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.