteful side of the Internet monthly and issue take-down notices for new instances, but it'll never be completely gone." That is why she finally decided to publicly discuss what

happened to her: talking about it herself was the only way she could meaningfully reassert agency over the situation, "to own it again," as she says, "without shame this time."

The rights exercised in Seida's story are those granted by copyright law, rather than privacy law, but they function in exactly the same way. All of these laws treat information about us as something that we *possess* and over which we have exclusive right of control. In Seida's case, copyright law grants her control over her picture, because she originally published it. In the case of a California teen or a European citizen who wishes to have information about them removed from the internet, the law grants them that right because they have a right to privacy. The normative foundation of the right to control the information is different in the two cases, but the means of asserting the right is the same: to petition individual websites and online services to take down the offending content. What Seida's case illustrates is that even if it can be demonstrated that we are owed a right to be forgotten (or to have our youthful indiscretions erased, or so on), it is not at all clear how that right could be meaningfully enforced in today's world. Defining information privacy in terms of information control consigns us to lives of take-down notice whack-a-mole.

There are competitors to the dominant, control approach. The main alternative is the idea that privacy is not about controlling access to information about us, but rather about a lack of access itself. For these theorists, we have privacy just insofar as others lack access to our bodies, knowledge about us, our personal space, and so on. Philosopher Ruth Gavison argues that "our interest in privacy […] is related to our concern over our accessibility to others: the extent to which we are known to others, the extent to which others have physical access to us, and the

extent to which we are the subject of others' attention" (1980, 423).[7] A close cousin of the access view, which focuses specifically on information, is the notion that privacy is *secrecy*. According to Judge Richard Posner, information privacy means "concealment of information" (1981, 272).[8] More recently, legal theorist Woodrow Hartzog and philosopher Evan Selinger have put forward a kind of updated secrecy view, adapted to the Information Age. In their view, privacy ought to be understood as a form of *obscurity.* "Many contemporary privacy disputes are probably better classified as concerns over losing obscurity," they write, which is "the idea that when information is hard to obtain or understand, it is, to some degree, safe" (2013a).[9]

Though intuitively plausible, defining information privacy in terms of limited access—or more specifically, as secrecy or obscurity—has strange implications. For instance, if privacy is limited access, then one would enjoy the most privacy in the complete absence of others, on a deserted island, perhaps, or in solitary confinement. In other words, privacy on this view is tantamount to seclusion. Yet, we normally think of privacy as a kind of relationship between individuals (and between individuals and governments). As legal theorist Daniel Solove writes, "in a world without others, claiming that one has privacy does not make much sense" (2008, 20). What's more, as the solitary confinement example shows, if privacy is seclusion then it is the kind of thing that can be forced upon us. This too seems wrong. We normally think of privacy as something we choose, not something we're forced to endure.

Furthermore, as philosopher Judith DeCew points out, the ideas of privacy and secrecy might overlap, but they clearly aren't coextensive. "First," she writes, "whatever is secret is

---

[7] Quoted in Solove (2008), p. 20.

[8] Quoted in Solove (2008), p. 21.

[9] See also Hartzog and Selinger (2013).

withheld from others, and it may not always be private. Thus secret treaties or military plans kept from the public are not private transactions or information. Second, privacy does not always imply secrecy, for private information about one's debts or odd behavior may be publicized. Although it is no longer concealed it is no less private" (1997, 48). The same could be said about obscurity. My elementary school yearbook is hard to obtain (and thus obscure, in Hartzog's and Selinger's sense), but one would be hard pressed to argue that the information contained in it is meaningfully private.

More importantly, what seems misguided about the access, secrecy, and obscurity views is that they envision privacy as a fundamentally antisocial value. If privacy is seclusion, then wanting privacy is wanting to be alone. If privacy is secrecy or obscurity, then wanting privacy is wanting to be unknown. For Posner, who makes the case most clearly, privacy is merely a "form of deception."[10] But I don't think that squares with common intuitions. I don't think the desire for privacy is at bottom the desire for secrecy or deception. I think we largely understand privacy to be a distinctly *prosocial* phenomenon, a value we cherish precisely because it regulates healthy social and interpersonal boundaries.

It is often pointed out, for instance, that most Americans willingly post photographs and other information about themselves on social media networks, while at the same time claiming that they desire privacy. This is usually taken as a sign of confusion, or worse, hypocrisy. But despite the paradox, I don't think the people who make such claims are confused or hypocritical. I think they recognize that having information privacy isn't tantamount to keeping secrets, that what having information privacy means is having agency over how others know us. We put

[10] See Solove (2008), footnote 42 on p. 205.

information about ourselves online because we want other people to know things about us. In doing so, however, we aren't relinquishing all say over how that information is interpreted and used. The control approach understands this. It understands that having information privacy is ultimately about agency, not anonymity. Yet, as I've suggested, we can't control information about ourselves. The question, then, is whether or not we can have agency over how others know us without controlling which particular pieces of information about us they have.

The central argument of this dissertation is that we can. I argue that trying to control information about ourselves is just one way we exercise agency over how others perceive and understand us. In addition to concealing and revealing information about ourselves, we work to shape or influence how that information is contextualized and interpreted. Information privacy thus involves more than control over particular pieces of information. It involves the entire process through which we negotiate our public identities—what I call the process of *social self-authorship.* If we want to protect information privacy in the Information Age, I argue, we ought to focus broadly on protecting our capacity to author our social selves.

My argument proceeds as follows. In chapter 1, I make good on the claim gestured at above, that defining information privacy in terms of information control is a dead end. This is true, I argue, for both conceptual and practical reasons. First, the very idea of "personal information" is suspect. It relies implicitly on the distinction between *information about someone* and *information* not *about them,* and I argue that distinction can't meaningfully be drawn. If we can't distinguish between personal information and non-personal information, then we can't specify, on a control theory of privacy, which information ought to be controlled. Second, even if we could somehow salvage the concept of personal information, I argue that effort would be in

vain. As a practical matter, information simply cannot be controlled. Too much of it is generated, collected, and stored, by too many different parties, for too many different reasons. And once information is generated, it is extremely difficult, if not impossible, to destroy.

Instead of understanding information privacy exclusively in terms of individual control over personal information, I argue that we ought to think more broadly about the ways we use information to shape how others perceive and understand who we are—what I call social self-authorship. In chapter 2, I draw from sociology and social psychology (especially the work of Erving Goffman) to develop an account of social self-authorship, and I demonstrate how thinking about information privacy in terms of social self-authorship reveals a different set of privacy problems than thinking about it in terms of control does. Not only has information technology made controlling information exceedingly difficult, it threatens to undermine social self-authorship entirely.

In chapter 3, I point to what's at stake in protecting our capacity for social self-authorship, by examining its relationship to social and political agency. I argue that many important activities we engage in necessarily require the willing cooperation of other actors, and that those actors decide whether or not to cooperate with us in large part based upon how they perceive and understand who we are. Activities like taking out a loan or testifying in court are things we simply can't do on our own. They are intrinsically social endeavors, the success of which hinges on how others decide to treat us. Undermining our capacity for social self-authorship thus undermines our ability to engage successfully in crucial social, political, and economic processes. Furthermore, many of the decisions about how to treat us are today made not by other human actors, but by computers—decisions about how much money to lend us or

how much to charge us for insurance. At the end of chapter 3, I explore what it means to be perceived and understood not by another person, but by an algorithm.

Following the descriptive work of the first three chapters, I turn in chapter 4 to the normative implications of my view. I argue that shifting from a control approach to privacy, which focuses on our relationship to particular pieces of information, to an authorship model, which focuses more broadly on the process of negotiating our public identities, requires that we also shift from worrying about norms of consent to worrying about norms of fairness and due representation. On this view, respecting our information privacy is not about getting our permission to collect information about us; it's about ensuring that the process through which others come to know us is one in which we get to participate. I describe what that demands exactly, and consider how those demands could be actualized through technology design, technology education, and the law.

Contrary to the way many people talk about privacy, I assume in this dissertation that privacy is a *practice*. It is not a state that we sometimes find ourselves in, but rather a set of norms we sometimes abide by. We *give* each other privacy. We give people privacy when we leave them alone. We give people privacy when we let them make decisions about their own lives and their own bodies. We give people privacy when we keep our prying eyes away from their personal affairs. And, I will argue, we give people privacy when we let them influence how we perceive and understand them.

Privacy theorists have argued for decades about whether or not the various kinds of practices I just mentioned are all of one piece. They have tried to discern whether privacy—understood in relation to private spaces, private decisions, private expressions, and private

information—represents a coherent concept and a single value, or if it represents a cluster of concepts and a cluster of values. My position is somewhere in the middle. I think we value privacy for a number of different reasons, but that all of the reasons we value it bear a significant family resemblance. Namely, we value privacy in all of its forms because it allows us to draw boundaries between ourselves and other people. We live in a tightly-packed and raucous world of independence-minded individuals. Privacy is how we keep from falling all over each other.

As many privacy theorists have noted, different communities and different cultures have developed different privacy norms, which demand different privacy practices. In some cultures family and sex life are meant to be kept private, while in others information about them are shared without stigma. In some communities the home is the center of social life and doors are always left open, while in other communities the home is a castle and one is expected to knock before entering. In some places the government can intervene in the intimate affairs of its citizens, while in other places the government is required for the most part to leave consenting adults alone.

What makes all of these conflicting norms *privacy* norms is that they function in each context to define the boundaries between individuals (and between individuals and organizations and governments). The conflict is simply about where and how to draw the various lines. As these examples already make evident, we draw many different kinds of boundaries. Norms having to do with the privacy of one's home and personal spaces produce spatial boundaries. The norms of decisional privacy produce boundaries of influence or power. Information privacy involves norms which function to produce what I call *epistemic boundaries*—boundaries between what we should and shouldn't *know* about each other.

Of course, as I've already suggested, thinking about privacy is not only a descriptive project. In addition to identifying the privacy norms people actually adhere to in various contexts, we can argue about which privacy norms we *ought to* abide by. I argue in chapter 4 that there are a set of norms related to information privacy that are being undermined by information technology, which I call norms of *hermeneutic privacy*. Such norms have to do with our capacity to shape or influence how others interpret information about us, with the way the information about us out in the world becomes meaningful to its possessors.

Like much philosophy, the discussion that follows treats commonplace issues in ways that may, at first, seem strange. But the goal is to lend voice and clarity to intuitions which most of us hopefully share. The ideas in this dissertation aren't new or groundbreaking. They can be found, in bits and pieces, spread out across the sprawling privacy literature. They can be glimpsed in discussions of the social dimensions of privacy, in writing on privacy and reputation, and in work on the relationship between privacy and freedom. What is lacking, and what I hope to offer here, is a single coherent argument about the relationship between identity, agency, and information. To protect privacy in the Information Age, these are the issues we have to get straight.

References

Alexander, Kurtis, and Vivian Ho. 2013. "New Law Lets Teens Delete Digital Skeletons." *SFGate,* September 24. http://www.sfgate.com/news/article/New-law-lets-teens-delete-digital-skeletons-4837309.php

DeCew, Judith Wagner. 1997. *In Pursuit of Privacy: Law Ethics, and the Rise of Technology.* Ithaca, NY: Cornell University Press.

Ferenstein, Gregory. 2013. “On California’s Bizarre Internet Eraser Law For Teenagers.” *TechCrunch,* September 24. http://techcrunch.com/2013/09/24/on-californias-bizarre-internet-eraser-law-for-teenagers/

Gavison, Ruth. 1980. “Privacy and the Limits of the Law.” *The Yale Law Journal* 89 (3): 421-471.

Hartzog, Woodrow, and Evan Selinger. 2013. “Big Data in Small Hands.” *Stanford Law Review Online.* 66:81-88. http://www.stanfordlawreview.org/sites/default/files/online/topics/66_StanLRevOnline_81_HartzogSelinger.pdf

Hartzog, Woodrow, and Evan Selinger. 2013a. “Obscurity: A Better Way to Think About Your Data Than ‘Privacy.’” *The Atlantic,* January 17. http://www.theatlantic.com/technology/archive/2013/01/obscurity-a-better-way-to-think-about-your-data-than-privacy/267283/

Kozinski, Alex. 2012. “The Dead Past.” *Stanford Law Review Online* 64: 117-124. http://www.stanfordlawreview.org/online/privacy-paradox/dead-past

Martino, Paul. 2013. “Inside California’s New Online Privacy Law for Minors.” *Law 360*, October 11. http://www.law360.com/articles/479853/inside-calif-s-new-online-privacy-law-for-minors

Mayer-Shönberger, Viktor. 2011. *Delete: The Virtue of Forgetting in the Digital Age.* Princeton, NJ: Princeton University Press.

Posner, Richard. 1981. *The Economics of Justice.* Cambridge, MA: Harvard University Press.

Powles, Julia. 2014. “What We Can Salvage from 'Right to Be Forgotten' Ruling.” *Wired UK,* May 15. http://www.wired.co.uk/news/archive/2014-05/15/google-vs-spain

Rosen, Jeffrey. 2012. “The Right to Be Forgotten.” *Stanford Law Review Online* 64:88-92. http://www.stanfordlawreview.org/online/privacy-paradox/right-to-be-forgotten

Seida, Caitlin. 2013. “My Embarrassing Picture Went Viral.” *Salon,* October 2. http://www.salon.com/2013/10/02/my_embarrassing_picture_went_viral/

Solove, Daniel. 2008. *Understanding Privacy.* Cambridge, MA: Harvard University Press.

Southwell, Alexander. 2013. “California’s New ‘Digital Eraser’ Evaporates Embarrassment.” *Law Technology News,* November 19. http://www.legaltechnews.com/id=1202628537209

Streitfeld, David. 2014. “European Court Lets Users Erase Records On Web.” *New York Times,* May 13. http://www.nytimes.com/2014/05/14/technology/google-should-erase-web-links-to-some-personal-data-europes-highest-court-says.html

Tamò, Aurelia, and Damian George. 2014. “Oblivion, Erasure and Forgetting in the Digital Age.” *The Journal of Intellectual Property, Information Technology and Electronic Commerce Law* 5 (2): 71-87.

Chapter 1

Out of Control:
On Privacy and Personal Information

> "[I]n a networked age, a reasonable amount of control is not enough; control has to be absolute control. One slip-up or data leakage and whatever was once protected can easily enter into a networked public where it may enter broader databases, be aggregated with other data, and circulate. In a networked world, data is more persistent, replicable, searchable, and scalable than ever before. Trying to achieve perfect control will only lead to frustration."
>
> - danah boyd, "Networked Privacy"

Privacy means many things to many people. It is, as legal scholar and privacy theorist Daniel Solove says, "a sweeping concept, encompassing (among other things) freedom of thought, control over one's body, solitude in one's home, control over personal information, freedom from surveillance, protection of one's reputation, and protection from searches and interrogations" (2008, 1).[11] Because of the huge scope and many facets of privacy, discussions about it are usually restricted to some or other specific domain. People speak of decisional privacy, privacy from government intervention, expressive privacy, and the privacy of one's home. The subject of what follows is *information privacy*, which—as Solove's description above suggests—is often understood to mean control over personal information. This idea has its roots in legal discussions, but quickly made its way into philosophy, sociology, and policymaking debates. It takes other forms of privacy as analogous to information privacy: since decisional privacy has to do with control over one's body and self, and since privacy in one's home is a

[11] See also DeCew (1997). Privacy, she says, is "a broad and multifaceted cluster concept…" (4).

matter of controlling who can enter it and when, then information privacy, it is assumed, must amount to control over who has access to one's personal information.

As I aim to show, this assumption is flawed in two ways. The first problem is conceptual and has to do with the very idea of "personal information." As I will show, the concept of personal information relies implicitly on the distinction between *information about someone* and *information* not *about them*. And I will argue that distinction can't meaningfully be drawn. Information about one person can be *inferred* from information about another, and thus the line between information that's personal and information that isn't turns out to be no line at all. If no information is truly personal information (or if *all* information is, which amounts to the same thing), then the idea of personal information is meaningless. And if that's true, then the idea of *controlling* personal information is surely meaningless too.

The second problem is practical. Even if we could somehow salvage the concept of personal information, the effort, I'll argue, would be in vain. For we have reached the point where information about us simply cannot be controlled. There is too much of it being generated, and it's being collected by too many different parties, for too many different reasons (both good and bad). Moreover, once that information is generated it is too difficult (perhaps impossible) to destroy. Therefore theories which rely on the concept of controlling personal information, even if coherent, turn out to be impracticable.

Furthermore, I argue that such theories tend to understand privacy and privacy norms in exclusively negative terms. They assume, that is, that one has full information privacy absent violations or trespasses against it. They assume that we can control information about ourselves *by default*, and thus the norms we need to develop and justify are those that correctly identify

violations of our control. That, I argue, is only half the picture. While it's true that we need to be able to identify and challenge privacy violations, we must also consider what is required in order to produce information privacy in the first place.

If all of that is right, then we seem to be faced with two possibilities. Either those who claim that "privacy is dead" are right (because so much information about us is now being generated, collected, stored, and used), or information privacy means something other than control over personal information. The constructive side of this project is to make a case for the latter. I argue that information privacy isn't dead; it's just something other than what we thought it was. To get a better grasp on information privacy, I suggest that instead of asking what information privacy *is*, we should ask what information privacy is *for.* Why do we value it? What interests does it protect? What loss, precisely, are those who claim that privacy is dead lamenting? And though I think there are many plausible answers to these questions, I will focus on one. Namely, our interest in being able to shape how others perceive and understand who we are—what I call our capacity for *social self-authorship*. This interest is central to our tacit understanding of information privacy, or so I will argue, as well as to our anxieties about losing it.

Finally, and most importantly, this interest can still be protected, even without being able to control our personal information. Thus information privacy, despite worries to the contrary, is still within our grasp.

1.1. Control Theories of Privacy

The idea that information privacy means having control over personal information is pervasive. It is implicit in the way we talk about information privacy colloquially—for example, when we complain that Facebook has invaded our privacy by selling information about our preferences and interests to third party advertisers. It is implicit in the way information privacy is talked about on TV and in the news. And importantly, it is one of the principal ways that scholars frame theoretical discussions of information privacy.

What I call "control theories of privacy" have their roots in a famous (1890) *Harvard Law Review* article by Samuel Warren and Louis Brandeis, considered by many to be one of the most influential essays in the history of American jurisprudence.[12] In it the authors offer the first systematic defense of a Constitutional right to privacy. All agree, they claim, that each individual has a "right of inviolate personality" (82) from which stems a "right to be let alone" (76). From the latter follows "the right of determining, ordinarily, to what extent [one's] thoughts, sentiments, and emotions shall be communicated to others" (78). Since then-recent inventions, such as photography and sound-recording, threatened that right by allowing others to broadcast one's "thoughts, sentiments, and emotions" to a wide audience, a broader right to privacy needed to be recognized in order to protect it. As philosopher Judith DeCew puts it, Warren and Brandeis "emphasized the invasion of privacy brought about by public dissemination of details relating to a person's private life [… and] thus laid the foundation for a concept of privacy that has come to be known as control over information about oneself" (2013, 4).[13]

---

[12] See DeCew (1997), p. 14.

[13] See also DeCew (1997), chapter 1.

Though the idea originates in legal discussions, control theories of privacy exist in the philosophical literature as well. The theorist most often credited with importing it into philosophy is Alan Westin, who defines privacy in his (1967) book *Privacy and Freedom* as: "the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others" (7). Since then countless authors have taken up Westin's basic premise and incorporated it into their own individual approaches. Charles Fried, for instance, claims that privacy "is not simply an absence of information about us in the minds of others, rather it is the control we have over information about ourselves" (1968, 482).[14] David Meeler writes, "privacy is best understood in terms of protecting information about us from being known by others" (2008, 152). It is implicit in philosopher and legal theorist Anita Allen's claim that we are in the midst of a "Great Privacy Give-Away," because "People are giving away more and more personal data to intimates and strangers […]." (2013, 846-7). DeCew says that "control over information about oneself" is "the classic notion or core of privacy" (1997, 24). And while philosopher of information Luciano Floridi formulates his understanding of information privacy in the somewhat cryptic terminology of "ontological friction in the infosphere," it is clear that at bottom his is a control theory as well.[15]

Making what's at stake in viewing privacy as control over personal information even clearer is the fact that it is the primary way in which privacy activists frame their social and political agendas. The American Civil Liberties Union's (ACLU) Speech, Privacy and

---

[14] Cited in Nissenbaum (2010).

[15] Floridi writes: "Given a certain amount of personal information available in (a region of) the infosphere I, the lower the ontological friction in I, the higher the accessibility of personal information about the agents embedded in I, the smaller the informational gap among them, and the lower the level of informational privacy implementable about each of them. Put simply, informational privacy is a function of the ontological friction in the infosphere" (2006, 187).

Technology Project mission statement declares its dedication to "increasing the control that individuals have over their personal information."[16] Privacy International (PI) states "that everyone's personal information and communications must be carefully safeguarded, regardless of nationality, religion, personal or economic status."[17] The Freedom of Information Protection and Privacy Association (FIPA) advocates for "the right to control or limit the collection, use, and disclosure of one's own personal information."[18] "The overall policy goal in every country," writes Colin Bennett in *The Privacy Advocates*, "has been to provide individuals greater control of the information that is collected, stored, processed, and disseminated about them by public and private organizations" (2008, 6).

American privacy law, debates about privacy in the academy, and political organizations advocating for privacy reform all conceive of information privacy in terms of control over personal information. Obviously, such control theories of privacy merit some scrutiny.

1.2. On the Very Idea of Personal Information

What *is* personal information? As we've seen, this idea is thrown around a lot in discussions about information privacy, but it is rarely examined very closely.[19] Whatever else people mean by the term "personal information," I think we can say two things about it with certainty: (1) it is

---

[16] http://www.aclu.org/technology-and-liberty

[17] https://www.privacyinternational.org/about-us

[18] Cited in Bennett (2008), p18.

[19] I have found two exceptions to the rule. One is a (2007) National Research Council report, which defines personal information as "the set of all data that is associated with a specific individual, e.g., date of birth, gender, address, name of first pet, favorite chocolate, high school of graduation, geographical location at 3:14 p.m. on March 30, 2005, and on and on and on." (63). The other is philosopher Helen Nissenbaum, who writes in the introduction to *Privacy In Context* (2010): "Here and throughout the book, following use practices in the policy community, I use ["personal information"] to mean information about an identifiable person…" (4).

information about some person or persons, and (2) it is the personal information *of* the person or persons whom it is about. Thus information about the distance between Munich and Berlin certainly isn't personal information. And information about Henry and Adam's affair *is* (plausibly), but it isn't *my* personal information; it's theirs.

Now, a lot more needs to be said, since there exists a great deal of information about some person or persons which wouldn't obviously be thought of as personal. That I am an American citizen, a resident of New York City, and a registered Democrat, for example, are all information specifically about me that seem evidently *non*-personal. Fortunately, however, we needn't worry ourselves with such nuances, since the conceptual issues I want to describe arise already with the simple criteria laid out above. Whatever personal information is, if it is always, at the very least, information about some person or persons, and if it is the personal information *of* the person or persons it is about, then it will invariably run up against the following (related) problems. First, information is *generative*: by way of inference, having some information can lead to having other information. If one knows, for instance, that Berlin is a six-hour drive from Munich, and that Jena is half-way between the two, then one has access to the further information that Jena is about three-hours from either of them. Or to take another example, if one knows that Henry and Adam are having an affair and that Henry swears they aren't, then one also knows that Henry is a liar.

Second, having access to a greater *quantity* of information can produce a change in the *quality* of that information. Which is to say, the more one knows the more meaningful what one knows becomes. Certain information—which absent a larger context would not be considered personal—can *become personal* if that context (i.e. more information) is provided. Thus

information can be both personal and non-personal, depending on the presence or absence of other, related information. For example, an abortion clinic's phone records certainly aren't anyone's personal information. Nor is one's phone number. Yet if one had access to an abortion clinic's phone records, and knew that the clinic's policy is to call each woman scheduled for an abortion every day for the three days leading up to the procedure, one could quite easily gain access to extremely personal information—namely, that certain women are scheduled for abortions.

A helpful way of thinking about the point I'm trying to make is in terms of what in phenomenology and hermeneutics is called the *figure-ground* model or the *horizonal* model of meaning, or what analytic philosophers of language call *semantic holism*. The insight behind all of these theories is that the parts of some thing are always defined in some sense by the whole in which they are situated. For phenomenologists, the individual parts of a visual field (the figures) are bounded, and thus made coherent and cognizable by the rest of the field around them (the ground). In hermeneutics, the interpretation of a text or work of art is shaped by the history and culture (the "horizons") in which the interpreter is rooted. For semantic holists, linguistic expressions are only meaningful because they are part of larger networks of meaning—words in sentences, sentences in languages, and so on. All of these positions developed in response to arguments or assumptions that meaning inheres in discrete, atomic data—raw sensory input, individual phonemes, single words—and that the process of understanding entire visual scenes, soundscapes, or texts is simply a matter of breaking them down into their constituent parts. What phenomenology, hermeneutics, and semantic holism teach us in various different ways is that this commonsense idea is too simplistic. The meaning of something is not merely a function of the

meaning of its parts; rather, things are meaningful because they are embedded in meaningful contexts.

What I am arguing is that the same is true of information. One can't determine the full meaning of individual data in isolation; one must always factor in the context, which of course can change. Consequently, it isn't possible to classify particular pieces of information as personal or non-personal *in advance,* because the meaning of that information will change depending on the greater informational contexts in which it is situated. To make this more concrete, consider an example borrowed from danah boyd:

> Curious about what secrets might be hidden in my DNA, I decided to spit in a tube and turn my DNA over to the genetic testing service 23andMe. What came back was fascinating: hints that my ancestors might have origins that differ from the family narrative, and disease probabilities that suggest that family medical stories are either inaccurate or statistically curious. Through this test, I learned information about myself, but I also learned information about members of my family. Furthermore, by choosing to subject my DNA to this testing process, I didn't just reveal data about myself; I gave away data that provides insights into my mother, brother, grandparents, and even children that I don't yet have. I never asked my future grandchildren for permission to offer their data to a scientific database. I made a decision about the privacy of my data that affects numerous people who are implicated but who have no say. And, in doing so, I learned information about them that they may not wish to know, let alone have me know. (2012, 1)

What are we to make of boyd's DNA information? It seems obviously personal and obviously hers. But what about her relatives who were also implicated in its analysis? In certain contexts—namely, those in which it is known that they are related—boyd's DNA information seems to be both her personal information *and her relatives'* personal information. Absent such context, it's only hers. If boyd and some of her relatives have policies with the same insurance company, then the fact of their relation could be incredibly sensitive information. It could reveal, for instance, a

disposition to develop certain conditions or illnesses which could cause their insurance premiums to rise.

To take another example, consider recent revelations that the US National Security Agency (NSA) has for a number of years been systematically collecting telephone metadata from all of the major US telephone carriers.[20] Metadata is information *about* phone calls, such as when they took place, between whom, for how long, and so on. It is *not* information about the specific contents of the calls. In the wake of these revelations government officials have claimed that Americans' privacy rights are not violated by this NSA program because metadata isn't sensitive personal information. If the contents of the calls were collected, the argument goes, it would be another story, but collecting metadata shouldn't raise privacy concerns. As many have pointed out, however, metadata is in many cases far more revealing than the contents of phone calls. As Jane Mayer reported in the *New Yorker*:

> 'The public doesn't understand,' [mathematician and former Sun Microsystems engineer Susan Landau] told me, speaking about so-called metadata. 'It's much more intrusive than content.' She explained that the government can learn immense amounts of proprietary information by studying 'who you call, and who they call. If you can track that, you know exactly what is happening—you don't need the content.'
>
> For example, she said, in the world of business, a pattern of phone calls from key executives can reveal impending corporate takeovers. Personal phone calls can also reveal sensitive medical information: 'You can see a call to a gynecologist, and then a call to an oncologist, and then a call to close family members.' And information from cell-phone towers can reveal the caller's location. Metadata, she pointed out, can be so revelatory about whom reporters talk to in order to get sensitive stories that it can make more traditional tools in leak investigations, like search warrants and subpoenas, look quaint. (June 6, 2006)

[20] See Greenwald (2013).

In other words, having a great deal of information about the context in which something occurs is often sufficient for making reliable inferences about the thing itself. When the ground is visible enough the figure is bound to emerge.[21]

Though perfectly obvious once you think about them, these observations illustrate an important problem with the idea of controlling personal information. If personal information can be inferred from non-personal information, and if non-personal information can become personal information when placed in the right context, then it isn't possible to define in advance what information is personal and what information isn't. And if that isn't possible, then it isn't possible to determine what information, on a control theory of privacy, ought to be controlled. In the DNA example, control theories of privacy would seem to demand that in some contexts boyd's relatives control access to *boyd's* DNA information (which, presumably, would mean controlling access to boyd herself). In the case of phone call metadata, controlling personal information seems to demand controlling access to every conceivable detail about the calls one makes. Conceptually, control theories of privacy seem therefore to be either too vague or too demanding. Either they can't distinguish between information that ought to be controlled and information that ought not to be, or they demand that we control *all* information. Either way, such theories are not particularly useful for guiding our practices and policies.

The central argument of this dissertation is that what we are really interested in when we claim to be interested in controlling our personal information is not, in fact, that information (the

[21] Nissenbaum (1997) makes a similar point: "A single fact about someone takes on a new dimension when it is combined with other facts about the individual, or when it is compared with similar facts about other individuals. Applying ingenuity to one-dimensional bits of information can transform mere 'noise' and statistical data into rich portraits of people. Through the powers of information technology we acquire the capability not only to collect and store vast amounts of information, but to bring order to it, to manipulate it and to draw meaningful inferences from it. By these actions we are able to inject shape and also value into a riot of formless data" (217).

figure), but rather the overall picture that information paints about us (the ground). And I will argue in what follows that while managing the former is not possible, managing the latter is. Before we get to that, however, we need to examine the second problem with control theories of privacy—namely, that they are *impracticable.* Even if we could salvage the concept of personal information by coming up with a reliable way of distinguishing between information that is and is not personal, such efforts would be in vain, for information simply cannot be controlled.

1.3. Our Global Information Society

We generate a lot of information. From the beginning of written history to the year 2003, humans recorded something like five billion gigabytes of information. Today, we produce that much information every ten minutes.[22] Sometimes we mean to do it. I mean to generate information when I send my boyfriend a list of what's in the fridge, so that he knows what to buy at the grocery store. I mean to do it when I calculate my students' grades and when I record how much weight I lifted at the gym. At other times we generate information inadvertently, such as when I purchase a book online and the time, date, sale price, and other information about my purchase is recorded by the retailer website. Whether we mean to or not, we generate a lot of information, and a lot of the information we generate is information about ourselves.

Indeed, as information theorist danah boyd suggests in the epigraph to this chapter, we are fast approaching a point where so much information about us exists that we can in no meaningful sense *control* it. Luciano Floridi calls this new state of affairs a "global information society" (2010, 8). Consider how much information the US government alone collects about its

[22] See Turke (2012).

citizens: information about each birth and each death, about who enters and exits the country, when, and for what reasons, who can drive a car, how much income each person earns, who is married to whom, to which organizations one makes charitable contributions, each person's race and gender, hair and eye color, height and weight, religious and political affiliations. None of this seems consequential in isolation, but as we saw in the previous section it can become more meaningful in the aggregate. And that's just a small slice of the information we give over knowingly and willingly, and which in all likelihood is collected for entirely reasonable and legitimate purposes.

However, according to several National Security Agency whistleblowers, the US government is also collecting an enormous amount of information about its citizens in secret. Phone calls, emails, web searches, credit card purchases, travel information, and anything else that leaves a data trail is liable to be intercepted, without a warrant, and stored.[23] As discussed above, documents released by NSA analyst and whistleblower Edward Snowden in June 2013 show that the agency routinely gathers phone call metadata en masse*,* directly from major phone company computer systems.[24] Other documents reveal that the agency collected as many as three billion "pieces of information" from US computer networks in a single, presumably perfectly ordinary, month.[25]

What's more, one needn't engage in activity online to become the subject of such information collection. According to the ACLU, automatic license plate readers are being used by law enforcement to track where people go in their cars:

---

[23] See Bamford (2012).

[24] See Greenwald (2013).

[25] See Greenwald and MacAskill (2013).

> The devices have been proliferating around the country at worrying speed. Mounted on patrol cars or placed on bridges or overpasses, license plate readers combine high-speed cameras that capture photographs of every passing license plate with software that analyzes those photographs to identify the plate number. […] More and more cameras, longer retention periods, and widespread sharing allow law enforcement agents to assemble the individual puzzle pieces of where we have been over time into a single, high-resolution image of our lives. (2013, 2)

Individual states and municipalities are deploying unmanned aerial drones (UADs) to patrol the skies, recording information about citizens' whereabouts, travel patterns, and more.[26] In New York City, a program developed by the police department, in conjunction with Microsoft, unifies and streamlines the city's thousands of closed-circuit television cameras, license plate readers, 9-1-1 call centers, radiation detectors, and other surveillance systems, into one centralized "command-and-control center" in downtown Manhattan, so that the police are able to track and monitor people as they move around town. Microsoft hopes to sell the system to other cities interested in more efficient surveillance.[27]

Now think about how much information corporations collect. When I visit a website, say Amazon's, and buy a book, my name and address are obviously recorded, as well as my credit card number, the date and time of purchase, and so on. What many people don't realize, however, is that much more is recorded too. In order to predict what other books I might be interested in Amazon likely records the website I visit immediately prior to its own, as well as the one I go on to visit after. They record the books I browse but choose not to buy, the amount of time I spend considering them, the frequency with which I shop on their site, and the type of computer I use to access it. Google, in order to offer more accurate search results, as well as

---

[26] See the ACLU's (2011) *Protecting Privacy from Aerial Surveillance: Recommendations for Government Use of Drone Aircraft*.

[27] See Ungerleider (2012).

more accurately targeted advertising, scans the contents of my emails that I send and receive through its email service. The company stores and indexes every search query I enter, every result I select, the length of time between when I select one result and return to select another, every advertisement I click on, and where I'm located when I search.

Furthermore, governments and corporations work in tandem. The Snowden leaks revealed a remarkable amount of cooperation between the NSA and several of the largest information technology corporations in the world—companies such as Google, Apple, Facebook, Twitter, and Microsoft. The documents Snowden disclosed detail the ways in which these corporations provide the NSA with direct access to the information generated by and about their users. In a report titled "The Surveillance-Industrial Complex," the ACLU explains why this is so important:

> The Privacy Act of 1974, although riddled with exceptions and loopholes, does restrict the ability of law enforcement agencies to maintain dossiers on individuals who are not suspected of involvement in wrongdoing. But the government is increasingly circumventing those restrictions simply by turning to private companies, which are not subject to the law, and buying or compelling the transfer of private data that it could not collect itself. (8)

In other words, since the government (at least in the United States) is barred from conducting surveillance of its own citizens without a warrant, it now simply outsources that work to corporations not subject to the same restrictions.

Of course, one might want to argue that the government and corporations ought not to be collecting all of that information. But even if they were to stop, think about how much information we willingly generate and make available ourselves. Many of us continuously publish extensive information about our whereabouts, the activities we're engaging in, our tastes and preferences, and anything else we might think of *in real time* on social networking platforms

like Facebook and Twitter. We publish photographs of the food we're eating on Instagram and moan about its quality on Yelp.

Not all of the information we produce about ourselves is produced in the name of frivolous self-promotion either. A huge amount of information has to be generated about us in order to provide us with the technology-based services we want. If you want a good email spam filter, for instance, algorithms will need to collect information about who you email, how often, and what sort of thing you usually write about. If you want your phone to tell you the weather it has to collect information about where you are. If you want your computer to remind you about your appointments it has to know where you're supposed to be and what you're supposed to be doing. In other words, behind all of the wonderful technologies many of us have grown accustomed to using are complex algorithms. And those algorithms work by processing information. Unless we are going to radically de-technologize our lives, the amount of information we generate about ourselves is only going to grow.[28]

Making controlling that information more difficult still is the fact that once it is generated information is difficult, if not impossible, to destroy. That is because modern information processing is massively *distributed*. Cloud computing allows for the documents I'm working on at home on my computer to be stored and managed on servers thousands of miles away. Co-location services keep multiple copies of information in different physical locations to protect against hardware malfunction, human error, and natural disasters. Which is to say, Google doesn't keep just one copy of my search history in a place where it could be easily deleted; it

[28] As the Center for Democracy and Technology's Erica Newland put it in a 2012 speech to the DC Superior Court, "To disconnect from all of the services and technologies that collect personal, sensitive data about us would be to disconnect from society. The on-the-ground reality is that to 'opt out' of the data collection, correlation, and/or use that takes place when we go about the activities described above would be analogous to 'opting out' of electricity a mere thirty years ago."

keeps many copies all over the world. Moreover, once information is made accessible online, individuals can download and maintain copies of that information. And if that information is ever made *in*accessible, they can easily republish it. In the words of one technology writer, "The Internet never forgets" (Bradley 2010).

Finally, as discussed in the previous section, even if you somehow abstained from all information generating activities—the internet, credit card use, basic services, you never drove a car, and so on—one could still infer things about you from information they have access to about others. Recall the example of danah boyd's DNA information. Or imagine you are a student in one of my classes. One needn't know anything more about you than that to deduce with a fair degree of certainty your whereabouts from mine. Unless *everyone* abstains from all of the various modes of information production, then, no one can.[29]

With this being the state of things (or its near-term trajectory), it seems clear that there is no meaningful sense in which we can control information about ourselves. Save for one. The one way in which we can control information about ourselves is to prevent unnecessary information from being generated and collected in the first place. That means putting in place policies which limit the information governments and corporations seek out. It means requiring that individuals "opt-in" to providing information about themselves, rather than expecting them to know to "opt-out." It means demanding that corporations devise methods to purge their servers of information after it is no longer necessary for achieving the ends for which it was provided. Doing all of these things is eminently worthwhile, but one must recognize that it would provide for only a

[29] As Daniel Solove puts it, "Life today is fueled by information, and it is virtually impossible to live as an Information Age ghost, leaving no trail or residue" (2004, 8).

very limited form of control. And to quote boyd once again, "[I]n a networked age, a reasonable amount of control is not enough; control has to be absolute control."

Still, it is worth noting that I am *not* suggesting that we shouldn't care about who has information about us or what they do with it. As the chapters that follow make clear, I think such issues are of central importance. What I am arguing is that framing concerns about the dissemination and use of information about individuals in terms of their capacity to control it dooms privacy advocacy to failure. Since distinguishing information about individuals from information not about them is, at best, extremely difficult, and since controlling information of any kind is effectively impossible, information privacy worries need to be conceptualized in a way that doesn't make use of those terms. For many, the loss of control over personal information means a loss of privacy. Prevailing opinion would have it that we have "given up" our privacy, that we've reached the "end of privacy," and so on. The goal of what follows is to show that prevailing opinion is wrong. We have only reached the end of privacy if privacy means control over personal information. Fortunately, it doesn't.

1.4. Privacy in the First Place

A number of theorists have recognized problems with control theories of privacy and have sought to develop alternative approaches. One of the most influential—if influence is measured in terms of impact on policymaking—is philosopher and media theorist Helen Nissenbaum, whose conceptual framework was adopted most recently in President Barack Obama's 2012

proposal for a US "Consumer Privacy Bill of Rights."[30] For Nissenbaum, information privacy isn't about controlling personal information per se, but rather about the contexts in which personal information is provided and the purposes to which it is put. "Many [privacy advocates] argue that protecting privacy means strictly limiting access to personal information or assuring people's right to control information about themselves," Nissenbaum writes, "I disagree. What people care about most is not simply *restricting* the flow of information but ensuring that it flows *appropriately*..." (2010, 2, emphasis in original). Information "flows appropriately" when it is sent and received by the "actors," and according to the "transmission principles," defined by the norms inherent in the context in which the information was originally provided. For example, if I confide deeply personal information to a friend and that friend turns around and blogs about it, then the information I provided has flowed inappropriately. Inherent in the context of speaking to a friend in confidence are implicit norms which dictate that the only person authorized to receive that information is the friend, and thus that its transmission elsewhere is improper. In such cases, when information does not flow appropriately with respect to the context in which it was provided, *contextual integrity* has been violated. Information privacy, Nissenbaum argues, is the maintenance of contextual integrity.

There are many things to like about this approach. First of all, as we have seen, it is crucial to understand information in context. That Nissenbaum puts such a strong emphasis on information contexts is thus already a huge improvement over theories which neglect it. Second, her emphasis on appropriate information flows, rather than appropriate information, makes sense

[30] See the Obama Administration's (2012) *Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy.* http://www.whitehouse.gov/sites/default/files/privacy-final.pdf

of something which I discuss at length in the next chapter—namely, that it is perfectly consistent to be concerned with one's information privacy and still publish a great deal of information about oneself online. That is because, on Nissenbaum's account, the issue isn't the presence or absence of information, but rather where and how it flows.[31] Finally, the concept of violating contextual integrity accurately captures an essential part of the indignation we feel when we think our information privacy has been violated, and Nissenbaum's analysis of what it means to violate contextual integrity is extremely precise. Her theory is therefore invaluable as a heuristic tool for identifying *prima facie* privacy violations.

But despite Nissenbaum's shift of emphasis away from controlling personal information onto protecting contextual integrity, I think her approach is at bottom still about control. Rather than control specific personal information, protecting contextual integrity requires that we control information *flows*. To some extent, Nissenbaum's approach thus skirts the first problem with control theories—the conceptual issues having to do with delineating between personal and non-personal information. I'm not convinced, however, that it gets around the second.

Nissenbaum writes that her theory is "a justificatory framework for establishing whether socio-technical devices, systems, and practices affecting the flow of personal information in society are morally and politically legitimate" (236). "In applying contextual integrity to this question," she says, "I call on it to serve as a decision heuristic, a framework for determining, detecting, or recognizing when a [privacy] violation has occurred" (148). Again, I think Nissenbaum's approach is extremely useful as just this sort of heuristic device. Using contextual integrity to determine if an information privacy violation has occurred, however, raises the

[31] See Nissenbaum (2010), 187.

question of what to do in cases where it has. On this question, Nissenbaum is vague. She writes: "In light of contextual integrity, an evaluation that finds a given system or practice to be morally or politically problematic, in my view, is grounds for resistance and protest, for challenge, and for advocating redesign or even abandonment" (191). Thus it seems that the kinds of cases she envisions are (1) cases where certain information flows are explicitly proposed and evaluated *in advance* of their implementation, (2) cases in which an inappropriate flow of information, once identified, could be easily righted, and (3) cases where the person or persons responsible for an inappropriate information flow could be identified and held liable for violating information privacy norms.[32] In all of these cases it is assumed that information flows can be controlled.

Yet it strikes me that such cases encompass only a tiny fraction of those we might worry about in relation to information privacy. Again, in such cases where we *can* exercise some modicum of control over information it is eminently worthwhile to have a set of principles for determining how to do so legitimately, and Nissenbaum's theory is by far the most nuanced approach to elaborating them. Apart from those cases, however, I think my argument in the previous section about the difficulty of controlling information applies equally to controlling information flows.

Take one of Nissenbaum's own examples: data brokers, or what she calls "omnibus information providers." Companies like ChoicePoint, Acxiom, and First Advantage aggregate information about individuals and groups from myriad sources. They mine public records and buy data from private companies in order to produce detailed dossiers that can be used by businesses to tailor products to specific customers or customer demographics, by law

[32] See Nissenbaum (2010), 231-243.

enforcement to identify and target suspects, or by anyone else willing to pay for it to do with it what they wish (Nissenbaum 2010, 45-9). On Nissenbaum's analysis, such services may be ethically problematic because they likely violate the norms inherent in the contexts in which the information was originally provided. When one provides information about one's health to an insurance company, for instance, it is expected that the company will keep that information in confidence and use it only to set one's insurance rates. For the insurance company to sell that information to a company like ChoicePoint, and for ChoicePoint to sell it to other "actors" to use in whatever way they like is probably a breach of the "transmission principles" implicit in the original information context (204-16).

I think Nissenbaum's analysis is right. Data brokers violate contextual integrity when they buy and sell information intended for a specific audience and for a specific purpose. But what does Nissenbaum's approach recommend as a *solution* to the problem? "One of the most important contributions contextual integrity can make," she says,

> is to debunk the logic once and for all in the claim that information shared with *anyone* (any *one*) is, consequently 'up for grabs' […] What this reasoning fails to recognize is how critical it is to spell out the actual and potential recipients of information. Whether the information is transmitted in raw form or assembled, digested, and mined, it makes an enormous difference whether the recipient of this information is your neighbor, your boss, a potential employer, an insurance company, a law enforcement officer, your spouse, or your business competitor. […] Sensitivity to these differences means resisting, once and for all, the idea of public information; that is, resisting the idea that recipients do not need to be explicitly spelled out. (2010, 216, emphasis in original)

I agree with Nissenbaum that we should resist the idea that once we share information it is "up for grabs." But her solution, it seems, is to make explicit who one's information is intended for and what uses it is intended for—or to infer those intentions from the nature of the relevant

"context-dependent informational norms"—and to *challenge each and every breach of those intentions/norms.* To see why this solution is unsatisfying, consider again the case of data brokers. Imagine I provide information about my health to my insurance company and the insurance company sells that information to ChoicePoint. The theory of contextual integrity tells us (rightly) that the insurance company has violated my information privacy by transmitting information about me that was provided in confidence to an unintended third party. Consequently, I might sue my insurance company. Meanwhile, however, ChoicePoint turns around and sells that information to myriad other parties—pharmaceutical companies, advertisers, and the Centers for Disease Control. The theory of contextual integrity tells us that my information privacy has now been violated again. So, I might sue ChoicePoint. At the same time, the advertisers use my information to produce manipulative advertisements that entice me to buy their drugs. The theory of contextual integrity tells us that my information privacy has been violated *again*, and I might want to sue the advertiser. As we saw in the previous section, so much information is being generated and collected about us by so many different parties for so many different reasons that this process is likely to go on *ad infinitum*. On Nissenbaum's approach, then, having information privacy means constantly working to track breaches of contextual integrity and challenging those breaches in order to control information flows. That work, as I've suggested, is futile.

This problem with Nissenbaum's approach is nevertheless instructive, for it highlights a deeper problem with control theories. Namely, they consider privacy and privacy norms in exclusively negative terms. They assume, that is, that one has full information privacy absent violations or trespasses against it. They assume that we can control information about ourselves

(or information flows) *by default*, and thus the norms we need to develop and justify are those that correctly identify violations of our control. But that is only half the picture. While it's true that we need to be able to identify and challenge privacy violations, we must also consider what is required in order to bring about information privacy in the first place.

Consider the privacy of one's home as an analogy.[33] It doesn't make sense to worry about trespasses against that sort of privacy unless one already has a home and privacy in it. The privacy of one's home thus relies on antecedent social and political structures having to do with private property, material structures that realize them, and so on. (By contrast, the homeless don't have the privilege of worrying about privacy violations of this kind.) The same can be said about information privacy. Social and political (as well as material and technological) structures must exist in order to *produce* information privacy. And only once those structures are in place does it make sense to worry about trespasses against them. Obviously, the structures required for information privacy are rather different from those needed to produce privacy in one's home. The work of the chapters that follow is to explain what such structures might look like.

In order to do so, I want to suggest that we first ought to ask what information privacy is *for*. Why do we so desire information privacy? What do we think we lose when it is violated? We might call this a *functional* or *instrumental* approach to information privacy, one which asks after the function of information privacy with respect to other goods or goals. There are two important advantages to this kind of approach. First, it accommodates the widely accepted view mentioned at the start of this chapter that privacy generally, and information privacy specifically, is something of an "umbrella concept." That is to say, privacy seems to mean many different things,

[33] As I've already suggested, the privacy of one's home and information privacy are not in all respects meaningful analogues. But they are similar in relevant ways in this particular context.

and we seem to value it for many different reasons. A functional approach to theorizing privacy suits such a view, because it allows for the possibility that privacy has any number of different functions. The privacy of one's home might function to create a space for one to live free from government interference (which we value for one set of reasons), while at the same time it also functions to create a space where one can think and express oneself free from the judgmental gazes of one's neighbors (which we value for a different set of reasons). Similarly, information privacy might function, as some have argued, to make intimacy possible (by allowing people to choose who has and who doesn't have access to sensitive information about them). At the same time it might also function, as others have argued, to protect one's "inviolate personality" (as we saw with Warren and Brandeis, above).

Second, a functional approach brings into sharper relief the sort of positive structures that make information privacy possible, which I'm suggesting privacy theorists have largely neglected. That is because if you know what something is meant to do or accomplish it is easier to see what conditions are required for it do so. If, for example, we understood the privacy of one's home simply as the right or ability to control who could enter it and when, we wouldn't understand very much. We might be able to define all kinds of violations against such control, but we wouldn't have any way of explaining what sort of home has privacy in the first place. We wouldn't be able to explain, for instance, why living in an all-glass structure, with no curtains or shades of any kind, would mean living without privacy in one's home. As long as no one was violating one's control over who could enter the all-glass house and when, it would seem as though one's privacy was perfectly intact. On the other hand, if the privacy of one's home were described in functional terms, as I describe it above, then it might be understood that one of its

functions is to create a space where one can think and express oneself free from the judgmental gazes of one's neighbors. In which case, it would be perfectly clear why having privacy in one's home means, at the very least, having curtains to close. This is more or less the state of information privacy theory today. It is assumed that we have information privacy (or that we had it until very recently), and virtually all work on the subject is aimed at identifying violations against it. Few have stopped to ask how we came to have information privacy in the first place, or what the social, political, and technological conditions are for bringing it back. Thinking about information privacy in functional terms will go some way to helping right the course.

1.5. Information Privacy Without Control

So what is information privacy for? There are many plausible answers to this question. As others have argued, information privacy is a condition for the possibility of intimacy, autonomy, and individual liberty. I want to argue that, in addition to these things, we desire information privacy because it is necessary for what I call *social self-authorship*.

Social self-authorship is the process by which individuals negotiate their public identities with others. By public identities I mean our personas, our projected selves, who others think we are. Authoring our public identities is what we are doing when we dress one way rather than another, attenuate our local accents, put bumper stickers on our cars, and post updates on Facebook. It is also what we are doing when we clarify the intentions behind those Facebook posts, publicly dispel rumors about us, explain our behavior, and correct someone's recollection about our pasts. We spend a great deal of time and energy trying to shape the way others perceive us, and for good reason: the way others perceive us has an enormous impact on our ability to act

effectively in society. It plays an important role in how seriously we're taken on the job (or if we get the job in the first place), whether or not our testimony is believed in court, who is physically and emotionally attracted to us, how we're treated by law enforcement, whether or not we're granted tenure, and so on. Everything we do in concert with others—which is to say, nearly everything we do—depends in some part on who those others think we are.

This process is central to intuitions about information privacy. Indeed, I think this is really what the intuitions underlying control theories of privacy are all about. We aren't interested in control over personal information *for its own sake*. We are interested, rather, in controlling the way that information portrays us to others. Of course, we can't entirely control how people perceive us. But we can (and do) participate in the process by which they arrive at those perceptions.

Information is an integral part of this activity. We add, retract, clarify, and complicate information about us and the contexts in which that information is understood. We offer other perspectives. If we detect that we are being deeply misunderstood we try to clarify differences of principle. Moreover, we gauge the reactions of those around us to the way we present ourselves and adjust that presentation accordingly.[34] Thinking about information privacy in terms of social self-authorship thus offers a way of heeding the intuitions motivating our desire for control over personal information, while conceding that control is impossible. It shifts the focus away from defending against violations of our (elusive) control, onto creating the conditions for being able to actively take part in a fundamental social and political activity.

---

[34] Note that this picture does not entail that the public identity or persona being authored is the perfect expression of a fully known transparent self, nor the expression of a self independent of its surroundings and social relations. On the contrary, I describe social self-authorship as a process of identity *negotiation* precisely to highlight its complexity and multi-directionality. The way others perceive and react to one's public identities certainly has an impact on one's self-understanding.

In the chapters that follow, I carefully articulate the relationship between privacy and social self-authorship, the effects of information technology on our capacity for social self-authorship, its social and political value, and the privacy norms its protection demands. This is not an exhaustive account of information privacy—for if I am right, and privacy ought to be conceptualized in functional terms, then a complete account would describe each of its many functions and is well outside the scope of this dissertation. Nevertheless, the function of privacy I deal with here is, I think, the one about which people are most concerned. Even if they don't take us all the way, the theoretical and practical implications of my account therefore move us well along toward the goal of understanding how to have robust privacy in the Information Age.

References

Allen, Anita. 2013. "An Ethical Duty to Protect One's Own Information Privacy?" *Alabama Law Review* 64 (4).

American Civil Liberties Union. 2004. *The Surveillance-Industrial Complex: How the American Government is Conscripting Businesses and Individuals in the Construction of a Surveillance Society.* https://www.aclu.org/national-security/surveillance-industrial-complex

American Civil Liberties Union. 2011. *Protecting Privacy from Aerial Surveillance: Recommendations for Government Use of Drone Aircraft.* http://www.aclu.org/files/assets/protectingprivacyfromaerialsurveillance.pdf

American Civil Liberties Union. 2013. *You are Being Tracked: How License Plate Readers are Being Used to Record Americans' Movements.* https://www.aclu.org/technology-and-liberty/you-are-being-tracked-how-license-plate-readers-are-being-used-record

Bamford, James. 2012. "The NSA Is Building the Country's Biggest Spy Center (Watch What You Say)." *Wired,* March 15. http://www.wired.com/threatlevel/2012/03/ff_nsadatacenter/all/

Bennett, Colin. 2008. *The Privacy Advocates: Resisting the Spread of Surveillance*. Cambridge, MA: MIT Press.

boyd, danah. 2012. “Networked Privacy.” *Surveillance & Society* 10 (3/4).

Bradley, Tony. 2010. “Erasing your Digital Tracks on the Web.” *PC Magazine,* May 2. http://www.pcworld.com/article/195270/xxx.html

DeCew, Judith Wagner. 1997. *In Pursuit of Privacy: Law, Ethics, and the Rise of Information Technology.* Ithaca, NY: Cornell University Press.

DeCew, Judith Wagner. 2000. “Privacy and Information Technology.” In *Privacy and Data Protection: Theory and Practice,* ed. M. J. van den Hoven. Kluwer Academic Publishers.

DeCew, Judith Wagner. 2013. “Privacy.” In *The Stanford Encyclopedia of Philosophy,* ed. Edward N. Zalta. http://plato.stanford.edu/archives/fall2013/entries/privacy/

Floridi, Luciano. 2010. *Information: A Very Short Introduction*. Oxford: Oxford University Press.

Floridi, Luciano. 2006. “The Ontological Interpretation of Information Privacy.” *Ethics and Information Technology* 7 (4): 185-200.

Fried, Charles. 1968. “Privacy.” *Yale Law Journal* 77 (3): 475-493.

Greenwald, Glenn. 2013. “NSA Collecting Phone Records of Millions of Verizon Customers Daily.” *The Guardian,* June 5. http://www.guardian.co.uk/world/2013/jun/06/nsa-phone-records-verizon-court-order

Greenwald, Glenn, and Ewen MacAskill. 2013. “Boundless Informant: The NSA's Secret Tool to Track Global Surveillance Data.” *The Guardian,* June 11. http://www.guardian.co.uk/world/2013/jun/08/nsa-boundless-informant-global-datamining

Mayer, Jane. 2006. “What’s the Matter with Metadata?” *New Yorker* (blog), June 6. http://www.newyorker.com/online/blogs/newsdesk/2013/06/verizon-nsa-metadata-surveillance-problem.html

Meeler, David. 2008. “Is Information All We Need to Protect?” *The Monist* 91 (1): 151-169.

National Research Council. 2007. *Engaging Privacy and Information Technology in a Digital Age*. Washington, DC: The National Academies Press.

Nissenbaum, Helen. 1997. “Toward an Approach to Privacy in Public: Challenges of Information Technology.” *Ethics and Behavior* 7 (3): 207-219.

Nissenbaum, Helen. 2010. *Privacy in Context: Technology, Policy, and the Integrity of Social Life* Stanford: Stanford University Press.

Newland, Erica. 2012. “Disappearing Phone Booths: Privacy in the Digital Age.” *Speech to the DC Superior Court judges*. Washington, DC. https://www.cdt.org/files/pdfs/Privacy-In-Digital-Age.pdf

Solove, Daniel. 2004. *The Digital Person: Technology and Privacy in the Information Age.* New York: New York University Press.

Solove, Daniel. 2008. *Understanding Privacy.* Cambridge, MA: Harvard University Press.

Turek, Dave. 2012. “The Case Against Digital Sprawl.” *Business Week,* May 2. http://www.businessweek.com/articles/2012-05-02/the-case-against-digital-sprawl

Ungerleider, Neal. 2012. “NYPD, Microsoft Launch All-Seeing ‘Domain Awareness System’ With Real-Time CCTV, License Plate Monitoring [Updated].” *Fast Company,* August 8. http://www.fastcompany.com/3000272/nypd-microsoft-launch-all-seeing-domain-awareness-system-real-time-cctv-license-plate-monito

Warren, Samuel, and Louis Brandeis. 1890. “The Right to Privacy.” *Harvard Law Review* 4 (5): 193-220.

Westin, Alan. 1967. *Privacy and Freedom.* New York, NY: Atheneum.

The White House. 2012. *Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy.* http://www.whitehouse.gov/sites/default/files/privacy-final.pdf

Chapter 2

Between You and Me: Epistemic Boundaries and the Work of Social Self-Authorship

---

> "All interpersonal contact goes through the visible surface, even if it penetrates fairly deep, and managing what appears on the surface—both positively and negatively—is the constant work of human life."
>
> - Thomas Nagel, "Concealment and Exposure"

> "[T]he imaginations which people have of one another are the solid facts of society…"
>
> - Charles Horton Cooley, *Human Nature and the Social Order*

Privacy, on the whole, has to do with negotiating boundaries between oneself and others.[35] The privacy of one's home has to do with negotiating physical boundaries. Decisional privacy is about boundaries of influence or power. And information privacy involves negotiating what we might call *epistemic* boundaries: boundaries between what others should and shouldn't *know* about us. Control theories of information privacy understand this, only they misunderstand the nature of the negotiation. They assume that epistemic boundaries are drawn merely by limiting access to information one desires to keep secret, when in fact the process of *revealing* information is equally, if not more, important. As Thomas Nagel remarks in the epigraph above, "managing what appears on the surface—*both positively and negatively*—is the constant work of human life" (2002, 5, emphasis mine).

Furthermore, the meaning of information is not self-evident; it must be *interpreted*. Thus the work of negotiating epistemic boundaries doesn't end with concealment and exposure, but also involves helping to shape the way information about us (or a lack of information) is

[35] See Cohen (1996).

contextualized and understood. Put another way, control theories assume that there is a one-to-one correspondence between information and knowledge, that withholding information means withholding knowledge and accessing information means acquiring it. In fact, information must be interpreted in order for it to lead to knowledge. Drawing epistemic boundaries—determining what people do and don't know about us—is not, therefore, a function of simply concealing and revealing information, but also working to influence how that information is interpreted and understood.[36]

To see why this is so important, imagine, for example, a student who always falls asleep in class. Having watched this happen day in and day out, the instructor might plausibly conclude that the student is simply immature and irresponsible, someone who parties too much to get enough rest. If asked, however, the student could explain to the instructor that she has to work nights and watch her young son in the mornings when her husband goes to work, and only gets to sleep for a few hours in the afternoon before going to class and then to work again. After learning about all that, having re-contextualized and thus re-interpreted the fact that the student always falls asleep in class, the instructor might have a very different understanding of who the student is and why she behaves the way she does. He might decide to cut her some slack, to offer help during office hours and an alternative time to take the exam. And in doing so, the instructor might significantly reshape the options available to that student, and therefore strengthen her capacity to achieve her goals.

---

[36] In *The Unwanted Gaze: The Destruction of Privacy in America* (2000), legal theorist Jeffrey Rosen also warns against confusing information with knowledge. Yet from similar premises he and I arrive at different conclusions. For Rosen, the fact that information taken out of context can mislead others about who we are indicates that we should be extra-vigilant in protecting information about ourselves from getting out. As we will see in what follows, the same fact indicates to me that instead of focusing all of our energies on protecting information about us from getting out, we should be equally concerned about the processes by which the possessors of information about us interpret that information and put it to use.

This example demonstrates two things. First, it shows that having information about someone (e.g., that they fall asleep in class) underdetermines how that person will be perceived and understood. Information must be interpreted. Second, it points to why being able to take part in the process by which others interpret information about us is central to our capacity to act as effective agents in society. I take up the latter issue (the connection between privacy and agency) in the next chapter. My goal in this chapter is to describe how we shape the way others perceive and understand who we are—what I call *social self-authorship*—not only by choosing what to reveal about ourselves and what to conceal, but also by shaping how that information (or its absence) is interpreted. I aim to show that the process of negotiating epistemic boundaries between oneself and others is not merely about protecting secrets, but rather involves the constant work of producing and managing public identities—the work of social self-authorship.

In the second part of the chapter, I examine how information technology affects our ability to do that work. I argue that while it's true that information technology undermines our ability to negotiate epistemic boundaries, it does so for different reasons than control theories suggest. As we saw in chapter 1, it is certainly true that information technology has made it more difficult to control information about ourselves. But being able to control information about us is only one of several tools we use to author our social selves. The deeper and more insidious threats to information privacy wrought by the advent of information technology are that it renders social self-authorship *invisible* and *unnecessary*. In the global information society, we are often unaware that others are forming impressions of who we are, and we are rarely allowed to participate in the processes by which they form them. Thus, in addition to removing one of the

tools we use to negotiate epistemic boundaries, information technology undermines our ability to participate in the negotiation at all.

Now, some will object that worrying about how we are perceived is nothing more than vanity, and that what I am describing is merely a species of self-promotion, "personal branding" or "self-marketing." They will argue that information privacy could not, at bottom, have to do with such shallow interests; that it is, rather, about loftier concerns—intimacy, autonomy, etc. A third aim of this chapter is to show how this objection gets things backwards. Social self-authorship isn't merely self-promotion. On the contrary, self-marketing, personal branding, and so on are merely ways of understanding social self-authorship through a market-oriented, commodifying lens. The work of creating and managing public identities is a basic feature of social life, something we must do on account of the simple fact that other people can't read our minds. Since they can't know our personal identities the way we know them, they have to make judgments about who we are. And we influence those judgments by presenting ourselves in various different ways, sometimes consciously and sometimes unconsciously, often differently to different people.

As we will see in chapters 3 and 4, the way we understand how we negotiate epistemic boundaries has profound implications. It shapes the discourse around information privacy and is reflected and codified in privacy law. As such, correcting our conception of how epistemic boundaries are drawn is more than an exercise in philosophical reflection; it is the first step toward improving information privacy policy.

2.1. Social Selves (and How We Author Them)

"A man's Social Self," writes William James, "is the recognition which he gets from his mates. […] Properly speaking, a man has as many social selves as there are individuals who recognize him and carry an image of him in their mind" (1890, 293-4). In other words, my social selves are the various different ways other people understand who I am.[37] And while it's true *properly speaking* that there are as many such understandings as there are understanders, James is right to point out that "as the individuals who carry the images fall naturally into classes, we may practically say that he has as many different social selves as there are distinct groups of persons about whose opinion he cares" (294). Thus I present myself in one way to my parents and in a somewhat different way to my closest friends. I project a third version of myself to my students and a fourth to my colleagues. Of course, in most cases these projections are not wholly different from one another. They are all variations on a theme. But different aspects of myself are salient to different people, and my ability to create and sustain meaningful relationships relies in part on my ability to emphasize some of them and de-emphasize others.[38] In some cases, it is necessary to conceal things about myself entirely.

This need not be understood as a form of dishonesty or dissimulation ("lying by omission"). A teacher could hardly be accused of dishonesty for neglecting to tell his students about his sex life. Indeed, he would likely be accused of gross impropriety if he *did* share details about such things with them. Likewise, maintaining a healthy relationship with extended family

---

[37] For an excellent history of the concept of the social self, from William James through Charles Horton Cooley, George Herbert Mead, Erving Goffman, and others, see James A. Holstein and Jaber F. Gubrium's (2000) *The Self We Live By: Narrative Identity in a Postmodern World,* especially chapter 2, "Formulating a Social Self."

[38] For a discussion of the relationship between privacy and the ability to create and sustain different kinds of relationships, see Rachels (1975).

members often requires withholding one's political views from them (as we are reminded each Thanksgiving); maintaining a professional relationship with one's boss might mean being more reserved in her presence than one is otherwise disposed; and maintaining a loving relationship with one's children probably demands on occasion that one refrain from expressing how one truly feels about them in that moment. In all of these cases the work that is being done is that of shaping the "images," as James calls them, of ourselves in other people's minds.[39] Obviously, we can't *control* those images entirely. But we try, consciously and unconsciously, to influence them.

Moreover, we shape the images others have of us not only by concealing things from them but also by *revealing* things. We reveal our tastes by dressing a certain way and by decorating our homes. We reveal our political beliefs in conversation and with bumper stickers on our cars. We reveal our strengths when we endure adversity, and our vulnerabilities when we ask for help. We also reveal more specific things—facts about ourselves, about our families, and about our pasts. And again, we do so selectively. I might reveal that I am gay to my friends and family, but depending on where I work I may choose not to reveal that fact about myself to my coworkers. Conversely, I might share my professional goals with my coworkers but not my closest friends.

In all of these examples context is essential. The way others perceive and understand who we are isn't determined simply by the information they have about us. If it were, then you could "get to know" someone just by reading a well-curated dossier about them. On the contrary, we shape the way others perceive and understand who we are by revealing information about

[39] Of course, that is not *all* one is doing in those situations. There are, for instance, any number of reasons one might refrain from telling one's children how they really feel about them in a given moment. To shape the image of oneself in their eyes is one of them.

ourselves in very particular circumstances, and by embedding that information in larger informational contexts. For example, it often matters *who* reveals some information, as anyone who has ever said "I want to be the one who tells them" knows. That is because information learned second-hand can take on a different significance than information heard first-hand. To take another example, the *timing* of a revelation is often a crucial factor in determining how the information revealed is perceived and understood.[40] Many women carefully time the announcement that they are pregnant. Children choose the right moment to show their parents their report cards. Managers decide when to tell employees that they have been promoted or laid off.

What's more, and what control theories of privacy miss entirely, is that revealing information is only the beginning of a process in which that information is contextualized and understood. As I argued in chapter 1, information is only meaningful by virtue of the larger informational contexts in which it is embedded. Thus what someone knows about us can always, given new information, take on a new or altered significance, as it did in the example of the sleeping student, above. The images of us others carry in their minds are not static images; they change constantly in light of new revelations.[41] Shaping those images is therefore not a one-time task, but rather, to quote Nagel again, "the constant work of human life."

---

[40] I am grateful to Serene Khader for this example.

[41] That is not to say that those perceptions aren't *stable*—for the most part they probably are. But they are nonetheless ever susceptible to change, sometimes in large and sometimes in small ways.

Though it has been ignored for the most part by philosophers, the nature and function of this work has long been an object of interest to sociologists and social psychologists, who call it "impression management."[42] Sociologist Barry Schlenker writes:

> *Impression management is the conscious or unconscious attempt to control images that are projected in real or imagined social interactions.* When these images are self-relevant, the behavior is termed *self-presentation*. We attempt to influence how other people—real or imagined—perceive our personality traits, abilities, intentions, behaviors, attitudes, values, physical characteristics, social characteristics, family, friends, job, and possessions. In so doing, we often influence how we see ourselves. […] Impression management is a central part of the very nature of social interaction; it is inconceivable to discuss interpersonal relations without employing the concept. (1980, 6-7, emphasis in original)

The reason impression management is "a central part of the very nature of social interaction" is that, again, we can't read each other's minds. We can't know the totality of someone else's identity. Consequently, the impressions we have of one another are necessarily partial and perspectival. When we interact with other people we know this, intuitively, and so we work (sometimes intentionally, sometimes through force of habit) to form impressions of ourselves that are socially appropriate, relevant, strategically useful, genuine, and so on.

The idea of impression management was first and most influentially elaborated by Erving Goffman, who described it using metaphors of stage acting. We are each, for Goffman, both actor and audience, performing various roles and observing the performances of others. We act one way when we are "front stage" (when we know we are being observed and judged), and another way when we are "back stage" (free from observation and judgment). We act in "teams" when our performances are coordinated, such as when the staff of a restaurant works together to

[42] I explain why I prefer the term "social self-authorship" to "impression management" below.

impress a visiting critic. And we act according to various "scripts"—social conventions which dictate what to say and do in given situations.[43]

"Every person lives in a world of social encounters, involving him either in face-to-face or mediated contact with other participants," Goffman writes in his famous essay "On Face-Work,"

> In each of these contacts, he tends to act out what is sometimes called a *line*—that is, a pattern of verbal and nonverbal acts by which he expresses his view of the situation and through this his evaluation of the participants, especially himself. Regardless of whether a person intends to take a line, he will find that he has done so in effect. The other participants will assume that he has more or less willfully taken a stand, so that if he is to deal with their response to him he must take into consideration the impression they have possibly formed of him. (1967, 5, emphasis in original)

To "take a line" is to adopt a specific understanding of the social situation one is in—an understanding of one's role in it and the roles of others, the social norms which dictate appropriate behavior in that particular situation, the situation's plausible outcomes, etc.—and to act accordingly. For instance, when I go to the gym I understand that my role is that of "participant" or "trainee," and that the coaches at the gym can therefore make certain demands of me. It is entirely appropriate for a coach to yell, "run around the block!" And it is entirely appropriate for me to do so. In this situation we have each taken the correct line—we have acted out our parts. If, by contrast, I walked into class and one of my students yelled at me to run around the block, it would be clear to everyone else that the student had taken the wrong line, had misjudged his role, and had therefore acted according to the wrong set of behavioral norms.

As Goffman makes clear, however, we are not restricted in any situation to merely identifying the correct line or choosing from a set of pre-defined options. Rather, the impressions

[43] For more on Goffman's elaborate stage analogy see his (1959) *The Presentation of Self in Everyday Life.*

we create (and the lines implicit in them) help to *define* the situation. When I walk into class on the first day of the semester and take my place at the front of the room my behavior indicates to the students that I am the instructor. Not only where I stand in the room, but what I say, what I do, and the authority with which I speak, contribute to the impressions they have of me, of how I should act, how they should act, and of the nature of the activity about to take place. If I entered the room in sweatpants, sat in the back, and played on my cell phone, the situation would be defined very differently. "[T]he initial definition of the situation projected by an individual tends to provide a plan for the co-operative activity that follows" (1959, 10), writes Goffman, and "the others, however passive their role may seem to be, will themselves effectively project a definition of the situation by virtue of their response to the individual and by virtue of any lines of action they initiate to him" (9). If the students in my class refused to acknowledge me as the instructor, if they ignored me and continued to talk amongst themselves despite my pleas to pay attention, then my role as the instructor could be threatened and I might be forced to take another line.

Because this back and forth between the various actors helps to articulate the bounds of socially acceptable behavior in a given situation, it takes on a normative quality—what Goffman calls a "moral character":

> Society is organized on the principle that any individual who possesses certain social characteristics has a moral right to expect that others will value and treat him in an appropriate way. Connected with this principle is a second, namely that an individual who implicitly or explicitly signifies that he has certain social characteristics ought in fact to be what he claims he is. In consequence, when an individual projects a definition of the situation and thereby makes an implicit or explicit claim to be a person of a particular kind, he automatically exerts a moral

> demand upon the others, obliging them to value and treat him in the manner that persons of his kind have a right to expect. (1959, 13)[44]

By projecting a certain public identity we establish the various roles each person is meant to play in a particular situation, define what is expected of each other, contextualize each other's behavior and thereby render it intelligible. Shaping the impressions we make on other people is thus a central part of what it means to be a social creature, a creature engaged in cooperative endeavors and guided by the rules that make social order possible.[45] As social beings, we can't help but behave in a way which takes into account how others will perceive and understand us. "To live effectively as human beings," writes Schlenker, "our actions can't be simply random; actions must follow some pattern or plan that establishes who we are, how we see ourselves and desire others to see us, and how we see the world and wish the world to treat us" (1980, 6). Or as Goffman puts it, "[T]he very obligation and profitability of appearing always in a steady moral light, of being a socialized character, forces one to be the sort of person who is practiced in the ways of the stage" (1959, 251).

Of course, the "ways of the stage" to which Goffman refers are largely ways of managing information. First, in order to foster good impressions we have to collect information about the situations we are in. We have to know who we are dealing with, what their goals are, what sorts of power dynamics exist between us, and what expectations they have (both of us and of our interaction). That information makes it possible to calibrate one's self-presentation to fit the

---

[44] While we have every reason to believe that Goffman, writing in the late 1950s, is pointing here toward the "morality" of a kind of inegalitarian social order that most of us today reject, his claims about the normative quality of the impressions we foster don't rely on it. As the classroom example above shows, one needn't be a "person of a particular kind" in order for him or her to be due a certain kind of treatment. One merely needs to occupy a certain role in relation to others. I use the term "normative" instead of "moral" to suggest that the norms invoked in a given situation are not necessarily right or good ones.

[45] "Although some people seem to regard concerns with others' impressions as as sign of vanity, manipulativeness, or insecurity, self-presentation is an essential and unavoidable aspect of everyday interaction" (Leary 1996, 15).

situation.[46] If, for instance, a teenager is in a restaurant with his friends and a group of adults are seated at the table next to them, it could be important for him to notice that one of those adults is his girlfriend's father. Having learned that information he can change the way he is presenting himself (from one oriented toward impressing and entertaining other teenagers to one oriented toward cultivating an air of maturity around adults).

Second, we have to manage the information we make available about ourselves. Goffman writes:

> One over-all objective of any team [of actors] is to sustain the definition of the situation that its performance fosters. This will involve the over-communication of some facts and the under-communication of others. Given the fragility and the required expressive coherence of the reality that is dramatized by a performance, there are usually facts which, if attention is drawn to them during the performance, would discredit, disrupt, or make useless the impression that the performance fosters. These facts may be said to provide 'destructive information.' A basic problem for many performances, then, is that of information control; the audience must not acquire destructive information about the situation that is being defined for them. In other words, a team must be able to keep its secrets and have its secrets kept. (1959, 141)

In some respects, this idea is fairly straightforward. Our ability to produce and sustain impressions of who we are is clearly, to some extent, a function of how well we can control what information about us others have. How people perceive us is, after all, a function of what they know about us. But as Goffman points out, communicating information about ourselves directly to others is not the only way they acquire it. In addition to giving information about ourselves to others, we give information *off*:

> The expressiveness of the individual (and therefore his capacity to give impressions) appears to involve two radically different kinds of sign activity: the expression that he *gives*, and the expression that he *gives off*. The first involves verbal symbols or their substitutes which he uses admittedly and solely to convey

[46] See Goffman (1959), p. 294.

> information that he and the others are known to attach to these symbols. This is communication in the traditional and narrow sense. The second involves a wide range of action that others can treat as symptomatic of the actor, the expectation being that the action was performed for reasons other than the information conveyed in this way. (1959, 2, emphasis in original)

When a job candidate's hand shakes during an interview he gives off information about his being nervous. When a father misses his child's soccer game he gives off information about his priorities. And when someone pronounces "Houston Street" in Manhattan like they would the city in Texas, they give off information about how well they know New York. In all of these cases, the people in question are communicating information about themselves to others, only probably without knowing it.

The distinction between expressions given and given off is important, because it highlights the fact that our ability to control information about ourselves is (and always has been) partial at best. We try, of course. We think about where to go, how to dress, and what to say. But any time we appear before others (in person, or, as we shall see, online) the information we convey to them about ourselves utterly exceeds our capacity to manage it. As psychologist Mark Leary writes, "Virtually every aspect of our behavior provides information from which other people can draw inferences about us. Whenever we are in the presence of other people, they have ready access to a wealth of information from which they can form impressions of our personalities, abilities, attitudes, moods, and so on" (1996, 16).[47] Indeed, even appearing as though one is trying too hard to control one's appearance conveys revealing information.

[47] Or as Goffman puts it, "Whatever an individual does and however he appears, he knowingly and unknowingly makes information available concerning the attributes that might be imputed to him and hence the categories in which he might be placed" (1961, 90). Also see Leary's discussion of "secondary impressions" (1996, 11-13).

This should not be taken to suggest that our attempts to affect the way others perceive us are doomed to fail. Instead, it should remind us that the work of impression management is ultimately about *shaping* or *influencing* others' impressions of us, not controlling them, and that doing so is not merely a performance *for* others, but is rather an ongoing, interpersonal process.[48] That is why I prefer the term "social self-authorship" to "impression management." Management implies unilateral control over relevant variables. (A competent manager is one who keeps everything in order and lets nothing fall through the cracks.) Social self-authorship, on the other hand, involves the ongoing work of building and maintaining relationships in which one feels fully and accurately represented.

Shifting from a management lens to an authorship lens also helps us see that controlling access to information is but one strategy for shaping the way others perceive and understand who we are, but one means for drawing the right epistemic boundaries between oneself and others. As I have been suggesting, the way in which we influence how others *interpret* information about us is equally important. Sociologists refer to this as "packaging" information:

> Effective communication involves packaging information to have a desired impact on an audience. To communicate effectively, one must put oneself in the place of the audience; take into account their perspective, including their competencies, interests, and attitudes; gauge how they are likely to interpret and react to alternative message possibilities; and then edit, package, and transmit the information in a way that leads the audience to draw the desired conclusion. (Schlenker and Pontari 2000, 211)

Psychologists and sociologists have catalogued a whole host of strategies we use to color the way information about us affects how others perceive and understand us. We use different tones of voice to package bits of information, depending on the audience. We use different facial

[48] Despite Goffman and other sociologists' insistence on the term control, he and the rest readily admit that such control is partial at best. See Branaman (1997, lii) and Schlenker (1980, 71-2).

expressions, different gestures, and different phraseologies. We use verbal and gestural "attitude statements" to convey that we are certain *kinds* of people, such as when a teenager tries to appear detached from and "above" the concerns of his family. We make "public attributions" about ourselves and the events that take place in our lives in order to cast them in a particular light to others—when, for instance, we attribute our successes to our own competence and hard work, and our failures to outside forces. "Emotional expressions" (both voluntary and involuntary) tend to influence how others contextualize and interpret information about us, as any student with a sob story about his many-times-dead grandmother knows. And we "associate" or "disassociate" ourselves from others in order to benefit from their social capital or distance ourselves from their lack of it.[49]

Furthermore, as I've said, our public identities change and evolve. As such, beyond working to package the initial disclosure of information, we constantly work to update, modify, correct, and re-interpret it. Anyone who has had an embarrassing fact about them revealed to the public knows that, despite how it might feel in the moments immediately following the disclosure, life goes on. We explain ourselves, put facts about us in greater context. We correct falsehoods and misunderstandings, and we ask others to vouch for us. Sometimes other people have to change before the images of us in their minds can—so, we wait it out. The thing that most enables us to draw healthy epistemic boundaries—to effectively author our social selves—is not control over information about us, but simply that the process remains open-ended, that those who are trying to perceive and understand us are always open to revision and

[49] I put these terms in quotes because they are technical terms in sociology and psychology. For a detailed overview of these and other strategies, see the second chapter, "Tactics," in Leary (1996, 16-38).

reinterpretation, that they never assume that they have understood us in our totality, once and for all.

Control theories of privacy miss all of this. By focusing exclusively on the binary condition of having or lacking information they fail to see all of the many ways that information is made *meaningful* to its possessors. I should note, again, that I am not arguing that having control over information about ourselves is unimportant. Rather, on the picture I am trying to offer, such control is merely one amongst many means by which people author their social selves. It is a *component* of social self-authorship, not its totality. Moreover, this is not a new condition wrought by the development of information technology. One need only think for a moment about the presumably timeless phenomenon of *gossip* to see that we have never had complete (or even particularly robust) control over information about ourselves, and that we have therefore always relied on these other strategies to shape how other people perceive and understand who we are. This point is important because it reveals the way in which control theories of privacy rely on a romanticized notion of pre-technological social reality. Since the development of information technology has made it extremely difficult to control information about ourselves, control theorists assume that prior to its development our control was robust. With this false picture in mind, they conclude that the problem technology poses to privacy is that it erodes our ability to control information.

If my account is right, however, a different picture emerges. If information privacy is the drawing of epistemic boundaries, and if we draw epistemic boundaries not only by controlling information about ourselves, but also through a whole host of other strategies, then the fact that information technology has made it more difficult to control information about ourselves is

merely one of several problems technology might pose to privacy. And as I argued in the previous chapter, it is the problem we are least likely to solve. Control theories thus focus all of our attention on a single, un-winnable battle, and tell us it's the war. To see the other privacy problems we ought to be confronting, we must consider how information technology affects not only our ability to control information about ourselves, but how it affects social self-authorship more broadly. When we do that, we see that the most important problem confronting our capacity to draw epistemic boundaries is not that information technology has made one of our means of doing so more difficult, but that it undermines the process of social self-authorship in its entirety.

2.2. On Facebook-Work

On one level, things remain much the same as they were before the advent of information technology. We still can't read each other's minds. Others' conceptions of who we are thus remain partial and perspectival, and we remain ever trying to shape how they think about and understand us. On this level, all that information technology has changed with respect to negotiating epistemic boundaries is that we now do the work of constructing and maintaining our public identities with new tools and through new media.

Chief amongst them, of course, is the internet. In addition to revealing information about ourselves through face-to-face interaction, we now email and instant message, have personal websites and blogs, and maintain social media presences on Facebook, Twitter, Instagram, and Tumblr. We put out personal ads on dating sites like OkCupid and create professional networks through services like LinkedIn. As I discussed in chapter 1, the enormous amount of information we make available about ourselves on the Internet (and that our friends and acquaintances make

available about us) is difficult, if not impossible, to control. But the argument I have been making in this chapter—that controlling information is but one of many means by which social self-authorship takes place—is as applicable to the parts of our lives we live online as it is to the parts we live offline.

Indeed, we employ many of the same strategies for shaping how information about us is interpreted and understood on- and offline. We reveal different information to different people depending on the nature of our relationships with them, either by maintaining different kinds of social media profiles on different kinds of social media sites, or by maintaining multiple profiles on the same site. Some services, like Facebook, even allow users to determine exactly what information specific other people will see when they visit their profiles. We create email addresses intended to reflect just the right aspects of our personalities to just the right audiences. For example, while most people use a simple combination of their first and last names for the email addresses they use for professional correspondence, many maintain a second email address that uses a nickname or other informal name for personal correspondence. Likewise with usernames or handles on social networks, online dating sites, forums and message boards, and so on. As psychologist Patricia Wallace notes, we put a great deal of attention and care into crafting our online monikers, because they frame our online self-presentations.[50]

Beyond the names and usernames that designate our online personas, we do a great deal of work creating and curating the content that comprises them in order to shape how others perceive who we are. On Facebook people comment approvingly or ironically about politicians or celebrities, and post literary quotations meant to reflect their tastes and interests. On Instagram

[50] See Wallace (1999), especially chapter 2.

people post photos of food, to show how well they cook or how finely they dine. Academics often publish their resumes or curricula vitae, as well as links to their publications, in their faculty profiles on department websites. Indeed, while information technology has made it extremely difficult to control information about ourselves, it has at the same time made it much easier to author our social selves in these other ways. It is far easier now than it was fifty years ago to convey information about ourselves to large groups of friends and acquaintances, to quickly and loudly dispel rumors and respond to personal attacks (before they have a chance to fester and take on lives of their own), to profess solidarity and association with others, to express one's attitude toward public events and institutions, and so on. Concealing information is more difficult in a global information society, but revealing it is easier than ever.

What's more, in the same way that I argued revealing and concealing information are not the only means by which we author our social selves offline, that we also work to shape the way that information is interpreted and understood, so too do we work to shape the way information about us is interpreted and understood online. We carefully choose the angle from which to take an Instagram "selfie" and meticulously select the filter with which to color it. We time our Facebook posts for maximum impact, and clarify and contextualize their meaning in the comments beneath them. In an early episode of the television show *Girls*, the main character Hannah struggles to find exactly the right way to express her fear that she might have contracted a sexually transmitted disease. After going through several iterations she finally tweets, "All adventurous women do"—an intentionally vague formulation meant to express her literary character and cool nonchalance. The scene perfectly captures the work we do to author our social selves online. These new media offer the opportunity to somewhat more carefully weigh how we

present ourselves, to take a little bit more time deciding how to frame a picture or phrase an admission. But none of this is different in any fundamental way from the kind of social self-authorship that we undertake face-to-face. We reveal and conceal information, and work to shape how that information is interpreted; only now we do so with new tools and through new media.

On a deeper level, however, things have changed. To understand what about the process of social self-authorship has been altered fundamentally by the advent of information technology, we need to return to an idea of Goffman's I discussed in the previous section—namely, the distinction between information *given* and information *given off*. Information given is that which we offer up intentionally, while information given off is any other information that can be gleaned or inferred about us apart from what we communicate directly. Just as we give off information about ourselves, oftentimes unknowingly, by speaking or dressing a certain way, we unknowingly give off an enormous amount of information about ourselves online when we visit websites, make purchases with a credit card, read e-books, and so on.

Offline, there is only so much information we can plausibly give off. Unless one is a public figure who is watched and reported about, giving off information about oneself face-to-face requires actually coming face-to-face with the potential recipients of that information. Because we only personally interact with a limited number of people, only a limited number of people have access to the information we give off. And since those people can only notice so many things about us, and as they will remember only a portion of what they notice, there exists a natural limit to the amount of information we give off face-to-face.[51]

[51] Of course, there is an infinite amount of information we could *potentially* give off face-to-face. What I am claiming here, however, is that there is only a limited amount of information that we *actually* give off.

Online, the situation is entirely different. To give off information about oneself online one obviously need not personally interact with anyone. Rather, one simply needs to use information technology or be subject to digital surveillance. As we saw in the previous chapter, nearly everything we do today leaves a trail of data behind us. When we send and receive email we give off information about our personal and professional acquaintances. When we shop online we give off information about our tastes and preferences, our buying habits, and our finances. When we check the weather online we give off information about our whereabouts, and when we research medical issues online we give off information about our health and our bodies.

Moreover, one need not "go online" at all—in the traditional sense of browsing websites—to give off information about oneself online. To be "online" just means that one is connected in some way to the Internet, and we are today connected to the Internet in myriad ways. We give off information about ourselves online when we make purchases with a credit card and they are recorded in a store's online records system. We give off information about ourselves online when we use our smartphone's GPS software to navigate to a destination. We give off information about ourselves online when we drive past an electronic license plate reader and when we walk past surveillance cameras. The breadth and depth of information about ourselves that we give off online is limitless. It can be accessed by a limitless number of people. And it can be stored for a limitless period of time.

As we saw in the previous chapter, everything from what we buy to where we go to who we know is susceptible to tracking. And as these great repositories of information grow, a vast number of government and private sector organizations are putting considerable time, money, and human effort into developing ways of turning this riot of information into usable caricatures

of who we are. As the legal theorist Daniel Solove puts it, we are all the subjects of increasingly complex and precise "digital dossiers":

> Digital technology enables the preservation of the minutia of our everyday comings and goings, of our likes and dislikes, of who we are and what we own. It is ever more possible to create an electronic collage that covers much of a person's life—a life captured in records, a digital person composed in the collective computer networks of the world. (2004, 1)

The metaphor people often turn to for help thinking through this phenomenon is "Big Brother," the all-seeing totalitarian government of George Orwell's *1984.* But as Solove points out, that metaphor is not particularly apt. Big Brother is a centralized power that keeps its subjects in submission by constantly reminding them that "Big Brother is watching," while the wide variety of government and private sector organizations that track us today are just the opposite. They are manifold and diffuse. They have different interests and purposes. And the last thing they want is for us to notice that they are watching. Democratic governments collect information largely for benign, bureaucratic reasons—to monitor the economy, make predictions about the demand for services, regulate potentially dangerous industries, and so on—and businesses collect information mostly to help them better understand how to sell us things. "[B]usinesses don't punish us so long as we keep on buying," writes Solove, "and they don't make us feel as though we are being watched. To the contrary, they try to gather information as inconspicuously as possible. Making us feel threatened would undermine rather than advance the goal of unencumbered information collection" (2004, 7).[52]

[52] Of course, not *all* of the reasons governments and businesses collect and analyze information about us are benign. And Solove doesn't claim that the problems identified by the Big Brother metaphor are absent entirely. His argument is, rather, that there are problems the Big Brother metaphor misses, and that they are in fact the ones about which we ought to be the most concerned.

Instead of looking to Orwell for help, Solove suggests that we turn to Kafka. For the state of affairs we find ourselves in now, with so much information being collected by so many parties for so many different reasons, has far more in common with the world of Kafka's *The Trial* than it does with Orwell's *1984*. The protagonist of *The Trial*, Joseph K., learns one morning that he is being investigated under suspicion of having committed some unspecified crime. And though he tries for the length of the novel to discover what the Court suspects he has done and why, he is never able to see the dossier about him that the judges are considering. At the end of the story, Joseph K. is executed, still clueless about the crimes he was alleged to have commit or the evidence used to make the case against him.

Though obviously exaggerated, Kafka's story is like our present situation in that we don't know what information about us is being collected, who is collecting it, how it is being evaluated, or to what ends. Of course, most of us needn't worry about being suddenly hauled off and executed on account of information collected about us. But a wide variety of much smaller-scale decisions *are* being made about how to treat us, based on information collected without our knowledge. To take a commonly cited example, many online retailers offer their products and services at different prices to different customers, depending on what they know about them.[53] Amazon, for instance, once charged members of its own "Prime" service—which, for $79 per year, offers unlimited two-day shipping on most of its products—more for books and other goods than it did non-Prime customers, on the assumption that anyone willing to spend that much money each year for quicker shipping would also spend more to buy the items themselves. Similarly, the online travel-booking company Orbitz offered more expensive hotel options to

[53] See Valentino-DeVries, et. al. (2012).

customers who accessed its website from Apple computers than it did to customers who accessed it from non-Apple computers, reasoning that someone who buys an upscale computer would want an upscale hotel room too.[54] In neither case were the customers aware that the goods and services being offered to them were tailored to the particular conceptions these businesses had of who they were.

This sort of price discrimination—or what in the industry is euphemistically called "dynamic pricing" or "price customization"—is but one of a hundred ways information we give off online is used to make decisions about how to treat us. Companies frequently use information about how their customers make use of their services to divide customers into "angel" and "demon" groups. "Angel customers" are those most profitable to the company, while "demon customers" are those who cost the company money (for instance, by frequently calling customer service but rarely making large purchases). The best service is then reserved for "angel customers," while "demon customers" are relegated to second-tier service.[55] Advertisers use the search terms people enter into health-focused websites to determine what medical conditions they have, and then use that information to target them with pharmaceutical ads.[56] And government surveillance agencies monitor the kind of language we use in social media posts to determine if we ought to be treated as suspected terrorists.[57]

Of course, none of this is *inherently* problematic. What makes these examples problematic is that, in each of them, the people involved are generally unaware that they are being monitored

---

[54] See Mattioli (2012).

[55] See Solove (2004), p. 50.

[56] See Franzen (2013).

[57] See Stone (2012).

and evaluated. Thus, like Joseph K., we are increasingly subject to decision-making processes we know nothing about, which function according to information we aren't aware is being collected. And this kind of Kafkaesque situation, where we don't know exactly what information we are giving off online, who has access to it, how they evaluate it, or to what ends, undermines social self-authorship in two ways.

The first is epistemic. As we saw in the previous section, shaping how others perceive and understand who we are requires taking into account who *they* are, how we are related to them, what power dynamics exist between us, what they already think or know about us, and so on. Without knowing those things it is difficult to know how we are being perceived or how we should intervene to change it. To play a part well, as Goffman might say, one needs to know one's audience. By obscuring *what* information we give off and to *whom*, information technology thus undermines our ability to effectively shape the way others perceive and understand who we are. If we don't know who has information about us, what information they have, how they evaluate it, or to what ends, we have no way of knowing how to positively influence their conceptions of who we are. Indeed, we may not know that we ought to be trying to do so in the first place.

The second problem is structural. As Solove points out, having access to so much information about individuals makes governments and businesses more reliant on it, and therefore more reliant on processes for coming to understand us in which we aren't able to participate.[58] It is, after all, much cheaper to digitally monitor behavior and make educated guesses about the actors' intentions and desires than it is to solicit each actor's own account.

---

[58] See Solove (2004), p. 49.

In this way, information technology undermines social self-authorship by *obviating the need for it.* It disposes many of the individuals and organizations with whom we interact to rely for their understanding of who we are on information we unknowingly give off, rather than information we provide intentionally. And it makes it difficult for us to help shape the way that information is interpreted and understood.

In sum, information technology undermines social self-authorship by making it both *invisible* and *unnecessary*. Social self-authorship is invisible in the Information Age, because we don't know who is forming opinions about who we are. It is unnecessary because they need not involve us in the processing of forming them.

There are, I think, two important things to notice about this. First, these problems posed by the advent of information technology are much more serious than the problem of diminishing information control. Again, that is not to say that the latter isn't a problem; it is. But whereas losing control over information about ourselves means losing one tool we use to negotiate epistemic boundaries, the problems described above mean an end to the negotiation itself. For if we don't know who is perceiving us and how, and if we have no opportunity to engage in the processes by which they arrive at those perceptions, then the boundaries between ourselves and others are no longer ours to draw.

On the other hand, we should also notice that, unlike the problem of diminishing information control, the problems described above are problems we might actually solve. As I argued in chapter 1, the only way we could regain meaningful control over information about ourselves is by radically de-technologizing our lives. For all of the information technologies we enjoy and rely upon require a mass of information about us to provide us with the services we

want. By contrast, the problems described above—though deeper and more insidious—could be addressed without curtailing technological development. Information systems could be built in such a way that their users are aware of the processes in which they are implicated, and importantly, are able to contribute meaningfully to them.

References

Branaman, Ann. 1997. "Goffman's Social Theory." In *The Goffman Reader,* eds. Charlers Lemert and Ann Branaman. Oxford: Blackwell Publishing Ltd.

Cohen, Jean L. 1996. "Democracy, Difference, and the Right of Privacy." In *Democracy and Difference: Contesting the Boundaries of the Political,* ed. Seyla Behnabib. Princeton, NJ: Princeton University Press.

Cooley, Charles Horton. 1902. *Human Nature and the Social Order.* New York: Scribner.

Franzen, Carl. 2013. "Advertisers Can Learn Your Health Conditions From Your Web Activity, Study Claims." *The Verge,* July 8. http://www.theverge.com/2013/7/8/4505164/medical-data-isnt-anonymous-in-ad-tracking-study-finds

Goffman, Erving. 1959. *The Presentation of Self in Everyday Life.* New York: Anchor Books.

Goffman, Erving. 1961. *Encounters: Two Studies in the Sociology of Interaction*. Penguin University Books.

Goffman, Erving. 1967. *Interaction Ritual: Essays on Face-to-Face Behavior.* New York: Pantheon Books.

Holstein, James, and Jaber Gubrium. 2000. *The Self We Live By: Narrative Identity in a Postmodern World.* Oxford: Oxford University Press.

James, William. 1890. *The Principles of Psychology.* New York: Dover Publications.

Leary, Mark. 1996. *Self-Presentation: Impression Management and Interpersonal Behavior.* Boulder, CO: Westview Press.

Mattioli, Dana. 2012. “On Orbitz, Mac Users Steered to Pricier Hotels.” *The Wall Street Journal,* August 23. http://online.wsj.com/news/articles SB10001424052702304458604577488822667325882

Nagel, Thomas. 2002. *Concealment and Exposure: And Other Essays.* Oxford: Oxford University Press.

Rachels, James. 1975. “Why is Privacy Important?” *Philosophy and Public Affairs* 4 (Summer): 323-333.

Rosen, Jeffrey. 2000. *The Unwanted Gaze: The Destruction of Privacy in America*. New York: Vintage Books.

Schlenker, Barry. 1980. *Impression Management: The Self-Concept, Social Identity, and Interpersonal Relations.* Monterey, California: Brooks/Cole Publishing Company.

Schlenker, Barry, and Beth Pontari. 2000. “The Strategic Control of Information: Impression Management and Self-Presentation in Daily Life.” In *Psychological Perspectives on Self and Identity,* eds. Abraham Tesser, Richard B. Felson, and Jerry M. Suls. Washington, DC: American Psychological Association.

Solove, Daniel. 2004. *The Digital Person: Technology and Privacy in the Information Age*. New York: New York University Press.

Stone, Andrea. 2012. “Homeland Security Manual Lists Government Key Words for Monitoring Social Media, News.” *The Huffington Post,* February 24. http://www.huffingtonpost.com/ 2012/02/24/homeland-security-manual_n_1299908.html

Valentino-DeVries, Jennifer, Jeremy Singer-Vine, and Ashkan Soltani. 2012. “Websites Vary Prices, Deals Based on Users' Information.” *The Wall Street Journal,* December 24. http:// online.wsj.com/news/articles/SB10001424127887323777204578189391813881534

Wallace, Patricia. 1999. *The Psychology of the Internet.* Cambridge: Cambridge University Press.

Chapter 3

Acting with Others:
Agency and Social Self-Authorship

"One might hold that when I am identified, it is my horizon of agency that is identified."

- Linda Martín Alcoff, *Visible Identities*

Our lives are deeply intertwined. Our actions, indeed our very ability to act, are intertwined with the actions and the abilities of others. This fact is everywhere apparent in a big, urban metropolis like New York City. Just getting to work or school involves navigating vast seas of other human beings, each of whom has it in their power to make your journey a little bit easier or a little bit harder. They can gather closer together so you can fit into the subway car. Or not. They can hold the door open for you when your hands are full. Or not. They can jumpstart your car for you when it dies, or hold your place in line for you at the store, or feed your cats for you while you're away for the weekend. Or not.

This is true not only in our personal lives, but in our professional lives too. Many of us, acting in our professional capacities, have it in our power to either help or hinder others as they go about their lives. We are tasked with making decisions about how the companies or organizations that employ us should treat the people they have power over. A professor can cut a struggling student some slack and let them turn an assignment in late. A manager can decide to read a less experienced job candidate's resumé charitably and give them a shot at the job. A judge can have sympathy for someone convicted of a crime and give them a lenient sentence. Or not. The vast majority of decisions like this, everyday, commonplace decisions, but decisions which

can nonetheless have enormous impacts on the people wrapped up in them, are decisions about which there is no one right answer. They are decisions which require judgments.

In today's world, many of these decisions are made by computers. Which is to say, our actions and our ability to act are intertwined not only with other people, and with businesses and governments, and other organizations, but also with algorithms. Computers determine whether or not banks will extend individuals new lines of credit. They decide if insurance companies will offer people policies. They choose who in the airport security line will be singled out for extra scrutiny, and in some cases, who will be allowed to fly at all. Much like the questions of whether or not to cut a student some slack or whether or not to hire the young, inexperienced job candidate, these questions don't have clear-cut answers. They too require judgments. We have merely automated them.

In all of these cases, the judgments being made are based, at least in part, on what the judger thinks or knows about the judged. From the simplest decisions, like whether or not to make room for someone in a crowded subway car, to the most complicated, like whether to give a convicted criminal a lenient sentence or a harsh one, the judgment turns to some extent on the information the decider has about the person being decided about, and on how that information is interpreted and understood. If the person approaching the subway car looks rude or pushy, people are less likely to go out of their way to let them on. If the criminal appears contrite and understanding, the judge or jury might be more inclined to let them off. Figuring out whether or not to let the student turn in the assignment late involves determining if he is struggling or just lazy. Deciding if you should hire the newbie hinges on whether you perceive her as an amateur or a rising star.

If we are to have any agency in these decisions made about us, we therefore have to work to shape or influence the way those making decisions about us perceive and understand who we are. As we saw in the last chapter, we have to conceal some pieces of information and reveal others. We have to time those revelations just right. We have to emphasize certain things about us and de-emphasize others, and we have to work to package information about ourselves so that it has the right impact. But as we also saw, the facts about us can tell more than one story. A lack of job experience can mean the burden of training or the opportunity of mentorship. Poorly written papers are sometimes the product of misplaced priorities and other times the product of a student being in over his head. Being able to exercise agency in the processes by which people (or computers) make important decisions about us requires not only that we be able to reveal some information and conceal other information, but that we are able to help shape or influence how that information is interpreted and understood.

My arguments in the first two chapters were concerned with the *nature* of information privacy. I claimed that privacy, generally, is the drawing of boundaries of various kinds between persons, and that information privacy specifically is the drawing of epistemic boundaries. I argued that in order to draw epistemic boundaries we engage in what I call social self-authorship —not merely controlling information about ourselves, but actively working to shape how others perceive and understand who we are. In this chapter, my concern is the *value* of information privacy. Privacy theorists have proposed a number of possibilities. Some, like philosopher James Rachels, argue that information privacy is valuable because it makes intimacy possible.[59] Others, like Stanley Benn, think we value information privacy because it reflects respect for personhood

[59] See Rachels (1975) and Gerstein (1978).

and autonomy.[60] I don't dispute these views; I think we value information privacy for many different reasons. As I argued in chapter 1, privacy is best understood in functional terms—it serves a number of different purposes. Likewise, we value those different purposes in different ways. That we value privacy for so many different reasons does not suggest to me that privacy is a concept "in disarray," as others have argued[61], but rather that it is something that matters a great deal to us, that it plays a number of different roles in our lives. The more reasons we can marshall to defend it the better.

My goal in this chapter is to put forward the idea that, in addition to the reasons mentioned above, we value information privacy because it strengthens individual social and political agency. The vast majority of what we do, we do in concert with others, and those others can either help or hinder us. Which they choose to do is in large part a function of who they think we are. As such, our perceived identities act as a kind of litmus test for determining what sort of treatment we are owed. As philosopher and social theorist Linda Alcoff writes, our perceived identities are "profoundly significant in determining the state of the 'world' (or worlds) that each individual inhabits: whether they experience that world as hospitable, friendly, judgmental, skeptical, intrusive, or cold" (2006, 90-91). Acting effectively is far easier in a friendly and hospitable world than a skeptical and cold one.

In the first section, I spell out more clearly what I mean by social and political agency. It isn't just that others can make things go more or less smoothly for us; there are certain kinds of activities we engage in that are entirely dependent on the cooperation of other people. To take out

---

[60] See Benn (1984).

[61] See Solove (2008), especially chapter 1.

a loan, or testify in court, or organize people politically, others have to decide to treat us in certain ways. And whether or not they decide to do that is contingent on who they think we are. Our very freedom is thus tied in an important sense to our capacity for social self-authorship.

In section 2, I describe a range of examples of what's at stake in the kinds of decisions I've been referring to. I argue that our ability to access the things we want and need—from material things to people and places to valuable information—hinges on gatekeepers deciding that we ought to be granted that access. Furthermore, our credibility, whether or not people believe what we say, is a function of what else they know or believe about us. Even our ability to struggle for social and political recognition, to secure equal rights and equal opportunities in a democratic society, is, I argue, in part dependent on our capacity to shape how our fellow citizens perceive and understand us. Of course, these examples of the connection between agency and social self-authorship are not exhaustive. My goal is not to chart all of the different ways that our perceived identities and our ability to shape them impact social life. I aim, simply, to demonstrate that the connection exists, and that it pertains to aspects of our lives that are deeply important. For if that is the case, then information privacy is valuable, and worth defending, for that reason alone.

Finally, in the last section, I consider what it means for us that so many of these kinds of decisions are today being relegated to computers. Some have heralded the automation of day-to-day judgments as not only a victory for efficiency, but also a triumph of algorithmic objectivity over the deep prejudices of human deciders. I argue that we ought not to celebrate too quickly. Algorithms are just as prone to bias and partiality as human beings are, it's just buried in code rather than consciousness. Thus to whatever extent privacy demands that we intervene in these

decision-making processes, we must not forget to look to those places where the processes have been digitized too.

3.1. Social and Political Agency

Philosophers largely agree that agency means, in some sense, the capacity to act. Acting, in contrast to doing or reacting, means choosing to behave in one way rather than another—behaving, as it were, *freely*. The paradigm case of an agent is an adult human being. And while other "higher" mammals might also be agents, plants, for example, are not. Though plants do things—absorb water, grow toward sunlight—they can't choose whether or not to do them, so they haven't the capacity to act.

Persons, then, are agents, and things like plants are not. Beyond that the agreement amongst philosophers mostly ends. Important for present purposes, there is little consensus about what it means to behave freely. As Isaiah Berlin argues in his famous (1969) essay "Two Concepts of Liberty," there are two main proposals. On one side, proponents of what Berlin calls "negative" freedom argue that freedom (or liberty—the terms are for our purposes interchangeable) is the absence of external obstacles to action. "By being free in this sense," he writes, "I mean not being interfered with by others" (123). Or to put things the other way around, "If I am prevented by others from doing what I could otherwise do, I am to that degree unfree" (122). Berlin calls this conception negative, because it attributes freedom to the *absence* of something—namely, external obstacles to action—and traces its roots to the classic liberal tradition, from the work of Hobbes and Locke to that of Spencer and Mill.

Proponents of a "positive" conception of freedom argue, by contrast, that freedom and unfreedom are internal matters. The obstacles to freedom aren't other people, but rather one's own irrational desires and dispositions. "The 'positive' sense of the word 'liberty'," Berlin says, "derives from the wish on the part of the individual to be his own master. I wish my life decisions to depend on myself, not on external forces of whatever kind. I wish to be the instrument of my own, not of other men's, acts of will" (131). Thus, to take a standard example, consider a person who wants badly to quit smoking but can't resist the urge. If she gets up, walks to the corner bodega, buys a pack of cigarettes and blissfully smokes them, she is to the proponent of negative freedom a paradigm case of a free agent. To the proponent of positive freedom, however, she is the picture of unfreedom: a slave to her addiction. The dispute between negative and positive conceptions of freedom, as Berlin paints it, is about where the constraints on an individual's capacity to act originate. For negative theorists, such constraints originate from other actors. For positive theorists, they originate from within the actor herself. For negative theorists, freedom is the ability to act without encountering human-imposed barriers. For positive theorists, freedom is self-realization—the capacity to resist distraction and temptation, see through ideology and overcome false-consciousness, and actualize one's deepest potential.

When I refer to "social and political agency" I am pointing to something outside of this Berlinian picture. I am neither talking about the absence of external barriers nor the presence of self-mastery or self-realization. What I mean by social and political agency is *the freedom to engage in social and political practices as a full and equal participant.* Consider, for example, what we mean when we say that someone is "free to take out a loan." In order to evaluate

whether that claim is true or false in a particular case, the proponent of negative freedom would ask questions like: "Is someone preventing them from entering the loan office?" Or, "Is someone barring them from signing the loan contract?" By contrast, the proponent of positive freedom would be interested in whether or not some competing desire or disposition was preventing them from seeking the loan in the first place.

Neither approach captures what we're really interested in when we ask if someone is free to take out a loan, because they both assume that the actions a free person carries out are carried out by them alone.[62] They miss the fact that some actions are necessarily carried out in concert with other people. I say "necessarily," because I'm not talking about actions that are usually carried out by a group of people but could in principle be carried out by someone alone, activities like having a picnic or moving a sofa. Nor am I talking about collective actions, like strikes or political protests.[63] I'm talking about actions individuals undertake, which necessarily involve other parties. To take out a loan, one needs to do more than drive to the bank or fill out a form online. Those are just the steps one takes to initiate the loan-request process. The real activity we're interested in takes place after one makes the loan request—namely, the process of the bank determining whether or not they will grant it.

What we're really interested in when we ask if someone is "free to take out a loan" is whether or not they can request a loan from a bank and then be subject to a fair process according to which the bank evaluates their default risk. That is to say, we're interested in whether or not they will be recognized as a potential loan-recipient—treated as someone who

---

[62] I have no doubt that staunch defenders of either approach could shoehorn my intuitions into their theory. But both approaches seem to exclude what I am trying to point to here from the core of what they understand freedom to mean.

[63] For an overview of the sort of agency involved in collective or shared action, see Bratman (2009).

ought to be considered for a loan—and whether or not they will be given the same consideration as anyone else. Note, the question of whether or not someone is free to take out a loan is *not* about whether or not they are actually granted the loan, since it is perfectly right for a bank to deny a loan to someone who has defaulted on many previous ones. Rather the question is whether or not they have the opportunity to petition for one and are granted a fair hearing.

We would say that someone lacked the freedom to take out a loan if the banks they approached dismissed their applications out of hand. Undocumented immigrants in the United States, for instance, sometimes lack the freedom to take out loans because they don't have Social Security numbers.[64] They can drive to the bank, fill out the form, and so on. And they can overcome whatever internal fears or inhibitions might prevent them from trying. Yet if they do so, the bank will most likely reject their application without even reviewing it. They will not see it as a legitimate application, because it was submitted by someone who is not seen as a potential loan-recipient.

Even if the bank does treat them as a potential loan-recipient, its involvement in their freedom to take out a loan does not end there. For they need more than to be granted access to the loan request process; they also need to be fully and fairly represented in it. Imagine, for example, that my identity was stolen.[65] Someone opened credit cards in my name, charged tens of thousands of dollars to them, and then vanished. All of this would be reflected on my credit report. If I then went and applied for a loan, the bank would likely deny it, pointing to this

---

[64] Some banks will issue loans to undocumented immigrants who lack social security numbers, but have been assigned Individual Tax Identification Numbers (ITINs). Such loans, however, often carry much higher than normal interest rates. See Huseman (2014).

[65] This happened to 16.6 million people in the United States in 2012, resulting in nearly $25 billion in financial losses, according to the most recent U.S. Department of Justice report. See Bureau of Justice Statistics (2013).

history of unpaid credit card bills as evidence that I posed a high risk of defaulting. In a fair process, the bank would give me the opportunity to explain and prove that the information they had about me was false or misleading, that it should not be read to mean that I pose a high default risk, that it could be explained away, that I am, in fact, a safe candidate for a loan. And if the bank refused to do that—if it refused to let me contest its interpretation of the information about me—then we would probably say that I was treated unfairly. We would say that the institutional process used to evaluate loan applications malfunctioned, since it failed to distribute loans according to the real default risks posed by loan applicants.

In this example, then, the threat to my agency (my freedom to take out a loan) has little to do with internal or external barriers to individual action. Rather, it has to do with my ability to engage in a social process as a full and equal participant. Positive and negative conceptions of freedom miss what is at stake in this process, because they fail to understand that for some kinds of activities to take place, the institutions or social contexts in which they are embedded must function properly. Taking out a loan requires more than simply being free from external obstacles, and more than being free from psychological barriers to requesting it. It requires that other actors treat you in a particular way.

To take another example, consider what it means to give testimony, to be "free to testify." Testifying involves more than merely speaking; it requires that one's speech be heard, and in a particular way—namely, as asserting the truth about something concerning which one has first-hand or expert knowledge. In court, for instance, at least in the United States, not everything uttered is taken as testimony, and not everyone present is permitted to testify. Members of a jury

are not permitted to testify during a case they are deciding.[66] And others may be disqualified from testifying if it is determined that they lack first-hand or expert knowledge about the issues under consideration by the court.[67]

Furthermore, even if one is given the opportunity to testify, one's testimony must be heard and considered. If everyone on the jury assumed in advance that the person testifying lacked all credibility, and if they therefore refused to even think about and consider their testimony, then we would only say that the witness was free to testify in a nominal sense. And if, like in the loan example, they were not given an opportunity to contest the jury's perception of their credibility, we would say that they were treated unfairly. Thus again, one's freedom in this case hinges not on internal or external barriers to action, but rather on whether or not one is recognized and treated by others in a particular way.

The above two examples—taking out a loan and giving testimony—involve formal institutional or legal contexts, with clearly defined roles and norms. Finally, consider an example which relies instead on an informal social context and looser roles and norms: community organizing. In order to organize people one must be seen and treated in a particular way. Indeed, much of the work community organizers do involves building trust and solidarity with the people they are trying to organize. That is because getting people to follow and work with you requires that those people see and treat you as a leader—as capable, visionary, and committed to the same political agenda as they are. As soon as they stop seeing and treating you that way, as soon as they stop following and working with you, you cease to be organizing them. Like taking out a

---

[66] *United States Federal Rules of Evidence,* Rule 606.

[67] Ibid., Rules 602 and 702.

loan or giving testimony, acting as a political organizer is thus something one simply can't do on one's own. It requires that one be seen and treated by others in a particular way.

If acting in the kinds of cases I've been describing requires more than self-mastery and a lack of obstacles, if it requires that other agents behave in a certain way too, then one's own agency with regard to those kinds of actions can be undermined by the other agents refusing or failing to act in the way required. If the bank does not treat you as a potential loan-recipient, or if it does not subject you to a fair loan application process, then you lack the freedom to take out a loan. If the jury does not treat you as a legitimate witness, or if it makes unfounded assumptions about your credibility, then you lack the freedom to testify. If a community does not treat you as a leader, then you lack the freedom to organize in it politically. In other words, the set of all possible actions open to us at any given moment is both produced and constrained by the willingness of other people to cooperate with us.

Crucially, part of what determines whether they help or hinder us is how they perceive and understand who we are. It is here that social self-authorship enters the picture. If one's agency in certain circumstances is a function of how one is perceived, then one's ability to shape or influence that perception partially determines the extent of one's agency. My ability to influence how the bank perceives my default risk affects my ability to take out a loan. My ability to influence how a jury perceives my credibility affects my ability to effectively testify, and so on. In what follows, I describe a series of concrete examples of how this works—how our ability to influence the way others perceive us is connected to our social and political agency. I aim to demonstrate not only that the connection exists, but that these examples are not strange or special cases. The kind of agency I am pointing to is a basic feature of everyday life. It involves our

ability to engage in the kinds of social and political processes that everyone needs to engage in to flourish. Daniel Solove writes:

> We depend upon others to engage in transactions with us, to employ us, to befriend us, and to listen to us. Without the cooperation of others in society, we often are unable to do what we want to do. Without the respect of others, our actions and accomplishments can lose their purpose and meaning. Without the appropriate reputation, our speech, though free, may fall on deaf ears. Our freedom, in short, depends in part upon how others in society judge us. (2007, 31)

Social self-authorship is not a minor or peripheral phenomenon; it is central to our ability to act effectively in society, to realize our life plans and achieve our goals, and it deserves careful protection.

3.2. The Stakes of Self-Authorship

Access

The most tangible way our public identities shape our ability to act effectively in society involves our dealing with gatekeepers—people who stand between us and the things we want or need. For whether or not they grant us access to those things is in part a function of how they perceive us. Consider, for instance, access to jobs. Who gets a particular job has, one hopes, something to do with who has the best credentials, most experience, clearest vision, and greatest promise. In most cases, however, which candidate meets those criteria is not self-evident. Information about the various candidates, gleaned from resumés or interviews, is not enough, since information about us can tell more than one story. One candidate might not have worked in the field as long as the others, which could suggest that she would require more on-the-job training. But her experience might be the most relevant to the job, which might suggest otherwise. Another candidate might

have a gap in his employment record, which could suggest to the employer that he isn't reliable. During that period, however, he might have been caring for a sick relative, which could indicate instead that he is dedicated and loyal. Strengthening or diminishing a person's ability to influence how information about them is interpreted in a hiring situation thus strengthens or diminishes their ability to get a job.

To take another example, consider what is required in order to access certain people and certain places. Meeting powerful people often requires that their handlers perceive us as useful or important ourselves. Think of trying to meet a wealthy investor or a member of Congress. Likewise, to visit closely guarded places the guards must perceive us as safe and orderly. Think of going to the Vatican or the Louvre. Sometimes problems arise when we try to access certain people *in* certain places. In some US states, for instance, one might not be allowed to visit their partner in a hospital. People in same-sex couples are routinely barred from visiting their partners, because they are not understood to be "real" family members. Even in states where hospitals are legally required to grant same-sex partners access, hospital staff and administrators are often unfamiliar with the law and thus perceive them to be ineligible.[68] Being able to contribute to the narratives gatekeepers use to understand us thus impacts our ability to access the people we want to see and the places we want to go.

How others perceive us also affects our ability to access information. We only confide in friends if we perceive them to be trustworthy and discreet. Citizens only convey information to the police if they perceive them to be honest and working in their best interest.[69] The US

[68] See Riou (2014).

[69] Or, perhaps, threatening.

government only allows civil servants to see sensitive information if special investigators decide that they are eligible for security clearance. And whether or not they are deemed eligible is contingent upon the investigator's interpretation of facts gathered about the person: are they trustworthy, loyal, discreet, and dependable? Are they sympathetic to people or groups whose interests conflict with U.S. interests? Are they vulnerable or susceptible to being bribed?

What all of these examples show is that the way we are perceived and understood has an enormous impact on our ability to access a wide variety of important goods—from resources such as jobs, to people and places, to valuable information. Since how we are perceived and understood is not determined in advance, but is rather the product of ongoing interpretive negotiations, our ability to influence those negotiations affects whether or not we are granted access to the things we want and need.

Credibility

A second way that our perceived identities are connected with our social and political agency has to do with their impact on our credibility. The extent to which other people believe us is determined in part by how they perceive us, and that they do believe us is enormously important for our ability to act effectively in society. As we saw above, the extent to which others believe we are credible impacts the strength of our testimony in court. It affects how seriously our colleagues and clients take us, and thus how successful we are on the job. It determines whether or not our friends seek out our advice, and when we offer advice it impacts their decisions about whether to follow it. Moreover, because credibility is the kind of thing that one person can

transfer to another, our credibility determines whether or not we can meaningfully vouch for the credibility of others.

Judgments about credibility are based on a number of factors. For someone to make a truly informed decision about how much to believe us, they would have to consider whether or not we have all of the credentials required to make accurate claims about the particular issues at hand, whether or not claims we have made in the past have turned out to be true, and whether or not our claims cohere with their other beliefs. In our day-to-day lives, though, we often don't have time to thoroughly vet each person we meet and research the veracity of their claims. We can't access and review all of their credentials, and thereby make completely informed judgments about their credibility. Instead, we rely on heuristics—most importantly, our perceptions of who they are and the groups to which they belong.[70]

We constantly and often unthinkingly categorize people. My hairdresser is a hipster. The old man who waits outside my apartment building for his grandson to come home from school is retired. The President is black. Rush Limbaugh is an ideologue. Attached to many of these groups or categories of people are attributes: hipsters are associated with superficiality, ideologues with intransigence. Our initial impressions of which groups people belong to thus help to form the overall impressions we have of them, by associating them with the attributes of the groups. Put another way, we judge people in large part by stereotyping them. "Whenever we

---

[70] See Fricker (2007), p. 30-32; Quinn, et. al. (2007), "To simplify the demands of daily interaction, mere exposure of a stimulus to a target is sufficient to stimulate categorical thinking and promote the emergence of its associated judgmental, memorial, and behavioral products (i.e., stereotyped reactions). According to this account, then, categorical thinking is an unavoidable aspect of the person perception process" (69); and Alcoff (1999), "We cannot often directly assess the processes by which the other upon whom we are relying has obtained their knowledge; we cannot know with certainty *how* they obtained their knowledge nor do we necessarily have the expertise to know *what a reliable procedure would be* for obtaining certain kinds of knowledge. Therefore, we must assess the person in a more general way before we can afford them an authority in any epistemic matters" (75-6).

encounter someone and categorize her as a member of a particular group," write social psychologists Kimberly Quinn, C. Neil Macrae, and Galen Bodenhausen, "stereotypes about this group will exert an influence on the interpretive process involved in forming an impression of the person" (2007, 69).

Philosopher Miranda Fricker points out that the attributes or valences different stereotypes carry can be positive, negative, or neutral (2007, 30-31). In terms of credibility, then, being perceived as belonging to this or that group can increase one's credibility, decrease it, or have no effect at all. Doctors, for instance, are generally believed to be intelligent and trustworthy; stereotypes about doctors carry positive credibility associations. If I put on a lab coat and walked around a nearby hospital, giving off the impression that I was a doctor, that perception of me would likely increase my credibility in the eyes of those I met along the way. Stereotypes about criminals, by contrast, indicate that they are untrustworthy, and being perceived as a criminal therefore detracts from one's credibility. If I put on an orange jumpsuit and hung out at my local prison, visitors I came in contact with would likely be less inclined to believe me than if I were wearing plain clothes. Finally, stereotypes about retired people do not generally suggest anything about their credibility. One might assume that a retired person is older and has a lot of time on their hands. But perceptions of retired people in the collective imagination (as Fricker calls it) don't associate them with any particular level of credibility. If I were perceived as retired, the stereotypes associated with that perception would not be likely to impinge on judgments about my credibility at all.

Given the discussions of the previous chapters, about the way the meaning of information can change depending on the contexts in which it is embedded, it should come as no surprise that

the valence a stereotype carries can also vary by context. Fricker gives as an example of this the stereotype that women are intuitive:

> Some stereotypes may resist any definitive categorization because they can carry either a positive or negative valence, depending on the context. The stereotype of women as intuitive is a case in point. In contexts where it is assumed that 'intuitive' suggests irrationality, the stereotype is derogatory; but in contexts where intuition is regarded as a cognitive asset, the stereotype is complimentary. (2007, 31)

An example of the former kind of context might be a philosophy seminar room, in which rationality and argumentation are considered paramount, and feeling and intuition are often thought to be antithetical to the work at hand. A woman subject to the stereotype that women are intuitive would suffer what Fricker calls a *credibility deficit* in such situations—those who hold that stereotype would be less inclined to believe what she says than those who don't. An example of the latter kind of context, where intuition is regarded as a cognitive asset, might be caring for an infant. When a baby cries and nothing is obviously causing it any discomfort, what is usually required to figure out what it needs is not argument and logic, but empathy and intuition. In such situations, if women are perceived to be intuitive, a woman's claims about the baby's needs might be taken more seriously than conflicting claims put forward by a man.

It is also important to note that stereotypes can be more or less reliable. The word "stereotype" has taken on a negative connotation in common parlance and is often used to mean not only a generalization, but a false and demeaning one. While it's true that some stereotypes are false and demeaning, many are not. In the social psychology literature, and for present purposes, stereotype is a generic term meaning a "widely held association between a given social group and one or more attributes" (Fricker 2007, 30). The reason stereotypes are so valuable, cognitively, as a heuristic for making judgments about other people, is that more often than not

they lead to the correct judgment. Fricker is for the most part interested in false and demeaning stereotypes—what she calls *negative identity prejudices*—because her goal is to reveal a form of injustice (epistemic injustice). For our purposes, however, we should be equally interested in reliable stereotypes and unreliable ones, since our goal is to understand the whole range of effects that perceptions about us can have on our credibility, regardless of whether or not they fairly represent us.

If stereotyping people is one of the main ways we judge their credibility, then trying to influence how we are stereotyped is one of the main ways we affect others' judgments of our credibility. Of course, when others perceive us, especially face-to-face, there are certain groups to which they will almost invariably assign us—our very appearance activates certain stereotypes. For many people, for instance, race, gender, and age group can be easily "read" off their bodies, and all of the credibility-relevant attributes those categories are associated with thus get attached to them as soon as they meet someone new. One need not look too carefully to see that I am white, male, and an adult. However, other categories or groups I belong to are less obvious—for instance, I am also Jewish, an academic, a fan of pop-music, gay, a resident of Brooklyn, and originally from Tennessee. Aspects of my appearance might subtly suggest my membership in some of those groups to some people, but I am not assigned to them automatically. As such, I can more effectively influence other people's perceptions of my membership in such groups. Though I have certain stereotypically Jewish features, if I talked a lot about Christmas and Easter and never mentioned anything having to do with Judaism, I could give someone who didn't know me well the impression that I belonged to a different religious community. Likewise with my vocation and sexual orientation.

Being able to influence how we are categorized affects our perceived credibility, not only because the groups we belong to are associated with various credibility-relevant attributes (like trustworthiness, expertise, and so on), but also because the attributes associated with the different groups we belong to often conflict. I look like I live in Brooklyn, which could signal to others that I lack credibility when talking about, for instance, the political culture of the southern United States. Given the opportunity, however, I could explain that I was born and raised in the south, and that my experiences there make up for the credibility deficit that the stereotypes associated with my now being a New Yorker impose on perceptions of me.

What's more, the relationship between stereotypes and identity perceptions runs both ways. Stereotypes color how individual members of stereotyped groups are perceived, and individuals who don't conform to stereotypes can change which attributes are associated with the groups of which they are members.[71] Social psychologists call this *stereotype correction.* Someone who holds the stereotype that Jewish people are greedy, might, upon meeting sufficiently many non-greedy Jews, revise their stereotype. Someone who holds the stereotype that gay men are promiscuous, might, upon meeting sufficiently many monogamous gay men, revise their stereotype. Someone who holds the stereotype that women are more intuitive than rational, might, upon meeting sufficiently many eminently rational women, revise their stereotype.

Our ability to do that, however, to revise others' stereotypes about the groups to which we belong, requires being able to shape the way they perceive us. That is because stereotypes are self-reinforcing—they undermine evidence that they are wrong. Psychologists call this

---

[71] See Brown (2010) *Prejudice: Its Social Psychology*, especially chapter 9.

confirmation bias. We can only attend to so much information, and we are inclined to attend most carefully to information that confirms our assumptions. When different pieces of information conflict, we usually give more weight to the information that doesn't require us to revise our other beliefs. Seeing past or through stereotypes therefore requires a concerted effort, either on the part of the perceiver (to challenge her own assumptions) or on the part of the perceived (to challenge the assumptions of others). If we weren't able to contribute to the processes by which others interpret information about us—to author our social selves—it would be extremely difficult to induce others to revise their stereotypes. Our capacity for social self-authorship thus contributes not only to our ability to access the things we want and need, but also to our ability to establish credibility for ourselves. Without it we are beholden to others' assumptions, unable to influence the judgments they make about us.

Recognition

Finally, being able to change the way other people understand certain stereotypes is important for both individuals and the groups to which they belong. Obviously, it's important for me to be able to overcome false stereotypes about myself, so that I am better situated to act effectively in society. But it's also important for the groups associated with those stereotypes to detach from them. Consider stereotypes about gay men in the United States. For a long time gay men were associated with sexual deviancy. Indeed, homosexuality was *defined* as an "inversion" of the normative sexual order. As a result, gay rights were hardly considered. Who would think of enshrining in law the right to be a deviant? The political position of gay people in the United States has recently changed in large part because gay people came out of the closet and dispelled

stereotypes about themselves. We demonstrated to friends and loved ones that being gay isn't harmful or deviant; there just happen to exist more ways of being a healthy, loving, sexual being than people previously admitted.

Social self-authorship played a crucial role in this process. Until gay people came out in large numbers—both interpersonally and publicly—there was no meaningful force pushing back against ignorance and homophobia, no challenge being made against false and demeaning stereotypes.[72] Overcoming negative identity prejudices, as Fricker calls them, requires actively working to shape the way others perceive who we are, offering alternative contexts and narratives for interpreting our behaviors and intentions. And insofar as struggles like the gay right struggle, which is to say struggles for social and political recognition, rely for their success on changing the way others in a democratic society think about and feel toward members of the struggling groups, their ability to author their social selves is crucial for achieving meaningful progress. Social self-authorship therefore contributes to agency, not only socially and interpersonally, but also politically.

The above are just a few illustrations of what is at stake in our capacity for social self-authorship. One could identify many more. What's important for present purposes is that our capacity to influence how others perceive and understand us is deeply connected to our ability to act effectively in society—to things like our ability to access what we want and need, to establish credibility, and to secure political recognition. In the next section, I argue that it isn't only other people who make important decisions about how to treat us, much of that work is today delegated to computers. Moreover, much like human perceivers, computers are prone to

[72] See, for example, Herek and Glunt (1993): "Personal contact with a gay man or lesbian is a powerful predictor of heterosexuals' attitudes toward gay men" (242).

prejudice and error, to making unfounded assumptions about who we are, what we want, and why we act the way we do. To safeguard social and political agency in the Information Age, we therefore need not only to be able to influence how other people understand us, but to be able to influence the way that algorithms understand us too.

3.3. Identity and Algorithms

Until the late 1970s, if someone requested a loan, the bank would dispatch a loan investigator to determine whether or not the applicant was likely to repay it. A human being would go out into the world, collect information, and assess it. He or she (probably he) would make a judgment, an informed guess as to whether or not the applicant would default on the loan. In those days, loan investigators collected information pertaining to the "5 Cs":

> the character of the person (do you know the person or their family?); the capital (how much is being asked for?); the collateral (what is the applicant willing to put up from their own resources?); the capacity (what is their repaying ability. How much free income do they have?); the condition (what are the conditions in the market?). (Thomas 2000, qtd. in Malheiros, et. al. 2013)

In other words, investigators gathered what information they could about a person's life and interpreted it to mean that the person was or wasn't creditworthy. Needless to say, this process was subject to prejudice and prone to error. As legal theorist Frank Pasquale writes, the reports "included attributes like messiness, poorly kept yards, and 'effeminate gestures.' The surveillance could be creepy and unfair—virtually everyone has some habit that could be seized on as evidence of unreliability or worse" (2015, 21-22).

Passage of the Equal Credit Opportunity Acts in 1975 and 1976, along with the introduction of credit cards and the computerization of much of the banking industry, led in the

late 1970s and early 1980s to a shift from credit investigations to personal credit *scoring* (Pasquale 2015; Thomas 2000). Today, if I apply for a loan, whether or not I receive it has mostly to do with my credit (or "FICO") score—a three-digit number, issued by one of the three major U.S. credit bureaus, based on information gathered by computers, and decided entirely by algorithms. The score, like the old investigator's report, represents the credit bureau's best guess at the likelihood I will pay back any money I borrow. It is based on a number of factors, including my record of payments, the ratio of available credit to money I owe, the length of my credit history, and the types of credit I use (Kelly 2010).

Because a credit score is indicated so clearly and concisely, and because it is based on hard *facts*, it may seem as though it represents something more true or more real than the subjective, prejudiced reports the banking industry once used. That, however, is not the case. While it's true that credit scores somewhat more accurately predict default risk[73], they still represent a guess—a judgment, a perception, an interpretation of available information—and one set of facts can tell more than one story. Imagine a man who has had a mortgage for a few years and has always paid his bills on time. On the day his latest bill is due, he gets hit by a car on the way to mail in his payment. Knocked unconscious, he doesn't wake up for two weeks. As soon as he gets home from the hospital he sends in his check, but it's too late. One late mortgage payment can decrease a credit score by as much as seventy-five points (Berger 2013). That lower score, in turn, can significantly raise the overall cost of a future mortgage, perhaps by so much as to make it too expensive to take on.

[73] See Thomas (2000).

A credit score is thus like any other perception one might have of who another person is. It is partial and perspectival. The difference is that the perspective is coded in an algorithm, rather than an individual's singular history and prejudice. The late payment in the previous example is an anomaly. If the man were given the opportunity to contextualize and explain it, he could show that it doesn't indicate anything significant about the likelihood that he'll default on future loans. "[Credit scores] may *feel* as objective and real as the score on a math test," writes Pasquale, "But a critical mass of complaints over the past twenty years has eroded credit assessors' claims to objectivity and reliability" (2015, 25, emphasis in original).

What's more, credit scores are only the beginning. An episode of the British science fiction show *Black Mirror* depicts a world in which everything a person does is scored: whether they brush their teeth in the morning, how many calories they consume at each meal, the time they spend exercising, and how much sleep they get.[74] Each person watches their score fluctuate throughout the day and sees how it compares to the scores of those around them. According to a report from the World Privacy Forum, that fictional world is very nearly the world we live in; most people just don't know it. The report lists dozens of different ways we are scored. Rating agencies assign us "consumer profitability" scores, which gauge how much money a lender is likely to make from extending us credit, and "charitable donor" scores, which predict how much money we're likely to give away. They give us "churn" scores, which tell them whether or not we might move our accounts to competitors, and "fraud" scores that determine if our accounts have been hacked. One data broker developed a "pregnancy predictor" score to figure out who to target with advertising for baby products. The Affordable Care Act (ACA) in the United States

[74] The episode is titled "15 Million Merits."

requires that each person covered under it be assigned a "health risk" score to predict how sick they might get and how much care they might require. "Frailty" scores guess whether elderly people can handle various medical treatments, and "social" scores measure how much influence individuals have over members of their social networks (Dixon and Gellman 2014).

Like credit scores, these other scores are guesses. They do not represent what is known about individuals, but rather what can be surmised. They are interpretations of the information available about us, inferences made using sophisticated calculations and complex algorithms. Also like credit scores, these guesses about who we are and what we are likely to do or be can have an enormous impact on us. "After a consumer is scored, ranked, described, or classified," write the authors of the World Privacy Forum report, "companies, governments, private enterprises, health care entities, and others including law enforcement, can then use the resulting score to make decisions about an individual or group" (Dixon and Gellman 2014, 19). They can decide whether or not we are eligible for trivial things like advertisements and discounts, or if we should have access to deeply important things like loans, insurance, and healthcare.[75]

When we negotiate our public identities today, we are thus negotiating not only with other people, but also with computers. Cultural theorist John Cheney-Lippold calls these efforts to digitally score and classify us the production of "new algorithmic identities":

> The networked infrastructure of the internet, with its technological capacity to track user movements across different web sites and servers, has given rise to an industry of web analytics firms that are actively amassing information on individuals and fine-tuning computer algorithms to make sense of that data. The product of many of these firms is a 'new algorithmic identity', an identity formation that works through mathematical algorithms to infer categories of identity on otherwise anonymous beings. (2011, 165)

[75] See Graham (2005).

Crucially, these algorithms and the categories they sort us into are not neutral (boyd 2014). People often assume that algorithms are "objective," and that the activities which have been delegated to computers are therefore entirely free from bias and error. What this overlooks is that algorithms don't fall from the sky; they are designed, by humans. An algorithm is just an abstract set of instructions for how to solve a problem, and any given problem can be solved in a number of different ways.[76] Moreover, computers don't know anything more about the world than we tell them, so programmers have to specify for them what information is relevant to solving the problem at hand. Settling both of these questions—how the problem ought to be solved and what information is relevant to solving it—involves making a number of decisions.

To take the standard, Computer Science 101 example, the algorithm for making a peanut butter and jelly sandwich is to spread peanut butter on one slice of bread, jelly on another, and then put the slices together so that the spread sides meet. In this trivial example, at least one decision has already been made that could have been made differently—namely, to spread the peanut butter side first. Perfectly good PB&Js have been made by spreading the sides in the opposite order. Furthermore, for a computer to run this algorithm it would have to be told what counts as "peanut butter," "jelly," and "bread."[77] Could it use white bread *or* whole wheat? What about naan or a bagel? Does the bread's expiration date matter? What about the thickness of the slices? When programmers make these kinds of decisions—which again, they *have* to make—they build various assumptions into their programs. Different assumptions can produce vastly different outputs.

---

[76] See Goffey (2008).

[77] Indeed, it would have to be told a great deal more than that. But these variables will suffice to make my point.

Assumptions about data, which is to say, assumptions about which information is relevant (and which isn't relevant) to solving a particular problem, are especially important. As sociologist Tarleton Gillespie writes,

> There is a premeditated order necessary for algorithms to even work. More than anything, algorithms are designed to be and prized for being functionally automatic, to act when triggered without any regular human intervention or oversight. This means that the information included in the database must be rendered into data, formalized so that algorithms can act on it automatically. […] Recognizing the ways in which data must be 'cleaned up' is an important counter to the seeming automaticity of algorithms. (2014, 4-5)

In order to produce working information systems, technologists need not only to construct clever algorithms, but also to organize the data that's to be fed into those algorithms into discrete, processable kinds. They have to pluck relevant information from out of the "great blooming buzzing confusion," as James said, and then encode that information in a form that the algorithms can understand and manipulate. "Evaluations performed by algorithms always depend on inscribed assumptions about what matters," says Gillespie, "and how what matters can be identified" (12).

It's not hard to see how important decisions about "what matters" can become when the objects of computer analysis are human beings and their identities. If our 1970s loan investigator had been tasked with articulating what matters about loan applicants, he would likely have trained a computer to evaluate creditworthiness based on the appearance of applicants' lawns or the effeminacy of their gestures. As Cynthia Dwork and Deirdre Mulligan write:

> While many companies and government agencies foster an illusion that classification is (or should be) an area of absolute algorithmic rule—that decisions are neutral, organic, and even automatically rendered without human intervention—reality is a far messier mix of technical and human curating. Both the datasets and the algorithms reflect choices, among others, about data, connections, inferences,

> interpretation, and thresholds for inclusion that advance a specific purpose. […] They reflect the explicit and implicit values of their designers. (2013, 35)

Deciding which classifications or attributes are relevant for determining who is eligible for a loan is the computer equivalent of interpreting the available information about someone; it is the interpretive lens through which computers make decisions about us. And we human beings hard-code those lenses into the logics of computer algorithms and the structures of datasets. The fact that different criteria are used today to decide who is creditworthy than were used in the 1970s is not only the result of the shift from human investigators to computer scoring. It is also a product of conscious efforts on the part of data scientists, regulators, law makers, and the public to change the assumptions involved in making those decisions.

Much like the decisions that the human beings we interact with make about us, the decisions computers make about us are value-laden and subject to bias. Furthermore, the range of decisions being delegated to computers is constantly expanding. Algorithms are tasked with deciding who is a security risk, and therefore ought to be the subject of extra surveillance, and sometimes barred from traveling. They decide who is entitled to social services and resources like Medicaid and food stamps (Kerr and Earle 2013). One company, called Gild, has developed an algorithm which promises to automate the process of choosing who to hire. For now, it is limited to software developer jobs:

> As with any grading scheme, Gild's algorithm—which scores from 1 to 100—is making judgments about what makes a good developer. The software grades the quality of someone's code by checking for basic errors and also gauging its complexity. It also looks at how extensively programmers' open-source code has been taken up by other projects. (Leber 2013)

Some look at a tool like Gild's and see the promise of perfect meritocracy. They see a system in which computers impartially choose the best possible job candidate, and that candidate is hired

on that basis alone, rather than on the basis of who they know or how well they interview. As we've seen, however, no judgment made by an algorithm is truly impartial. They are subject to the beliefs and values of the people who program them. As one user of Gild's software points out, "It tends to reward developers who know a small number of programming languages really well and dabble in several others" (Leber 2013). That someone who meets that description is a more desirable software developer than someone who doesn't reflects a certain perspective, which I'm sure not all technology firms share.

Having agency over the kinds of decisions that computers make about us thus requires that we be able to correct mistakes, challenge assumptions, and offer our own interpretations of the facts. We need to be recognized as the authors of our social selves, not only by the people who make judgments and decisions about us, but also the algorithms. This might seem like a strange or perplexing claim, that computers ought to treat us a certain way. What it means is that computers and information systems ought to be *designed* in such a way that affords such treatment. In the next chapter I examine what exactly that demands, and how those demands might be realized.

References

Alcoff, Linda Martín. 1999. "On Judging Epistemic Credibility: Is Social Identity Relevant?," *Philosophic Exchange* 29 (1): Article 1.

Alcoff, Linda Martín. 2006. *Visible Identities: Race, Gender, and the Self.* Oxford: Oxford University Press.

Benn, Stanley I. 1984. "Privacy, Freedom, and Respect for Persons." In *Philosophical Dimensions of Privacy: An Anthology,* ed. Ferdinand Schoeman. Cambridge: Cambridge University Press.

Berger, Rob. 2013. “5 Myths About Late Payments & Your FICO Scores.” *Credit.com* (blog), December 5. http://blog.credit.com/2013/12/5-myths-about-late-payments-your-fico-scores-71720/

Berlin, Isaiah. 1969. *Four Essays on Liberty.* Oxford: Oxford University Press.

boyd, danah. 2014. *It’s Complicated: The Social Lives of Networked Teens.* New Haven: Yale University Press.

Bratman, Michael. 2009. “Shared Agency.” In *Philosophy of the Social Sciences: Philosophical Theory and Scientific Practice,* ed. Chrysostomos Mantzavinos. Cambridge, UK: Cambridge University Press.

Brown, Rupert. 2010. *Prejudice: Its Social Psychology*. *2nd Edition.* West Sussex, UK: Wiley-Blackwell.

Bureau of Justice Statistics. 2013. “16.6 Million People Experienced Identity Theft in 2012.” http://www.bjs.gov/index.cfm?ty=pbdetail&iid=4911

Cheney-Lippold, John. 2011. “A New Algorithmic Identity: Soft Biopolitics and the Modulation of Control.” *Theory, Culture & Society* 28 (6): 164-181.

Dixon, Pam, and Robert Gellman. 2014. *The Scoring of America: How Secret Consumer Scores Threaten Your Privacy and Your Future.* Report by the World Privacy Forum. http://www.worldprivacyforum.org/2014/04/wpf-report-the-scoring-of-america-how-secret-consumer-scores-threaten-your-privacy-and-your-future/

Dwork, Cynthia, and Deirdre K. Mulligan. 2013. “It’s Not Privacy, and It’s Not Fair.” *Stanford Law Review Online* 66: 35-40. http://www.stanfordlawreview.org/online/privacy-and-big-data/its-not-privacy-and-its-not-fair

Fricker, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing.* Oxford: Oxford University Press.

Gerstein, Robert. 1978. “Intimacy and Privacy.” *Ethics* (89): 76-81.

Gillespie, Tarleton. 2014. “The Relevance of Algorithms.” In *Media Technologies,* eds. Tarleton Gillespie, Pablo Boczkowski, and Kirsten Foot. Cambridge, MA: MIT Press. Pre-publication version accessed online at http://culturedigitally.org/2012/11/the-relevance-of-algorithms/

Goffey, Andrew. 2008. “Algorithm.” In *Software Studies: A Lexicon,* ed. Matthew Fuller. Cambridge, MA: MIT Press.

Goffman, Erving. 1967. *Interaction Ritual: Essays on Face-to-Face Behavior.* New York: Pantheon Books.

Graham, Stephen D.N. 2005. "Software-Sorted Geographies." *Progress in Human Geography* 29 (5): 562-580.

Herek, Gregory and Eric Glunt. 1993. "Interpersonal Contract and Heterosexual's Attitudes Toward Gay Men: Results from a National Survey." *The Journal of Sex Research* 30 (3): 239-44.

Huseman, Jessica. 2014. "Setting the Record Straight on Mortgages for Undocumented Immigrants." *National Mortgage News,* October 17. http://www.nationalmortgagenews.com/news/risk-management/setting-the-record-straight-on-mortgages-for-undocumented-immigrants-1042907-1.html

Kelly, Jeanne. 2010. "Recipe for a High FICO Credit Score." *The Huffington Post,* November 15. http://www.huffingtonpost.com/jeanne-kelly/recipe-for-a-high-fico-cr_b_777627.html

Kerr, Ian, and Jessica Earle. 2013. "Prediction, Preemption, Presumption: How Big Data Threatens the Big Picture." *Stanford Law Review Online* 66: 65-72. http://www.stanfordlawreview.org/online/privacy-and-big-data/prediction-preemption-presumption

Leber, Jessica. 2013. "A Startup That Scores Job Seekers, Whether They Know It or Not." *MIT Technology Review,* March 7. http://www.technologyreview.com/news/511896/a-startup-that-scores-job-seekers-whether-they-know-it-or-not/

Malheiros, Miguel, Sacha Brostoff, Charlene Jennett, and M. Angela Sasse. 2013. "Would You Sell Your Mother's Data? Personal Data Disclosure in a Simulated Credit Card Application." In *The Economics of Information Security and Privacy,* ed. Rainer Böhme. Berlin: Springer-Verlag.

Pasquale, Frank. 2015. *The Black Box Society: The Secret Algorithms That Control Money and Information.* Cambridge, MA: Harvard University Press.

Quinn, Kimberly, C. Neil Macrae, and Galen Bodenhausen. 2007. "Stereotyping and Impression Formation: How Categorical Thinking Shapes Person Perception." In *Sage Handbook of Social Psychology: Concise Student Edition,* eds. Michael A. Hogg and Joel Cooper. Thousand Oaks, CA: Sage Publications.

Rachels, James. 1975. "Why is Privacy Important?" *Philosophy and Public Affairs* 4 (Summer): 323-333.

Riou, Garrett. 2014. "Hospital Visition and Medical Decision Making for Same-Sex Couples." *Center for American Progress Blog*, April 15. https://www.americanprogress.org/issues/lgbt/news/2014/04/15/88015/hospital-visitation-and-medical-decision-making-for-same-sex-couples/

Solove, Daniel. 2008. *Understanding Privacy.* Cambridge, MA: Harvard University Press.

Thomas, Lyn. 2000. "A Survey of Credit and Behavioural Scoring: Forecasting Financial Risk of Lending to Consumers." *International Journal of Forecasting* 16: 149-172.

Chapter 4

Hermeneutic Privacy:
Toward Effective Privacy Policy

> "Privacy covers many things. It protects the solitude necessary for creative thought. It allows us the independence that is part of raising a family. It protects our right to be secure in our own homes and possessions, assured that the government cannot come barging in. Privacy also encompasses our right to self-determination and to define who we are. Although we live in a world of noisy self-confession, privacy allows us to keep certain facts to ourselves if we so choose. The right to privacy, it seems, is what makes us civilized."
>
> - Ellen Alderman and Caroline Kennedy, *The Right to Privacy*

> "We need more shades and more blinds and more virtual curtains."
>
> - Jeffrey Rosen, *The Unwanted Gaze*

In 2010, the founder and CEO of Facebook, Mark Zuckerberg, claimed during an interview that privacy is no longer a social norm. "People have really gotten comfortable," he said, "not only sharing more information and different kinds, but more openly and with more people. That social norm is just something that has evolved over time."[78] This was obviously a self-serving message: Facebook makes its money by collecting and analyzing information about its users. Were we to decide all of a sudden that we didn't want to share information about ourselves "more openly and with more people," Facebook's business model would be upended. Nevertheless, Zuckerberg's statement still rings true. Facebook has over 1 billion active users, who share more than 5 billion pieces of content, including 350 million photos, *every day.*[79] That

---

[78] See Johnson (2010).

[79] As of February 2014. See Zimmer (2014).

kind of zeal for self-publicity would seem to suggest that, at least for the world's young people, privacy is no longer a top concern.

Yet social media scholar danah boyd cautions us not to jump too quickly to any conclusions. It may *seem,* she says, that young people don't care about privacy, but that's only because we don't know what to look for. "Journalists, parents, and technologists seem to believe that a willingness to share in public spaces—and, most certainly, any act of exhibitionism and publicity—is incompatible with a desire for personal privacy," boyd writes. "The teens that I met genuinely care about their privacy, but how they understand and enact it may not immediately resonate or appear logical to adults" (2014, 56). In fact, boyd finds, teens go to great lengths to ensure their privacy. They maintain multiple profiles across a variety of social media networks, and as soon as a network becomes too familiar to adults, they abandon it. In their posts and online comments, teens engage in what boyd calls "social steganography"—they speak elliptically and in slang, so that even though their messages are visible to everyone, their meanings are only accessible to the target audience. Some teens disable their social media profiles each afternoon, so their parents can't browse them when they get home from work, and then they revive them in the morning before they go to school.[80]

In other words, adults look at teen behavior online and see a troublesome disregard for how they will appear to future college admissions officers, potential employers, loan officers, and the police. What the adults fail to realize is that, for teens, it is they themselves—parents, teachers, and all the other adults in teens' lives—that teens are worried about. For teenagers, parents and teachers are the principal characters who control them, who constrain their behavior

[80] See boyd (2014), especially chapter 2.

and limit their agency. Thus the epistemic boundaries teens are concerned about drawing are the ones between themselves and their parents, their teachers and principals, and other authority figures. And when it comes to drawing *those* boundaries, shaping how *those* people access and interpret information about them, teens are extremely creative.

Moreover, adults are not only looking in the wrong places to find young people caring about privacy, many also fail to recognize privacy-preserving behavior when they see it. That is, at least in part, because privacy is often thought of as a *state* that one can be in, rather than, as philosopher Jeffrey Reiman, argues, a set of *practices.* Privacy, for Reiman, "involves a complex of behaviors that stretches from refraining from asking questions about what is none of one's business to refraining from looking into open windows one passes on the street, from refraining from entering a closed door without knocking to refraining from knocking down a locked door without a warrant. Privacy can in this sense be looked at as a very complicated social ritual" (1976, 38-9). Specifically, it is a social ritual aimed at creating and respecting interpersonal boundaries.[81] If privacy is understood to be a state, then changes in the world which threaten that state would seem to threaten the very possibility of privacy. For example, if privacy is understood to be the state in which information about us is inaccessible to others, and if the world has become the kind of place where it is increasingly difficult to enter that state, then we would seem to be living in a world of diminished privacy.

On the other hand, if privacy is understood as a practice, as a social ritual or set of behaviors aimed at drawing boundaries between people, then we could imagine adapting those behaviors to new circumstances. Today, we may not be able to keep information about us from

[81] Privacy "articulates intersubjectively recognized personal boundaries that are the sine qua non for the establishment and maintenance of autonomous identities" (Cohen 1996, 204-5).

others, but as I've argued in previous chapters, we can draw epistemic boundaries in other ways. Many teens, boyd argues, intuitively understand this. They know that they live in a world where information about them is easily accessible. After all, they are the ones creating and distributing a great deal of it. They understand that their lives are intensely public. And yet that doesn't induce them, as Zuckerberg claims, to abandon their desire for privacy. It induces them, rather, to devise ways of creating what Helen Nissenbaum calls "privacy in public."[82]

Nissenbaum argues that privacy norms can apply even to information that is publicly accessible. For example, she quotes Ferdinand Schoeman: "just because something happens in public does not mean it becomes a public fact: the Central Park rape occurred in public as did the trial of the accused, but the victim maintains a measure of privacy as to her identity" (1994, 81).[83] In such situations, despite the fact that any number of people might have gained access to sensitive information by perfectly innocent means, privacy norms nevertheless bar them from using that information in certain ways. The information, in other words, despite being in one sense "public," is not therefore, as Nissenbaum puts it, "up for grabs." Prior to the advent of information technology, the range of situations where one could reasonably expect privacy around public facts was limited. Today, however, privacy in public is tremendously important. Another example Nissenbaum gives is that of information in the public record. As we've seen, the government and other public institutions collect all kinds of information for all kinds of reasons. Prior to the advent of information technology, accessing and analyzing such information was time and resource intensive. Now public records about private individuals can be accessed

---

[82] See Nissenbaum (1997).

[83] Quoted in Nissenbaum (1997), p. 214.

cheaply and easily, aggregated, and subjected to rigorous, automated analysis. Thus, where we were once able to rest assured that the public record didn't reveal anything sensitive about us, now we cannot.[84] And so, Nissenbaum argues, it raises considerable privacy issues. "Just because people are able to learn these facts by referring to public records," she says, "does not imply a right to distribute and use the information in any way they choose" (1997, 215).

"Many teens," boyd finds, "are developing innovative solutions to achieve privacy in public. To get there, they must grapple with the tools that are available to them, the norms that shape social practices, and their own agency" (2014, 59). That means developing the kinds of intricate privacy practices described above. "For the teens that I interviewed," writes boyd,

> privacy isn't necessarily something that they have; rather it is something they are actively and continuously trying to achieve in spite of structural or social barriers that make it difficult to do so. Achieving privacy requires more than simply having the levers to control information, access, or visibility. Instead, achieving privacy requires the ability to control the social situation by navigating complex contextual cues, technical affordances, and social dynamics. Achieving privacy is an ongoing process because social situations are never static. (2014, 60)

Privacy, for teens, is just as it is for adults: the desire for interpersonal boundaries. Teens are just better equipped than adults are to understand how we might create such boundaries in the Information Age. "Instead of signaling the end of privacy as we know it," boyd writes, "teens' engagement with social media highlights the complex interplay between privacy and publicity in the networked world we all live in now" (2014, 57).

That networked world, as described in previous chapters, is one in which information about us abounds, persists, and is widely accessible. It is, in other words, a world in which we live far more publicly than people once did. It's a world in which, as I have been arguing, we

---

[84] As Anita Allen says, "Data once resigned to the dustbin of history is now at anyone's fingertips" (2011, 162).

must learn to rely less on tools of concealment and exposure to shape how others perceive us, and more on tools of interpretation. Rather than thinking of information as something we possess and control, we must learn to think of it as a way of relating to other people, as our means for authoring the identities through which others come to know us. Teens understand this. They understand that in an information-saturated world the key to drawing epistemic boundaries has less to do with controlling information than it does with shaping *meaning.* "They recognize," boyd writes, "that limiting access to meaning can be a much more powerful tool for achieving privacy than trying to limit access to the content itself" (2014, 69).[85]

This is what I call *hermeneutic privacy.* It is the notion that successfully drawing epistemic boundaries, and thus achieving information privacy, has as much to do with the processes by which information about us is interpreted and made meaningful as it does with the processes by which information about us is obtained in the first place. I should note, again, that I am not claiming that hermeneutic privacy is *all there is* to information privacy. My account is not exhaustive. My argument is rather that hermeneutic privacy is one dimension of a broader phenomenon, one of the many "social rituals" or sets of practices we engage (or ought to engage) in in order to draw healthy interpersonal boundaries. It is an important part of privacy overall—especially in the Information Age—and a part not yet given its due by privacy theorists.

As one can see in the first epigraph to this chapter, privacy theorists have long recognized that a central component of information privacy is what Ellen Alderman and Caroline Kennedy call the right to "define who we are." We are all, as Jeffrey Rosen puts it, "entitled to be regarded as self-defining individuals […] The ideal of privacy, similarly, insists that individuals should be

[85] For more on the notion that privacy as a value must be realized differently in different social and technological contexts, see Lessig (1999).

allowed to define themselves" (2000, 223). Legally, at least in the American context, protecting such a right has meant recognizing torts of defamation, public disclosure of embarrassing facts, appropriation of name or likeness, and publicity which places a person in a false light.[86] Each of these torts recognizes in one way or another the harms that come from interfering with a person's ability to shape how others perceive and understand who they are. But these protections are insufficient.[87] They recognize our interest in not having others lie or reveal embarrassing information about us. They fail, however, to recognize that others can interfere with or undermine our ability to influence how we are perceived and understood in other ways. As we've seen in previous chapters, truths can tell more than one story. We have an interest not only in preventing lies about us and embarrassing facts from getting out, but also in being able to shape the way the truth about us is interpreted and understood. As I argued at the end of chapter 2, information technology undermines that ability in two ways: by making social self-authorship invisible and unnecessary. That is to say, information technology makes it difficult for us to know when we are being perceived, and it makes it easy for others to perceive who we are without our input. Protecting information privacy in the Information Age requires solving both of these problems.

In what follows, I suggest how we might do that. In the first section, I argue that the privacy norms entailed by control theories are norms of consent, and that such norms are inadequate. Information is too difficult to control, and as a result, the only meaningful consent

---

[86] Defamation is technically not a privacy tort, though it is for our purposes conceptually tied with the other three, in that each has to do with protecting an individual's ability to define who they are. For the classic delineation of privacy torts, see Prosser (1960). For a critique of Prosser's view, see Richards and Solove (2010).

[87] For a more thorough discussion of the difficulty of using privacy and defamation torts to protect oneself online, see Abril (2007).

decisions we can make regarding information about us is whether or not to divulge it. For reasons I put forward below, all-or-nothing consent of this kind cannot provide for robust information privacy. Thinking about information privacy in terms of social self-authorship, by contrast, as a process in which multiple parties negotiate their public identities, entails norms of fairness and due representation, rather than consent. I articulate the specific normative demands of what I call hermeneutic privacy, and explain why they are better suited to the task of protecting information privacy in the Information Age. In the subsequent sections, I put the norms articulated in the first to work, showing what they can do, concretely, to protect hermeneutic privacy. I discuss their implications for technology design, technology education, and technology law.

4.1. From Consent to Due Representation

I have argued that information privacy has to do with more than merely concealing information about ourselves; it has to do with influencing the way information about ourselves is contextualized and understood—the work, as I've called it, of social self-authorship. I have argued that our capacity for social self-authorship is valuable because it contributes to our social and political agency, and that it's being undermined by information technology, though in different ways than competing theories of privacy would have us believe. What remains to be explained is how we can right the course, how we can protect our capacity for social self-authorship from the forces that threaten it, and thus maintain information privacy in the Information Age. In other words, we need to specify the norms of hermeneutic privacy, the set of practices (and constraints on our practices) that enable social self-authorship to function.

Existing information privacy norms are built for the most part upon the notion of individual *consent*. On a control model this makes perfect sense: if information about me is mine to control, then if someone wants to collect it and use it, they have to ask me first and get my blessing. However, for reasons that I hope are becoming increasingly clear, respecting consent norms in the Information Age is extremely difficult. Such a vast amount of information about us is being collected, and so much of that information—*on its own*—is not particularly sensitive, that asking for our permission to collect each and every bit of it would be ludicrous. Imagine, for instance, getting a pop-up on your computer or phone each time Google wanted to record some piece of information about you: "Can we record your search term?" "Can we record the time of your search?" "Can we record the brand computer you use?"

This situation has created enormous problems in privacy law. As the head attorney of a large privacy law practice told the attendees of a conference on privacy and technology law, "the sheer quantity of data involved in this field makes giving notice and choice to individual subjects virtually impossible, with the result that the concept underpinning privacy law is already 'out the window'" (Chapman 2015).[88] In other words, if it's impossible for us to control information about ourselves, it's equally impossible for others to respect our right to do so. Rather than admit that, though, businesses generally deal with this situation by requiring customers to accept Terms of Service (TOS) or End User License Agreements (EULAs) that grant the companies the right to do with their information what they will. The legal function of such agreements is, of course, to shield those companies from liability should they fail to secure the information they are given or should they use it in a way that upsets people. Symbolically, however, the agreements serve as a

[88] The specific field the attorney is referring to is the field of Big Data, which I discuss in the conclusion.

kind of tacit admission that the most control these organizations can offer users over the information about themselves that they provide is the choice of not divulging it in the first place.[89]

All-or-nothing consent of this kind—which is to say, granting the people, businesses, and organizations with whom we interact either no rights over the information that we provide them or the right to do anything at all with it—is unsatisfactory for a number of reasons. First, it suggests that information privacy is an extremely blunt instrument, a sort of spigot we get to control, either putting information about ourselves out into the world or not. Even those who think that information privacy is indeed about information control imagine the control being more fine-grained than that. They think that we ought to be able to control who has information about us, what they are allowed to do with it, under which circumstances, for how long, and so on.

Second, it leaves people with only two, mutually-exclusive options: either enjoy information privacy or enjoy services that require information about you, but not both. In today's world the former is not a real option. Opting out of all of the various technologies and technology-dependent services we encounter on a day-to-day basis, choosing, as it were, to keep the spigot of information about ourselves closed entirely, would mean decamping to a Thoreauvian cabin in the woods. Yet relinquishing information privacy is not an appealing option either. As we saw in the previous chapter, our ability to influence how others perceive and understand who we are is a necessary condition for social and political agency. It determines, in

[89] In addition to providing few legal rights to users and blanket protections to service providers, such agreements are famously filled with legal jargon and are so difficult for lay people to understand that websites like "Terms of Service; Didn't Read" (http://www.tosdr.org) have sprouted up to explain to people in readable language what various companies' user agreements actually say. For more on why consent agreements aren't particularly useful in information privacy contexts, see Cate (2006).

part, how other people and organizations decide to treat us. Giving up information privacy means giving up our say in those decisions.

Third, all-or-nothing consent suggests that once we release information about ourselves to others that information is "up for grabs." If companies cannot (or will not) guarantee that the information we give them about ourselves will be protected, and that it will only be put to use in ways that we have explicitly sanctioned, then when we choose to open the spigot of information about ourselves we have to assume that it's being released not just to that specific company, but to anyone at all. Indeed, American privacy law makes this explicit. In order to have an actionable claim against someone for publicly disclosing private facts about oneself, they must be facts that are nowhere else available in public. "Certainly no one can complain," writes William Prosser, "when publicity is given to information about him which he himself leaves open to the public eye" (1984, 110). For all of these reasons, the normative dimension of control theories of privacy —i.e., the consent model—is as incapable of protecting information privacy as its conceptual dimension is incapable of describing it.

Thinking about information privacy in terms of social self-authorship, rather than information control, suggests a way out of these problems. On a control model, the normative issue is whether or not an individual has exclusive rights over information about them. The problem is that it is nearly impossible to enforce such a right in any meaningful way. On an authorship model, however, the normative issue is whether or not the process by which an individual shapes their public identity functions correctly. The norms of hermeneutic privacy are thus regulative ideals: standards which will never be perfectly achieved, but indicate how close or how far from the ideal process a particular instantiation of it is. The question this raises, of

course, is what the ideal process of social self-authorship looks like, what it means for social self-authorship to function correctly. One might be tempted to say here that social self-authorship functions correctly when a person is perceived and understood in just the way they want to be. But such an approach falls back into control-oriented thinking. Social self-authorship aims, again, not to control how others perceive and understand who we are, but to shape or influence it. The process of social self-authorship is, as I've said, a *negotiation*. Negotiations function well not when one or another party is in control, but when each party and their interests are adequately represented. The relevant norms are not, therefore, those of consent, but rather of fairness and due representation.

Before moving on, I want to point out how shifting from a control model of privacy to an authorship model reflects an understanding of privacy overall as a *social value* rather than an individual one. The interests privacy protects are not (or are not only) the interests of individuals against other people, organizations, and the state. They are, instead, the interests of a well-functioning society. Privacy theorists have traditionally emphasized the value of privacy for promoting autonomy, arguing that respect for privacy is an expression of respect for persons.[90] And while that is no doubt true in part, it neglects the fact that respect for privacy is also an expression of respect for a well-ordered society. In order for human beings to live together in relative harmony we must be able to draw boundaries. In order for us to interact and coordinate our behavior in such a way that we are each able to exercise agency over ourselves and our lives, despite the tangle of our conflicting and overlapping efforts, we must be able to contribute to the

[90] For the classic formulation of this view, see S. I. Benn (1984).

processes by which others come to understand our interests and our intentions. Hermeneutic privacy helps to ensure that we are able to do that.

The question which now faces us, then, is what a fair process of social self-authorship looks like, what it means to enjoy due representation when negotiating our identities with others. To answer that question, I propose we turn to that old and true tool in the privacy theorist's toolkit—namely, the imagined state of surveillance. Philosophers and legal theorists have long looked for inspiration and insight to the surveillance states of Orwell's *1984*, Yevgeny Zamyatin's *We*, Aldous Huxley's *A Brave New World*, and even, as we saw in chapter 2, Kafka's *The Trial*. They exaggerate the most fearsome and problematic aspects of surveillance in order to highlight the harms that, even in its more modest forms, surveillance can cause. They examine those worlds in order to find what's missing—which values are absent, which human needs aren't being met. In our case, we already know the answer to those questions. As we saw in chapter 3, the value that social self-authorship promotes is social and political agency; its absence means potentially being unable to access important resources, to assert credibility, and to secure political recognition. Instead of looking to a surveillance scenario in order to discover what's wrong with it, I propose that we look at a form of surveillance which we take to be legitimate and ask why we think intuitively that such surveillance is alright. For surveillance would seem to be the antithesis of well-functioning social self-authorship. It is, almost by definition, a situation in which there exists an asymmetry of power between two parties in the drawing of epistemic boundaries between them. And yet if we find certain forms of surveillance acceptable, then the reasons we do—the constraints on that surveillance which make it acceptable—ought to tell us something important about what is required in order to produce and

protect social self-authorship. If social self-authorship can function even in those situations most antithetical to it, then whatever enables it to do that is what we ought to be promoting in order to protect it elsewhere.

In chapters 1 and 2, I described some of the uncountably many forms of surveillance we are subject to today. Nearly everything we do, both online and offline, leaves a trail of data behind it, and that data is collected, stored, and analyzed by myriad actors for reasons both good and bad. Some forms of surveillance seem perfectly justifiable and unproblematic; others do not. A paradigm case of *prima facie* acceptable surveillance is a closed-circuit television (CCTV) camera monitoring the retail floor of a local bank. When we step into a bank and see a CCTV camera watching us, few of us feel outrage or indignation. We accept that banks are targets for attack, and that we as much as the banks are being protected by those cameras. Yet there are limits. If we learned that our local bank was monitoring citizens' phones and email accounts, in order again to prevent attacks upon the bank and its customers, we would be outraged indeed. Thus it's not the *aim* of the surveillance that justifies it, at least not on its own. There is something about the *means* of the surveillance that we find acceptable.

I want to suggest that we intuitively accept the kind of surveillance we experience at our local bank for the following reasons: (1) we know that we are being monitored; (2) we know how we are being monitored, which is to say, we know what information about us is being collected; (3) we know why we are being monitored; and (4) we know how we are being evaluated. In other words, at our local bank we know that we are being perceived and judged, we know why, and we know how we ought to act in order to be judged positively (i.e., we know not to look like we're going to rob the bank). Furthermore, we know that if we are judged negatively,

if the bank believes that we are going to rob it, then we will be handed over to the police. And at that point, if we live in a democracy, we know that the government will accord us formal due process, meaning (5) we will have the opportunity to explain, to give our side of the story, our interpretation of the facts.

Now, just because we find this kind of surveillance acceptable does not mean that we would accept being subjected to it at all times. Many theorists have examined the perils of total surveillance, and I will not rehearse their findings here.[91] What I want to point out is that if any of the above caveats were removed, we would have cause to worry about even the kind of surveillance our local bank was undertaking. If it were monitoring us in secret, we would not accept it. (Indeed, we would be unable to in principle.) If we learned that it had begun to monitor us in ways we weren't aware of, we would cease to accept it. If the surveillance took place somewhere the public did not have reason to believe was at risk, and it was therefore unclear why we were being monitored or how we were being evaluated, we would not accept the surveillance. Finally, if we didn't know what would happen in the case that we were judged negatively—if, that is, we lacked reason to believe that we would be able to contest the judgment —we would not think the surveillance was legitimate.

The reason these constraints on the kind of surveillance we experience at the bank lead us to accept it is, I think, that we intuitively recognize that they prevent it from undermining our social self-authorship. They ensure that our social and political agency is protected. We would happily, of course, not experience any surveillance at all. But provided that a large, complex society requires it in certain cases, this is the kind that we'd like to have. We accept being

[91] See, most famously, Foucault (1977). Also see Lyons (2001) and Cohen (2010).

monitored if there is a good reason for it, and just as long as it doesn't undermine our ability to draw epistemic boundaries. Surveillance is a limit case. Again, we would not put up with being monitored everywhere, all the time. The fact, however, that the above constraints on surveillance ensure that it doesn't quash social self-authorship entirely indicates that they are important means of protecting it.

Let us then restate the constraints as positive normative requirements:

(1) We ought to know when others are collecting information about us.
(2) We ought to know what information they are collecting.
(3) We ought to know why they are collecting information about us.
(4) We ought to know how the information about us is being evaluated.

These four requirements (which correspond to the constraints above) ensure that social self-authorship is *visible*, that we are aware that others are forming opinions about us, why they are forming such opinions, and how. Constraint (5) above—that we have the opportunity to explain, to respond to the perceptions others have of us, to offer our own interpretations of the facts—can be spelled out more carefully in two further requirements:

(5) We ought to be able to correct misinformation about us.
(6) We ought to be able to contextualize and offer our own interpretations of the information held about us (by providing more information).

Finally, as I've argued, the process of social self-authorship is necessarily ongoing. New information can, at any time, lend new significance to old facts. Thus, as I've said, one of the most important means of protecting our capacity for social self-authorship is,

(7) The process ought to remain open-ended.

As long as others are making decisions that will affect our lives, and as long as they are making them based on assumptions about who we are, they ought to give us the opportunity to contest

those assumptions. It ought never to be assumed that we have been understood in our totality, once and for all.

Having identified the norms which enable and protect social self-authorship, we can now imagine a situation in which they are perfectly realized, and we can ask if the parties in that situation are duly represented in the process of negotiating their public identities. In a situation where we know others are collecting information about us (1), what information they are collecting (2) , why they are collecting it (3), and how the information is being evaluated (4), we can say that social self-authorship is perfectly visible. We have at our disposal everything we need to know in order to calibrate our self-presentation to the situation at hand, to know what about ourselves to emphasize or de-emphasize, how to package certain information for optimal effect, which errors need to be corrected, and so on. As Goffman would say, we know our audience. And as long as we have the opportunity to correct misinformation about us (5), contextualize and offer our own interpretations of the information the other party has (6), and as a long as the process as a whole remains open to contestation and revision (7), we have the full capacity to shape the other party's understanding of who we are. We still can't *control* how they perceive us. The other party still gets to interpret the information we provide, to judge its veracity and weight its relevance, to decide if the information we provide is significant, and to see for themselves how it fits into the larger picture of us that they have. But under these conditions we have every opportunity to exert our influence.

Of course, no actual situation will ever perfectly realize this ideal. The world we live in is a messy and imperfect one. We don't have time to tell the people with whom we interact exactly what information about them we are noticing and attending to, what it tells us, or why we care.

We can't always ask for their input before acting on the information about them that we have. If we think, though, that the scenario above accurately captures what it means for the process in which we negotiate our public identities to be fair, for each party to be duly represented, and for the process to thus protect each party's social and political agency, then we can use it as a regulative ideal, as a measuring stick for gauging the relative fairness of any real situation we find ourselves in. We can ask about such situations, to what extent do they meet the normative criteria laid out above? And if we want to improve them, to make them more protective of hermeneutic privacy, we can ask how they can be made to better approximate our ideal.

I argued in chapter 2 that information technology makes social self-authorship invisible and unnecessary. We can now spell out what that means in more detail. To make social self-authorship invisible is to create a situation in which norms (1) - (4), above, are not met. It is to create a situation in which we don't know that information is being collected about us, what information is being collected, why it is being collected, or how it is being evaluated. To say that information technology has made social self-authorship unnecessary is to say that it has created a world in which people are able to form opinions about who we are without asking for our input —i.e., without meeting requirements (5) and (6).

I also claimed that while these problems wrought by information technology are deeper and more insidious than the problem of diminishing information control, they, unlike the control problem, are problems we might actually solve.[92] We are now in a position to make good on that claim. In what remains of this chapter, I describe how the abstract norms enumerated above might actually be implemented. First, I discuss how hermeneutic privacy can be protected

[92] As I argued earlier, regaining control over information about ourselves would require radically de-technologizing our lives. And even then, the control we might regain would not be particularly robust.

through careful and concerned technology design. Second, I examine the implications of the above discussion for technology education. These two sections show what hermeneutic privacy might look like in practice, and how its demands differ from those of control models. Finally, while a full discussion of what the above view means for privacy and technology law is outside the scope of this project, I end by gesturing at its legal implications.

4.2. Hermeneutic Privacy by Design

In *The Whale and the Reactor* (1986), Langdon Winner argues that technologies aren't value-neutral, as is often assumed. Values can be embedded in technologies, intentionally or unintentionally, and those values structure the way the technologies are used. "If the experience of modern society shows us anything," he says, "it is that technologies are not merely aids to human activity, but also powerful forces acting to reshape that activity and its meaning" (6). Winner's argument is that we don't just use technologies as means to ends we've chosen in advance; the possibilities for action that different technologies open up to us, the affordances they provide, the activities they suggest we ought to engage in or refrain from, imply ends or purposes which we may or may not realize we are working toward. For this reason, when we design new technologies we are doing more than solving problems—we are, as he says, "making worlds" (11). Such world-making comes with certain responsibilities. If the tools we create aren't value-neutral, then we have to think about and assess the values we're building into them. It is especially important that we do this, Winner notes, early in the process of creating a new technology, since once it is adopted and becomes a regular part of its users' lives it is very difficult to change it. "By far the greatest latitude of choice exists the very first time a particular

instrument, system, or technique is introduced. Because choices tend to become strongly fixed in material equipment, economic investment, and social habit, the original flexibility vanishes for all practical purposes once the initial commitments are made" (29).

Those concerned about privacy in the Information Age have realized this too. In the early 1990s, a number of privacy advocates began arguing that technologists ought to build respect for information privacy into the very technologies they created—an approach known as "Privacy by Design" (PbD).[93] The idea is that privacy shouldn't be an afterthought, something which is worried about and tacked on at the end of the design process. If technologists are building values into their products, then privacy ought to be one of them. Every stage of developing new technologies should be guided by a set of agreed-upon privacy principles, PbD advocates argue. What's more, such principles shouldn't only be used to guide the development of individual, concrete tools, but should apply equally to "organizational practices and processes, and to broader information eco-systems and architectures" (Cavoukian 2012, 170).

PbD is thus a commitment to and a set of strategies for designing technologies to comply, "from the ground up," with privacy norms. The question this raises is what exactly those privacy norms are. In the US and elsewhere, the principles that privacy advocates generally recommend are the Fair Information Practices, or FIPs. These are information privacy norms developed originally in a 1973 report by the US Department of Health, Education, and Welfare (HEW), titled *Records, Computers, and the Rights of Citizens*, which set out to address "the problems

[93] For an overview of the history and development of Privacy by Design, see Hustinx (2010), Cavoukian (2009), and Cavoukian (2012).

arising from the application of computer technology to record keeping."[94] The HEW report made five broad recommendations:

- There must be no personal data record-keeping systems whose very existence is secret.
- There must be a way for an individual to find out what information about him is in a record and how it is used.
- There must be a way for an individual to prevent information about him that was obtained for one purpose from being used or made available for other purposes without his consent.
- There must be a way for an individual to correct or amend a record of identifiable information about him.
- Any organization creating, maintaining, using, or disseminating records of identifiable personal data must assure the reliability of the data for their intended use and must take precautions to prevent misuse of the data. (xx-xxi)

Following the report's widespread adoption, these recommendations became the basis for both US information privacy law, in the form of the Privacy Act of 1974, and international information privacy norms, agreed upon by the Organization for Economic Cooperation and Development (OECD) in its 1980 report, *Guidelines on the Protection of Privacy and Transborder Flows of Personal Data* (Dixon 2006; Cavoukian 2012). There is, of course, considerable disagreement about whether or not the Fair Information Practices have been a useful tool for safeguarding information privacy.[95] There is no disagreement, however, about the fact that they represent an international consensus among government and inter-governmental regulators regarding what protecting information privacy means. The FIPs can therefore serve as a touchpoint for our examination into whether and how things ought to change in order to protect hermeneutic privacy.

---

[94] A scanned copy of the original report is available here: http://www.justice.gov/opcl/docs/rec-com-rights.pdf

[95] For representative critiques of the both PbD and the FIPs, see Rubinstein (2011) and Mulligan and King (2014).

The first thing to notice about the FIPs is that they share much in common with the norms I presented in the previous section. Specifically, they agree with (1) that we ought to know when others are collecting information about us, (2) that we ought to know what information they are collecting, and (3) that we ought to know why—for what purpose—the information is being collected. It should come as no surprise that the view I have been arguing for and the one codified in the FIPs have this overlap, since the above requirements are a pre-condition for *any* approach to protecting information privacy: taking an informed approach toward something requires knowing what one is approaching. What's missing from the FIPs, in relation to my view, is what I presented as requirement (4) that we ought to know how the information about us is being evaluated, and requirement (6) that we ought to be able to contextualize and offer our own interpretations of the information held about us (by providing more information).

The Fair Information Practices are focused singularly on *the information itself*—its existence, who has access to it, how it's stored, whether or not it's true. They are focused, in other words, on giving individuals control over information about themselves. The conceptual foundation of the FIPs is the control theory of privacy (Mulligan and King 2012, 993). They pay no attention at all to what the information *means* to its possessors, what it says about the real people it identifies, nor do they pay sufficient attention to what sorts of decisions are made on account of its interpretation. In a world where vast amounts of information are collected about us each moment of every day, focusing on the information itself is a dead-end. We lack the resources, the time, to worry about each individual piece of information about us—who has it, whether or not it's secure, and so on.

By the same token, the design principles that this approach recommends, while admirable, are inadequate. The FIPs suggest that new technologies be built in such a way that

minimizes the amount of data collected. They recommend that users be given notice that information is being collected about them and have the opportunity to consent. They advocate robust data security, and mechanisms by which users can correct errors (Wang and Kobsa 2008). All of this is to the good. It is insufficient for protecting hermeneutic privacy, however, because it doesn't provide users with any way of understanding how the information that is being collected about them is interpreted or evaluated. And although it recommends that users be able to correct misinformation about them, it fails to recognize that true information about someone can tell more than one story about them, and that users therefore ought to be able to contextualize information about themselves and offer their own alternative interpretations of it.

We can do that, we can design principles of hermeneutic privacy into our technologies and information systems, not by informing someone each time a piece of information about them is collected and stored, but rather by giving them a running image of the impressions the technologies have of them. We can design technologies to give users a sort of gestalt view of what the information they have about them *means* to the system, and how it might affect decisions made about them. An institution that has already begun to move in this direction is one we've seen before: credit rating.[96] It used to be the case that no one had any idea what went into their "FICO" scores. They were, as Frank Pasquale puts it, "black boxes." But thanks to the introduction of the Fair Credit Reporting Act and other institutional changes, we now have a rough sense of what goes into our scores, and the credit rating agencies are required to let us see them without penalty.

Now imagine if this model were advanced further. Instead of just providing us with dry reports of the information collected about us—bank statements, loan statements, bills paid and

[96] I discuss credit scoring as an example of the way we negotiate our public identities with businesses and other organizations in chapter 3.

payments missed—the credit rating agencies could occasionally inform us when something we did affected our scores. A missed payment could trigger a notification that one's score had decreased. A year of on-time payments could trigger a notification that one's score had improved. Or to make this idea less intrusive and potentially annoying, imagine that we received monthly emails that summarized all of things we did that helped or hurt our scores. In the case of credit scoring, individual data points like missed or made payments might not singularly affect one's score anyway. They are factored into more complex calculations. A monthly digest approach would thus be of more use to us in negotiating our public identities than detailed reports about which information was collected, since they would inform us about how the credit rating agencies were perceiving and understanding us on the whole.

Furthermore, in addition to being able to dispute facts about ourselves, such as whether or not we really did miss a payment, we could design credit rating tools that allowed us to contextualize that information and suggest alternative interpretations of it. It might be true that someone missed a payment. But it might also be true, as in the example I gave in the previous chapter, that he missed the payment because he was hit by a car on the way to mail it in. As things stand, there is no way for someone in that situation to explain to the credit rating agencies that the missed payment doesn't actually indicate anything significant about the likelihood that they will miss more payments in the future. Yet we could imagine a scenario in which the credit rating agencies had avenues through which we could not only contest but contextualize information about ourselves.

To take another example we've seen before, consider the issue of "price customization," discussed in chapter 2. Companies like Amazon and Orbitz use the information they collect or purchase about their customers to make a range of decisions about how to structure their

customers' user experiences. Their websites recommend different products and product options, different deals and coupons, and as we've seen, they sometimes even offer their products at different prices, depending on who the customer is and what the company thinks it knows about them. Even if customers are notified that such websites are collecting information about them, indeed even if they are notified about the *purpose* of collecting that data, that knowledge isn't useful to them unless they know what the information *means* to the company. If Orbitz were to ask its users if it could collect information about the types of computers they use, in order to help "customize" their pricing, it likely would not be obvious to the customers how that information would be helpful.

But Orbitz could do things differently. It could display a message at the top of its search results page, which said something to the effect of:

> *Based on what we know about you, we think that you prefer high-end products, and that you are willing to spend a little bit more on luxury amenities. So we've selected the following hotel options for you to choose from. If you want to know why we think that about you, or if we've got you wrong, click here.*

If the customer chose to pursue more information about Orbitz's understanding of them, the site could display some of the customer data which led it to its conclusions, and could offer the customer the option of disputing Orbitz's interpretation of that data.

Of course, nothing would compel the website to adopt the user's interpretation. It would merely serve as more information about the customer for it to consider and calculate. What structuring things in this way would accomplish, however, is that it would turn the process of learning about who the company's customers are back into a *negotiation*. Instead of merely surveilling its users, the company would promote a kind of dialogue. Users would have the opportunity to do the work of shaping their own public identities, to learn how they are being

understood by the company and to modify their self-presentation accordingly. And companies would in all likelihood end up with a net increase in meaningful information about their users.

I should point out that I realize some people will read the above imagined scenarios and find the kind of future I am advocating for to be utterly dread-inducing. *What is he thinking?*, they'll ask. *I do not want to live in a world where all the technologies I interact with are constantly judging me, making decisions about me, and interrupting me to ask if they've correctly interpreted the information that they've collected about me*. It is important, therefore, to remember, that save for the last part of that scenario—where the technologies ask for our input—we already live in that world. Our technologies are already collecting information about us, using that information to judge us, and making decisions based on those judgments that impact our lives. Only in the world as it is, our technologies are just so many Peeping Toms. They observe us quietly, collecting information without intruding. It seems clear to me that we would be better off in a world where we knew we were being watched and were in conversation with our watchers.

Yet in order for that conversation to be a meaningful and fair one, we have to be told not only that information is being collected about us, but also what that information means to those collecting it. This will only become increasingly important as new information processing technologies are developed, which make it possible to extract ever more meaning from the masses of seemingly trivial data collected about us. Technologies like the computer processors and algorithms associated with the field of Big Data, which I discuss in the conclusion, already make it possible to discern correlations and patterns in the minutiae of our everyday comings and goings, calls, purchases, likes, and dislikes, and so on. We have reached the point where we can (and ought to) safely assume that more information is being collected about us at each moment

than we could possibly bring under our own control. And there is no way of knowing in advance what this or that piece of information about ourselves might mean to those who collect it.

Protecting hermeneutic privacy in the face of these challenges will require building technologies that include us in the process of getting to know us, that give us the opportunity to author our own social selves. It will also require that we be competent authors. If people lack a basic understanding of the tools they use and the systems in which they are implicated, then making those systems open to us and inclusive of our input will only get us so far. To protect and promote social self-authorship in the Information Age we will have to do more than build the principles of hermeneutic privacy into our technologies and information systems. We will have to educate people about how they work.

4.3. Privacy and Technology Literacy

Throughout this dissertation I have referred to social self-authorship as a *negotiation.* I have argued that hermeneutic privacy demands we be given access to the negotiation, that we be given the opportunity to influence how others perceive and understand who we are. And in the previous section, I suggested how technologies and technological systems could be designed to provide such opportunities. Having the opportunity to participate in the negotiation—having, as they say, a seat at the table—is not, however, sufficient for successfully authoring our social selves. We also have to be competent negotiators. Public policy aimed at promoting privacy in the Information Age must therefore include a technology education component.

It's often assumed that children are naturally technologically competent. "Digital natives," they're called, by contrast with adults who are mere "digital immigrants."[97] Putting

[97] See Barlow (1996) and Prensky (2001), cited in boyd (2014).

aside the problematic language of natives and immigrants, danah boyd argues that the assumption isn't even correct.[98] While it's true that those who have grown up entirely in a world of information technology are more comfortable with it, understand themselves in relation to it more intuitively, and so on, it isn't true that all such people are capable computer users.[99] While I began this chapter by arguing that many young people have a better intuitive understanding of the new privacy landscape than many adults, that should not therefore be taken to mean that young people are necessarily more technologically competent. Technology literacy and competence vary widely both within and across age groups, and track primarily with computer experience and access.[100] As boyd puts it, "there is no magical relation between skills and age" (2014, 177). Despite the fact that many teens are well situated to understand how privacy ought to be conceptualized in the Information Age, it isn't necessarily the case that they all have the tools or know-how to enact it.[101]

Interestingly, this may be attributable in part to how easy-to-use much personal technology has become. boyd points out that teens growing up in the 1980s and 90s, just as the internet became widely accessible, were often goaded into learning technology skills by their frustration with the tools on offer:

> In the early days of MySpace's popularity, a few teens learned that they could modify the looks and feel of their profiles by inserting code in the form of HTML, CSS, or JavaScript. This was the result of a bug in MySpace's development code. After watching teens explore self-expression through code, the company decided

---

[98] To her credit, boyd also flags the problematic language.

[99] See chapter 7 in boyd (2014).

[100] Experience and access in turn track significantly with race and class. See boyd (2014), p. 192-6. See also Hargittai (2008).

[101] "Many of today's teens are indeed deeply engaged with social media, but this does not mean that they inherently have the knowledge or skills to make the most of their online experiences" (boyd 2014, 176).

> not to patch the bug in order to see how users would personalize their pages. Excited by the ability to create "layouts" and "backgrounds," teens started learning enough code to modify their profiles. Some teens became quite sophisticated technically as they sought to build extensive, creative profiles. (boyd 2014, 182)

Indeed, this echoes my own experience. While I wasn't active on MySpace, I learned basic web design and programming skills, like many in my generation, in order to build GeoCities and Tripod webpages. I taught myself how to use graphics programs like Adobe Photoshop so that I could personalize them. I took programming classes in high school and over summers, and I learned how to build computers from parts by taking my computers apart and putting them back together again.

Now, hardly two decades later, the general movement in technology design and production is toward concealing from users the "guts" of the tools. This is both because technology companies want to provide their customers with seamless and streamlined user experiences, and also because they want to maintain tight control over their intellectual property. New computers today are very difficult to take apart, and doing so frequently voids the manufacturer's warranty. Software EULAs prohibit tinkering with the code. And social networking sites like Facebook and Twitter thwart any attempts to alter or customize their products in ways that aren't explicitly sanctioned. As a result, kids aren't given as many incentives to learn how the technologies they use work. Unless they happen to have an aptitude for math and science and are exposed to computer programming directly, they aren't as likely to find their way to those skills themselves. "When technologies are designed to make everyday use as easy as possible," writes boyd, "it is not necessary for users to learn the technical skills that early internet adoption required. Although it is not necessary to be technically literate to

participate, those with limited technical literacy aren't necessarily equipped to be powerful citizens of the digital world" (2014, 183).

The situation may in fact be even more pressing than boyd suggests. If what I have been arguing about information technology and social self-authorship is true, then the issue isn't merely one of being a *powerful* citizen of the digital world; the issue is whether or not one is capable of exerting one's agency online at all. Doing so requires being able to create and maintain public identities, and in the Information Age that means understanding how one's public identities are constituted in and through information technology. Even if all of the technology design suggestions I offered above were implemented, one would still need to have a basic understanding of the social, informational, and technological contexts in which one is operating in order to be a capable agent in the digital world.

There are, I think, three domains of knowledge or types of literacy relevant to protecting one's privacy online. First, as I've been suggesting, one must have a rudimentary technical understanding of computers and information systems. Second, one also needs what is often called information or media literacy. Third, one needs to understand specifically what it means to value privacy—what we might call privacy literacy. Roughly speaking, privacy literacy is an understanding of why one should care about protecting one's privacy and what it looks like to do so; media and information literacy provide an understanding of when—in what circumstances—one's privacy needs protecting; and technical competence or computer literacy involves the skills or know-how required to realize one's privacy in such circumstances. I discuss each briefly in turn.

We think it is important for children to learn about "how the world works." So, from a very young age we teach them about the abstract and concrete structures of reality. Not every child will grow up to be a mathematician or geologist, yet we nevertheless think that they should all have rudimentary math skills and a basic understanding of how mountains form. It is an odd thing, then, that the same logic has not been applied to the digital world. Not every child will grow up to be a programmer or systems administrator, but they will be at a disadvantage if they don't even know what programmers and sysadmins do. They should know what, at bottom, a computer is, how it works, and what its limits are. They should learn about what the internet is, concretely and materially, and how it is different from websites and the World Wide Web. They should learn what it means to "put something in the cloud." They should learn how an email travels from one computer or smartphone to another, what happens to your credit card information when you swipe it at the vending machine, what it means to encrypt data. We live in a world profoundly and increasingly structured by digital technologies. One can't know how our world works without knowing, at least at a basic level, how computers do.

Of course, I am not the first to suggest this. Organizations have long recognized this gap in basic education and have started to try and fill it. Non-profit organizations like Khan Academy[102] and Girls Who Code[103] offer free, online, introductory lessons in computer programming. Codeacademy says in its mission statement that "Education is broken. Come help us build the education the world deserves."[104] Even President Obama, in a now-famous 2013 speech, pointed out that the vast majority of US states do not allow high school computer science

[102] http://www.khanacademy.org

[103] http://www.girlswhocode.com

[104] http://www.codecademy.com/about

courses to fulfill math or science graduation requirements.[105] And Advanced Placement (AP) Computer Science is only taught in 10% of American high schools.[106] "Don't just play on your phone," Obama urged, "Program it."[107]

I don't think we should interpret the President's message to mean that everyone needs to become *proficient* at computer programming. Rather, kids ought to have some exposure to it. They should get a sense for the kind of work that goes into building the digital world they move through so fluidly, the kinds of problems constructing that world poses, the trade-offs one has to make in constructing it, and the values implicit in making them. They should learn, in other words, what Tasneem Raja and others call *computational thinking.*[108] Like scientific or mathematical thinking, computational thinking is a kind of reasoning—the kind required to solve problems algorithmically:

> Much like cooking, computational thinking begins with a feat of imagination, the ability to envision how digitized information—ticket sales, customer addresses, the temperature in your fridge, the sequence of events to start a car engine, anything that can be sorted, counted, or tracked—could be combined and changed into something new by applying various computational techniques. From there, it's all about "decomposing" big tasks into a logical series of smaller steps, just like a recipe.
>
> Those techniques include a lot of testing along the way to make sure things are working. The culinary principle of *mise en place* is akin to the computational principle of sorting: organize your data first, and you'll cut down on search time later. Abstraction is like the concept of "mother sauces" in French cooking (béchamel, tomato, hollandaise), building blocks to develop and reuse in hundreds of dishes. There's iteration: running a process over and over until you get a desired result. The principle of parallel processing makes use of all available

[105] See Sheehy (2012).

[106] See Wagstaff (2012).

[107] https://www.whitehouse.gov/blog/2013/12/09/don-t-just-play-your-phone-program-it

[108] In addition to Raja (2014), see Chalmers and Watts (2014), Cuthbertson (2014), and Paul (2014).

> downtime (think: making the salad while the roast is cooking). Like a good recipe, good software is really clear about what you can tweak and what you can't. It's explicit. Computers don't get nuance; they need everything spelled out for them. (Raja 2014)

These metaphors are glosses, to be sure, but they are powerful. Understanding that "computers don't get nuance," and more importantly *why* they don't, goes a long way toward understanding why a computer system doesn't always produce the results one expects, why it can't solve every problem, and why it can only solve the problems it does in a certain way. "The happy truth," says Raja, "is, if you get the fundamentals about how computers think, and how humans can talk to them in a language the machines understand, you can imagine a project that a computer could do, and discuss it in a way that will make sense to an actual programmer" (2014).

Again, none of this is new. People have long argued that the American education system is woefully lacking in the technology department. What is important for our purposes is that the knowledge and skills described above are crucial for creating and protecting information privacy. In the second epigraph to this chapter, Jeffrey Rosen writes, "We need more shades and more blinds and more virtual curtains" (2000, 224). The principles I put forward above for guiding technology design have to do with how we can build a digital world with more shades and more curtains. But we also have to know enough about how our digital house works to be able to draw them shut.[109]

---

[109] This is perhaps a strained metaphor. Rosen operates on a control theory of privacy, so by building "virtual curtains" he means building ways of operating online without being subject to surveillance and data collection. That, as I've said, is an admirable goal. But, I think, it's an implausible one. By adopting his metaphor I mean that it's true that we should build privacy principles into our technologies, as described in the previous section. The principles I've put forward are not about obscuring us online, but rather about making the processes we're implicated in open to our participation. When I say here that we have to know enough about how our digital house works in order to draw the virtual curtains shut, I mean that we need to understand our technologies sufficiently well in order to take advantage of the privacy mechanisms we build into them.

In addition to basic *computer literacy*, we also need what scholars and activists call *media* or *information literacy*. Instead of dealing with how computers work or the kind of reasoning involved in programming them (the *architecture* of the digital world), information literacy deals with how we treat the information the digital world confronts us with. It is "the ability to access, analyze, evaluate, and create messages in a variety of forms" (Aufderheide 1993, qtd. in Livingstone 2004). Or, as information theorist Michael B. Eisenberg writes, information literacy "is the set of skills and knowledge that not only allows us to find, evaluate, and use the information we need, but perhaps more importantly, allows us to filter out the information we don't need" (2008, 40). Being information literate means being able to find specific information one needs, having the tools to determine its provenance and credibility, and understanding how it is situated in greater informational contexts. Furthermore, information literacy involves having the skills to create new content, to determine the best or most impactful form that content can take, to distribute it to all and only those who ought to see it, and to understand the potential consequences of others gaining access to it.

Like computer literacy, information literacy is essential for creating and protecting information privacy. In order to successfully negotiate our public identities with others, to exert influence over the way they perceive and understand who we are, we have to be at least minimally savvy about how the information we gather and the information about ourselves that we produce is situated in greater informational contexts. We have to know whether or not the information we are relying on as a guide for how we present ourselves is reliable, and we have to know how the information we put out into the world will appear to those who access it. Information and education theorists have produced a variety of frameworks for teaching these

skills[110], and it is already commonplace for universities to identify information literacy as a teaching objective in college-level curricula. Like computer literacy, though, it should be something which is taught to students from the very beginning of formal education, so that skills and knowledge-sets can be refined and scaffolded as students mature and grow.

Finally, once one knows how computers and computer systems work, and one knows how to find, evaluate, and create information using them, there is still one domain of knowledge or know-how left that is necessary for competent social self-authorship in the Information Age. This is a form of know-how which is usually passed down culturally, often without explicit mention. What we might call *privacy literacy* is inculcated in children when parents instruct them to close the blinds before going to sleep, to speak quietly when discussing sensitive matter in public, to be careful not to leave important documents out on the table when visitors are coming over, and so on. As the discussion of danah boyd's work at the beginning of this chapter shows, children for the most part understand the value of privacy, and many intuitively understand, better than most adults, how actualizing that value in the Information Age might look different that it once did. Still, as boyd points out, children are for the most part concerned with drawing epistemic boundaries between themselves and their parents, their teachers, and other authority figures who are present in their everyday life. They need to be taught to worry about negotiating their public identities with people and organizations who are less present, like those discussed in chapter 3.

---

[110] Some prominent models include the "Big6" approach, Carol Kulthau's "information search process, the AASL/ AECT IL Standards, and the ACRL IL Competency Standards for Higher Education. For an overview of the differences and commonalities between these various approaches, see Eisenberg (2008).

One scholar who has recognized this need is sociologist Eszter Hargittai, who teaches a course at Northwestern University called "Managing Your Online Reputation." Having discovered in her own research that, as boyd also argues, assumptions about young people being tech-savvy "digital natives" are false[111], Hargittai and a colleague, Brayden King, designed a course to teach college students how to think about and constructively engage in what I have been calling social self-authorship. Prior to the start of the semester, Hargittai researches her students online, and when they arrive on the first day of class, she presents them with her findings. For many students, this exercise immediately reorients the way they understand how the information they put out into the world is situated. Absent proper safeguards, information intended for friends and other college students is available for any and all to see. Over the course of the semester, students are required to:

> search for themselves on various search engines, including Bing, DuckDuckGo, and Dogpile, and take note of what the Internet turns up. Other assignments include creating profiles and platforms such as Google Plus and Tumblr, and engaging with people of professional interest on Twitter. Each student must create and present to the class a brief case study of someone who has successfully managed and benefited from a positive online reputation. (O'Neil 2014)

These exercises provide what I call privacy literacy. They teach students how information about themselves is situated with respect to the people, organizations, businesses, and institutions, which often remain in the background of our our everyday lives, but nonetheless have enormous impacts on us. It teaches them that they have to draw epistemic boundaries not only between themselves and their parents, but also between them and all of the other actors and groups with whom one interacts in contemporary life. It trains students "to build robust, productive online

[111] See Hargittai (2008).

identities through which they can engage topics of interest, command audiences, and advance their careers" (O'Neil 2014).

Hermeneutic privacy demands not only that we be able to participate in the processes by which others come to perceive and understand who we are, but also that we are capable participants. Thus, in addition to designing and building information technology in a way that is inclusive of users' input into how information about them should be contextualized and interpreted, users must be empowered to offer that input effectively. Such empowerment requires computer, information, and privacy literacies, and adequate privacy policy must therefore involve a means of providing them.

4.4. Information Privacy Law

Finally, what remains to be discussed are the implications of the above arguments for information privacy law. Although a full treatment of the legal ramifications of hermeneutic privacy is outside the scope of this project, a few preliminary gestures are warranted. I argued above that the consent model, which is currently the normative foundation of American information privacy law, is not up to the task of safeguarding our interests in being able to shape how others perceive and understand who we are. Something must be said about what ought to supplement it or take its place. Furthermore, a noteworthy implication of the view I have put forward, which up to this point I have neglected to address, is the fact that protecting information privacy in the ways I have suggested would require demanding positive obligations of others. Since privacy is usually thought to be a quintessentially negative right, this implication of my view bears discussion.

American privacy law, in contrast to the privacy statutes of Canada and the European Union, is almost singularly focused on protecting citizens from the reach of government, rather than from the reach of other citizens and private sector organizations.[112] When it comes to dealing with the federal government, we enjoy a long list of protections regarding what it is allowed to do with information about us. In addition to our well-known Fourth Amendment protections against illegal search and seizure, the Privacy Act of 1974 requires that any federal government agency collecting information about US citizens adhere to the fair information practices, or FIPs, described above.[113] Such practices include: publicly revealing the existence of databases containing information about US citizens, only collecting as much data as is needed, allowing citizens to review the information collected about them and correct mistakes, requiring the government to secure the data collected, limiting its right to disclose any information collected, and so on.[114] In other words, apart from cases where there are competing and overriding interests (such as law enforcement and national security), the US government is legally barred from surveilling US citizens in secret.

What's more, if the government concludes that someone has done something wrong, or that they are about to do something wrong, based on information it has collected about them, it cannot simply act on that information with no further worry about its veracity. The Fifth and Fourteenth Amendments guarantee US citizens the right to due process. If the government intends to coerce us in some way, we are entitled to a trial, during which we have the right to be confronted with the evidence against us, to challenge its accuracy, to offer our own

---

[112] See Levin and Nicholson (2005) and DeVries (2003).

[113] 5 U.S.C § 552a (2000).

[114] See https://epic.org/privacy/1974act/

interpretations of the facts, and to have the government's case and our own adjudicated impartially.[115] Which is to say, the government is required to include us in the process by which it perceives and understands who we are, before it can make any decisions that will impact our lives. As Ian Kerr and Jessica Earle write, "privacy and due process values seek to limit what the government […] is permitted to presume about individuals absent evidence that is tested in the individuals' presence, with their participation. As such, these values aim to provide fair and equal treatment to all by setting boundaries around the kinds of assumptions that can and cannot be made about people" (2013, 70-71).

With respect to the government, then, I think American privacy law is more or less sufficiently protective of our hermeneutic privacy interests.[116] The problem is that, unlike in Canada and the EU, the protections described above do not extend to our dealings with the private sector (DeVries 2003). The FIPs delineated in the Privacy Act of 1974 are mandatory only for Federal Government agencies. For state and local governments, and for private businesses, the Fair Information Practices are merely suggestions. As long as private organizations obtain information about us consentingly, or more commonly nowadays, if they obtain it from a third party, there is nothing in US privacy law which prevents them from interpreting that information any way they want, nor from using those interpretations as the basis for making decisions that impact our lives. As I hope to have shown convincingly in the previous chapter, there is a great deal at stake in our dealings with businesses and other actors in the

[115] See Friendly (1975).

[116] At least in theory, as a general framework. As many have noted, the laws described above are "riddled with exceptions" (DeVries 2003). What's more, it is not at all clear that the US government actually abides by US privacy laws, especially since 9/11. Unfortunately, treating these more concrete, contingent issues having to do with the actual application of US privacy law is outside the scope of this project.

private sector. And as I hope to have shown above, obtaining our consent in order to collect information about us is not an adequate mechanism for legally protecting our privacy.

The simplest solution would obviously be to extend the scope of privacy and due process law to cover the private sector. This is the approach that Canada and the EU have taken. As Avner Levin and Mary Jo Nicholson write, "Unlike the US, the EU imposes controls over business processing and use of personal data, both before and after the data are collected. […It] grants individuals the right to challenge any decision significantly affecting them that is based on an automatic processing of data, including decisions involving credit worthiness or employment" (2005, 376). Recognizing that there is just as much opportunity for coercion and abuse of power in our dealings with others in the private realm as there is in our dealings with the government, Canada and the European Union have acted to ensure that the processes through which individuals negotiate their identities with businesses and other private sector organizations are bound by the same rules of fairness and due representation as those that bind the processes through which they negotiate them with the government. The United States would do well to follow suit.

Doing so, however, would raise at least one problem. It would institute positive obligations on the part of businesses and other private organizations, requiring them to create organizational and technological means for their customers and users to participate in the processes through which they perceive and understand them. Rather than merely refraining from collecting or using information about someone without their consent, it would demand that they build certain features into their tools and expend resources to gather and incorporate users' own accounts of how information about them ought to be contextualized and interpreted. Since

privacy is normally thought of as a quintessentially negative right[117]—one which imposes on others the obligation only to refrain from doing certain things, like prying—the suggestion that American privacy law ought to impose positive obligations on private businesses and other organizations might strike some as odd.

While I can't treat this issue fully here, there are two things worth saying about it. First, despite its usual characterization as a negative right, commonsense intuitions about what privacy demands undoubtedly include a range of positive obligations, apart from those described above. One can see hints of those obligations in the very language used to describe expectations of privacy, such as when we say that we are required to "give" someone their privacy. This idea is clearest in cases where privacy is owed in public. If I walk down the street and a man's pants fall down or a woman's dress billows up, most would say that I am obligated to look away. Goffman calls this the demand of "civil inattention."[118] Or to take a technology-related example, consider the recent incident in which hackers stole nude photographs of a number of celebrities from cloud storage and released them to the public.[119] Not only do the privacy rights of those celebrities demand that we refrain from seeking the photos out, should we happen to find ourselves in possession of them (say, because a friend emailed them to us), their right to privacy would seem to demand that we proactively do the work of ridding ourselves of them. The notion that the right to privacy creates positive obligations on others, though often neglected, is thus not totally foreign to our commonsense intuitions about it.

---

[117] See for instance Velasquez, Andre, Shanks, and Meyer (1990).

[118] See Goffman (1971).

[119] See Dewey (2014).

Still, some who accept all of that will nonetheless reject the specific hermeneutic privacy-related obligations I've outlined. We may have positive privacy obligations, they'll say, but new ones can't simply arise out of nowhere, just because we have developed new technologies. To this I would respond that they can and they must, and indeed, that they have in the past. As Judith DeCew notes, Warren and Brandeis themselves wrote their famous "The Right to Privacy" (1890) in response to the development of then-new technologies, which they believed threatened privacy (DeCew 1997, 16). At the turn of the 20th Century, those technologies had to do with high-speed photography and mass-distributed newspapers. Warren and Brandeis worried that such technologies made previously private information accessible to the public at large, and thus achieving privacy in this new circumstance required creating new protections. "Instantaneous photographs and newspaper enterprise have invaded the sacred precincts of private and domestic life," they wrote, "and numerous mechanical devices threaten to make good the prediction that 'what is whispered in the closet shall be proclaimed from the house-tops'" (1890, 195). In their article, they endeavored to show that we needed new means for realizing privacy, and that the law was adaptable enough to offer such means. Moreover, they recognized that we would almost certainly face this situation again. Warren and Brandeis, writes DeCew, "pointed out that the right to privacy was incapable of being given an exhaustive and wholly accurate definition; it would instead be worked out through opinions in a vast number of cases. That process would be possible, they felt, due to the law's elasticity, capacity for growth, and adaptability to new conditions in the face of modern devices" (1997, 16).[120]

[120] See also Richards and Solove (2010).

And that is precisely what happened. As Richard Posner points out, the trend—at least in American jurisprudence—has been and continues to be the expansion of legal privacy protections (1978). Today, as our lives are becoming ever more enmeshed in information technology, and as the vast majority of that technology is owned and controlled by private interests, the new privacy protections we need are in the private realm. In order to have agency over the important decisions that impact our lives, business and other private organizations will have to create ways for us to participate in the decision-making processes. If our public identities are going to be ours to author, we'll need robust privacy laws that obligate others to treat them that way.

References

Abril, Patricia Sanchez. 2007. "A (My)Space of One's Own: On Privacy and Online Social Networks." *Northwestern Journal of Technology and Intellectual Property* 6 (1): 73-88.

Alderman, Ellen, and Caroline Kennedy. 1995. *The Right to Privacy.* New York: Alfred A. Knopf, Inc.

Allen, Anita. 2011. *Unpopular Privacy: What Must We Hide?* Oxford: Oxford University Press.

Aufderheide, Patricia. 1993. *Media Literacy: A Report of the National Leadership Conference on Media Literacy.* Aspen, CO: Aspen Institute.

Barlow, John Perry. 1996. "A Declaration of the Independence of Cyberspace." https://projects.eff.org/~barlow/Declaration-Final.html.

Benn, Stanley I. 1984. "Privacy, Freedom, and Respect for Persons." In *Philosophical Dimensions of Privacy: An Anthology,* ed. Ferdinand Schoeman. Cambridge: Cambridge University Press.

boyd, danah. 2014. *It's Complicated: The Social Lives of Networked Teens.* New Haven: Yale University Press.

Cate, Fred H. 2006. “The Failure of Fair Information Practice Principles.” In *Consumer Protection in the Age of the ‘Information Economy’,* ed. Jane K. Winn. Hampshire, UK: Ashgate Publishing Limited.

Cavoukian, Ann. 2009. “Privacy by Design.” Report by the Office of the Information and Privacy Commissioner of Ontario. https://www.ipc.on.ca/images/Resources/privacybydesign.pdf

Cavoukian, Ann. 2012. “Privacy by Design: Origins, Meaning, and Prospects for Assuring Privacy and Trust in the Information Era.” In *Privacy Protection Measures and Technologies in Business Organizations: Aspects and Standards,* ed. George Yee. Hershey, PA: Information Science Reference.

Cavoukian, Ann, Stuart Shapiro, and R. Jason Cronk. 2014. “Privacy Engineering: Proactively Embedding Privacy, by Design.” Report by the Office of the Information and Privacy Commissioner of Ontario. https://www.privacybydesign.ca/content/uploads/2014/01/pbd-priv-engineering.pdf

Chalmers, Jim, and Tim Watts. 2014. “Kids Should Code: Why ‘Computational Thinking Needs to Be Taught in Schools.” *The Guardian,* December 18. http://www.theguardian.com/commentisfree/2014/dec/19/kids-should-code-why-computational-thinking-needs-to-be-taught-in-schools

Chapman, Cate. 2015. “Evaluating Privacy and Legal Concerns Related to Big Data.” *Cyber Risk Network,* January 15. http://www.cyberrisknetwork.com/2015/01/15/evaluating-privacy-legal-concerns-related-big-data/

Cohen, Elliot D. 2010. *Mass Surveillance and State Control.* New York: Palgrave Macmillan.

Cohen, Jean L. 1996. “Democracy, Difference, and the Right of Privacy.” In *Democracy and Difference: Contesting the Boundaries of the Political*, ed. Seyla Behnabib. Princeton, NJ: Princeton University Press.

Cuthbertson, Anthony. 2014. “Coding in the Classroom: Computational Thinking Will Allow Children to ‘Change the World.’” *International Business Times,* September 2. http://www.ibtimes.co.uk/coding-classroom-computational-thinking-will-allow-children-change-world-1463493

DeVries, Will Thomas. 2003. “Protecting Privacy in the Digital Age.” *Berkeley Technology Law Journal* 18 (1): 283-311.

Dewey, Caitlin. 2014. “A Comprehensive, Jargon-free Guide to the Celebrity Nude-photo Scandal and the Shadowy Web Sites Behind It.” *The Washington Post,* September 2. http://www.washingtonpost.com/news/the-intersect/wp/2014/09/02/a-comprehensive-jargon-free-guide-to-the-celebrity-nude-photo-scandal-and-the-shadowy-web-sites-behind-it/

Dixon, Pam. 2006. “A Brief Introduction to Fair Information Practices.” *World Privacy Forum,* June 5. https://www.worldprivacyforum.org/2008/01/report-a-brief-introduction-to-fair-information-practices/

Eisenberg, Michael B. 2008. “Information Literacy: Essential Skills for the Information Age.” *Journal of Library & Information Technology* 28 (2): 39-47.

Foucault, Michel. 1977. *Discipline and Punish: The Birth of the Prison.* New York: Vintage Books.

Friendly, Henry. 1975. “Some Kind of Hearing.” *University of Pennsylvania Law Review* 123: 1267-1317.

Goffman, Erving. 1971. *Relations in Public: Microstudies of the Public Order.* New York: Basic Books.

Hargittai, Eszter. 2008. “The Digital Reproduction of Inequality.” In *Social Stratification*, ed. David Grusky. Boulder, CO: Westview.

Hustinx, Peter. 2010. “Privacy by Design: Delivering the Promises.” *Identity in the Information Society* 3 (2): 253-255.

Johnson, Bobbie. 2010. “Privacy no longer a social norm, says Facebook founder.” *The Guardian,* January 10. http://www.theguardian.com/technology/2010/jan/11/facebook-privacy

Kerr, Ian, and Jessica Earle. 2013. “Prediction, Preemption, Presumption: How Big Data Threatens Big Picture Privacy.” *Stanford Law Review Online* 66: 65-72. http://www.stanfordlawreview.org/online/privacy-and-big-data/prediction-preemption-presumption

Lessig, Lawrence. 1999. “The Architecture of Privacy.” *Vanderbilt Journal of Entertainment Law and Practice* 1 (1): 56-65.

Levin, Avner, and Mary Jo Nicholson. 2005. “Privacy Law in the United States, the EU, and Canada: The Allure of the Middle Ground.” *University of Ottawa Law & Technology Journal* 2 (2): 357-395.

Livingstone, Sonia. 2004. “Media Literacy and the Challenge of New Information and Communication Technologies.” *Communication Review* 1 (7): 3-14.

Lyons, David. 2001. *Surveillance Society: Monitoring Everyday Life.* Buckingham, UK: Open University Press.

Mulligan, Deirdre, and Jennifer King. 2012. "Bridging the Gap Between Privacy and Design." *Journal of Constitutional Law* 14 (4): 989-1034.

Nissenbaum, Helen. 1997. "Toward an Approach to Privacy in Public: Challenges of Information Technology." *Ethics & Behavior* 7 (3): 207-219.

O'Neil, Megan. 2014. "Confronting the Myth of the 'Digital Native.'" *The Chronicle of Higher Education,* April 21. http://chronicle.com/article/Confronting-the-Myth-of-the/145949/

Organization for Economic Cooperation and Development. 1973. *OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data.* http://www.oecd.org/internet/ieconomy/oecdguidelinesontheprotectionofprivacyandtransborderflowsofpersonaldata.htm

Paul, Annie Murphy. 2014. "Interrupt Kids' Regular Scheduled Programming: Teaching Computer Science Without Touching a Computer." *Slate,* August 27. http://www.slate.com/articles/technology/future_tense/2014/08/computer_science_unplugged_teaching_computational_thinking_without_computers.html

Posner, Richard. 1978. "An Economic Theory of Privacy." *Regulation*. May/June Issue.

Prensky, Marc. 2001. "Digital Natives, Digital Immigrants." *On the Horizon* 9.http://www.marcprensky.com/writing/Prensky%20-%20Digital%20Natives,%20Digital%20Immigrants%20-%20-Part1.pdf

Prosser, William. 1960. "Privacy." *California Law Review* 48: 383–423.

Prosser, William. 1984. "Privacy: A Legal Analysis." In *Philosophical Dimensions of Privacy: An Anthology,* ed. Ferdinand Schoeman. Cambridge: Cambridge University Press.

Raja, Tasneem. 2014. "Is Coding the New Literacy? Why America's Schools Need to Train a New Generation of Hackers." *Mother Jones,* June 16. http://www.motherjones.com/print/253891

Reiman, Jeffrey. 1976. "Privacy, Intimacy, and Personhood." *Philosophy & Public Affairs* 6 (1): 26-44.

Richards, Neil, and Daniel Solove. 2010. "Prosser's Privacy Law: A Mixed Legacy." *California Law Review* 98 (6): 1887-1924.

Rosen, Jeffrey. 2000. *The Unwanted Gaze: The Destruction of Privacy in America.* New York: Vintage Books.

Rubinstein, Ira. 2011. “Regulating Privacy by Design.” *Berkeley Technology Law Journal* 26 (3): 1409-1456.

Schoeman, Ferdinand. 1994. “Gossip and Privacy.” In *Good Gossip,* ed. Robert F. Goodman and Aaron B. Ze’ev. Lawrence: University of Kansas Press.

Sheehy, Kelsey. 2012. “High Schools Not Meeting STEM Demand.” *US News & World Report,* October 1. http://www.usnews.com/education/blogs/high-school-notes/2012/10/01/high-schools-not-meeting-stem-demand

U.S. Department of Health, Education, and Welfare. 1973. *Records, Computers, and the Rights of Citizens.* DHEW Publication No. (OS) 73-94. http://www.justice.gov/opcl/docs/rec-com-rights.pdf

Velasquez, Manuel, Claire Andre, Thomas Shanks, and Michael Meyer. 1990. “Rights.” *Issues in Ethics* 3 (1).

Wagstaff, Keith. 2012. “Can We Fix Computer Science Education in America?” *Time,* July 16. http://techland.time.com/2012/07/16/can-we-fix-computer-science-education-in-america/

Wang, Yang, and Alfred Kobsa. 2009. “Privacy-Enhancing Technologies.” In *Handbook of Research on Social and Organizational Liabilities in Information Security,* eds. Manish Gupta and Raj Sharman. Hershey, PA: IGI Global.

Warren, Samuel, and Louis Brandeis. 1890. “The Right to Privacy.” *Harvard Law Review* 4 (5): 193-220.

Winner, Langdon. 1986. *The Whale and the Reactor: A Search for Limits in an Age of High Technology.* Chicago: The University of Chicago Press.

Zimmer, Michael. 2014. “Mark Zuckerberg’s theory of privacy.” *The Washington Post,* February 3. http://www.washingtonpost.com/lifestyle/style/mark-zuckerbergs-theory-of-privacy/2014/02/03/2c1d780a-8cea-11e3-95dd-36ff657a4dae_story.html

Conclusion

The Future of Identity

> "Speechless action would no longer be action because there would no longer be an actor, and the actor, the doer of deeds, is possible only if he is at the same time the speaker of words. The action he begins is humanly disclosed by the word, and though his deed can be perceived in its brute physical appearance without verbal accompaniment, it becomes relevant only through the spoken word in which he identifies himself as the actor, announcing what he does, has done, and intends to do. […] In acting and speaking, men show who they are, reveal actively their unique personal identities and thus make their appearance in the human world, while their physical identities appear without any activity of their own in the unique shape of the body and sound of the voice."
>
> - Hannah Arendt, *The Human Condition*

> "Big data's predictive benefits belie an important insight historically represented in the presumption of innocence and associated privacy and due process values—namely, that there is wisdom in setting boundaries around the kinds of assumptions that can and cannot be made about people."
>
> - Ian Kerr and Jessica Earle, "Prediction, Preemption, Presumption: How Big Data Threatens Big Picture Privacy"

The central argument of this dissertation has been that information privacy involves more than control over particular pieces of information, that it involves the entire process through which information about us is interpreted and acted upon. I argued that beyond the mere concealment and exposure of information, we work to contextualize and guide the interpretation of information about us, a process I've called social self-authorship. I argued that our capacity for social self-authorship is central to our ability to draw interpersonal boundaries, and that our ability to draw such boundaries is a necessary condition for social and political agency. Finally, I argued that in order to protect information privacy in the Information Age, we need to design new technologies in a way that empowers users to author their social selves, we need to help

people become technically competent enough to take advantage of those opportunities, and we need to use the law to ensure that the private sector respects privacy norms.

An important underlying theme of these arguments has been the need to shift from thinking about information as something we possess and (ideally) control, to thinking about information as a way of relating to other people, as the medium through which others come to know us. When information is understood as a possession, we think of our rights over it as rights of ownership and control. But if it's understood as the means for carrying out an important social and interpersonal process, then our rights and our interests have to do with being guaranteed fair and equal participation in that process. I want to conclude by illustrating why it's so important that we make this conceptual shift *now.* Up to this point, I've discussed technology mostly in generalities. I've described the kinds of information collected about us, the way it is processed, and the uses to which it is put. I now want to discuss a particular technology, or more precisely a particular field of technological development, which has enormous privacy implications. I want to talk about *big data.*

Big data is a concept from computer science, which has moved into the mainstream technology industry. Originally, it referred to datasets that were too large for existing computers and existing algorithms to sort through and extract meaningful information. Now we *have* computers that are sufficiently powerful and algorithms that are sufficiently complex to process those huge datasets, so Big Data has come to represent the subfield within computer science and the computer industry that works to extract as much meaningful information as possible from the depths of vast datasets that would otherwise have gone un-plumbed.

As I've described throughout the preceding chapters, nearly everything we do today, in the course of our day-to-day lives, leaves an enormous trail of data behind us. Using cell phones produces data about calls and text messages sent and received, apps used, websites accessed, and the location of the phone. Buying things with a credit or debit card produces data about what was purchased, for how much, when, where, and by whom. Watching Netflix or Hulu produces data about which shows or movies were watched and which were considered and ultimately passed over. Researchers estimate that by the year 2020 1.7 megabytes of data will be created for each human on the planet each second (Bansal 2014). That's the same as around 850 double-spaced pages of plain text *per person per second*.

It is on the basis of all this data that companies like Netflix and Hulu are able to predict which movies you'll want to watch next, how Amazon predicts which books you'll want to read, how companies like Target can figure out who is likely to get pregnant, and therefore ought to be sent maternity-related advertisements.[121] It is on the basis of all this data, as we've seen, that companies are able to "customize" the prices at which they offer the same goods to different customers, and the government is able to pinpoint suspected terrorists.[122] It is on the basis of all this data that researchers at the University of Rochester and Microsoft Research were able to predict where one of their research subjects would be 80 weeks into the future, with 80 percent accuracy (Tucker 2013).[123] It is also on the basis of all this data that many very welcome innovations are being made possible. Researchers are using cell phone location data in Ivory Coast to optimize bus routes (Tablot 2013). Google has used search results to predict outbreaks

---

[121] Much to some people's chagrin. See Duhigg (2012).

[122] See chapter 2.

[123] By analyzing GPS data.

of the flu.[124] And by analyzing tens of thousands of clinical records, drug regulators are discovering hidden risks in ordinary drugs and pulling them from the market.[125]

Yet incredibly, despite all that has been learned from the troves of data we produce, only one half of one percent of the data that currently exists has ever been analyzed (Bansal 2014). As governments, private corporations, and research institutions all continue to spend more money developing new big data tools, both the amount of data that is analyzed and the amount of meaningful information that can be extracted from it will continue to grow. It is impossible to predict now what we will soon be able to know, about ourselves, each other, and the world. "If the last century was marked by the ability to observe the interactions of physical matter," says a recent report on big data, "—think of technologies like x-ray and radar—this century […] is going to be defined by the ability to observe people through the data they share" (Regalado 2013).

There are several important things to note about this. First, the vast majority of the data that we produce and that is produced about us is about *trivial* things. It is information about what kind of computer we use, which website we visit to check the weather, how long we spend watching TV, and the time of day we purchased socks. Even if we could control information about ourselves, this is not the kind of information most people would spend time and effort trying to control. Second, a lot of the data being collected about us isn't being collected directly from us, but rather through what are called machine-to-machine transactions. In order to work properly, and to coordinate their functionality, the devices we use "talk" to each other by

---

[124] Though it has been argued that Google's predictions are not yet terribly accurate. See Lohr (2014).

[125] This is what happened, for instance, in the case of Vioxx. See Tene and Polonetsky (2012).

transmitting data between them. Our phones communicate with the Bluetooth systems in our cars, which communicate with event data recorders (EDRs) in the case of an accident. This "communication" means the transmission of data about the state of each device. In an increasingly networked world, the tools we use and rely on need to create and transmit data about us (or data from which information about us can be inferred) to provide us with the services we want.

Third, big data works by detecting often unpredictable patterns—sets of correlations—in enormous datasets. Which is to say, the businesses and research organizations using big data technology to analyze the sea of information available to them don't know in advance what they're going to find. What's more, the kinds of advanced artificial intelligence and machine learning techniques used to analyze big datasets often rely on algorithms which are "trained," rather than programmed in the conventional sense. Consequently, legal theorist Ira Rubinstein argues that "the newly discovered information is not only unintuitive and unpredictable, but also results from a fairly opaque process" (2013, 76). Opaque, that is, not only to end users, but also to the technologists who build the tools themselves.

Finally, and crucial for our purposes, is the fact that when it comes to big data research, the information being analyzed doesn't even have to be about *us* to say meaningful things about us. Much of the information that is collected about us is anonymized and aggregated with other information from people who share certain characteristics with us. The group data is then mined to see if there are any interesting correlations between membership in the group and some other characteristic. When businesses and other organizations want to make decisions about us, they are then able to use wholly non-private information about our group memberships to make

educated inferences as to what is likely true about us. Which is to say, they profile and stereotype us (Vedder 1999). Consider one well-known case, reported a few years ago in the New York *Times*:

> When an Atlanta man returned from his honeymoon, he found that his credit limit had been lowered to $3,800 from $10,800. The switch was not based on anything he had done but on aggregate data. A letter from the company told him, 'Other customers who have used their card at establishments where you recently shopped have a poor repayment history with American Express.' (Andrews 2012)[126]

In other words, big data makes it easy for others to make judgments and assumptions about us without even collecting information specifically about us. All they need are one or two data points about entirely public aspects of who we are—where we live or somewhere we've recently visited—and they're good to go. "In the end," says legal theorist Joseph Jerome, "the worry may not be so much about having information gathered about us, but rather being sorted into the wrong or disfavored bucket" (2013, 51).

To recap: in the world of big data, (1) the information collected about us mostly pertains to things we wouldn't normally care that others knew about us; (2) it is in significant part a byproduct of machine-to-machine transactions, rather than direct surveillance; (3) it is extremely difficult to predict what will be learned from analyzing it; and (4) the information need not even be *about us* to affect us. We are entering a world in which we stand to be defined by the minutia of our lives—the information we *give off,* as Goffman would say—rather than what we say about ourselves and what we set out intentionally to do. What's more, in this world we will be defined more by the aggregate interests, preferences, successes and faults of our demographic groups than by our own faults, strengths, and idiosyncrasies. In such a world, our public identities will

[126] Quoted in Jerome (2013).

no longer be ours to author. They will be assumed, inferred, predicted identities, *algorithmic* identities, understood not in terms of what matters to us, or even to society, but rather by what matters to whichever public or private entities happen to be analyzing information about us.

It isn't hard to see why control theories aren't up to the task of protecting our information privacy interests in this technological landscape. Even if we could control information about ourselves, we probably wouldn't bother to control most of *this* information, and even if we did it probably wouldn't do us much good. If we can't predict what kinds of inferences could be drawn from particular pieces of information about us, we can't know whether we ought to reveal or conceal them. And if decisions are going to be made about us on account of information about other people, what's the use of controlling information about ourselves?

Thinking about information privacy in terms of authorship, rather than control, offers a way out of these problems. First, it refocuses our attention away from particular pieces of information and toward the overall understanding others have about us. Second, it aims to ensure that we are included in the process of reaching that understanding. Ian Kerr and Jessica Earle write:

> Big data enables a universalizable strategy of preemptive social decision-making. Such a strategy renders individuals unable to observe, understand, participate in, or respond to information gathered or assumptions made about them. When one considers that big data can be used to make important decisions that implicate us without our even knowing it, preemptive social decision making is antithetical to privacy and due process values. (2013, 71)

The hermeneutic privacy norms I presented in chapter 4 work precisely to counter this tendency to make decisions about us based on predictions and assumptions. They demand that we be included in the processes through which others come to perceive and understand us, that we be recognized as the authors of our social selves. If a bank wants to lower our line of credit, we

ought to know why it suddenly thinks our risk of defaulting has gone up. And if it has misinterpreted the facts, we ought to be able to explain why it has misunderstood us. Again, that doesn't mean that others are required to treat what we say about ourselves as overriding truths. It merely means that we ought to have a place in the interpretive negotiations.

Big data offers a particularly stark illustration of why control theories of privacy ought to be abandoned, and why an authorship approach better equips us for the challenges information technology poses to privacy in the Information Age. It shows why are our interests are better served by looking not at the details, but at the big picture, not at particular pieces of information, but at the stories information tells. The threat technology poses to privacy is not that it undermines our ability to control information about ourselves, but rather that it cuts us out of the process of defining ourselves at all. Protecting information privacy in the Information Age means building ourselves back into that process, demanding that we be recognized as the authors of our own social selves.

References

Andrews, Lori. 2012. "Facebook Is Using You." *The New York Times,* February 4. http://www.nytimes.com/2012/02/05/opinion/sunday/facebook-is-using-you.html

Bansal, Manju. 2014. "Big Data: Creating the Power to Move Heaven and Earth." *MIT Technology Review,* September 2. http://www.technologyreview.com/view/530371/big-data-creating-the-power-to-move-heaven-and-earth/

Duhigg, Charles. 2012. "How Companies Learn Your Secrets." *The New York Times Magazine,* February 16. http://www.nytimes.com/2012/02/19/magazine/shopping-habits.html

Jerome, Joseph W. 2013. "Buying and Selling Privacy: Big Data's Different Burdens and Benefits." *Stanford Law Review Online* 66: 47-53. http://www.stanfordlawreview.org/online/privacy-and-big-data/buying-and-selling-privacy

Kerr, Ian, and Jessica Earle. 2013. “Prediction, Preemption, Presumption: How Big Data Threatens the Big Picture.” *Stanford Law Review Online* 66: 65-72. http://www.stanfordlawreview.org/online/privacy-and-big-data/prediction-preemption-presumption

Lohr, Steve. 2014. “Google Flu Trends: The Limits of Big Data.” *The New York Times,* March 28. http://bits.blogs.nytimes.com/2014/03/28/google-flu-trends-the-limits-of-big-data/

Regalado, Antonio. 2013. “The Data Made Me Do It.” *MIT Technology Review*, May 3. http://www.technologyreview.com/news/514346/the-data-made-me-do-it/

Rubinstein, Ira S. 2013. “Big Data: The End of Privacy or a New Beginning?” *International Data Privacy Law* 3 (2): 74-87.

Talbot, David. 2013. “African Bus Routes Redrawn Using Cell-Phone Data.” *MIT Technology Review,* April 30. http://www.technologyreview.com/news/514211/african-bus-routes-redrawn-using-cell-phone-data/

Tene, Omer, and Jules Polonetsky. 2012. “Privacy in the Age of Big Data: A Time for Big Decisions.” *Stanford Law Review Online* 64: 63-69. http://www.stanfordlawreview.org/online/privacy-paradox/big-data

Tucker, Patrick. 2013. “Has Big Data Made Anonymity Impossible?” *MIT Technology Review,* May 7. http://www.technologyreview.com/news/514351/has-big-data-made-anonymity-impossible/

Vedder, Anton. 1999. “KDD: The Challenge to Individualism.” *Ethics and Information Technology* 1 (4): 275-281.

References

Abril, Patricia Sanchez. 2007. “A (My)Space of One’s Own: On Privacy and Online Social Networks.” *Northwestern Journal of Technology and Intellectual Property* 6 (1): 73-88.

Alcoff, Linda Martín. 1999. “On Judging Epistemic Credibility: Is Social Identity Relevant?,” *Philosophic Exchange* 29 (1): Article 1.

Alcoff, Linda Martín. 2006. *Visible Identities: Race, Gender, and the Self.* Oxford: Oxford University Press.

Alderman, Ellen, and Caroline Kennedy. 1995. *The Right to Privacy.* New York: Alfred A. Knopf, Inc.

Alexander, Kurtis, and Vivian Ho. 2013. “New Law Lets Teens Delete Digital Skeletons.” *SFGate,* September 24. http://www.sfgate.com/news/article/New-law-lets-teens-delete-digital-skeletons-4837309.php

Allen, Anita. 2011. *Unpopular Privacy: What Must We Hide?* Oxford: Oxford University Press.

Allen, Anita. 2013. “An Ethical Duty to Protect One’s Own Information Privacy?” *Alabama Law Review* 64 (4).

American Civil Liberties Union. 2004. *The Surveillance-Industrial Complex: How the American Government is Conscripting Businesses and Individuals in the Construction of a Surveillance Society.* https://www.aclu.org/national-security/surveillance-industrial-complex

American Civil Liberties Union. 2011. *Protecting Privacy from Aerial Surveillance: Recommendations for Government Use of Drone Aircraft.* http://www.aclu.org/files/assets/protectingprivacyfromaerialsurveillance.pdf

American Civil Liberties Union. 2013. *You are Being Tracked: How License Plate Readers are Being Used to Record Americans’ Movements.* https://www.aclu.org/technology-and-liberty/you-are-being-tracked-how-license-plate-readers-are-being-used-record

Andrews, Lori. 2012. “Facebook Is Using You.” *The New York Times,* February 4. http://www.nytimes.com/2012/02/05/opinion/sunday/facebook-is-using-you.html

Aufderheide, Patricia. 1993. *Media Literacy: A Report of the National Leadership Conference on Media Literacy.* Aspen, CO: Aspen Institute.

Bamford, James. 2012. “The NSA Is Building the Country’s Biggest Spy Center (Watch What You Say).” *Wired,* March 15. http://www.wired.com/threatlevel/2012/03/ff_nsadatacenter/all/

Bansal, Manju. 2014. “Big Data: Creating the Power to Move Heaven and Earth.” *MIT Technology Review,* September 2. http://www.technologyreview.com/view/530371/big-data-creating-the-power-to-move-heaven-and-earth/

Barlow, John Perry. 1996. “A Declaration of the Independence of Cyberspace.” https://projects.eff.org/~barlow/Declaration-Final.html.

Benn, Stanley I. 1984. “Privacy, Freedom, and Respect for Persons.” In *Philosophical Dimensions of Privacy: An Anthology,* ed. Ferdinand Schoeman. Cambridge: Cambridge University Press.

Bennett, Colin. 2008. *The Privacy Advocates: Resisting the Spread of Surveillance*. Cambridge, MA: MIT Press.

Berger, Rob. 2013. “5 Myths About Late Payments & Your FICO Scores.” *Credit.com* (blog), December 5. http://blog.credit.com/2013/12/5-myths-about-late-payments-your-fico-scores-71720/

Berlin, Isaiah. 1969. *Four Essays on Liberty.* Oxford: Oxford University Press.

boyd, danah. 2012. “Networked Privacy.” *Surveillance & Society* 10 (3/4).

boyd, danah. 2014. *It’s Complicated: The Social Lives of Networked Teens.* New Haven: Yale University Press.

Bradley, Tony. 2010. “Erasing your Digital Tracks on the Web.” *PC Magazine,* May 2. http://www.pcworld.com/article/195270/xxx.html

Branaman, Ann. 1997. “Goffman’s Social Theory.” In *The Goffman Reader,* eds. Charlers Lemert and Ann Branaman. Oxford: Blackwell Publishing Ltd.

Bratman, Michael. 2009. “Shared Agency.” In *Philosophy of the Social Sciences: Philosophical Theory and Scientific Practice,* ed. Chrysostomos Mantzavinos. Cambridge, UK: Cambridge University Press.

Brown, Rupert. 2010. *Prejudice: Its Social Psychology*. *2nd Edition.* West Sussex, UK: Wiley-Blackwell.

Bureau of Justice Statistics. 2013. “16.6 Million People Experienced Identity Theft in 2012.” http://www.bjs.gov/index.cfm?ty=pbdetail&iid=4911

Cate, Fred H. 2006. “The Failure of Fair Information Practice Principles.” In *Consumer Protection in the Age of the ‘Information Economy’,* ed. Jane K. Winn. Hampshire, UK: Ashgate Publishing Limited.

Cavoukian, Ann. 2009. “Privacy by Design.” Report by the Office of the Information and Privacy Commissioner of Ontario. https://www.ipc.on.ca/images/Resources/privacybydesign.pdf

Cavoukian, Ann. 2012. “Privacy by Design: Origins, Meaning, and Prospects for Assuring Privacy and Trust in the Information Era.” In *Privacy Protection Measures and Technologies in Business Organizations: Aspects and Standards,* ed. George Yee. Hershey, PA: Information Science Reference.

Cavoukian, Ann, Stuart Shapiro, and R. Jason Cronk. 2014. “Privacy Engineering: Proactively Embedding Privacy, by Design.” Report by the Office of the Information and Privacy Commissioner of Ontario. https://www.privacybydesign.ca/content/uploads/2014/01/pbd-priv-engineering.pdf

Chalmers, Jim, and Tim Watts. 2014. “Kids Should Code: Why ‘Computational Thinking Needs to Be Taught in Schools.” *The Guardian,* December 18. http://www.theguardian.com/commentisfree/2014/dec/19/kids-should-code-why-computational-thinking-needs-to-be-taught-in-schools

Chapman, Cate. 2015. “Evaluating Privacy and Legal Concerns Related to Big Data.” *Cyber Risk Network,* January 15. http://www.cyberrisknetwork.com/2015/01/15/evaluating-privacy-legal-concerns-related-big-data/

Cheney-Lippold, John. 2011. “A New Algorithmic Identity: Soft Biopolitics and the Modulation of Control.” *Theory, Culture & Society* 28 (6): 164-181.

Cohen, Elliot D. 2010. *Mass Surveillance and State Control.* New York: Palgrave Macmillan.

Cohen, Jean L. 1996. “Democracy, Difference, and the Right of Privacy.” In *Democracy and Difference: Contesting the Boundaries of the Political,* ed. Seyla Behnabib. Princeton, NJ: Princeton University Press.

Cooley, Charles Horton. 1902. *Human Nature and the Social Order.* New York: Scribner.

Cuthbertson, Anthony. 2014. “Coding in the Classroom: Computational Thinking Will Allow Children to ‘Change the World.’” *International Business Times,* September 2. http://www.ibtimes.co.uk/coding-classroom-computational-thinking-will-allow-children-change-world-1463493

DeCew, Judith Wagner. 1997. *In Pursuit of Privacy: Law Ethics, and the Rise of Technology.* Ithaca, NY: Cornell University Press.

DeCew, Judith Wagner. 2000. "Privacy and Information Technology." In *Privacy and Data Protection: Theory and Practice,* ed. M. J. van den Hoven. Kluwer Academic Publishers.

DeCew, Judith Wagner. 2013. "Privacy." In *The Stanford Encyclopedia of Philosophy,* ed. Edward N. Zalta. http://plato.stanford.edu/archives/fall2013/entries/privacy/

DeVries, Will Thomas. 2003. "Protecting Privacy in the Digital Age." *Berkeley Technology Law Journal* 18 (1): 283-311.

Dewey, Caitlin. 2014. "A Comprehensive, Jargon-free Guide to the Celebrity Nude-photo Scandal and the Shadowy Web Sites Behind It." *The Washington Post,* September 2. http://www.washingtonpost.com/news/the-intersect/wp/2014/09/02/a-comprehensive-jargon-free-guide-to-the-celebrity-nude-photo-scandal-and-the-shadowy-web-sites-behind-it/

Dixon, Pam. 2006. "A Brief Introduction to Fair Information Practices." *World Privacy Forum,* June 5. https://www.worldprivacyforum.org/2008/01/report-a-brief-introduction-to-fair-information-practices/

Dixon, Pam, and Robert Gellman. 2014. *The Scoring of America: How Secret Consumer Scores Threaten Your Privacy and Your Future.* Report by the World Privacy Forum. http://www.worldprivacyforum.org/2014/04/wpf-report-the-scoring-of-america-how-secret-consumer-scores-threaten-your-privacy-and-your-future/

Duhigg, Charles. 2012. "How Companies Learn Your Secrets." *The New York Times Magazine,* February 16. http://www.nytimes.com/2012/02/19/magazine/shopping-habits.html

Dwork, Cynthia, and Deirdre K. Mulligan. 2013. "It's Not Privacy, and It's Not Fair." *Stanford Law Review Online* 66: 35-40. http://www.stanfordlawreview.org/online/privacy-and-big-data/its-not-privacy-and-its-not-fair

Eisenberg, Michael B. 2008. "Information Literacy: Essential Skills for the Information Age." *Journal of Library & Information Technology* 28 (2): 39-47.

Ferenstein, Gregory. 2013. "On California's Bizarre Internet Eraser Law For Teenagers." *TechCrunch,* September 24. http://techcrunch.com/2013/09/24/on-californias-bizarre-internet-eraser-law-for-teenagers/

Floridi, Luciano. 2010. *Information: A Very Short Introduction*. Oxford: Oxford University Press.

Floridi, Luciano. 2006. “The Ontological Interpretation of Information Privacy.” *Ethics and Information Technology* 7 (4): 185-200.

Foucault, Michel. 1977. *Discipline and Punish: The Birth of the Prison.* New York: Vintage Books.

Franzen, Carl. 2013. “Advertisers Can Learn Your Health Conditions From Your Web Activity, Study Claims.” *The Verge,* July 8. http://www.theverge.com/2013/7/8/4505164/medical-data-isnt-anonymous-in-ad-tracking-study-finds

Fricker, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing.* Oxford: Oxford University Press.

Fried, Charles. 1968. “Privacy.” *Yale Law Journal* 77 (3): 475-493.

Friendly, Henry. 1975. “Some Kind of Hearing.” *University of Pennsylvania Law Review* 123: 1267-1317.

Gavison, Ruth. 1980. “Privacy and the Limits of the Law.” *The Yale Law Journal* 89 (3): 421-471.

Gerstein, Robert. 1978. “Intimacy and Privacy.” *Ethics* (89): 76-81.

Gillespie, Tarleton. 2014. “The Relevance of Algorithms.” In *Media Technologies,* eds. Tarleton Gillespie, Pablo Boczkowski, and Kirsten Foot. Cambridge, MA: MIT Press. Pre-publication version accessed online at http://culturedigitally.org/2012/11/the-relevance-of-algorithms/

Goffey, Andrew. 2008. “Algorithm.” In *Software Studies: A Lexicon,* ed. Matthew Fuller. Cambridge, MA: MIT Press.

Goffman, Erving. 1959. *The Presentation of Self in Everyday Life.* New York: Anchor Books.

Goffman, Erving. 1961. *Encounters: Two Studies in the Sociology of Interaction*. Penguin University Books.

Goffman, Erving. 1967. *Interaction Ritual: Essays on Face-to-Face Behavior.* New York: Pantheon Books.

Goffman, Erving. 1971. *Relations in Public: Microstudies of the Public Order.* New York: Basic Books.

Graham, Stephen D.N. 2005. “Software-Sorted Geographies.” *Progress in Human Geography* 29 (5): 562-580.

Greenwald, Glenn. 2013. “NSA Collecting Phone Records of Millions of Verizon Customers Daily.” *The Guardian,* June 5. http://www.guardian.co.uk/world/2013/jun/06/nsa-phone-records-verizon-court-order

Greenwald, Glenn, and Ewen MacAskill. 2013. “Boundless Informant: The NSA's Secret Tool to Track Global Surveillance Data.” *The Guardian,* June 11. http://www.guardian.co.uk/world/2013/jun/08/nsa-boundless-informant-global-datamining

Hargittai, Eszter. 2008. “The Digital Reproduction of Inequality.” In *Social Stratification*, ed. David Grusky. Boulder, CO: Westview.

Hartzog, Woodrow, and Evan Selinger. 2013. “Big Data in Small Hands.” *Stanford Law Review Online.* 66:81-88. http://www.stanfordlawreview.org/sites/default/files/online/topics/66_StanLRevOnline_81_HartzogSelinger.pdf

Hartzog, Woodrow, and Evan Selinger. 2013a. “Obscurity: A Better Way to Think About Your Data Than ‘Privacy.’” *The Atlantic,* January 17. http://www.theatlantic.com/technology/archive/2013/01/obscurity-a-better-way-to-think-about-your-data-than-privacy/267283/

Herek, Gregory and Eric Glunt. 1993. “Interpersonal Contract and Heterosexual’s Attitudes Toward Gay Men: Results from a National Survey.” *The Journal of Sex Research* 30 (3): 239-44.

Holstein, James and Jaber Gubrium. 2000. *The Self We Live By: Narrative Identity in a Postmodern World.* Oxford: Oxford University Press.

Huseman, Jessica. 2014. “Setting the Record Straight on Mortgages for Undocumented Immigrants.” *National Mortgage News,* October 17. http://www.nationalmortgagenews.com/news/risk-management/setting-the-record-straight-on-mortgages-for-undocumented-immigrants-1042907-1.html

Hustinx, Peter. 2010. “Privacy by Design: Delivering the Promises.” *Identity in the Information Society* 3 (2): 253-255.

James, William. 1890. *The Principles of Psychology.* New York: Dover Publications.

Jerome, Joseph W. 2013. “Buying and Selling Privacy: Big Data’s Different Burdens and Benefits.” *Stanford Law Review Online* 66: 47-53. http://www.stanfordlawreview.org/online/privacy-and-big-data/buying-and-selling-privacy

Johnson, Bobbie. 2010. “Privacy no longer a social norm, says Facebook founder.” *The Guardian,* January 10. http://www.theguardian.com/technology/2010/jan/11/facebook-privac

Kelly, Jeanne. 2010. “Recipe for a High FICO Credit Score.” *The Huffington Post,* November 15. http://www.huffingtonpost.com/jeanne-kelly/recipe-for-a-high-fico-cr_b_777627.html

Kerr, Ian, and Jessica Earle. 2013. “Prediction, Preemption, Presumption: How Big Data Threatens the Big Picture.” *Stanford Law Review Online* 66: 65-72. http://www.stanfordlawreview.org/online/privacy-and-big-data/prediction-preemption-presumption

Kozinski, Alex. 2012. “The Dead Past.” *Stanford Law Review Online* 64: 117-124. http://www.stanfordlawreview.org/online/privacy-paradox/dead-past

Leary, Mark. 1996. *Self-Presentation: Impression Management and Interpersonal Behavior.* Boulder, CO: Westview Press.

Leber, Jessica. 2013. “A Startup That Scores Job Seekers, Whether They Know It or Not.” *MIT Technology Review,* March 7. http://www.technologyreview.com/news/511896/a-startup-that-scores-job-seekers-whether-they-know-it-or-not/

Lessig, Lawrence. 1999. “The Architecture of Privacy.” *Vanderbilt Journal of Entertainment Law and Practice* 1 (1): 56-65.

Levin, Avner, and Mary Jo Nicholson. 2005. “Privacy Law in the United States, the EU, and Canada: The Allure of the Middle Ground.” *University of Ottawa Law & Technology Journal* 2 (2): 357-395.

Livingstone, Sonia. 2004. “Media Literacy and the Challenge of New Information and Communication Technologies.” *Communication Review* 1 (7): 3-14.

Lohr, Steve. 2014. “Google Flu Trends: The Limits of Big Data.” *The New York Times,* March 28. http://bits.blogs.nytimes.com/2014/03/28/google-flu-trends-the-limits-of-big-data/

Lyons, David. 2001. *Surveillance Society: Monitoring Everyday Life.* Buckingham, UK: Open University Press.

Malheiros, Miguel, Sacha Brostoff, Charlene Jennett, and M. Angela Sasse. 2013. “Would You Sell Your Mother’s Data? Personal Data Disclosure in a Simulated Credit Card Application.” In *The Economics of Information Security and Privacy,* ed. Rainer Böhme. Berlin: Springer-Verlag.

Martino, Paul. 2013. “Inside California’s New Online Privacy Law for Minors.” *Law 360*, October 11. http://www.law360.com/articles/479853/inside-calif-s-new-online-privacy-law-for-minors

Mattioli, Dana. 2012. “On Orbitz, Mac Users Steered to Pricier Hotels.” *The Wall Street Journal,* August 23. http://online.wsj.com/news/articles SB10001424052702304458604577488822667325882

Mayer, Jane. 2006. “What’s the Matter with Metadata?” *New Yorker* (blog), June 6. http://www.newyorker.com/online/blogs/newsdesk/2013/06/verizon-nsa-metadata-surveillance-problem.html

Mayer-Shönberger, Viktor. 2011. *Delete: The Virtue of Forgetting in the Digital Age.* Princeton, NJ: Princeton University Press.

Meeler, David. 2008. “Is Information All We Need to Protect?” *The Monist* 91 (1): 151-169.

Mulligan, Deirdre, and Jennifer King. 2012. “Bridging the Gap Between Privacy and Design.” *Journal of Constitutional Law* 14 (4): 989-1034.

Nagel, Thomas. 2002. *Concealment and Exposure: And Other Essays.* Oxford: Oxford University Press.

National Research Council. 2007. *Engaging Privacy and Information Technology in a Digital Age*. Washington, DC: The National Academies Press.

Nissenbaum, Helen. 1997. “Toward an Approach to Privacy in Public: Challenges of Information Technology.” *Ethics and Behavior* 7 (3): 207-219.

Nissenbaum, Helen. 2010. *Privacy in Context: Technology, Policy, and the Integrity of Social Life* Stanford: Stanford University Press.

Newland, Erica. 2012. “Disappearing Phone Booths: Privacy in the Digital Age.” *Speech to the DC Superior Court judges*. Washington, DC. https://www.cdt.org/files/pdfs/Privacy-In-Digital-Age.pdf

O’Neil, Megan. 2014. “Confronting the Myth of the ‘Digital Native.’” *The Chronicle of Higher Education,* April 21. http://chronicle.com/article/Confronting-the-Myth-of-the/145949/

Organization for Economic Cooperation and Development. 1973. *OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data.* http://www.oecd.org/internet/ieconomy/oecdguidelinesontheprotectionofprivacyandtransborderflowsofpersonaldata.htm

Pasquale, Frank. 2015. *The Black Box Society: The Secret Algorithms That Control Money and Information.* Cambridge, MA: Harvard University Press.

Paul, Annie Murphy. 2014. “Interrupt Kids’ Regular Scheduled Programming: Teaching Computer Science Without Touching a Computer.” *Slate,* August 27. http://www.slate.com/articles/technology/future_tense/2014/08/computer_science_unplugged_teaching_computational_thinking_without_computers.html

Posner, Richard. 1978. “An Economic Theory of Privacy.” *Regulation*. May/June Issue.

Posner, Richard. 1981. *The Economics of Justice.* Cambridge, MA: Harvard University Press.

Powles, Julia. 2014. “What We Can Salvage from 'Right to Be Forgotten' Ruling.” *Wired UK,* May 15. http://www.wired.co.uk/news/archive/2014-05/15/google-vs-spain

Prensky, Marc. 2001. “Digital Natives, Digital Immigrants.” *On the Horizon* 9.http://www.marcprensky.com/writing/Prensky%20-%20Digital%20Natives,%20Digital%20Immigrants%20-%20-Part1.pdf

Prosser, William. 1960. “Privacy.” *California Law Review* 48: 383–423.

Prosser, William. 1984. “Privacy: A Legal Analysis.” In *Philosophical Dimensions of Privacy: An Anthology,* ed. Ferdinand Schoeman. Cambridge: Cambridge University Press.

Quinn, Kimberly, C. Neil Macrae, and Galen Bodenhausen. 2007. “Stereotyping and Impression Formation: How Categorical Thinking Shapes Person Perception.” In *Sage Handbook of Social Psychology: Concise Student Edition,* eds. Michael A. Hogg and Joel Cooper. Thousand Oaks, CA: Sage Publications.

Rachels, James. 1975. “Why is Privacy Important?” *Philosophy and Public Affairs* 4 (Summer): 323-333.

Raja, Tasneem. 2014. “Is Coding the New Literacy? Why America’s Schools Need to Train a New Generation of Hackers.” *Mother Jones,* June 16. http://www.motherjones.com/print/253891

Regalado, Antonio. 2013. “The Data Made Me Do It.” *MIT Technology Review*, May 3. http://www.technologyreview.com/news/514346/the-data-made-me-do-it/

Reiman, Jeffrey. 1976. “Privacy, Intimacy, and Personhood.” *Philosophy & Public Affairs* 6 (1): 26-44.

Richards, Neil, and Daniel Solove. 2010. “Prosser’s Privacy Law: A Mixed Legacy.” *California Law Review* 98 (6): 1887-1924.

Riou, Garrett. 2014. “Hospital Visition and Medical Decision Making for Same-Sex Couples.” *Center for American Progress Blog*, April 15. https://www.americanprogress.org/issues/lgbt/news/2014/04/15/88015/hospital-visitation-and-medical-decision-making-for-same-sex-couples/

Rosen, Jeffrey. 2000. *The Unwanted Gaze: The Destruction of Privacy in America*. New York: Vintage Books.

Rosen, Jeffrey. 2012. “The Right to Be Forgotten.” *Stanford Law Review Online* 64:88-92. http://www.stanfordlawreview.org/online/privacy-paradox/right-to-be-forgotten

Rubinstein, Ira. 2011. “Regulating Privacy by Design.” *Berkeley Technology Law Journal* 26 (3): 1409-1456.

Rubinstein, Ira. 2013. “Big Data: The End of Privacy or a New Beginning?” *International Data Privacy Law* 3 (2): 74-87.

Schlenker, Barry. 1980. *Impression Management: The Self-Concept, Social Identity, and Interpersonal Relations.* Monterey, California: Brooks/Cole Publishing Company.

Schlenker, Barry, and Beth Pontari. 2000. “The Strategic Control of Information: Impression Management and Self-Presentation in Daily Life.” In *Psychological Perspectives on Self and Identity,* eds. Abraham Tesser, Richard B. Felson, and Jerry M. Suls. Washington, DC: American Psychological Association.

Schoeman, Ferdinand. 1994. “Gossip and Privacy.” In *Good Gossip,* ed. Robert F. Goodman and Aaron B. Ze’ev. Lawrence: University of Kansas Press.

Seida, Caitlin. 2013. “My Embarrassing Picture Went Viral.” *Salon,* October 2. http://www.salon.com/2013/10/02/my_embarrassing_picture_went_viral/

Sheehy, Kelsey. 2012. “High Schools Not Meeting STEM Demand.” *US News & World Report,* October 1. http://www.usnews.com/education/blogs/high-school-notes/2012/10/01/high-schools-not-meeting-stem-demand

Solove, Daniel. 2004. *The Digital Person: Technology and Privacy in the Information Age.* New York: New York University Press.

Solove, Daniel. 2008. *Understanding Privacy.* Cambridge, MA: Harvard University Press.

Southwell, Alexander. 2013. “California’s New ‘Digital Eraser’ Evaporates Embarrassment.” *Law Technology News,* November 19. http://www.legaltechnews.com/id=1202628537209

Stone, Andrea. 2012. “Homeland Security Manual Lists Government Key Words for Monitoring Social Media, News.” *The Huffington Post,* February 24. http://www.huffingtonpost.com/2012/02/24/homeland-security-manual_n_1299908.html

Streitfeld, David. 2014. “European Court Lets Users Erase Records On Web.” *New York Times,* May 13. http://www.nytimes.com/2014/05/14/technology/google-should-erase-web-links-to-some-personal-data-europes-highest-court-says.html

Talbot, David. 2013. “African Bus Routes Redrawn Using Cell-Phone Data.” *MIT Technology Review,* April 30. http://www.technologyreview.com/news/514211/african-bus-routes-redrawn-using-cell-phone-data/

Tamò, Aurelia, and Damian George. 2014. “Oblivion, Erasure and Forgetting in the Digital Age.” *The Journal of Intellectual Property, Information Technology and Electronic Commerce Law* 5 (2): 71-87.

Tene, Omer, and Jules Polonetsky. 2012. “Privacy in the Age of Big Data: A Time for Big Decisions.” *Stanford Law Review Online* 64: 63-69. http://www.stanfordlawreview.org/online/privacy-paradox/big-data

Thomas, Lyn. 2000. “A Survey of Credit and Behavioural Scoring: Forecasting Financial Risk of Lending to Consumers.” *International Journal of Forecasting* 16: 149-172.

Tucker, Patrick. 2013. “Has Big Data Made Anonymity Impossible?” *MIT Technology Review,* May 7. http://www.technologyreview.com/news/514351/has-big-data-made-anonymity-impossible/

Turek, Dave. 2012. “The Case Against Digital Sprawl.” *Business Week,* May 2. http://www.businessweek.com/articles/2012-05-02/the-case-against-digital-sprawl

Ungerleider, Neal. 2012. “NYPD, Microsoft Launch All-Seeing ‘Domain Awareness System’ With Real-Time CCTV, License Plate Monitoring [Updated].” *Fast Company,* August 8. http://www.fastcompany.com/3000272/nypd-microsoft-launch-all-seeing-domain-awareness-system-real-time-cctv-license-plate-monito

U.S. Department of Health, Education, and Welfare. 1973. *Records, Computers, and the Rights of Citizens.* DHEW Publication No. (OS) 73-94. http://www.justice.gov/opcl/docs/rec-com-rights.pdf

Valentino-DeVries, Jennifer, Jeremy Singer-Vine, and Ashkan Soltani. 2012. “Websites Vary Prices, Deals Based on Users' Information.” *The Wall Street Journal,* December 24. http://online.wsj.com/news/articles/SB10001424127887323777204578189391813881534

Vedder, Anton. 1999. “KDD: The Challenge to Individualism.” *Ethics and Information Technology* 1 (4): 275-281.

Velasquez, Manuel, Claire Andre, Thomas Shanks, and Michael Meyer. 1990. “Rights.” *Issues in Ethics* 3 (1).

Wagstaff, Keith. 2012. “Can We Fix Computer Science Education in America?” *Time,* July 16. http://techland.time.com/2012/07/16/can-we-fix-computer-science-education-in-america/

Wallace, Patricia. 1999. *The Psychology of the Internet.* Cambridge: Cambridge University Press.

Wang, Yang, and Alfred Kobsa. 2009. “Privacy-Enhancing Technologies.” In *Handbook of Research on Social and Organizational Liabilities in Information Security,* eds. Manish Gupta and Raj Sharman. Hershey, PA: IGI Global.

Warren, Samuel, and Louis Brandeis. 1890. “The Right to Privacy.” *Harvard Law Review* 4 (5): 193-220.

Westin, Alan. 1967. *Privacy and Freedom.* New York, NY: Atheneum.

The White House. 2012. *Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy.* http://www.whitehouse.gov/sites/default/files/privacy-final.pdf

Winner, Langdon. 1986. *The Whale and the Reactor: A Search for Limits in an Age of High Technology.* Chicago: The University of Chicago Press.

Zimmer, Michael. 2014. “Mark Zuckerberg’s theory of privacy.” *The Washington Post,* February 3. http://www.washingtonpost.com/lifestyle/style/mark-zuckerbergs-theory-of-privacy/2014/02/03/2c1d780a-8cea-11e3-95dd-36ff657a4dae_story.html