

## PAPER TITLE

“Information Privacy and Social Self-Authorship,” forthcoming in  
*Techné: Research in Philosophy & Technology*

## AUTHOR

Daniel Susser  
Research Fellow, New York University School of Law  
Assistant Professor, San Jose State University

## ABSTRACT

The dominant approach in privacy theory defines information privacy as some form of control over personal information. In this essay, I argue that the control approach is mistaken. I claim that information privacy involves the drawing of epistemic boundaries—boundaries between what others should and shouldn’t know about us. While controlling what information others have about us is one strategy we use to draw such boundaries, it is not the only one. We conceal information about ourselves and we reveal it. And since the meaning of information is not self-evident, we also work to shape how others contextualize and interpret the information about us that they have. Information privacy is thus about more than controlling information; it involves the constant work of producing and managing public identities, what I call “social self-authorship.” In the second part of the essay, I argue that thinking about information privacy in terms of social self-authorship helps us see ways that information technology threatens privacy, which the control approach misses. Namely, information technology makes social self-authorship invisible and unnecessary, by making it difficult for us to know when others are forming impressions about us, and by providing them with tools for making assumptions about who we are which obviate the need for our involvement in the process.

## KEYWORDS

privacy, information privacy, control, personal information, social self-authorship

## Information Privacy and 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 has to do with negotiating boundaries between oneself and others.<sup>1</sup> 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. The dominant approach in privacy theory defines information privacy as some form of control over personal information. Alan Westin famously argued that privacy is “the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others” (1967, 7). For Charles Fried information 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). This idea is central not only to debates in privacy theory, but also to debates in privacy law. According to philosopher and legal theorist Judith DeCew, the “crucial core of privacy” protected by American tort law is “having control over information about oneself” (1997, 14).<sup>2</sup> In other words, the dominant conception of information privacy holds that we draw epistemic boundaries by controlling the information others have about us.

I argue in what follows that this approach—exemplified in so-called “control theories” of privacy—is wrong.<sup>3</sup> Control theories understand correctly that information privacy is about negotiating

epistemic boundaries. Indeed, the basic moral intuition underlying such theories, that we value information privacy because we value having agency over how others know us, seems to me fundamentally right. The problem is that they misunderstand how we normally exercise that agency, which is to say, 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 contextualized and understood. Put differently, 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.

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.

This example demonstrates 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. My goal in the first part of this essay 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.<sup>4</sup>

In the second part of the essay, 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. It is true, of course, 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 Information Age, 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.

The way we understand how we negotiate epistemic boundaries has important 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 a step toward improving information privacy policy.

## 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” (293-4). In other words, my social selves are the various different ways other people understand who I am.<sup>5</sup> 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. 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 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.<sup>6</sup> 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 taste 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 strength when we endure through adversity, and our vulnerability 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 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.<sup>7</sup> 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. 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.<sup>8</sup> Shaping those images is therefore not a one-time task, but rather, to quote Nagel again, "the constant work of human life."

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.”<sup>9</sup> 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 impress a visiting critic. And we act according to various “scripts”—social conventions which dictate what to say and do in given situations.<sup>10</sup>

“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 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 professor. 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 professor 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)<sup>11</sup>

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.<sup>12</sup> 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 situation.<sup>13</sup> 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 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).<sup>14</sup> Indeed, even appearing as though one is trying too hard to control one's appearance conveys revealing information.

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.<sup>15</sup> 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 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. “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.<sup>16</sup>

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 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 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. What's more, it is the problem we are least likely to solve. A huge amount of information has to be generated about us in order to provide us with the technology-based services we want. If we want good email spam filters, for instance, algorithms will need to collect information about who we email, how often, and what sorts of things we usually write about. If we want our phones to tell us the weather they have to collect information about where we are. If we want our computers to remind us about our appointments they have to know where we're supposed to be and what we're supposed to be doing. 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 generated about us is only going to grow, and our ability to control that information is only going to further diminish.<sup>17</sup>

Control theories thus focus all of our attention on a single, unwinnable 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. 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. 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 advancing—that controlling information is but one of many means by which social self-authorship takes place—is as applicable to our activities online as it is to our activities 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 different 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.<sup>18</sup>

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 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 (though this sometimes does more harm than good), 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; now we simply do so in new ways.

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.<sup>19</sup>

Online, the situation is 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. 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.

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. 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.” (2004, 7).<sup>20</sup>

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 committed 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.<sup>21</sup> Amazon, for instance, once charged members of its own “Prime” service—which, for an annual fee, 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 extra 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 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.<sup>22</sup> 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.<sup>23</sup> Advertisers use the search terms people enter into health-focused websites to determine which medical conditions they have, and then use that information to target them with pharmaceutical ads.<sup>24</sup> 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.<sup>25</sup>

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 and evaluated, and they are therefore denied the opportunity to participate in the processes by which they are perceived and understood. Consider how this is different from the way things work offline. As I argued in the previous section, we give off all kinds of information face-to-face. But in face-to-face interactions we can make reliable, educated guesses about who is observing us, what kinds of things they are noticing about us, and how that information is affecting their perceptions of who we are. We can physically look around and see who is watching us, ask what their motives for watching us might be, reflect on what we’re doing and how it might make us appear to them, and so on. Online very little of that is possible. It’s extremely difficult to know what information the technologies we engage with are collecting, where that information might go, or how it might be used. 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 specific 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, then 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.

This problem is captured in part by the notion of "context collapse," an idea put forward by Alice Marwick and danah boyd (2011).<sup>26</sup> Like Goffman, they argue that whenever we communicate we construct real or imagined audiences that determine how we express and present our thoughts. However, certain forms of information technology interfere with this process. "The need for variable self-presentation," they write, "is complicated by increasingly mainstream social media technologies that collapse multiple contexts and bring together commonly distinct audiences" (115). Twitter, to take their example, requires that its users write messages to potentially unknown audiences, thus making it difficult to know how best to shape one's message in advance. "The potential diversity of readership on Twitter ruptures the ability to vary self-presentation based on audience, and thus manage discrete impressions" (116).

Yet the issue I'm pointing to involves more than mere context-collapse. The problem is not just that in online social media we can't anticipate who will receive and try to decode our communications. The problem is that when our activities are mediated by information technology we often don't realize that we're communicating at all. When we post a message on Facebook or Twitter we understand that we're communicating something to someone, although to whom exactly might be opaque. When we browse the Internet and websites track our comings and goings the situation is different. When we drive past

license plate readers or purchase something in a store using a credit card it is not immediately apparent that we're communicating information about ourselves. Information technology undermines social self-authorship not only by making it difficult to know our audiences, but by making it difficult to know when we even have one.

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.<sup>27</sup> 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. 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.<sup>28</sup>

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. 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

- Allen, Anita. 1988. *Uneasy Access: Privacy for Women in a Free Society*. Savage, MD: Rowman & Littlefield.
- Allen, Anita. 2013. "An Ethical Duty to Protect One's Own Information Privacy?" *Alabama Law Review*, 64(4).
- Altman, Irwin. 1975. *The Environment and Social Behavior: Privacy, Personal Space, Territory, Crowding*. Monterey, CA: Brooks/Cole Publishing Company, Inc.
- Bennett, Colin. 2008. *The Privacy Advocates: Resisting the Spread of Surveillance*. Cambridge, MA: MIT Press.
- Branaman, Ann. 1997. "Goffman's Social Theory." In *The Goffman Reader*, ed. Charles 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.
- DeCew, Judith. 1997. *In Pursuit of Privacy: Law, Ethics, and the Rise of Information Technology*. Ithaca, NY: Cornell University Press.
- DeCew, Judith. 2013. "Privacy." In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. Retrieved from <http://plato.stanford.edu/archives/fall2013/entries/privacy/>
- 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>.
- Fried, Charles. 1968. "Privacy." *Yale Law Journal* 77(3), 475-493.

- 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.
- Marwick, Alice and danah boyd. 2011. "I Tweet Honestly, I Tweet Passionately: Twitter Users, Context Collapse, and the Imagined Audience." *New Media Society* 13(1), 114-133.
- Mattioli, Dana. 2012. "On Orbitz, Mac Users Steered to Pricier Hotels." *The Wall Street Journal*, August 23.  
<http://online.wsj.com/news/articles/SB10001424052702304458604577488822667325882>
- Meyrowitz, Joshua. 1985. *No Sense of Place: The Impact of Electronic Media on Social Behavior*. Oxford: Oxford University Press.
- Nagel, Thomas. 2002. *Concealment and Exposure: And Other Essays*. Oxford: Oxford University Press.
- Newland, E. 2012. "Disappearing Phone Booths: Privacy in the Digital Age." *Speech to the DC Superior Court judges*. Washington, DC. Retrieved from <https://www.cdt.org/files/pdfs/Privacy-In-Digital-Age.pdf>
- Nissenbaum, Helen. 2010. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford: Stanford University Press.
- 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*, ed. Abraham Tesser, Richard Felson, and Jerry 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.
- Solove, Daniel. 2008. *Understanding Privacy*. Cambridge, MA: Harvard 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](http://www.huffingtonpost.com/2012/02/24/homeland-security-manual_n_1299908.html)
- Valentino-DeVries, Jennifer, Jeremey 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.
- Warren, Samuel and Louis Brandeis. 1890. "The Right to Privacy." Reprinted in *Philosophical Dimensions of Privacy: An Anthology*, ed. Ferdinand Schoeman. Cambridge: Cambridge University Press.
- Westin, Alan. 1967. *Privacy and Freedom*. New York, NY: Atheneum.

## Notes

---

<sup>1</sup> See, for example, Irwin Altman's (1975) *The Environment and Social Behavior: Privacy, Personal Space, Territory, Crowding*. "Privacy," writes Altman, "is an interpersonal boundary-control process, which paces and regulates interactions with others" (10). Also see Cohen (1996).

<sup>2</sup> For lengthier discussions of control theories of privacy and their place in the broader privacy literature, see chapter 4 in Nissenbaum (2010) and chapter 2 in Solove (2008).

<sup>3</sup> I am not the first to argue that control theories are misguided. For other arguments against the control approach, see Allen (1988), DeCew (1997), Solove (2008), and Nissenbaum (2010).

<sup>4</sup> 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. I think such an 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.

<sup>5</sup> 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."

<sup>6</sup> 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.

<sup>7</sup> I am grateful to Serene Khader for this example.

<sup>8</sup> 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.

<sup>9</sup> I explain why I prefer the term "social self-authorship" to "impression management" below.

<sup>10</sup> For more on Goffman's elaborate stage analogy see his (1959) *The Presentation of Self in Everyday Life*.

<sup>11</sup> 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.

<sup>12</sup> “Although some people seem to regard concerns with others’ impressions as a sign of vanity, manipulateness, or insecurity, self-presentation is an essential and unavoidable aspect of everyday interaction” (Leary 1996, 15).

<sup>13</sup> See Goffman (1959) p. 294.

<sup>14</sup> 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).

<sup>15</sup> 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).

<sup>16</sup> 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).

<sup>17</sup> 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.”

<sup>18</sup> See chapter 2 of Wallace’s (1999) *The Psychology of the Internet*.

<sup>19</sup> Of course, there is an infinite amount of information we could *potentially* give off face-to-face. What I am claiming here is that there is only a limited amount of information that we *actually* give off. Moreover, I don’t mean to suggest that the information we do in fact give off face-to-face isn’t significant. As I argued in the previous section, the fact that we give off so much information offline

---

undermines the notion that we ever had robust control over information ourselves, even prior to the introduction of information technology. What has changed is (1) that we give off orders of magnitude more information online than we do offline (meaning that much more information is *actually observed, collected, and stored*), and (2) most people have very little understanding of the kinds of information they give off online and therefore don't know how to modulate their behavior online to take that into account.

<sup>20</sup> 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.

<sup>21</sup> See Valentino-DeVries, et. al., "Websites Vary Prices, Deals Based on Users' Information."

<sup>22</sup> See Mattioli, D., "On Orbitz, Mac Users Steered to Pricier Hotels."

<sup>23</sup> See Solove (2004), p. 50.

<sup>24</sup> See Franzen, C., "Advertisers can learn your health conditions from your web activity, study claims."

<sup>25</sup> See Stone, A., "Homeland Security Manual Lists Government Key Words for Monitoring Social Media, News"

<sup>26</sup> See boyd and Marwick (2011) "I Tweet Honestly, I Tweet Passionately: Twitter Users, Context Collapse, and the Imagined Audience." Although boyd and Marwick coined the term, the idea stems, as they note in their article, from Joshua Meyrowitz's work. See Meyrowitz (1985).

<sup>27</sup> See Solove (2004), p49.

<sup>28</sup> As one anonymous reviewer of this paper helpfully pointed out, this is not an entirely new phenomenon. We have long made snap judgments about people on account of the information they give off, rather than explicitly soliciting information from them. What changes in the context of information technology, I think, is that governments, corporations, and institutions are incentivized to do that too. The decisions such entities make about us are likely to have greater

---

impacts on our lives than the decisions individuals make (indeed, governments are allowed to coerce us), and thus there is more at stake in ensuring that they interpret the information they have about us correctly.