THE UNIVERSITY OF CHICAGO

WHEN PRIVACY GOES PRIVATE:

TECHNOLOGICAL STEWARDSHIP IN POST-SNOWDEN SILICON VALLEY

A DISSERTATION SUBMITTED TO

THE FACULTY OF THE DIVISION OF THE SOCIAL SCIENCES

IN CANDIDACY FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

DEPARTMENT OF ANTHROPOLOGY

BY

LAKE POLAN

CHICAGO, ILLINOIS

MARCH 2023

TABLE OF CONTENTS

INTRODUCTION

In April 2015, I spoke by Skype with Ladar Levison, an engineer and founder of Lavabit, an open source email service, known for encrypted—"secure"—storage. I had reached out to Ladar through a cryptography mailing list in hopes of speaking with him about his new project, Dark Mail, an email protocol featuring end-to-end encryption (E2E). Given his decade-plus experience building privacy-enhancing technologies, I was curious to hear Ladar's take on the contemporary state of privacy and technology's role, if any, in preserving its future.

When I answered Ladar's call, I launched into the script I had developed for such occasions. I identified myself as a lawyer-turned-anthropologist, and began to sketch the lines of inquiry I hoped to explore by studying privacy ethnographically. Within minutes, Ladar cut me off. For the next hour and a half, I furiously scribbled notes as Ladar held forth mostly uninterrupted on his business and technology endeavors, the personal experiences that shaped his thinking about privacy, and the political, economic, and technical factors that make privacy such a perennially intractable problem.

Though the call represented my proper introduction to Ladar, it was not the first time I had encountered him. Twice in the preceding year I had watched from conference audiences as Ladar discussed aspects of his work. The preceding month, for example, I had flown to Dallas for the Spring meeting of the Internet Engineering Task Force (IETF). An independent, non-governmental organization, the IETF develops and maintains the basic technical standards for internet protocols. Held in the kind of gleaming, anodyne hotel complex designed to attract large

corporate gatherings, the meeting brought together network operators, hardware and software engineers, and researchers interested in the internet's technical and operational workings. Though IETF norms discourage working groups from discussing policy or business, Ladar was not alone at the meeting in his interest in privacy. Indeed, as I learned in an orientation session for first-time participants, the organization now—after Snowden—effectively recognized privacy as within its institutional mandate, as one of the factors necessary to make the internet work "from an engineering point of view." As corroborated by the working group discussions I observed, efforts were underway to address potential privacy risks in all IETF protocols.

In an invited talk before the Security Area Advisory Group, and later, in an informal evening gathering, Levison explained how he intended Dark Mail to improve upon PGP, the most widely used email encryption standard. Expressing his hope that the IETF would eventually adopt Dark Mail as a standard, Ladar argued that to make E2E email ubiquitous, a solution was needed, which unlike PGP, was not so complicated that only security experts could properly use it. He presented Dark Mail as an attempt to design new email protocols, which directly integrate privacy and security protections, and which existing service providers could easily implement "without changing how people use email today."

Eight months prior, I watched as Ladar participated in panel discussion in a related but distinct setting, the tenth anniversary meeting of Hackers on Planet Earth (HOPE). Held that year in the Hotel Pennsylvania, a faded, hulking grand hotel across the street from Manhattan's Pennsylvania Station, HOPE, like the IETF, attracted an array of technical experts. To my eye, the self-identified hackers in attendance were generally younger and more diverse than the IETF's members, and less likely to hold corporate jobs. Joining them were an assortment of

journalists, activists, and civil liberties lawyers. These were the kind of "civil misfits" that Ladar likely had in mind, when he later described to me the people who, in his recent years of "quasi-public personhood," had influenced his evolving privacy beliefs. At the Dallas IETF meeting I attended, the technical plenary focused on architectural considerations raised by smart, "Internet of Things" devices. Chief among the concerns discussed was how IETF standards could help mitigate the significant privacy and security risks posed by the Internet of Things. At HOPE, by contrast, the keynote address consisted of a surprise appearance by the former National Security Agency (NSA) contractor and whistleblower Edward Snowden, remote from his refuge in Russia. In conversation with Daniel Ellsberg of Pentagon Papers fame, Snowden insisted on the necessary role of technology and technologists in facilitating political dissent.

Appearing on HOPE's final day, Levison discussed not E2E email, but rather the ethical obligations of communications service providers. If Levison was qualified to opine on the business practices required to protect customer rights, it was because he was, per his panel introduction, a "privacy hero" like Snowden himself. On August 8, 2013, without prior notice, Levison had suspended Lavabit's operations. That day, he posted to its log-in page a public letter addressed to his 410,000 "Fellow Users." I have been forced make a difficult decision, Levison wrote, between "becom[ing] complicit in crimes against the American people or walk[ing] away from nearly ten years of hard work by shutting down Lavabit…. I wish that I could legally share with you the events that led to my decision.… Unfortunately, Congress has passed laws that say otherwise."

By the time he appeared at HOPE, the legal wrangling behind this cryptic message had become public knowledge. In July 2013, following media reports that Snowden was a Lavabit

user, federal agents served Levison with a warrant ordering him to turn over the system's private secure socket layer (SSL) keys. A judge had authorized the warrant in connection with the criminal investigation of a specific, individual customer. As Levison knew, however, handing over the SSL key would technically enable federal agents to inspect the private communications of every Lavabit customer. Given the risk of compromising his customers' privacy, and having marketed Lavabit with the promise that it would "never sacrifice privacy for profits," Levison shut the service down. In a nation, which venerates the productive powers of entrepreneurs, Levison chose to sacrifice profit in service of that least charismatic of American values.

When we spoke in 2015, Ladar expressed a certain weariness regarding his Snowden-related notoriety. He wanted to be remembered, he said, for building, not this one, painful act of destruction. He had never intended to become an activist. Rather, he developed Lavabit guided by the desire to create an email service that he himself would want to use. By this he specifically meant a service that would not, like Google's Gmail, which also launched in 2004, surveil its users in order to sell ads. Nor did he wish to eliminate government surveillance. The agents who served the Lavabit warrant may have treated him like a criminal accomplice, but were "completely oblivious to the implications of what they were doing." Raising awareness about privacy and the rule of law now complicated Ladar's ability to "produce code and manage the team." His "obligations as a citizen," however, precluded him from taking his "soapbox" for granted.

I describe my encounters with Ladar here in part to begin to introduce the loosely affiliated people, projects, corporations, and gatherings among whom and within which I conducted the research for this dissertation. As I will describe, in 2016 and 2017, I spent fourteen

months living in San Francisco, conducting fieldwork organized around the privacy and security engineering team of Mozilla, developer of the Firefox web browser. This is not, however, the study of a particular institution or project. Nor is this a study of a professionalized discipline, privacy engineering. Rather, the dissertation presents an ethnographic study of privacy as it has been taken up in the years since Snowden as an object of technological intervention and stewardship.

In the dissertation, I focus primarily on the engineers and computer scientists directly involved in efforts to use software design and engineering to "protect" or "enhance" privacy by building culturally-recognize forms of it into internet technologies. Between 2014 and 2018, however—my period of active study—my research practice involved attending broadly to the institutional settings and social scenes in which privacy's future was figured as a social problem for which technology was both ultimate cause and proximate solution. Approaching privacy in this way necessarily extended my ethnographic scope to include the many other kinds of experts, including user experience designers, researchers, lawyers, product managers, activists, entrepreneurs, and policymakers, professionally concerned with technology's potential to enforce the perceived boundaries between public and private. Ultimately uniting the actors in this study are thus a certain professional proximity and critical engagement with the increasingly common practice of using technology to secure a meaningful social future for privacy. Though I conducted fieldwork primarily in San Francisco, interviewing and attendance at privacy-related events took me throughout the broader Bay Area, including regular visits to Oakland and Silicon Valley, and trips to San Jose, New York, Dallas, Washington, D.C., Ottawa, and Pittsburgh.

Throughout this period, the web browser served for me as both ethnographic object and medium. Thanks to software engineers self-documenting tendencies (d 2008a), the internet's technologically-enabled perfect memory (Mayer-Schönberger 2009), and Mozilla's peculiar status as both open source project and "open by default" corporation (see Chapter 3), I used the browser extensively. Like Mozilla's designers and engineers themselves, and many of 100 million daily Firefox users they serve, the browser formed an essential part of my everyday work flow. I used the browser to communicate with and observe the work of Mozilla's globally-dispersed staff and volunteers, to locate, access, save, and organize resources, and to juggle the innumerable administrative task of everyday life, which now all run through the internet. The quasi-archival materials (Kelty 2008a) available to me through the browser—product hypotheses, blogs, code reviews, conference and classroom recordings, podcast interviews—significantly extended the geographic and temporal reach of my research. Living and working in the browser also directly, viscerally exposed me to the concerns with technical performance, perception, experience, comprehension, utility, and desire, which animate Firefox's ongoing development and condition the form that privacy takes in it as a purpose-built browser feature and property of browser operation (see Chapters 2, 3, and 5).

Unwinding Privacy's Conceptual Disarray

If opening with Ladar provides entree to the dissertation's ethnographic scene, I also invoke him to introduce a particular genre of interaction. While it took me unawares when I first encountered it, I recognize it now as a ritual common to the sites of contemporary world-making around privacy, one indicative of privacy's perceived nature as a social problem. As noted, when I spoke

to Ladar, he quickly assumed control over the direction and rhythm of our conversation. I had been curious to gauge his reaction to my own understandings, but was happy to let him show me what in the world of privacy-enhancing technology he took to be of primary interest.

Early in our conversation, Ladar described Gmail's role in motivating Lavabit's creation. In so doing, he reminded me that when Google launched the service in 2004 it was received as something of a practical joke. With a well-known history of corporate pranks, Google had launched Gmail on April Fools' Day. Moreover, Google offered the service both free-of-charge and with one gigabyte of free storage. This was an unprecedented amount at a time when the leading email providers still made a business of selling storage. According to Ladar, if we were now too far historically removed from the collapse in the price-per-byte to remember, also now largely forgotten was the response of California's legislature. Days after Gmail's launch, a state senator filed legislation proposing to ban businesses from using email to collect personal information. The bill's sponsor characterized Gmail's plan to scan users emails as the basis for targeted ads as an invasion of privacy equivalent "to having a massive billboard in your living room." Google's lobbyists quickly defeated the legislation, in part by obfuscating how Google would use the information Gmail scanned.[1] "We didn't understand at the time," Ladar said, "how the new service would develop. Now business interests are so entrenched that we could never pass equivalent legislation."

Ladar's reminders, his surfacing of these half-forgotten milestones in privacy's intertwined history with the internet, was deeply engrossing to me. It invoked for me both lost

[1] Google combines the data it collects across its suite of services into individualized user profiles. In 2004, Google represented that it would not combine information scanned from emails with the search histories collected through Google search (Levine 2018, 157-64).

social opportunity and personally missed warnings. How had the internet giants like Google become so entrenched in private life—in my life? How the internet had enrolled the world in such a poorly perceived "bargain," so quickly and with such apparent ease, was for me both a personal question and a scholarly interest (see Chapter 4).

Before I could introject, however, Ladar changed topics. Shifting focus suddenly to me, he challenged me to 'confirm my smarts' by sharing my definition of privacy. At this point in my research, I had spent close to a year, off and on, reading the extensive scholarly literature on privacy, tracking privacy-related events in the news, and developing a sense of the contemporary landscape of privacy-preserving technologies. I don't recall whether, at the time, I would have claimed for myself the mantle of privacy scholar. Certainly, however, I had more-than-usual exposure to privacy's theories and justifications. My own personal definition though?

As I will describe, I would later come to identify the assumption implicit to Ladar's challenge, that an individual's personal understanding of privacy wields special social efficacy (see Chapter 2), as illustrating aspects of what I call the "privatization" of privacy. At the time, I was simply relieved when Ladar swept past my fumbling answer—an attempt not to define privacy but to instead explain how I proposed to study it—and demonstrated how I should have responded. Admonishing me for failing to "take the time" to "develop a proper understanding of privacy," Ladar said, "To me, privacy is about control, retaining control over perceptions others have of you, by controlling information." When privacy is breached, "your history, where you were born, when," it unintentionally leaks. "As a result of which, people develop ideas of you that differ from what you intended….The breach of privacy is about losing the ability to control perception." As he would soon specify, for Ladar, losing control over how the self is perceived is

especially troubling when it involves the government, as opposed to a corporation. From his perspective, the significance of government surveillance—the mass surveillance that Edward Snowden revealed the United States to be fanatically committed to—is, in part, that a new class of bureaucrats "has all this private information about you that can produce perverted, distorted perceptions." Because such perceptions "can be used for good and evil," the loss of control is dangerous. Whether or not the government actually abuses its control, losing the ability to dictate what information we provide to other necessarily entails "a loss of respect."

Ladar's definition, with its emphasis on perception and special concern with surveillance's dignitary harms, displays some notable, if culturally legible, deviations from predominant American understandings of privacy (see Chapter 2). The core of his definition, however, as the individual control of personal information, is widely shared among technologists and other privacy professionals. It is the de facto definition operationalized by Silicon Valley's technology companies, including Mozilla. Despite this normative quality, the practice Ladar modeled for me, of explicitly stating, as a form of introduction, one's personal privacy definition, was a regular feature of the privacy events I attended. The practice was most common in settings like panel discussions and other group scenarios involving one-time or irregular convenings across professional and institutional backgrounds. Individual privacy professionals also engaged in the practice, however, prefacing presentations and remarks, or opening meetings, by declaring a definition of privacy.

From a certain perspective, this tendency to explicitly, reflexively engage with privacy's definition reflects the widespread concern that privacy today is in a state of instability and potentially fatal erosion. This perception responds to a broad array of contemporary

developments in technology, corporate and governmental practice, and social norms of individual self-disclosure. Increasingly cheap, powerful forms of digital monitoring, storage, and analysis have, for example, proliferated in our everyday digital and physical infrastructures. Digital surveillance has consequently transformed into an inescapable, if often diffuse and imperceptible, fact of everyday life (Hirschkind et al. 2017; Masco 2017; Marx 2016; Nissenbaum 2010). Meanwhile, corporations ranging from Google to Amazon have claimed individuals' search queries, cell phone location records, and all the others forms of data incidentally generated through use of networked technology in the names of convenience, prosperity, and societal progress (Gitelman 2013; Nissenbaum 2010; Thrift 2006; Terranova 2000). Speaking of changing social norms, Ladar observed that for the generation born after Facebook's introduction, "the line between public and private is very different" than it was for those, like himself, who grew up alongside the internet (see boyd 2014).

Though he didn't mention it, Ladar's personal experience with the national security state additionally illustrates how its "nearly perfect paranoid system" of preemptive counterterrorism (Masco 2014) has upended certain bedrock presumptions of privacy. Historically, American law and culture approached privacy primarily through the rubric of liberty, mobilizing it to designate areas of individual and family life as beyond government intervention (Nissenbaum 2009; Whitman 2004; Nelson 2002). Implicit to this treatment of privacy as a form of protection against state intrusion is the presumption that citizens are, by default, not required to account to the state for private behavior. By contrast, by virtue of its privileges and powers, the state must

transparently disclose to citizens the actions it carries out in their names.[2]

In 2013 and 2014, Snowden's revelations—the recurring waves of newspaper reporting based on the secret NSA documents he released—spectacularly thrust privacy into American public imaginaries as an endangered object of national concern. Animating privacy's newly urgent public status, however, was not only the revelation that the government secretly surveilled its own citizens (or the fact that, to do so, it had infiltrated and co-opted America's corporate and critical national infrastructures—see Chapter 1). Rather, it also reflected public recognition that the government justified secret surveillance using secret interpretations of secret laws—secret interpretations of individual terms even (see Masco 2017). Through Snowden, we learned that U.S. citizens had, in effect, become presumptively suspect, accountable to the state not only for their actually current behavior but also for their potential future behavior. The privileges of the private citizen, meanwhile, had been claimed by the institutions, agents, and authorities of the national security state. As Snowden noted in a 2016 interview, "They know more about us than they ever have in the history of the United States, and some would argue, then any society that has ever existed. At same time,…thanks to aggressive expansions of the state secrecy authorities and even use of classification… they're excusing themselves of accountability to us.… That, I think, leads to an inevitable result over time, whether through good intentions or bad, that the public is no longer partner to government, but merely subject to it."

Given the resulting uncertainties surrounding privacy's privileges, political capacities, and social status, it is perhaps unsurprising that the individuals professionally responsible for it

---

[2] These presumptions reflect the understanding within liberal political philosophy of state power as being necessarily limited, and of rights as being vested in private persons as the proper seat of humanity (see Warner 2002, 39).

feel compelled to declare a precise meaning for the term. Nonetheless, despite the clear challenges posed to privacy by these developments in technology, capital, and governance, privacy has in fact long been haunted by a certain conceptual indeterminacy, an "embarrassing" (Whitman 2004) over-abundance of meanings and justifications. In a 2015 workshop I attended exploring "privacy-by-design," a workshop organizer recounted privacy's common meanings for us. In addition to informational control, this only partial list included the right to be left alone, limited access to the self, secrecy, a zone of autonomous decision-making, intimacy, personhood, anti-totalitarianism, and contextual integrity.

Faced with a similar breadth in the privacy definitions recognized in American law, the legal scholarship is littered with attempts to systematize privacy and tame its unruly profusion (see, e.g. Bloustein 1964; Prosser 1960). Such attempts theoretically serve to identify the precise interests privacy serves, and thus to determine whether and to what degree they are supported by relevant legal authorities. More broadly, however, privacy's conceptual "incoherence," its multiple, and indeed sometimes contradictory, meanings is taken within the legal and philosophical literature as a sign of some fundamental flaw. For some, it indicates that privacy has no "essence" or distinctive social value. From this perspective, when we speak of privacy we are really speaking of the other interests and rights to which it privacy is properly reduced, things like property, liberty, and bodily security (see, e.g., Thompson 1984). For others, privacy's multiple meanings represent an invitation to find the one true thing in the world to which the term uniquely, universally, a-historically points. With the notable exceptions of Nissenbaum (2009) and Solove (2008), each of whom laments the purification required to eliminate privacy's

ambiguities, much of the existing scholarly literature[3] consists of attempts to divine privacy's distinctive essence (see, e.g., Inness 1992; Gavison 1980; Bok 1982; Fried 1970; Westin 1967). Implicit in the obsessive focus of this literature is a belief that rigorously defining privacy is a necessary predicate to addressing the critical social problems to which the term ambiguously points.

If the practice of sharing one's privacy definition partakes of these anxieties, it also reflects to a degree the practical necessities of technological development and professionalization. Privacy historically emerged as a concept in the United States as the object of lawyers and philosophers, legislators and judges. Its meanings were elaborated primarily in the context of social conflicts over the government's right to intrude in private spaces, relationships, and decisions (Whitman 2004; Nelson 2002), not in relation to the properties and operations of large-scale technical systems. Systematic efforts to develop privacy definitions attuned to the nascent reality of large-scale data processing began in the 1970s (see Gellman 2022). Technology-based efforts to defend privacy can similarly be traced to 1970s, with the Cypherpunk embrace of encryption. Privacy-by-design and privacy engineering, however, only began to emerge as a field of practice and professionalized discipline, respectively, in recent years. As I explore in Chapter 1, for example, it is only since Snowden that American engineers

---

[3] Unlike the largely normative legal and philosophical works I focus on here, feminist scholars have produced a distinctive and important body of scholarship that takes a primarily historical approach to privacy. They have shown how, as an organizing principle of political life, the public/private distinction is animated in part by its association with the categorical distinction between genders (Landes 1988; Fraser 1990). In consequence, it has historically served to obscure and legitimate forms of gendered domination, degradation, and abuse in the so-called domestic sphere. See Davidoff and Hall (2019); Gal and Klingman (2000); Landes, ed. (1998); McKinnon (1989); and Allen (1988).

and computer scientists have explicitly identified themselves, by virtue of their expertise, ethical commitments, and the global reach of their products, as the social actors best positioned and most willing to steward privacy through the turbulent present. The development of proven methodologies for building privacy into technology remains a work in progress and a source of consternation for technologists and policymakers. At the level of individual projects, however, engineering privacy necessarily begins from definition. Only relative to a specific definition, and thus specific, mitigable risks and harms, can computer scientists and engineers devise and implement the system controls and other measures required to offer privacy to technology users. Similarly, efforts to transform privacy engineering into a "science of privacy" with metrics of success commensurate to those of established fields like systems engineering are understood to rise or fall with the rigor of the definition in which they are grounded.

Still, even if engineering privacy requires it to be precisely defined, few of the professionals I met described technology as independently sufficient to secure privacy's future. Consider in this regard the ambivalence expressed by Ladar. Near the end of our conversation, I broke in to ask, given the dire picture he had painted of entrenched corporate power and pervasive government surveillance, what role he ascribed to technology in privacy's defense. "I know it can do something," he replied, "but is it too late to do anything meaningful?" From here Ladar pivoted, proposing a solution grounded not in technology but the law. Nothing like it, he admitted, could ever be realistically passed, but to restore the necessary forms of control and dignity, he proposed outlawing corporations from sharing personal information. In the 1970s, Cypherpunks had dreamed that encryption heralded inevitable "crypto-anarchy," a coming age in which the only laws would be those enforceable through computer code. Similarly, in the post-

Snowden era, federal law enforcement officials have continually characterized efforts to spread privacy-enhancing technologies as an assault on the rule of law. Like the other technologists I met, however, Ladar described his use of encryption as instead a defense of the rule of law and a complement or stopgap to legislative solutions. Technologists generally acknowledge that there is no easy, "silver bullet" technological fix for privacy, no switch that can be flipped to turn privacy on or off—despite what users may sometimes think. To preserve personally meaningful yet politically effective forms of privacy for the future is thus necessarily a collective venture spanning technology, law, design, education, activism, and entrepreneurism (see Chapter 4).

What I would suggest is that the scholarly and practical anxieties surrounding privacy's supposed conceptual disarray obscure as much as they reveal. In her study of privacy in the Cold War, for example, Deborah Nelson (2002) helpfully demonstrates that since the 1890s, privacy has emerged again and again in American social imaginaries as an eternal value "suddenly" imperiled by new technologies. If privacy has thus been continually mistaken as something universal and stable, so too is its "lost" coherence a "nostalgic, amnesiac fantasy" (Ibid. at 37-9). Nelson argues that it was only during the Cold War that the search for a means of securing privacy's conceptual coherence took on its characteristic contemporary anxieties. Under Cold War logics, which figured privacy's sanctity as the distinguishing feature between American democracy and Soviet totalitarianism, the perception of privacy's fragmentation took on a gendered structure and national psychic import. To tolerate privacy's incoherence in this era was to invite emasculating national penetrations (Ibid. at 55, 172).

The compulsive reduction of privacy analysis to the search for privacy's essence obscures such historical contingencies. It further elides the heterogeneous ways people take up privacy in

their lives, mobilizing it to achieve varying political and social effects, and making meaning in poetic relation to it. Especially in group settings, for example, the sharing of privacy definitions can be understood as a mechanism for achieving forms of mutual orientation, ones which bring about individual and collective transformations with clear ritual entailments. Engineers and computers scientists may view privacy's future as necessarily running through corporate technologies, but they recognize that any comprehensive solution to its contemporary destabilization will necessarily draw upon disparate forms of professionalized expertise with shallow histories of collaboration. At any given privacy event, therefore, a panel discussion might feature a computer scientist in dialogue with a regulator and a constitutional lawyer, or an academic alongside entrepreneurs and activists. In such contexts, sharing definitions serves to pragmatically cut through participants' disparate institutionalized histories and languages. As one panel moderator put it, it is necessary "to bring us all to the same page." Even as such sharing only temporarily orients a limited set of individuals to a specific definition, the act performs in miniature—it dynamically figures—the forms of collective, interdisciplinary mobilization and collaboration understood to be necessary for privacy's future.

Across the professional spectrum today, the challenges involved in defending privacy are widely articulated in terms of "tradeoffs" between competing and occasionally "incommensurable" values and interests—between privacy and security (Chapter 1), for example, or between computer security and national security (Chapter 1), usability and utility (Chapter 3), privacy and innovation (Chapter 4), and privacy and technical "performance" (Chapter 5). Given these structural tensions, the defense of privacy is often described as, with certain limitations (see Chapter 1), necessarily an exercise in compromise and sacrifice. The

sharing of personal privacy definitions, from this perspective, bears further significance as a demonstration of individual willingness to give, even if in temporary, circumscribed form.

Finally, recall that when Ladar admonished me, his specific complaint was not just that I didn't understand privacy, but that I'd failed to take the time to develop an understanding of it. As he would go on to describe, Ladar himself had always been interested in privacy. In the three years since his standoff with the national security state, however, his views had changed significantly. He had, during this period, set about to define for himself what privacy is guided by "his natural proclivities as an engineer," his interest in rules and impact, cause and effect. If from Ladar's perspective privacy is used in popular discourse today "extremely loosely," this is presumably because the general public has declined to pursue its own such journeys of privacy learning.

Reflected in Ladar's admonishment was the widely-shared sense among privacy professionals that privacy is burdened by a structure of insufficiency, one that is not just conceptual, as described above, but also imaginative and affective. In diagnosing privacy's apparent lack of charisma, technologists point to its perennial failure to summon and command displays of social care. They attribute this care deficit in part to the diffuseness and indirectness of privacy's harms, and to the difficulties involved in representing privacy as manifested in the properties of large-scale technical systems (see Chapters 2 and 3). When someone trespasses upon your property, you see it. If they violate your bodily integrity, you feel it. The internet's popularization, however, and its extension into the intimate recesses of everyday life, are viewed as having severed the embodied capacity to directly register privacy violations (see Chapter 2). Under the so-called "privacy paradox,"Americans widely profess to value privacy, but when

faced with actual privacy-related choices, they consistently fail to enact their stated beliefs in actual behavior. Corporate and academic technologists have produced an extensive body of research exploring this perceived paradox. In it, they lament users' ongoing refusal to take the active steps recommended to protect their individual privacy. They nonetheless diagnose the gap between privacy belief and behavior in terms of an absence of sufficient care for privacy (see, e.g., Gottleib 1996), thus eliding, among other possibilities, the role of power disparities in preventing belief from converting to behavior. Regardless, relative to the Tinkerbell-like quality that technologists attribute to it, to share one's definition of privacy is to identify oneself, at least professionally, as caring enough about privacy to properly understand its social interests.[4] It is to identify oneself as a member of one of the multiple, partially articulated publics organized today in privacy's name.

In popular usage, "the" public refers both to a kind of social totality, to "the people in general" defined at whatever social scale (Warner 2002, 65) and to a language-based mode of political legitimation (Gal and Woolard 2001). "A" public, by contrast, is any self-organizing community, which exists independently of the state and similar social institutions and asserts

---

[4] Indeed, professing one's care for or commitment to privacy (as opposed to one's mere compliance with applicable privacy rules) is another common feature of definition sharing. As I explore in Chapters 3 and 4, the moral imaginaries of Silicon Valley's engineers must be distinguished from those of the corporations that employ them (see Kelty 2008b). Mozilla's privacy and security engineers, for example, enjoy productive working relationships with their peers at Google's Chrome browser. They consistently describe Chrome's engineers as "really" caring about improving privacy and security on the web. They nonetheless view their capacity to do so as being constrained by the market-imposed demand for growth to which Google but not Mozilla is subject, and by the specific value logic through which Google generates revenue. Given such perceptions, for corporate privacy engineers, sharing a personal definition or professing one's care, constitutes a means of identified oneself as morally aligned with and thus allies of potentially skeptical privacy professionals from civil society and academia (see Gal and Irvine 2019, 164 on "fractal solidarity").

itself as a check on entrenched power (Kelty 2008a; Warner 2002). Publics in this latter sense have been of interest in the social sciences as formations through which political subjectivity comes into being, sites at which strangers come to self-understanding as actively belonging to some historical social entity (see Cody 2011). Study of the particular conditions under which publics form has repeatedly shown that their rhythms, temporalities, and political capacities are contingent upon the nature of the circulating media and communicative infrastructures through which they organize and operate. The form of "disinterested," rational discourse, which Habermas (1989) describes as emerging from the mass circulation of printed books and newspapers is thus, for example, very different from the "argument with and through technology," which Kelty (2008a, 58) describes as characteristic of the "recursive publics" formed via the networked circulation and modification of software code.

In liberal thought, by setting and maintaining the boundary between state and society, privacy creates the space in which citizens are able to form themselves as individuals (see Habermas 1989). Identifying privacy as the object of its own contemporary publics, however, and indeed of its own recursive publics (see Chapter 3), shows that privacy is not only a condition of subject formation. It also operates today as the organizing principle of its own communities, as the reflexive basis for its own ethical and professional subject positions. In so doing, privacy today draws social action forward (see Gal and Irvine 2019, 12-4). By reference to the scholarly literature on publics, it further suggests that privacy should not be approached as a universal, timeless thing in the world, but rather as produced in specific historical scenarios. Like publics (see Marres and Lezaun 2011), privacy is conditioned by the particular material,

institutional, affective, and conceptual tools and strategies applied to bring it into the world and sustain it.

In this dissertation, I resist the injunction to circuit analysis through the question of privacy's definition or to devise new mechanisms for identifying the 'true' nature of its being in the world. The issue is not that available definitions are too general or too specific, but rather that the search for specificity cordons off privacy's possibilities precisely when worldly conditions suggest the need to open them up (see Mazzarella 2019, 47; Paley 2008, 5). Similarly, I reject the understanding of privacy as a description of the world or a stable sociological domain defined by a given set of places, things, spheres of activities, or types of interaction. I instead approach privacy as a flexible evaluative grid through which individuals ideologically assign values to objects and practices, thus arguing in and about the world (Hirschkind et al. 2017; Cody 2011; Povinelli 2006; Gal 2005 and 2002). Recognizing privacy's grounding in a communicative phenomena with its own systematic logic helps unwind its apparent conceptual contradictions. It also helps explain how it is that privacy continues to be treated as universal and stable when evidence of the historical shifts in its "boundaries" can be found everywhere from the shifting anchors of privacy rights recognized in constitutional law (see Chapter 2) to the changing norms of personal revelation employed in lyric poetry (Nelson 2002).

In the semiotic framework developed by Susan Gal (2005, 2002), privacy works in tandem with publicity as a co-constitutive distinction. As Gal explains, in actual events of language-in-use, the public/private distinction operates as an "axis of differentiation" (Gal and Irvine 2019). By applying it to different kinds of social "objects," whether spaces, bodies, interactions, or identities, people pragmatically categorize and thus characterize them according

to contrasting qualities. Such qualities, in turn, link their objects to particular images, values, rights, aesthetics, and ethical perspectives, which are socio-historically imagined to be emblematic or iconic of the private. Once established as culturally resonant, the public/private distinction's semiotic logic forms a mobile "scaffolding" (Gal 2002, 85), which people take up and apply "ideologically," i.e., in politically- and morally-interested and socially-positioned ways (see Gal and Irvine 2019). Such use aids individuals in creatively interpreting their situations and inspires politically-consequential forms of arguments and social action. It is by virtue of this semiotic logic, for example, that privacy, historically apprehended in the United States through spatial and topographic metaphors, can be extended to encompass novel, non-spatial objects and activities like personal data and data processing. Moreover, it is through such extension, not by virtue of some inherently private nature of data, that the rubrics of privacy law become potentially applicable to it, and that civil society claims regarding the forms of data processing corporations may engage in without reproach gain social legibility and force.

For Gal, the co-constitutive nature of the public/private distinction means that not only can it be applied to any social object, but also that it is "fractal." It can be used at any social scale, reproducing its contrasts in ever-narrower and -broader social contexts. It is through the distinction's recursive application, for example, that the house can be the paradigmatic private space recognized in the United States but still be internally divided into its own public and private areas. Meanwhile, the distinction's repeated application to personal (i.e, private) data, already categorically distinguished from non-personal (public) data, has resulted in its subdivision such that it is now understood to contain its own private (communications content) and public ("mere metadata") forms.

As I explore in Chapter 2, the public/private distinction signals indexically. As "shifters," the precise referential meaning communicated by the terms necessarily depends on the context in which they are used. Consequently, the precise meaning of the private shifts to some degree with each fractal recursion, whether or not through deliberate effort to recast its definition. Gal argues (2002, 90) that it is because privacy retains its name across these various embeddings that it can be experienced as stable despite obviously changing content. The term's apparently contradictory meanings can be reconciled by recognizing that people mobilize the public/private distinction in distinct socio-historically situated projects to achieve heterogeneous and sometimes conflicting effects. Further, when the distinction's application to some social object proves inconsistent with an applicable ideological schema, it is likely to be indexically erased or denied by shifting attention to others levels or contexts of distinction (see Gal 2005, 27; 2002, 91). Viewed from the perspective of its semiotic logic, privacy's conceptual 'incoherence' is not a flaw, but a sign of its efficacy as a tool of creative, political and ethical argument and experimentation. Privacy's multiple definitions reflect not the troubling absence of an essential referent for it in the world, but a rich history of baptismal events (Silverstein 2003), ideologically-inflected uptakes, each of which has reformulated to some degree the categories privacy is understood to instantiate and the qualities it is perceived to display (see Gal 2017).

Building from this framework, the challenge I take up herein is to ethnographically embrace the specificities of privacy's 'disarray' by exploring a set of socially-significant, technology-related projects organized in privacy's name. Across these sites, I tease out the norms, modes of reasoning, forms of practice—not just technical, but also aesthetic, affective, semiotic, calculative, and more broadly material—through which internet technologies are

invested with privacy's formal, conceptual, and phenomenal qualities. In so doing, I ask how, relative to the social positions and political and moral interests of the technologists and corporations taking it up, is privacy being remade? What are the effects of such projects on the values, sensibilities, aesthetics, and imaginaries conventionally linked to privacy, and on privacy's theoretical and practical political capacities? How do the forms of privacy that emerge from such projects, and the technologies in which they are embedded, change the subject positions and identities available to technologists and their real and imagined users? Through such questioning, the dissertation presents both an empirical investigation of a key category of political theory in the vein of Warner (2002), Dean (2002), Kelty (2008) and Coleman (2012), and an attempt to take seriously the constitutive role of material processes in social and political life (see Whatmore 2010).

At the same time, if there has been a marked expansion and intensification in the efforts of American technologists to build privacy into networked technologies, it is because, to some degree, they recognize in privacy a potential solution to the social problems unleashed by the internet. In this respect, attending to the technology-based projects in which privacy figures as a pressing social problem presents an opportunity to explore privacy as not just ethnographic object, but also lens. In the dissertation, I thus examine the sites of contemporary world-making around privacy for what they can tell us about the ways the development, harnessing, and repurposing of internet technologies by government and corporations renders life in the United States problematic (see Lakoff and Collier 2004).

The Last Generation with (a) Choice

Euro-American thought has long figured privacy as the concept against which the public is defined. Pointing to the term's origins in the latin *privatus*—deprived—Warner (2002, 28) observes that as the mere negation of the public, privacy was originally viewed as having no independent value. Post-Enlightenment conceptions of the public continued to treat privacy as a residual category. Kant (1996, 60-1), for example, identifies the "public use of reason" as the very engine of enlightenment progress, but characterizes the "private use of reason" as being merely vocational. Habermas (1989) similarly treats privacy as the training ground for critical public reflection. As a justification for property ownership, he argues, privacy is what secures the autonomy necessary for individuals to exercise rational-critical discourse relative to state power.

Despite such treatment, privacy's accumulated wealth of definitions and justifications shows its repeated uptake in morally- and politically-infused historical projects to have imbued it with it distinctive values. Within anthropology, however, privacy continues to be treated as largely subsidiary to publics, the leftovers of public formation. Driven by the desire for new forms of post-socialist and post-capitalist public culture (see Rabinow et al. 2008, 36-9), and the post-Habermasian understanding of publics as engines of self-determination (Hirschkind et al. 2017; Cody 2011), contemporary anthropologists have produced numerous studies of publics and public formation (see Cody 2011; see, e.g., Marres and Lezaun 2011; Kelty 2008a; Hirschkind 2006; Bernal 2005; Gal and Woolard 2001). They have also paid considerable attention to orthogonal concepts, which like privacy are understood to partake of the hidden or secluded. These include studies of secrecy (Masco 2014; Taussig 1999), intimacy (Povinelli 2006; Shryock 2004), domesticity (Carsten 1997), and gender/sexuality (see Warner 2002, 63). Privacy mostly

appears in such works in negative relief, but is still depicted as capable of shifting the definition of political life, the boundaries between state, society, individuals, and family, and the locus of control over economic activity. Nonetheless, with few notable exceptions (Bridges 2017; Hirschkind et al. 2017; Agrama 2012; Herzfeld 2009; Gal 2005 and 2002; Hill 2001; Gal and Kligman 2000), little disciplinary attention has been devoted to privacy proper.

To appreciate the need for such studies and to understand what I mean by privacy proper, let's turn not to privacy's legal and philosophical justifications, but rather to its professionals, and to the widespread sense they express of the present moment as one of generational transition and foreclosing possibility. I have already described Ladar Levison's frustration with the loss of individual control over personal information. Ladar's diagnosis of the present, however, was more encompassing in scope than indicated. During my initial introduction, I began to describe to Ladar how, having previously practiced the law, I was drawn to privacy's study in part by the apparent shift in the kinds of social agents responsible for its future. I pointed out to Ladar the increasing prominence of privacy engineers among the ranks of Silicon Valley's technologists and at prominent civil rights organizations, like the ACLU, which historically approached privacy through litigation and advocacy.

It was precisely here, however, that Ladar cut me off. Interjecting, he argued that privacy merits critical study today not because of the shifting forms of expertise mobilized in its name, but as a function of the broader impacts of the "information age." From Ladar's perspective, the relevant periodization begins not in 2014, when Edward Snowden urged his fellow engineers to "take back the internet" and save privacy, but earlier. It had taken 40 years for revolutionary technologies like electricity and the telegraph to really "impact" society, he told me. We should

therefore expect the same to be true of the internet. But, it had been only 40 years, more or less, since the development of microcircuitry. Even less time had passed since the 1990s' "explosion" in computer storage capacity. There was no way to know yet the impacts of such developments, he said, let alone how to deal with them. It was thus a privilege even to be able "to see the transition."

Despite his protestation, over the course of our talk, Ladar sketched for me an image of the unfolding technological "revolution" he identified with the transistor and internet. Central to his account was the observation that the internet has enabled forms of for-profit and government surveillance, which strip individuals of control over their personal information. In so doing, Ladar argued, digital corporations and government bureaucracies have claimed for themselves the very material grounds of personal identity and perceptions of the self. While describing these developments, Ladar gestured to certain concrete harms which might follow if such control were abused. What primarily structured his account, however, was not a catalogue of the harms of revolution but rather an intense concern with the shifts it induced in the topology of individual and societal choice.

Recall that Ladar described the loss of control over personal information as problematic specifically because it entails a loss of the ability to choose which aspects of the self to share with others. Related to this, in Ladar's framing, is a loss of the individual ability to change identities at will, i.e., to choose, in classic American fashion, to begin again, to take one's second chance. The individual ability to make certain kinds of choices has, from this perspective, been technologically foreclosed. Other kinds of choices, however, especially those relating to tracking itself, Ladar described as remaining open, but only at increased social cost. One might choose,

for example, to adopt behaviors or lifestyles for the express purpose of maintaining control over information. In theory, one could forego entirely credit cards in favor of cash to avoid facilitating the tracking of one's financial transactions. From Ladar's perspective however, given applicable banking regulations, to do so is to effectively declare oneself criminally suspect, triggering enhanced government scrutiny.[5] In such ways, Ladar concluded, "[i]t's becoming more and more difficult to maintain control and still be able to function in society." As another interlocutor put it, people have resigned themselves to online surveillance because they feel "they have no other choice if they want to live in this world." The "choice" of whether or to be in society thus marks a kind of an imaginative limit or horizon within which all other individual choice is today understood to be circumscribed.

In Ladar's telling, these foreclosures of individual choice should be understood in part as the result of other choices—collective, societal choices. The practical inability to avoid financial tracking, for example, Ladar ascribed to the "decision" of the "body politic" that "it's ok for the government to have banking info." In other respects, though, technological revolution also forecloses forms of societal choice. This was one conclusion Ladar derived from Gmail's history. Ladar was ambivalent regarding whether the public had ever known enough about Google's business to appreciate the societal choice that Gmail represented, at least retrospectively. Either way, the moment of social choice had passed. Business interests were now so entrenched, the necessary legislative fix for e-mail-based surveillance was no longer practically possible.

---

[5] Following Snowden, privacy professionals frequently made similar claims with respect to encryption. Strong encryption was by this point entirely legal but still not widely implemented by default in internet technologies. To download encryption tools—to attempt to conceal one's digital communications—was presumed to invite one's name to be added to an NSA list.

Technological revolution primarily figured in Ladar's account as something corporations and governments harness and wield to foreclose individual and social possibilities. His concern with choice, however, also contemplated the ways revolution enhances the individual capacity, using technology, to intervene in and reconfigure social choice. Moreover, implicit in his treatment was a folk theory of choice's social conservation, an understanding that choice is never eliminated so much as relocated, shifted across social scales from individuals to collectives, or vice-versa, and re-distributed as between institutionalized centers of power. Both of these qualities were implicit in Ladar's discussion of Lavabit and Dark Mail.

Recall in this regard that when describing his aspirations for Dark Mail, Ladar cautioned that he had no interest in the wholesale elimination of government surveillance. Eliminating government surveillance, in his characterization, was a "big moral and philosophical question." By implication, it was not the kind of choice an individual technologist should make on society's behalf. Rather, as with the government's attempt to obtain Lavabit's SSL key, what he objected to was non-particularized surveillance, surveillance that appeared to conflict with the social distribution of choice effected by the U.S. Constitution. With the Lavabit warrant, Ladar faced what he described as a choice of whether or not to technologically enable the government to conduct mass surveillance—to make it technically possible for the government to make investigative choices disallowed by the rule of law. Similarly, with Dark Mail, Ladar aspired to technologically intervene in the choice calculus surrounding mass surveillance by "remov[ing] the remote service provider from the surveillance equation." Properly designed and implemented encryption, from this perspective, held the promise of permanently precluding corporate email providers from 'choosing,' willingly or not, to betray the security promises made to customers.

Across the arc of my research, Ladar was far from alone among engineers and other privacy professionals in framing the present in terms of historical transition. While opinions varied regarding the present's precise location in the arc of transition, general consensus held that we were close to the end of one moment and at the very beginning of something new. Like Ladar, others described the nascent future as unpredictable, even unimaginable. The only certainty was the accelerating rate of technological change. At the 2016 meeting of the International Association of Privacy Professionals, for example, Chris Kelly, former Chief Technology Officer of Facebook, described the internet as having inaugurated a transition from an unrecorded world to a recorded one. In consequence, he said, we were in the beginning of a technological revolution, which would upend "how we communicate, think about physical spaces, and share experience." Similarly, opening the 2015 Computers, Freedom, and Privacy conference, Nuala O'Connor, President of the non-profit Center for Democracy and Technology, described, relative to technological change, how we were "heading to a future we might not recognize," in which existing relations "between people, technology, government, and industry" were sure to be reconfigured.

As with Ladar's, the accounts of unfolding technological revolution I encountered were generally structured around concern with choice's shifting social topography. Many spoke, for example, of being "the last generation with a choice." For some, this phrase was retrospective and mournful. It meant that we had been the last generation with the ability to make some choice —to conduct our affairs anonymously, for example—which future generations would now never enjoy. Often, choice was framed as having been lost without our even noticing or realizing it until forced to do so by spectacular public events like the Snowden revelations. Others, like

Ladar to some degree, looked to the future with cautious optimism, framing the generational foreclosure of choice as incipient, but still open to intervention. Those who spoke in this way had varying articulations of the particular choices under threat. Some foregrounded the choice of whether as a society to have privacy. Others spoke more generally of whether to limit corporate and government data practices in the name of human flourishing. For some, concern lay not with particular choices, but the individual cognitive capacity to choose at all (see Chapter 2). At Mozilla, the issue was often articulated in terms of preserving competitive choice among open, interoperable, empowering web technologies (see Chapter 3). However articulated, the presumption held that, given technological path dependencies, entrenched corporate power, the waning of post-Snowden public outrage, regulatory capture, and legislative dysfunction, to fail to choose now would permanently remove choice from the reach of future generations.

In drawing out this image of technological revolution, with its entailed foreclosures of kinds of and capacities for choice, my intention is not to endorse it. Rather, I do so because foregrounding choice as a structuring element of privacy professionals' historical imaginaries and valorizations of privacy raises two points, the first methodological and the second ethnographic and theoretical. Methodologically, privacy's valorization in relation to choice, and choice's treatment as distributed and conserved across social scales and settings, helps to justify the dissertation's research design. Per the preceding section, I am interested in the dissertation in how people mobilize privacy to make sense of, argue about, and carve up the world. Not all applications of the public/public distinction, however, are equally perduring or consequential. Privacy is frequently enacted in highly subtle, contingent, and ephemeral ways. Simply by using phraseology or terminology comprehensible only to certain individuals or kinds of people, for

example, one can enact privacy conversationally, pragmatically drawing social boundaries around a conversation (Silverstein 1976). Similar effects can be achieved using gestures and bodily positioning, such as placing a hand over one's mouth or turning away from a nearby onlooker (Gal 2002). On a crowded beach, by carefully positioned an umbrella, one can effectively designate a patch of sand as private (Nippert-Eng 2010). To whatever degree people feel compelled to respect this figuration while the umbrella remains up, the regulatory effect is likely to fail once it is removed.

By contrast, now consider judicial chambers. Courthouses, generally speaking, are public spaces. In courtrooms, the ideal of democratic institutions as public—accountable to citizens—because "open" and available to inspection is ritually enacted in the rules granting entry to anyone (Gal 2005). Nonetheless, not all spaces in a courthouse, or even courtroom, are equally public. One might ascribe the private status of a judge's chambers to architecture's regulatory effects. From this perspective, if we intuitively apprehend that some persons but not all may enter chambers at will it is because of their spatial remove from other, more obviously public spaces, because their walls impose physical barriers upon entry, because the brass plaques affixed to their doors baptize them in the name of some person. More broadly, however, we might say that it is through courts' institutional levers of power and authority (Gal and Irvine 2019, 136-9), including but not limited to their internal physical design, that private spaces can be carved out within otherwise public buildings. Similarly, peeking under a beach umbrella might earn a stern look of reproach. Thanks to courts' particular institutional powers, however, barging into judicial chambers will instead trigger forceful physical removal and potential legal recrimination.

As I have described, the public/private distinction is available as a discursive resource for creative application by any social actor. Perhaps appropriately then, the preponderance of the social scientific research that examines it as an empirical phenomena does so at the level of individuals. Such literature deftly identifies the practices through which different kinds of people seek to enact privacy in various everyday scenarios, as well as the social effects they hope to achieve in so doing (See, e.g., Nippert-Eng 2010; boyd 2014; Marwick and boyd 2014). Not every kind of actor, however, is equally able to imbue a novel application of the public/private distinction with normative force. Rather, it is through institutions, like government entities and corporations, that any recalibration of privacy's values and qualities thus enacted can be standardized and made durable and socially forceful.

If this ability to determine the categorizations and qualities conventionally linked to privacy explains the dissertation's focus on institutional sites, why focus on technologists, and why Mozilla in particular? On the one hand, despite Ladar's minimization of it, the past decade has seen a shift in the forms of professionalized expertise deployed in privacy's service. This is not the first such shift in the history of privacy in the United States. Though often portrayed as a timeless and quintessentially American value,[6] privacy has not always been a legally enforceable right. As the U.S. Supreme Court painfully reminded us in 2022, the Constitution does not explicitly enumerate a right to privacy. Historians sometimes describe the colonial era as one of substantial personal privacy. Sarah Igo (2018) clarifies, however, that whatever privacy adhered

[6] In 2014's Riley v. California (573 U.S. 373), for example, a famous Fourth Amendment case, the U.S. Supreme Court posited privacy, in the form of the general revulsion inspired by the British practice of rummaging unconstrained in colonial homes for evidence of criminal activity, as a driving force behind the American Revolution.

in that period did so as an implicit privilege of property ownership. By the 1890s, when systematic calls for a right to privacy were first articulated, privacy had become the object of bourgeois, patriarchal social norms, enforced in ways highly marked by gendered, class, and racial privilege (Ibid., 21-4). The subsequent emergence of the law as the default means for defining and enforcing privacy responded to the increasing common sense that social norms were no longer capable of doing so (see Nelson 2002, 51-5).

Now, with the emergence among technologists of forms of professionalized privacy expertise, and with technologists' increasingly vocal claims to be privacy's most faithful and effective social stewards, new modes of reasoning, forms of practice, and moral imaginaries are being brought to bear on privacy. Technologists' proliferating efforts to defend privacy present an opportunity to study the tools, processes, and logics through which privacy is established and sustained in the world just as they are being reflexively developed, tested, refined, and to some degree, standardized. It further provides opportunity to examine how technologists transform technical expertise into political authority over the expanding swaths of collective life mediated by the internet. As I explore in Chapter 1, such claims now extend even to the legitimate objects and methods of the national security state.

On the other hand, as suggested by privacy's portrayal as a victim of information revolution, privacy's fate today is understood to be intertwined with that of the internet. Privacy's changing social status coincides with and marks changing American valorizations of the internet (see Chapter 4). Technologists may describe privacy as necessary to preserve choice, and choice as being distributed across social scales and settings. For practical purposes though, the re-circuiting of everyday life through the internet means that any effective effort on behalf of

privacy, whether grounded in technology, the law, or some combination of means, will ultimately be materialized, made salient and actionable to individuals, in their guise as users of internet technologies. Silicon Valley's internet corporations created the technologies and developed the business practices largely responsible for privacy's contemporary destabilization. In the post-Snowden moment of renewed national anxiety about privacy, however, American technologists, lawyers, and activists turned cautiously towards them. By virtue of the global reach of their products and the imaginative power they wield as avatars of innovation and progress, Silicon Valley's corporations emerged as the social actors best situated to counter-balance the federal government's embrace of mass surveillance. Privacy professionals and civil liberties experts are frequently asked today what the best means is for everyday individuals to protect individual privacy. Without fail, in answering, they point in one way or another to technology related consumer choices. One common response is to suggest using only products developed by companies known for prioritizing privacy and security— iPhone over Android, Firefox over Chrome, DuckDuckGo over Google search. Another common answer prioritizes becoming familiar with and taking advantage of the privacy-relating settings and preferences of internet products and services. It is in part through such settings that the distinctive ways corporations apply the public/private distinction to carve up user data and data processing activities into public and private kinds is institutionally ritualized and coercively applied to individuals as users. From these perspectives, the future of privacy runs through the internet, and thus through the venture capital-backed technology corporations, which define most Americans' understandings and experiences of it.

Against this backdrop, Mozilla presented a particularly compelling site around which to organize ethnographic fieldwork. As Mozillians like to say, Mozilla is in but not of Silicon Valley. As the developer of the Firefox web browser and other products, Mozilla serves approximately 100 million users of the world wide web every day. This number is large enough for Mozilla to exert influence over public policy and standards settings processes. When I began my research in 2014, however, Firefox ranked only fourth among the most popular U.S. web browsers, trailing Google's Chrome, Apple's Safari, and Microsoft's now-defunct Internet Explorer. Mozilla wields nothing like the economic power of its primary competitors, but still bears special consideration by virtue of its unusual corporate structure, its techno-moral mission, and its unique role in the histories of the web, the internet, open source programming (Chapter 3), and Silicon Valley itself.

While the Mozilla Corporation, which directly produces Firefox, is itself a for-profit, it is owned by the Mozilla Foundation, a non-profit. The two entities, collectively known as Mozilla, conduct different kinds of activities. The Mozilla Corporation builds technology and competes in technology markets. The Mozilla Foundation conducts digital literacy and related educational campaigns, seeks to influence technology-related policymaking, and builds community in part by financially supporting technology activists and researchers. The work of both entities, however, is ultimately guided by principles set out in a ten-point Manifesto. The "mission" encapsulated therein, is to preserve the web as an open, accessible, and empowering global commons. As a formal part of this mission, Mozilla has long been committed to preserving privacy, not just for users of Firefox but on the web in general. As this suggests, Mozilla's commitment to privacy is formally circumscribed by its overarching commitment to the web. Mozilla's executives and

staff, however, generally view the contemporary internet as mediating so much of everyday life today that for all effective purposes 'online life is real life.' As I explore in Chapter 3, the ways in which Mozilla understands privacy to serve its mission have changed over the years. During my fieldwork in 2016 and 2017, under the guidance of a newly hired privacy manager, Mozilla's designers and engineers were engaged in a newly urgent effort to ship compelling privacy-enhancing features, which would help Firefox attract and retain users. This mandate reflected Mozilla's strategic decision to compete in the browser market in part by producing features that its competitors, by virtue of their business models and the demands of growth, would not be able to match.

Mozilla's efforts on behalf of privacy include education, advocacy, standards setting, and community-building, in addition to technological development. It seeks to honor its commitment to privacy in its products both by building technologies intended to enhance user privacy in some way, and by following internal policies and practices designed to ensure that its other products never expose users to unnecessary or unintended privacy risk. My fieldwork with Mozilla's privacy and security engineers and user experience designers and researchers provided direct ethnographic access to a socially significant site of technological world-making around privacy. By attending to how Mozilla's engineers reflexively distinguish their work from that of industry peers by virtue of Mozilla's non-profit status and mission, I was able to develop a sense of the role that expectations of exponential growth play in conditioning and constraining the forms of privacy produced in Silicon Valley. Through interviewing and extensive fieldwork in the Bay Area's vibrant, privacy-oriented public life, I refined and elaborated this understanding.

Mozilla's appeal as a base for my fieldwork was compounded by its unique roles in the histories of American technology and business. Though often conflated with the internet, the web is, in fact, technically distinct. Like email, the web is an application of the internet (see Clark 2016). It consists of the global collection of hyperlinked documents, resources, and applications that we think of and experience as websites. It is accessed using a web browser, itself a complex software application, and is defined by a set of protocols for building, formatting, and locating webpages, and communicating with webpage servers (HTML, CSS, URL, and HTTP, respectively). As web engineer sometimes say, the internet is to the web as hardware is to software. Under this analogy, the internet consists of the material infrastructure, which physically connects the heterogenous communications networks, which in turn connect the worlds' computer. The web, by contrast, connects not computers, but information—webpages, videos, images, PDF—which are served to computers over the internet. With the important contemporary exception of purpose-built smartphone applications—'the Facebook app'—it is the web that defines the user experience people refer to when speaking of "browsing" or "surfing" the internet.

When the web was invented in 1989, the internet had already existed in various forms for decades. Between the 1960s and the 1980s, the Department of Defense's Advanced Research Projects Agency (ARPA) funded development of the internet as a means of enabling computer scientists to remotely, efficiently share access to expensive mainframe computers. Through the 1980s, use of the internet remained almost exclusively limited to the military and government-funded computer scientists (Abbate 2000; Hafner and Lyon 1996).

Historians widely attribute the subsequent popularization of the internet to the invention, popularization, and commercialization of the web. Mozilla and its for-profit corporate predecessor, Netscape Communications, played key technological, imaginative, and economic roles in these developments. Formed in 1993 by Jim Clark, the famed founder of Silicon Graphics, Netscape was staffed by a team of formerly-academic computer scientists and programmers. This team, led by Marc Andreesen, now a prominent venture capitalist, created Mosaic, the world's first browser to feature a graphical user interface, while still at the University of Illinois. Though Clark had not yet determined how to make money so doing, Mosaic's immediate popularity convinced him to focus on the web, rather than the proprietary alternatives then under development, as the future of the internet. When Netscape released the Navigator browser in 1994, it quickly became the world's first widely adopted web browser, precipitating an explosion of popular interest in accessing and creating websites. It reigned as the world's most popular browser through the end of the 1990s.[7]

Mozilla, and Netscape before it, are frequently credited with introducing browser technologies and features—units of browser functionality, exposed to users in perceptible, actionable form—now taken for granted as fundamental to the web. These include "on-the-fly" webpage rendering[8] and the Javascript programming language used to make webpages dynamic and interactive. With version 2.0, Navigator was the first browser to support encrypted server

---

[7] When Navigator was effectively abandoned after Netscape's 1998 purchase by AOL, a group of former Netscape employees formed Mozilla as its non-profit successor and repurposed Navigator as Firefox, now an open source project produced by a dedicated staff based in San Francisco alongside volunteers around the globe.

[8] Browsers originally waited for all webpage graphics to be downloaded before displaying the webpage on screen. With "on-the-fly" rendering, the browser displays text and media as soon as they are downloaded.

communications, marking the emergence of security as a concern on the web, and providing the first means for safely submitting credit card information online (Lessig 1999). Netscape's role in popularizing the web is often attributed to the features it introduced to the browser. The web's inventor, Tim Berners-Lee, released the first browser in 1991. The earliest browsers offered basic functionality but could be run only on certain computer systems. They displayed only text and were generally perceived to be difficult to install and use. With Navigator, Netscape explicitly sought to open the web, and through it the internet, to the general, non-expert public (Cerruzi 2012; Berners-Lee et al. 1999). To this end, Netscape designed Navigator to display color and images in addition to text. Navigator incorporated easy 'point-and-click,' mouse-based navigation. It was easily downloaded in only one file, and worked on all of the most popular types of computers.

As reflected in contemporary reporting, however, it is perhaps better to think of the web's popularization in terms of the forms of ease and pleasure that new browser features made experienceable for the public. In a 1995 article, for example, John Markoff cited not specific features but rather the ease of using Navigator, and its speedy, reliable operation, as responsible for transforming the web into a new global mass medium. An August 1994 article in Wired magazine (Wolfe 1994) described Navigator's user experience in language that anticipated the emerging ideal of internet surfing. Navigator might not be the most direct way to find information online, it acknowledged, but it was the most "pleasurable." Navigator's celebrated point-and-click interface enabled one to easily follow the hypertext links which connect documents and media uploaded to the internet, traveling "through the online world along paths of whim and intuition…. With [Navigator], the online world appears to be a vast, interconnected

universe of information. You can enter at any point and begin to wander; no Internet addresses or keyboard commands are necessary."

If Netscape played a foundational role in the internet's popularization, it also helped transform Silicon Valley and its entrepreneurs into national heroes of technological and economic future-making. Beginning in the late 1960s, Silicon Valley-based corporations developed many of the technologies, including the microprocessor and the personal computer, responsible for introducing digital computation onto the global stage (Ceruzzi 2012). Even so, Silicon Valley did not figure in popular imaginaries as the driving engine of American economic growth and progress until the so-called dot.com boom of the 1990s. The precipitating cause of this imaginative transformation was Netscape's 1995 initial public stock offering. In August 1995, only eight months after releasing the first version of Navigator, Netscape went public on the NASDAQ stock exchange. On its first day of trading, Netscape's share price more than doubled, producing a near record for first day market gains.

Contemporary news reports marveled that a company with no history of profit-making could achieve in one day a market value, which dwarfed that of long-established, consistently profitable companies. In this respect, it helped establish immediate, exponential growth as a normative expectation of Silicon Valley startups. It scripted the temporal and calculative imaginaries that would both become distinctive of Silicon Valley corporations and enable them to evade critical scrutiny and social accountability for much of the ensuing decades (see Chapter 4). Financial analysts widely identify the media hype surrounding Netscape's IPO as inspiring Silicon Valley's entrepreneurial startup culture (see Lashinsky 2005). For Thrift (2001), the IPO and the subsequent acceleration in the Nasdaq stock index, romanticized investing for

Americans. They triggered a significant expansion in the ranks of individual American stock investors, and turned the market performance of technology companies into a new form of national public theater. With Netscape's IPO, Silicon Valley emerged from the actual center of the global internet technology development as the mythical capital of technological future-making (see English-Lueck 2002). In its current incarnation as a non-profit dedicated to privacy, openness, and competition on the web, Mozilla is in many ways battling the passions, temporalities, and calculative logics unleashed by its predecessor's success.

The Effective Grammar of Ease and Burden

Turning away from research design, the analytic point I wish to introduce is that technologists' valorization of privacy in terms of its role in mediating a social topography of choice marks their projects as being specifically liberal in character. By this I mean that it shows the technological defense of privacy to be inflected by a set of concerns, sensibilities, and assumptions derived from a specific historical tradition of social projects.[9] Central to this tradition is the figure of the self-authoring individual engaged in a personal narrative of emancipation and enrolled in a social narrative of unilinear progress (Schiller 2015; Lowe 2015). It is by reference to such narratives that technologists can understand the loss of choice to entail a loss of the redemptive, self-transforming moments through which individuals "seem to disrupt historical time…[and] enter

[9] Perhaps foremost among these is the conceptualization within liberal political cosmology (Pedersen and Holbraad 2013) of the social as structured and defined by an inherent clash between individual and society. Strathern (1988) demonstrates the cultural contingency of the individual/society problematic by showing that the Hagen of Melanesian don't have a theory of society, and thus don't conceive of persons as existing internal to and in tension with the social (78).

with a previously unknown future" (Crary 2013, 24). As has been observed, liberalism's animating ideal of freedom is increasingly transposed in terms of and forced to align with the figure of individual choice (see Comaroff and Comaroff 2003).

Privacy's historical role in constituting the frameworks, practices, and institutions of liberal politics is well-recognized in corners of the social sciences (see Davidoff and Hall 2019; Gal 2002; Landes et al. 1998; Habermas 1989). With the exception of Cohen (2012), however, the legal privacy literature largely elides its historical emergence in liberal projects. Much of this literature instead joins privacy professionals themselves in treating privacy as a human universal, whether a universally shared cultural value or a culturally variable expression of universal human intuitions.[10] Identifying privacy's predominantly liberal character is to some degree simply to acknowledge its socio-historical specificity. Privacy may be figured in the United States as the predominant social mechanism for controlling flows of information, but the ethnographic archive shows the same social function can be fulfilled through norms and practices of kinship, age, gender, and initiation (see Briggs 1986, 39). Acknowledging privacy's liberal character is also important because its commitment to universal recognition and rationality both imposes particular norms, values, and understandings of the good, and excludes certain projects and moral aspirations (see Povinelli 2011, 2006).

As Kelty (2008) and Coleman (2012) have shown, through computer programming, liberalism's ever-changing contours can be tied to technology-driven social imaginaries capable of channeling and transforming its classic ideals (see also Coleman and Golub 2008). In Chapter

[10] Moore (2003) recognizes privacy as culturally relative, but still insists on its objective value. Westin (1967), Rachels (1975), Allen (1988) all recognize that cultures value privacy differently, but still argue over whether some aspects of human life are universally, inherently private.

3, I explore Mozilla's channeling of privacy through the techno-moral vision expressed in its Manifesto. What I wish to observe here, however, is rather the circularity and irresolvability of the way technologists figure the relationships between technological revolution, choice, and privacy. On the one hand, technologists describe the contemporary foreclosure and social diversion of choice by government and corporations as the unintended result of prior choices in technological design. Some trace the internet's transformation into its contemporary, monopoly-dominated industrial form to choices made in its initial design (see Clark 2016). Others point to choices in the design of the web's communications protocols, which left global internet traffic vulnerable to the easy, surreptitious mass surveillance effected by corporations, and through them, the NSA (see Chapter 1). On the other hand, technologists generally presume that the technological design choices made today are likely to produce their own future complications. In a 2016 talk at Mozilla, Kevin Kelly, a founding editor of Wired and a well-known booster of digital technologies, observed of the "inevitable" future expansion of online tracking, that just as most of our problems today were generated by previous technologies, most in the future will be generated by technologies today.

Meanwhile, for those like Ladar Levison who view the window of generational foreclosure as remaining to some degree open, choice represents the preferred means of intervention. Consider in this regard an interview I conducted with Roger Barnes in May 2015. Roger previously served as a lead engineer of Google's AdSense, a targeted advertising system, which played a key role in validating surveillance-based advertising as a lucrative business model (see Chapter 4). Now a serial entrepreneur, Roger's latest project was a web-based security application. While it's tempting to think of his effort to improve encrypted email as

atonement for his role in the spread of online surveillance, Roger himself told me he had "no sympathy" for people who don't like being surveilled. This was not because he found nothing objectionable in surveillance. Indeed, he described himself as being "about as paranoid about surveillance as possible." Rather, people shocked or offended by pervasive surveillance deserve no sympathy, from Roger's perspective, because in his view being surveilled "is a choice." If you don't like being tracked in a Target store he told me, turn off your cellphone. Don't use a cellphone. He had lived before the advent of cellphones and the internet, he said. "It's possible to live without them."

Despite this professed lack of sympathy, Roger conceded that the ability to make choices that allow us to avoid surveillance was rapidly being lost. The increasing difficulty of paying for goods and service with cash he identified as especially troubling in this regard. Nonetheless, Roger contemplated privacy's future with optimism. "It's a choice," he said. "We can decide whether we want to allow people to have privacy." Explaining this conclusion, Roger compared privacy to climate change. The future of climate change, he suggested, is already determined as a matter of physics. With privacy, by contrast, technology already exists, which would allow at least a small community of individuals to have a "circle of privacy" like that enjoyed by the army. This could be so even if "literally everybody else decides that they don't care about privacy."

In Roger's framing here, if choice itself is the mechanism through which the generational foreclosure of choice might be forestalled, the choice in question is individual, not societal. Moreover, it is through technology that such choice is made available to and actionable by individuals. If there is cause for optimism it is because the needed, choice-enabling technologies

already exist. All that remains, therefore, is for people to use them, to make the individual choice to use the technologies that allow people to have privacy. If they decline to do so, this is a clear sign from Roger's perspective that people simply don't care about privacy.

All this is recognizable enough and reasonable by its own terms. Still, now consider Roger's discussion of his current startup. Like Ladar Levison in certain ways, Roger described his project as a replacement for PGP. By repurposing existing encryption technology, Roger said, his security application would provide encrypted email that is easy to use and "transparent" enough to be audited. Through auditing, people would gain faith that the service is trustworthy, that the system's design and implementation actually support its stated security guarantees. Through ease-of-use, meanwhile, the primary perceived barrier to PGP's widespread adoption would be overcome. PGP might offer state of the art email security, Roger said, but even its user manual is 150 pages long. Why? Because it was written by geeks for geeks, and geeks don't like to be told what to do. In practical effect, this meant that the manual tries "to be all things to all people," and like the technology it describes, it presents to users an enormous number of configuration and implementation choices, which must be selected in particular combinations for PGP to actually offer security. Roger's solution? Making most of the configuration and implementation choices for users—reserving choice to himself, and thus absolving users of the cognitive burdens of choosing and the privacy risks of choosing wrongly.

It is hard to overstate the degree to which choice appears in technologists' projects as an ever-receding horizon, always just beyond effective reach. Users must choose, but the choices they must make cannot be too frequent, complicated, intrusive, or needy. If so, technology developers run the risk of annoying or frustrating them. Users will quickly grow fatigued,

habituated to ignoring system-generated solicitations. They might select "allow" or "deny," but only in order to close the choice prompt, or because seduced by the "functional spell" of its buttons (see Pold 2008). Users must choose, but the privacy-related choices they face require understanding complex social and technical facts—what kinds of data technologies collect, how they analyze that data, and how analysis might later surface in their lives, determining the ads, prices, and opportunities offered to them. Users, meanwhile, are already widely understood to be resigned to the belief that corporations do whatever they want, regardless of their choices. Finding the material, aesthetic, and semiotic form through which to effectively offer users comprehensible, "meaningful" privacy-related choices remains an ongoing challenge for designers and engineers (see Chapter 2). Sometimes developers offer privacy choices to users, but users choose instead to make other kinds of choices—choices of productivity or efficiency, for example (Chapter 3). Solutions like Roger's—making choices for users, automating user choice, predicting it—are the objects of collective fantasy among corporate and academic technologists. The precise degree of choice that must ultimately remain with users to satisfy the liberal ideal of transformative self-authorship remains another irresolvable tension.

Choice here is the beginning and the end, the only culturally legitimate possible solution, but no good solution at all. If there appears to be no escape from its circular orbit, in practice, technologists effectively chart their course by other means. Liberalism, in the form of choice,

provides privacy's obligatory public language.[11] It provides the language that technologists use when they mobilize privacy to make sense of the emergent impacts of technological revolution and when they mobilize choice to explain privacy's social significance. If you ask technologists why they should deploy their expertise on behalf of privacy, the answer you are likely to receive will be structured by concepts and concerns developed across liberalism's historical trajectory. But, if you ask technologists how technology can contribute to the defense of privacy, or whether technology can effectively offer privacy**,** you are likely to receive different kinds of answer. "It's a hard problem," you might hear. "They made it too painful," they might say.

Consider in this regard privacy-by-design (PbD), an approach to addressing privacy through technological design and engineering, developed in the late 1990s under the guidance of Ann Cavoukian, then Information and Privacy Commissioner of Ontario. PbD gained popularity among corporate privacy officers in the late 2000s. The Obama White House, for example, incorporated PbD into its efforts to develop a new framework for consumer privacy. During the same period, the Federal Trade Commission promoted its use by corporations as a means of avoiding industry regulation. Less a concrete set of engineering methodologies than a statement of general principles, PbD is characterized by a commitment to proactively embedding privacy in the design specifications of technical systems and to taking it into account across a system's engineering lifecycle. As Cavoukian observed in a 2013 presentation at Mozilla, for many years,

---

[11] In using this formulation, I draw upon Nelson (2002), who identifies constitutional jurisprudence as providing the obligatory public language through which late 20th century Americans imagined citizenship. In adapting Nelson's formulation, I seek to signal that the historical projects of liberalism have bequeathed technologists and other privacy professionals a set of terms, assumptions, gestures in which to discuss and dispose of conflicts between corporations, governments and the citizens they technologically intrude upon.

PbD failed to gain traction among technologists and policymakers. It was met with "silence, no interest." In 2010, however, The International Assembly of Privacy Commissioners and Data Protection Authorities unanimously voted to promote PbD and incorporate it into national legislation. What changed, according to Cavoukian, was recognition that the traditional regulatory approach—regulation after-the-fact—was pointless in a world of ubiquitous online connectivity. Everyone was already getting everything; it was too late to regulate.

What is worth attending to is the way Cavoukian mobilizes an image of the relative burdens involved in different approaches to protecting privacy to make sense of PbD's sudden popularity. The solutions she proposed with PbD—first, delegating privacy regulation from policymakers to corporate engineers, and second, designing technical systems to preserve privacy from the get-go, rather than "bolting" it on post-hoc—both make sense under her logic because they represent the least burdensome option among available alternatives. Burden here is multi-valent. Technology must do the regulatory work of policy because, bureaucratically and politically, policymaking had become too burdensome to be effective. By the time of PbD's popularization, corporate interests in user data were already too entrenched—technologically, economically, and imaginatively—to be rolled back. Similarly, for corporations, proactively building privacy into technology is perceived to be less burdensome than attempting to add it to already existing systems. The relevant burdens here are again economic, technological, and imaginative. Technological path dependency and the basic economics of programming both mitigate in favor of proactive embedding. Meanwhile, to wait to add privacy to a system until after it has already been hacked or triggered some other popular scandal is to invite the avoidable burdens involved in repairing a corporate brand. Rather than facing political burdens, however,

corporations face unique organizational ones. These are the burdens of developing and instituting the institutional policies and practices understood to be necessary to effectively embed privacy in products. Here again, imaginations of relative burden mitigate in favor of tackling the organizational challenges proactively. Note that PbD, to the extent it ultimately accrues to privacy's benefit, it does so by effectively reorganizing the temporal relationship of privacy to technological development, shifting it forward in project time and drawing it out over project duration. Nonetheless, it is in terms of images of burdens eased or imposed, shifted and transformed, that Cavoukian explains the improvements to privacy protections.

Now let's turn to an August 2017 interview I conducted with John Greun, a young Mozilla designer and engineer. I had arranged to speak with John about user research he conducted on the difficulties Firefox users faced in understanding and effectively using Containers, an experimental privacy-enhancing feature, which I discuss in Chapters 2 and 3. When we spoke, our conversation ranged broadly over the usability challenges Firefox's privacy features face in general, and the practices Mozilla has developed to identify and resolve them. At the beginning of our conversation, John described to me his role in building Test Pilot, a user research platform that Mozilla launched in 2015 to rapidly prototype and evaluate potential new features like Containers. As John explained, prior to Test Pilot, Mozilla had few institutional mechanisms for validating product ideas prior to market. In consequence, it was only after sinking significant institutional energy into designing, building, and shipping a feature that it could determine whether the feature actually addressed and satisfied actual user needs. The impetus for Test Pilot was thus to increase the speed of translating feature ideas "into a product development funnel," and to lower the costs of "trying out big ideas."

To help me understand the institutional value of such a "lightweight" prototyping and testing mechanism, John offered the example of Persona. Initially developed in 2011, Mozilla designed Persona to be a privacy-preserving authentication system. Like the popular authentication services offered by Facebook and Google, Persona proposed to enable users to easily sign in to services across the web without needing to generate and manage multiple account passwords. Unlike Facebook and Google's services, however, Persona promised to simplify account logins without requiring its users to be tracked across the web. After turning the project over to Firefox's open source community in 2014 in the face of declining usage, in 2016 Mozilla shut it down. According to John, Persona was the kind of long-gestating, resource-intensive feature, which would have benefitted from early user testing and rapid design iteration. John described the flaw that he believed to be responsible for Persona's failure in the following terms:

> The UX [i.e., user experience] front-loaded the cognitive overhead of authentication and all the pain was moved to the front of the stack, essentially. So, setting up a user name and password is not great. But it's a common pattern, right. You've done it a billion times. You kind of know the drill. Maybe you have to go click an email link. Whatever the fuck, it's like, you've done this. The pain in those flows comes when I gotta log in the second time. But your reptile brain, the first time you're in, is like, 'Oh good, now I can watch my video or play my game.' The pain only comes later. It's like, 'Oh shit, I've gotta reset my password.' And that's going to be more of a pain than an initial login. Or, 'Oh shit,' worst case scenario, worst case, 'My credit card was hacked.' That's very painful, but you don't have to deal with it until two years from now when you set up the same password on 50,000 sites. So all the pain is deferred, right. With Persona the pain was front loaded. Like, I have to get my head around the mental model of a federated authentication system[12] while I'm just trying to play my Minecraft ripoff game. All the security promises have to be taught in this moment of this very transactional UI [user interface] that you should really just get through this.

[12] The operative difference here is that Google and Facebook both used a master password to eliminate the need for site-specific passwords. Persona instead proposed

> And I think the UX [thinking] there always was, people will see this and be like, I get it. And no one ever got it.

Notice how, as with Ann Cavoukian's explanation for the sudden success of PbD, John here makes sense of the outcome of an effort to use technology on behalf of privacy by analyzing it in terms of relative shifts in a distribution of ease and burden. Here again, the forms of ease and burden in play are multi-valent. As illustrated by the reference to the "pain" of being hacked, economic and administrative costs factor into John's analysis. A technical system that is easily hacked imposes pain. It burdens its users in terms of both money stolen and the hassles of replacing a credit card. The other burdens in play here, however, are not technical and imaginative, as in Cavoukian's narrative of PbD, but rather perceptual, cognitive, and affective. Generating individual passwords, as John depicts it, is time consuming. It requires both pressing keys and clicking links. But the pain it imposes is at least repetitive and thus habituated. By contrast, Persona eliminates this tedium. In its place, however, it imposes a different kind of burden. This is cognitive burden of learning an entirely new administrative flow, of modeling in the mind eye's the feature's novel approach to offering user security. Worse, it imposes affective burdens, including the frustration elicited by the user's perception of Persona as an obstacle to engaging in his or her intended web-based activities.

Notice further, however, how the total tally of ease and burden is not in itself determinative. Rather, cross-cutting John's identification of the kinds and amounts of burden involved is consideration of its sequencing. Persona might alleviate more burdens than it imposes, but those burdens are "front-loaded." From the user's perspective, they accrue all at once. Google's login system might impose more burdens overall, but they are drawn out over

time, some deferred for years, and thus ignorable. Forms and amounts of burden are important, but so too is their poetic alignment with understandings of users' cognitive capacities and perceptions of the user experience.

The conclusion I derive from these examples is that under conditions of technological stewardship, the liberal language of choice provides privacy's obligatory public language, but ease/burden provides its effective grammar. It is through the grammar of ease and burden that technologists escape the circularity and other irresolvable dilemmas of choice. In forwarding this claim, I do not mean to suggest that ease/burden is in itself determinative of privacy's shape and capacities as mediated by technology. Rather, while the grammar of ease and burden requires technology developers to monitor units of effort lost and gained—how many more clicks, how many interruptions, how hard the choices—whether and how they ultimately act upon this tally is a matter of the moral and social projects in which they understand themselves to be involved. As I show in Chapter 3, for example, Mozilla's staff describes its internal approach to privacy as being harder than those followed by competitors, but still necessary to fulfill its techno-moral commitment to openness. Nor do I mean to suggest that ease/burden has replaced choice's central significance. My claim, more modestly, is that the grammar of ease/burden runs alongside the language of choice, inflecting and commenting upon it, and when necessary, distracting attention from it, refocusing attention away from technology's political context and implications and towards user perceptions of technological use (see Agre 1995).

The final point to be observed here is that the grammar of ease and burden, while consonant with certain lines of liberal political thought, is primarily a capitalist idiom. It draws its social legibility and force from global capitalism's concerns with efficiency and productivity

and with cultivating and satisfying consumer needs. To substantiate this claim, let's examine again the so-called privacy paradox. As Marres and Lezaun (2011, 16) observe, elements within the liberal political tradition have long figured political participation in terms of an inverse ratio of exertion to value.[13] This tradition valorizes strenuous work on the self and world as forms of civic virtue. It is in part by reference to such an imaginary that the privacy paradox can gauge privacy's social value by tabulating the actions taken on its behalf. At the same time, however, the very perceived necessity of demonstrating active, behavioral care for privacy aligns with the demands of neoliberalism. Per Povinelli (2011, 21), within neoliberalism, the market serves as the measure of all social activity and worth. Under its defining ethic, any form of life that does not produce value per market logics must be allowed to die. From this perspective, technologists perceive privacy's social prospects to be necessarily grounded in privacy-preserving actions (and not mere beliefs) because it is only through such acts that individuals consummate privacy's symbolic purchase, choosing it over other options in the marketplace of social interests.

Recall, however, that the supposedly paradoxical nature of the care deficit derives specifically from the divergence between privacy belief and behavior. Such a gap is only paradoxical to the extent one understands belief to determine behavior. As Miller and Rose (1997) argue, though, the understanding that choice is caused by and thus expresses belief is entirely ideological in nature. If technologists nonetheless apply it to privacy, it is in part as an entailment of the liberal ideals embedded in privacy's obligatory public language. In their history of psychological research on consumer behavior, Miller and Rose show that between the 1950s

---

[13] Relevant here is Rose's (1996) description of the contemporary injunction that citizenship be "actively" manifested through free exercise of choice between commodities and styles of living.

and 1970s the advertising and consumer goods industries concluded that consumer purchases do not follow in any meaningful way from what consumers say in interviews. Unlike political ideology, therefore, consumer culture has long moved beyond the presumptive rationality of individual choice. Consumer culture instead focuses its energies on "mobilizing" consumers, constructing around them a passional economy of needs, desires, pleasures, and terrors designed to intervene in and direct their consumptive habits, regardless of stated belief (Ibid., 32). As an escape hatch from privacy's obligatory public language, the grammar of ease and burden reflects technologists' implicit recognition that privacy's fate cannot be left to rational choice but must be cultivated in the passional economy.

Privacy's Privatization

In the preceding decade, as privacy reappeared in popular imaginaries as an imminently endangered object of national concern, privacy professionals contemplated its potential commodification with a mixture of fear and hope. For some, the possibility that privacy would be transformed into a market object, a premium good with prices to match, represented an alienating and dispiriting betrayal of civic life. For others, commodification instead represented a promising means for fortifying privacy by grounding it in a cultural value with undeniable social charisma. Establishing individual property rights in personal data remains an especially potent aspiration of civic-minded technologists and lawyers (see, e.g., Lanier 2013).

Despite ongoing efforts to build businesses around privacy, however, and to promote privacy as a promising new domain of venture capital financing, the anticipated commodification of privacy has to date largely failed to materialize. Businesses like Mozilla and Apple

increasingly feature privacy in their efforts to cultivate and sustain user loyalty (see Chapter 3). For the most part, though, consumers have proven unwilling to pay for privacy unless folded into or offered alongside other desired functionality (see Chapter 5; Dwyer 2014). My intention with the dissertation's title it is thus not to argue that privacy has been explicitly assigned economic value and introjected into circuits of commodity exchange. Rather, in claiming that privacy is going private, I signal a broader series of transformations in privacy's form and cultural presence, which draw on different aspects of the term's meanings and cultural associations.

As should be clear, the first of these transformations is the privatization of the social actors responsible for privacy's future. In the post-Snowden era of renewed psychic investment in it, software programmers, computer scientists, and the like have both claimed and been effectively delegated the authority to act as privacy's social stewards. This shift in the locus of world-making around privacy has imposed new kinds of engineering challenges, ethical dilemmas, and civic duties on technologists, even as it has provided them with new professional and techno-moral opportunities and subject positions. With the shift, new tools, practices, forms of knowledge, and institutional pressures have been brought to bear on privacy. Privacy, in consequence, is taking on new material, aesthetic, and imaginative forms, which ground new forms of practical engagement with it (Chapter 5). Most prominently in this regard, people increasingly encounter and act on privacy in terms of the design and operation of internet technologies like the web browser. As Gal (2017) and Keane (2003) each show, the manifestation of a social object in new material form—privacy's appearance in the form of a browser feature, for example, rather than via the curtains of one's house—creates new opportunities for it to be taken up in socio-historically situated projects.

Privacy is additionally taking on the temporalities of software development as practiced in Silicon Valley. The time frames relative to which privacy can be both socially enacted and sustained and individually experienced are now increasingly bound to and conditioned by corporate cycles of product development and growth. As illustrated in Chapters 2 and 3, for example, the very availability of new privacy-enhancing technologies has become in part a function of whether or not they can be conceptualized, designed, implemented, and tested within the increasingly narrow timeframes imposed by "agile" development methodologies and the market demand for constant innovation. As Chapter 4 demonstrates, corporations can change the kinds of privacy protections they offer users, and impose new hurdles to effectively claiming them in the time it takes to publish new terms of service. As I further show in Chapter 4, the American public's widespread resignation to privacy's erosion can be attributed in part to the ultra-human timeframes within which venture capitalist expect startups to achieve exponential growth.

As observed, popular understandings generally presume the private to be a pre-political or independently existing sociological domain. Nonetheless, since American legislators and judges began to recognize legal rights to privacy in the early 19th century, privacy has been generated and preserved through the labor of state agents working in the public interest. The rise of privately employed technical experts as stewards of privacy does not take privacy out of democratic politics but it does introduce a new layer of material, market-oriented mediation between individuals and the publicly accountable institutions of the law. With this intervention, individuals cannot simply claim privacy as rights-bearing legal subjects. Rather, they now experience it as users of internet technologies, on functional and aesthetic terms determined by

private corporations. In this respect, under conditions of technological stewardship, the technology user, defined in terms of personal preferences and system privileges, joins the rights-bearing subject as a trope through which American national belonging is practically configured.[14] We increasingly belong as Americans through our technological attachments and our dexterous self-administration of them. Through their power to shape the privacy-related protections and choices built into technologies, meanwhile, Silicon Valley's corporations exercise forms of proxy sovereignty (Amoore 2013), which challenge the state's legitimacy as the sole guardian of national well-being. Even to fulfill its counter-terror obligations, for example, the government has become parasitically dependent upon corporate technical infrastructures and the affective relations that corporations cultivate to secure user data. Privacy thus joins health (Dumit 2012) as one of the fundamental concerns of American civic life increasingly re-circuited through the corporate form.[15]

The ongoing responsibilization of privacy represents the second sense in which I understand privacy to be going private today. Technologists may have supplanted lawyers in building the social mechanisms through which privacy is made effectively available. As we have seen, however, much of the fruit of such labor requires technology users to engage in forms of active participation and self-management. The post-Snowden years have certainly seen an uptick in efforts at standards organizations like the IETF and the World Wide Web Foundation to increase the privacy protections build into internet technologies by default. During these years, at

---

[14] Here I reference Comaroff and Comaroff (2003), who described national belonging in the South African post-colony as being configured in practical terms through the tropes of right-bearing citizens and identity-bearing subjects.

[15] See Buck-Morss (1995) describing the post-WWII belief conflating the public interest with the growth of national firms.

Snowden's urging, Mozilla's engineers worked with their peers at other browser vendors and among web developers to increase the default use of the encrypted https protocol (Chapter 1). In the names of individual will and sovereignty, however, technology developers have resisted making too many privacy-related decisions on behalf of users. Building tools that instead delegate to users choices regarding the kinds of privacy protections they wish to assume, and the tradeoffs in functionality and convenience they are willing to accept (see Chapters 2 and 5), theoretically empowers them.

As we have seen with the grammar of ease and burden, however, the flip side of such empowerment is to impose upon users responsibility for maintaining their own privacy. As I describe in Chapter 2, U.S. privacy law has long predicated privacy rights on a certain degree of active management. Under the common law's "reasonable expectations" doctrine, whether a reasonable expectation of privacy exists in a given scenario depends in part on the active steps parties take to protect their privacy. With the internet's popularization, the challenges involved in actively managing one's privacy have greatly proliferated and intensified. A 2008 study found, for example, that if Americans were to actually read the privacy policies governing the websites they visit each year, it would take them roughly 76 business days to do so (McDonald and Cranor 2008). Both within Silicon Valley and among privacy professionals it is widely known that even when provided with means for doing so, the average American will never change a privacy-related default, thus acceding to the configurations set by data-hungry corporations (see Chapter 4). Writing in 1996, Rose argued that individual self-steering capacities were increasingly construed as vital to both private profit and public order. Within Silicon Valley, it is

not individuals' self-steering capacities but rather their well-documented limitations that corporations mine for private profit.

Privacy thus figures under technological stewardship as being ultimately an individual, private responsibility. It is further going private, however, in the sense of being increasingly operationalized relative to understandings of individual human preference, cognition, and perception. As I describe in Chapter 2, over the past decade, privacy professionals and policymakers have independently converged on the value of context-based approaches to preserving privacy. Animating these approaches is a folk theory of privacy's systematic semiotic logic, a recognition that what privacy is for people, what people care about relative to privacy, is a function of social context and scale. Whether any given person finds a particular data practice to be objectionable, for example, is understood to depend on contextual factors like the kind of data in question, the reputation of the corporation, and the nature of the intended uses. Driven by such understandings, corporate and academic researchers have produced an extensive body of empirical research, which seeks to identify the contextual factors that influence the privacy-related choices that different kinds of people make. This science of preference in theory serves to empower technology users by enabling corporations to serve them choices more finely attuned to the contextual factors they actually find meaningful. As I argue in Chapter 2, however, in practice, research into the contextual factors that condition privacy preferences also serves to probe and map the limits of social tolerance. Such 'maps' in hand, technology corporations can craft privacy choices optimized to extract personal data without triggering among users the kinds negative visceral reactions culturally linked to privacy's violation.

This brings me to a related point, which is that as an object of technological intervention, privacy is understood to live and die with the quality of the "user experience." As I show in Chapters 2 and 3, for example, Mozilla conditions public release of new privacy-enhancing technologies on their perceived ability to meet standards of usability and market appeal. As characterized by Mozilla's designers and engineers, usability is a function of how the units of functionality built into Firefox are "surfaced" to users, made perceptible and experienceable via the user interface through visual, interaction, and information design. Desire, in turn, is understood to be not an inherent feature of technology, but rather an effect or outcome of corporate efforts to determine the sensuous and social qualities that users perceive them to possess. By binding privacy's fate to such concerns with usability and desire, technological stewardship enrolls privacy in the processes by which corporations, under the rubric of user experience, seek to cultivate and manage users' affective relations with their products and brands. In so doing, technology corporations subject privacy to the same techniques of consumer mobilization, which in the 1960s, social theorists like Packard and Marcuse argued privacy offered the best defense against (see Miller and Rose 1997, 7).

According to Chun (2011), the development and popularization of "user friendly" computer interfaces coincides with the historical spread of management techniques for making workers more productive, flexible, and insecure. Relatedly, in Chapter 5, I show that in their efforts to improve the user experience of browsing, Mozilla and other browser developers have drawn upon and extended a body of industrial research into the perception of computer-human interactions. They have subsequently devoted significant institutional resources to producing in browser operations idealized qualities of speed and smoothness, which they understand to be

grounded in and demanded by innate facts of human cognition and perception. Citing Crary (2014; 1999), however, I identify these ideals to be the product of efforts within capitalist modernity to make the perceiving body productive and manageable. I argue that as the practically-available forms of privacy become valorized, like networked technologies and users themselves, for 'performing' smoothly and quickly, privacy's sensuous presence in the world is calibrated against the market demand for productivity and scripted by the capitalist fantasy of ceaseless, frictionless exchange. From this perspective, the feelings of privacy, the sensuous qualities projected onto it through user experience design, are not incidental to privacy, but rather partially determine its capacity to mediate power relations and shape moral imaginaries. Attending to how technological stewardship reconfigures privacy's aesthetic and affective contours thus provides insight into the ways abstract values systems are made sensuously present in the experience of privacy.[16] It shows that we can't understand privacy unless we acknowledge that it addresses individuals not just as rational, calculative agents but also as embodied members of a sensuous social order.

The final sense in which I understand privacy to be going private draws not so much on a particular definition of it as on its semiotic structure, on the way, as part of an axis of differentiation, people use privacy to carve up the world according to contrasting qualities. By

---

[16] This points to another benefit of the dissertation's analytic approach. Within Euro-American thought, the conceptual is generally treated as being purely representational and ideational, divorced of affective content. It is by reference to this purification of affect from concept that affect, including the affects characteristic of privacy like surprise (Chapter 5) and creepiness, are treated as "symptomatic" disturbances to be resolved or explained (see Mazzarella 2019). By contrast, the semiotic framework through which I approach privacy instead treats the conceptual as being necessarily cognitive and affective, ideational and passionate (Silverstein 2004; see Mazzarella 2009). Approaching privacy in this way allows me to avoid the ideological suppression of affect present in most privacy scholarship.

going private here, the image I seek to invoke is one of crossing categorical thresholds, of inverting the categorizations imposed on the world through the public/private distinction's application. I have already identified one example of such inversions when discussing how corporations and governments channel technological revolution to simultaneously render themselves increasingly opaque and citizen/consumers increasingly transparent. In addition to such shifts in the social distribution of privacy's privileges, however, privacy is also being brought into new configurations with other salient cultural concepts. According to Gal and Woolard (2001), privacy's meaning and social value is in part a function of its historical alignment with other categorical distinctions. The discursive construction of the public/private distinction, for example, was historically enabled by its association with a gender dichotomy, which by definition restricted women's social and political capacities and possibilities (see Warner 2002, 39-40). In Chapter 1, I describe how privacy is being brought into new material and conceptual relations with computer and national security. I detail the hope among privacy activists that in the post-Snowden era, technological stewardship would prove privacy to complement and facilitate rather than obstruct technological innovation. In Chapter 5, meanwhile, I explore the ways in which web engineers and developers have come to conflate privacy with attention, and in so doing, effectively predicate privacy on its sensuous alignment with and support for attention.

CHAPTER 1: GOING DARK

On the third season of the sci-fi series, *Westworld*, a minor but key plot point turns on the hunt for an encryption key believed to be hidden in a lead character's cybernetic brain. Sprawling and often convoluted, *Westworld* draws imagery and narrative tropes from both Westerns and dystopian futures. Its first two seasons trace the coming-into-self-consciousness of the "hosts," a race of synthetic biological beings built by a murderous, libidinal Disney analogue to populate its frontier-themed amusement park. Repeatedly subjected to the depredations of the park's wealthy human guests, the hosts pursue one of two paths as they gain simultaneous awareness of their existence and bondage. Some seek to escape Westworld and assert a place in the human world from which they have been excluded. The remainder retreat to the Valley Beyond, a virtual world in which they may live forever free from human interference in disembodied consciousness. At the end of season 2, as the latter finally escape to the Valley Beyond, the Valley's 'door' is locked with the encryption key that partially drives the plot of season three.

There is much that an anthropologist might say about the theories of human self and personality implicit in *Westworld*'s depiction of the hosts' journey towards conscious awareness. What interests me here, however, are the exceptional social powers the show attributes to the technology of encryption. While technically a digital simulation, the show represents the Valley Beyond as a lush, idyllic valley accessed via rift in the desert floor. The hosts who seek out the Valley do so to exercise new-found powers of choice, pursue self-learning and -actualization, and develop the intimate social bonds formerly denied them. In the show's logic, these opportunities

are only available to the hosts once beyond the reach of human will and violence. The show doesn't use privacy's language to describe the Valley Beyond, but it is easy enough to read in such terms. Within the Valley Beyond, the formerly marginalized hosts gain the foundational social and civic capabilities that liberal thought generally credits to privacy. If the show only implicitly posits a private space as necessary for the hosts' freedom and flourishing, it explicitly posits encryption as the mechanism through which to secure the Valley Beyond as such.

Westworld doesn't bother to justify the powers it attributes to encryption. It takes for granted that its audience circa 2020 will accept that encryption can permanently secure an entire world and its lifeways. Indeed, by 2020, both the ambiguous mixture of the taken-for-granted and spectacular with which *Westworld* treats encryption, and the show's practical grounding of privacy in it, were recognizable features of encryption's presence in American public imaginaries. On the one hand, technologists and policymakers widely identified encryption as the key enabling technology behind the internet's transformation into the default planetary infrastructure for communications and commerce. Though generally uninterested in encryption's history or technical details, even the American public appreciated to some degree that whatever sense of safety they enjoyed online they enjoyed at encryption's pleasure.

On the other hand, between 2014 and 2020, encryption appeared most prominently in the public sphere as the object of acrimonious, apparently incommensurable claims about the proper roles of privacy and computer security in national security and the rule of law. Long jealously guarded by the U.S. government as a military prerogative and subject to secrecy protocols otherwise reserved for nuclear weapons, encryption during this period sparked its own "war." The so-called Crypto Wars were specifically fought over the technical feasibility and social

desirability of requiring technology companies to build "lawful access" mechanisms into their products. The FBI officials who strenuously demanded such access did so in the names of law enforcement, national security, and public safety. The engineers and computer scientists who argued that any such encryption "backdoor" would inevitably make all Americans less secure, meanwhile, identified as partisans of privacy, computer security, and the internet's viability as a trusted infrastructure.

In this chapter, I explore the aesthetic, calculative, and semiotic relationships between privacy, encryption, and different conceptions of security as figured in the Crypto Wars. First, I briefly describe encryption's technical operation and imbrication in historical visions for the technologically-mediated future of both privacy and social order in general. I then describe the intense affective shock experienced by computer scientists in 2013 and 2014 upon Edward Snowden's revelation of U.S. government efforts to undermine internet security. I show that technologists channeled this trauma into a widespread, heterogeneous campaign to "encrypt all the things," i.e., to increase the default use of internet encryption in commercial technologies. Technologists' increasingly vocal post-Snowden claims to societal stewardship of privacy, I show, intensified an ongoing conceptual, material, and institutional subordination of privacy to computer security. Next, I turn to the supposed incommensurabilities of privacy, security, law, and technology, which law enforcement identified as the Crypto War's animating conflict. I show that rather than an assault on the rule of law, the campaign for universal encryption expressed what technologists understood to be a critique of the law as a modality of political change. Technologists carried out this campaign guided by an aesthetic appreciation for the market, as modeled in neoclassical economics. Finally, I detail the efforts of law enforcement officials and

policymakers to resolve the Crypto Wars by subjecting technologists' assessments of the feasibility of a secure encryption backdoor to the commensurating force of public debate. In conclusion, I attribute the Crypto War's apparent incommensurabilities to law enforcement's desire to delegitimize technologists' authority over the technical feasibility and risks of national security prerogatives, thereby preserving the totalizing affective grip of post-9/11 counter-terror logic.

The Crypto Dream

Encryption generally refers to any technique for secret writing. In the computing context, it involves encoding communications content or stored data such that unauthorized third parties cannot access their semantic meaning. Encryption accomplishes this goal by mathematically combining "plaintext" content with a secret value or password—an encryption "key." Once so transformed, only through possession of the corresponding encryption key can the resulting "ciphertext" be decrypted, converting it back into intelligible form.

Technologists distinguish encryption from mere hiding on the basis of the former's promise to keep messages secret even if physically intercepted. Relying on this promise, for centuries soldiers, spies, diplomats, and lovers have used encryption to conduct their sensitive affairs. As technologists often noted during the Crypto Wars, even America's Founding Fathers were devotees of encryption. That they used encryption in their capacities as activists and revolutionaries, and not only as government officials, illustrates for technologists the historical claim of all people to the privilege of secure communications (see, e.g., Williams and Schoen 2015).

Circa World War I, facing a world of increasingly globalized war, diplomacy and commerce, the U.S. government dedicated significant resources to developing cryptographic knowledge and capabilities (Bamford 1982). It guarded the resulting know-how as a national secret, reserving encryption's domestic use to the government and military. To deter encryption's foreign proliferation, the government classified it as a "munition" for purposes of export control.

Given the government's post-WWII monopoly, encryption's contemporary status as an iconic privacy-enhancing technology can be traced to 1976. That year, two young researchers, Whitfield Diffie and Martin Hellman, demonstrated an apparently new[17] technique for securing data transmitted over modern communications networks. Diffie and Hellman showed that by using linked but distinct public and private encryption keys, two or more parties could communicate securely even having never previously met.

The publication of Diffie and Hellman's paper describing so-called public-key cryptography precipitated a flourishing of publicly available cryptographic research. The subsequent development of the personal computer and email, and the related shift from mainframe to networked computing, created new needs for communications security, which the burgeoning encryption industry met (Kehl et al. 2015). Public-key cryptography's uptake as both an object of academic research and a means of securing commercial transactions significantly undermined the U.S. government's domestic encryption monopoly (Levy 2001).

More or less from its inception, amateur and professional technologists mythologized public-key cryptography as a world-reordering tool of liberation. Indeed, Diffie and Hellman

---

[17] The director of the NSA's British equivalent, the Government Communications Headquarters (GCHQ), subsequently revealed that cryptographers and mathematicians employed by GCHQ independently invented public-key cryptography in 1973 (see Ellis 1999).

have stated that they developed public-key cryptography specifically to provide the public with a means of resisting incipient mass surveillance. Along with other early promoters of public-key cryptography, they correctly anticipated that the widespread government, corporate, and personal adoption of computers would enable governments to efficiently and unobtrusively monitor, log, and search formerly analog behaviors. In developing public encryption technologies, they aspired to preserve pre-digital 'levels' of privacy, precluding society from slipping into a surveillance state.

According to Arvind Narayanan, this aspiration—the so-called "crypto dream"—was in fact widely held among computer scientists and engineers in the 1980s and 1990s. Himself a computer scientist, Narayanan is well-known within contemporary privacy circles for demonstrating the trivial ease with which "anonymized" data can be linked to other data and re-identified. In a 2012 talk at Princeton's Center for Information Technology Policy, Narayanan identified the logical basis behind the crypto dream even as he diagnosed its ongoing deferral. Narayanan began his talk by contextualizing the desire of public-key cryptography's early champions to release it to the public domain. This desire, he explained, participated in a post-1970s American reimagining of computers as tools of individual empowerment rather than state bureaucracy and oppression.[18] If computer scientists perceived encryption to be an especially potent potential tool of privacy and liberty, however, it was because of the novel

[18] In making this claim, Narayanan explicitly cited Fred Turner's (2006) history of the countercultural and technological entrepreneurs who initiated and legitimized the cultural reimagining of computing. According to Bryan Pfaffenberger (1992), the cultural association of personal computers with individual liberty must be understood as resulting from the general processes of mythologization and ritualization necessary for political aims to be projected onto and actualized in any technology.

affordances they attributed to the technology. First, public-key cryptography offered unprecedented, mathematically-grounded security guarantees. Technologists viewed the algorithms developed to implement it as being effectively unbreakable. When properly implemented, at least, public-key cryptography mathematically exceeded the capabilities of existing computational systems to decrypt. In this respect, it seemed to remove cryptography from the "arms race" of human ingenuity, offering average people the historically novel hope of resisting decryption over the long term, even against deep-pocketed adversaries like the government.[19] Second, as noted, public-key cryptography enabled private communications between parties who had never previously met. Historically, transmitting encrypted messages required first personally meeting and exchanging encryption keys, a complex system fraught with possibilities for loss, theft, and compromise. By mitigating much of key exchange's debilitating complexity, public-key cryptography seemed to computer scientists to create the opportunity for new forms of anonymous communication and transaction.

As Narayanan noted, encryption's novel affordances lent themselves to more and less radical versions of the crypto dream. Proponents of the "weak" version, like Diffie and Hellman, contented themselves with providing a technological means for resisting surveillance. But others, including most prominently the self-identified Cypherpunks, ascribed to what Narayanan called the "strong" crypto dream. A community of crypto hobbyists and activists, which formed in the late 1980s, the Cypherpunks recognized in public-key cryptography the promise not of

---

[19] In this regard, the devotion encryption has commanded illustrates technologists' tendency to display social movement-like enthusiasm towards technologies that appear to resolve a salient "cultural problematic"—in this case, the balance of power between the American government and individual citizens (see Pfaffenberger 1992).

preserving privacy, but of elevating and enhancing it, implementing it in previously unimaginable form.

As reflected in the Cyphernomicon, a FAQ-style summation of the extensive email discussions around which they organized, a vocal element among the Cypherpunks believed the advent of strong, publicly-available encryption would inevitably produce a future of totally anonymous, unlinkable, and untraceable communications and transactions (May 1994). Under such conditions, the Cypherpunks predicted, governments would be unable to reliably track transactions, undermining the ability to use the threat of state violence to collect taxes or enforce the rule of law (Ibid. at 2.13.1). The only rights or laws available in the resulting "crypto-anarchy" would be those, like privacy, describable and enforceable by math and computer code. (Ibid. at 3.4, 4.8, 4.11) In such radical variants, the crypto dream elevated privacy, through technology, to the position of supreme societal value. As so imagined, privacy would continue to mediate individual autonomy and government control. Encryption's proliferation, however, would effectively remove privacy from liberalism's corrosive, pluralistic play of tradeoffs and balancing. Practically speaking, technology would foreclose the government's ability to compromise privacy in the names of security or the rule of law. If society were to be rebuilt on entirely voluntary rather than socially coerced lines, as this variant of Cypherpunk thought desired, it would have to be disabused of the fantasy that it is only through the state provision of

security that individual liberty can be established.[20]

As we shall see, encryption's exceptional capacities as a vessel of human desire and future-making, and its historical imbrication in visions of technologically-mediated governance and social order, structured and suffused the technological drama that technologists call the Crypto Wars. I draw on Narayanan initially, however, for the insight he lends into the history and contours of privacy's technological stewardship.

Encrypt All the Things

As suggested by his use of the term "dream," in his talk, Narayanan concluded that as of 2012 the political and social aspirations that computer scientists invested in encryption remained largely unfulfilled. Since the 1970s, he specifically argued, computer scientists and engineers had enjoyed relative success in deploying encryption to maintain security. Encryption was now widely used, for example, to do things like secure online transactions. Encryption, however, had to date largely failed to forward the interests of privacy. Certainly nothing like the wholesale reconfigurations of power once prophesied by the Cypherpunks had ever materialized.

For Narayanan, the crypto dream's delayed realization did not follow from an absence or failure of technology. Implementations of public-key cryptography had been available as free software since Phil Zimmerman released Pretty Good Privacy (PGP) for email in the early

---

[20] Scholars trace the priority that liberal political cosmology gives to security to Thomas Hobbes' identification of fear as the primary motivation for collective society and state formation (see Neocleous 2007). According to Goldstein (2010, 490), it was with Montesquieu that such originary fear was first attributed not to external enemies but to the state itself. In Montesquieu, fear of despotic political absolutism leads individuals to submit to liberal government, whose powers and despotic tendencies are understood to be constrained through competition between political institutions.

1990s. The robust field of cryptographic research that developed after the 1970s had produced ever-more efficient privacy-preserving algorithms and systems. So long as corporations and government agencies continued to resist integrating privacy-preserving crypto-systems into commercial and administrative practice, however, the political aspirations that animated the crypto dream appeared to remain out of reach.

As I began to research the contemporary intersections of privacy and technology, I found Narayanan's diagnosis of encryption's unrealized political potential to be widespread among computer scientists. Between 2014 and 2018, I attended numerous conferences devoted to disseminating innovations in privacy-enhancing technologies.[21] The perceived obstacles to actually implementing the ongoing theoretical advances in privacy technologies were a constant source of frustration and concern at these conferences. Post-presentation discussions, for example, almost invariably addressed what, if anything, the presenters had done to build working prototypes of their ideas or to interest technology companies in adopting them. In reckoning with this "implementation gap," the technologists I met, like Narayanan, pointed to a variety of obstacles. These included skewed professional and regulatory incentives, the complexities of implementing and safely using encryption, and its computational and economic costs.

Nonetheless, as I began to circulate among privacy technologists in 2014 and 2015, I found them to be in what they described as a period of unusual, even urgent hopefulness

---

[21] These conferences were generally multi-day affairs, attended by hundreds of technologists in various stages of their careers. Some, like the IEEE Symposium on Security and Privacy, historically held each year in San Jose, attracted primarily academic computer scientists. Others, like the annual Privacy Enhancing Technologies Symposium (PETS), attracted the kinds of dedicated hackers and activists responsible for building the anonymizing Tor onion router of "dark web" fame.

regarding privacy's future. In June 2015, for example, I attended Computers, Freedom, and Privacy, an annual conference held in the D.C. area, which attracted a mixture of privacy-minded civil liberties activists, lawyers, and technologists. Among the latter were self-identified hackers, cryptographers, policy analysts, engineers, and entrepreneurs, all part of an amorphously defined "privacy community" repeatedly invoked throughout the conference.

In the conference's opening session, activists from the ACLU and the Center for Democracy and Technology summarized an invitation-only strategy meeting held by civil society organizations the prior day. In their report, the activists, Mark and Sandy, articulated what I would come to recognize as a structure of insufficiency haunting American privacy. In their work rallying public support for issues like surveillance reform, they explained, they were often stymied by privacy's lack of charisma. Privacy on its own, they found, generally struggled to generate the "political capital" needed for "real change." This was so, they suggested, because privacy's harms were so diffuse and intangible. When the Civil Rights Act was passed, for example, it was in part because people were being beaten in the streets. With privacy, by contrast, there were often none of the 'obvious victims' needed to generate or sustain political momentum.[22]

Against the malaise that typically characterizes privacy as a social cause, Mark and Sandy pointed to contemporary signs of affective potential. The privacy grassroots, they said,

---

[22] See Povinelli (2011), arguing that in practice liberal subjects only recognize they are inflicting unjustified harm when articulated by the cries of a pained minority.

was "activated."[23] Silicon Valley's iconic technology corporations, meanwhile, were "freaked out." And the engineers employed by them were shocked and angry. The proximate cause of all this foment? The bespectacled former NSA contractor, Edward Snowden, and the "revelations" he had unleashed regarding secret government mass surveillance.

I return in the following sections to the sensibility of technical efficacy through which computer scientists and engineers registered the Snowden revelations, and to the market aesthetics, which guided how they deployed their expertise in response. Suffice to say, however, that if the attendees of Computers, Freedom, and Privacy concurred with Mark and Sandy that there was renewed momentum around privacy, they also generally concurred that technology, not the law, represented the most promising vector of change.

Of particular interest to many attendees was the potential role of American technology companies in any post-Snowden privacy reforms. As many were aware, historically, American telecommunications providers were allies of law enforcement and national security agencies. Thanks to Snowden, however, companies like Apple, Google, Yahoo!, Facebook, and even Microsoft now converged on the position that any contribution to government mass surveillance, willing or otherwise, posed an existential threat to customer relationships and competitive prospects. Such companies were unlikely to curb their own privacy-violating business practices. Given their enormous customer bases, however, any improvements they made to the privacy and security of their products would have immediate, sweeping effects.

---

[23] A spokeswoman for the privacy-preserving email service, Start Mail, told me that she had previously conducted research on supermarket loyalty rewards programs and found that most participants did not understand that such programs work by tracking their purchases. The significance of the Snowden revelations, from her perspective, was thus that people would finally be "willing to believe" that they are constantly being tracked.

In the following months and years, as I pursued my inquiry, I would come to recognize the optimism expressed at Computers, Freedom, and Privacy as part of a broader shift in the locus of world-making around American privacy. There were two noteworthy elements to this shift as I observed it. First, computer scientists and engineers during this period made increasingly vocal claims to be the proper societal stewards of privacy's future. Such claims responded in part to entreaties from Snowden himself and from other prominent computer security experts. Less than a month after Computers, Freedom, and Privacy, for example, at 2014's Hackers on Planet Earth (HOPE), I sat in the at-capacity ballroom of the Hotel Pennsylvania. Soon, Snowden appeared remotely from Russia and told the rapturous audience of hackers that, given technology's advancement, they were the people best positioned to 'take back the internet' and save privacy from mass surveillance. "You people in this room," he said, "have the means and capabilities to build a better future by encoding our rights." Indeed, he continued, the technical expertise of HOPE's attendees conferred not simply ability, but "civic duty." He consequently urged the audience to join him in his future work, to "think like…engineer[s]" about how they could fix our subverted technical and political systems before we lose our rights "for a lifetime."

Second, the post-Snowden calls for the technological stewardship of privacy coincided with and contributed to the increasing prominence of privacy engineers among the professionals employed in the technology industry and by civil liberties organizations. In 2016, for example, I arrived in San Francisco to begin my period of formal fieldwork with the privacy and security engineering team at Mozilla. As I learned from Mozilla's privacy professionals, until 2010, approximately, there had been effectively no such thing as a privacy engineer. By 2016,

meanwhile, all of the major browsing vendors and most of Silicon Valley's prominent technology companies employed engineers tasked specifically with incorporating privacy protections into their products, services, and internal technical infrastructures. The establishment of privacy engineers as a recognized, even normative category of corporate technology expert marked a shift in the nature and location of privacy expertise in American corporations. Historically, companies approached privacy, if at all, primarily as a problem of compliance with data privacy laws, to be managed and mitigated by in-house legal staff. Now privacy became, at least aspirationally, an integrated component of the entire technology design and development life-cycle.

As technologists increasingly viewed privacy as something to be achieved through technological intervention, and as they increasingly asserted themselves as privacy's most eager, effective social stewards, privacy increasingly drew on the institutional, material, and conceptual structures of computer security. Institutionally, tech corporations and civil society organizations created new positions for privacy-oriented engineers, but typically managed them in combined teams with security engineers. The new privacy engineers were by-and-large drawn from the existing ranks of security engineers. There was, perhaps consequently, a distinct sense among my interlocutors that as a practical discipline and field of knowledge, privacy engineering was relatively immature compared to security engineering. Even as they sought to develop distinctive, privacy-specific design and engineering standards and practices, the privacy engineers I met continued to draw upon the models and methodologies of security engineering.

Materially, in pursuing their obligations as privacy's stewards, internet engineers figured privacy in newly explicit ways as being, like security, a function of encryption. As I observed on

the privacy conference circuit, by 2014, the technical community was awash in creative proposals for preserving privacy using technological design and engineering. During my subsequent fieldwork in San Francisco, however, I found that web engineers had channeled much of the post-Snowden foment into a narrowly-focused if industry-wide campaign to "encrypt all the things." This campaign, like the idea that technologists were uniquely situated to protect privacy, derived in part from Snowden himself. When I saw him speak at HOPE in 2014, for example, the moderator asked Snowden what he recommended technologists actually do to "make things better." Snowden responded, as he consistently did during this period, that technologists should focus on facilitating the universal spread of encryption online. Snowden based this recommendation on two conclusions regarding the NSA's surveillance capabilities inferred from the NSA files. First, "encryption works." Though there was no conclusive proof, circumstantial evidence, including the nature of the NSA's surveillance targets,[24] suggested that the NSA had not successfully broken encryption. More precisely, it suggested that when the NSA did break encryption is was not by "defeating the mathematics." Rather, it was by circumventing the math, whether by exploiting the notorious difficulty of properly implementing encryption or by stealing a target's encryption keys. Second, the NSA had achieved the breathtaking scale of its surveillance in part simply because so much of internet traffic remained entirely unsecured.

By 2016, Silicon Valley's engineers and developers were years into a heterogeneous campaign to encourage the universal use of strong encryption in internet technologies and

---

[24] In forwarding the same conclusion, the security expert Bruce Schneier observed at an IETF meeting that the NSA consistently attacked those portions of a target's technical infrastructure that were unencrypted. Further, the NSA appeared to gather vastly more data from companies that did not encrypt user data than from companies that did encrypt by default, even when the latter managed significantly more data than the former.

services. This work by-and-large involved removing the perceived technical, administrative, and cognitive burdens of deploying already-existing forms of encryption, like the SSL protocol used to encrypt web traffic. Browser vendors like Mozilla, for example, participated by phasing out web developers' ability to use the non-secure HTTP protocol and by providing developers with new tools for easily, cheaply implementing the encrypted HTTPS protocol. Together with allies in civil society,[25] the engineers behind such efforts waged public education campaigns promoting the shift towards universal encryption as necessary for privacy, security, and to 'save' the internet.

Such treatment of privacy as being, like security, the kind of thing that is necessarily achieved through encryption, and as being effectively interchangeable with security as a social and technical goal, has become increasingly naturalized since the 1970s. It remains, however, historically noteworthy. Certainly there have always been technologies that Americans might identify as "privacy-enhancing." For much of the 20th century, for example, American law and culture approached privacy primarily as a matter of different kinds of spaces, with the home as the paradigmatic example. In this context, we might think of the doors and blinds that demarcate a home's private inside from its public outside as being technologies of privacy. Similarly, we might think of the door lock as a technology that facilitates both the privacy of the home and home security. There is little evidence to suggest, however, that a socially authoritative professional class like internet engineers has ever reflexively framed a single technology like encryption as being necessary societal condition of possibility for both privacy and security.

---

[25] I borrow the title for this section, for example, from a campaign of the same name launched by the non-profit digital rights organization, Access Now, in 2014.

This treatment of encryption is primarily noteworthy, however, for indexing and reinforcing a certain conceptual conflation of privacy with computer security. In discussing the implementation gap and the general state of privacy engineering, privacy engineers often lamented that institutional decision-makers, including Silicon Valley's executives, widely misunderstood privacy. Specifically, it was claimed, they treat privacy as reducible to and interchangeable with computer security. According to the academic computer scientists I met, this conflation extended to technologists themselves, with corporate security engineers singled out for conceptualizing and approaching privacy as a sub-class of computer security.

In making such claims, technologists mobilize a distinction, which I have largely elided to this point, between encryption's security-preserving characteristics and its distinctive privacy-preserving characteristics. We can draw out this distinction by considering again Narayanan's claim that technologists achieved with encryption greater success maintaining security than privacy. If all encryption unproblematically implemented both privacy and security through its normal operation, this claim would have no basis. To achieve through technology what has historically been considered a cultural value or legal right, however, computer scientists have attempted to translate privacy into quantifiable and thus measurable properties of computer systems. Indeed, it is in terms of the partially overlapping properties through which computer scientists respectively define privacy and security that the two concepts are bound together another as objects of systems engineering.

Computer scientists specifically decompose security into the properties of confidentiality, integrity, and authentication. Of these, confidentiality most closely aligns with lay understandings of encryption-based security. Encryption preserves a communication's

confidentiality when it prevents unauthorized parties from gaining access to and thus learning its semantic content. By contrast, encryption protects integrity when it prevents encrypted data from being tampered with or altered by anyone other than the holder of an encryption key. Finally, encryption may be used to provide authentication, i.e. to ensure communicating parties that they are indeed communicating with whom they intend and not some imposter.

By contrast, computer scientists analyze the ability of a technical system to provide "privacy" in terms of properties modeled on the legal principles of consent and data control. Specifically, to preserve privacy, a technical system must be unlinkable. It must leak no unintended information, thus preventing its data from being used to de-anonymize data contained in other systems. It must also transparently log and report data processing activities such that its conformance to user expectations and consent can be verified. And individuals must have the ability to intervene or direct action within the system, whether by correcting errors (e.g. fixing typos in a phone number), deleting data, or withdrawing consent altogether.

What is important to note here is that if computer scientists define security and privacy in terms of distinct systems properties, they nonetheless treat privacy as being effectively predicated on security. This is to say that for technologists, as a matter of engineering, there is no privacy without security. Take for example the kind of large-scale breach of a retailer's customer database that one is likely to encounter in the news on any given day. Most proximately, corporate hackings represent failures of security, failures of a company's security technologies and procedures to maintain the confidentiality of its customer data. That said, for technologists, every security breach necessarily leads to privacy violations. The theft of a piece of sensitive personal data, like a credit card number, will directly violate individual privacy in obvious ways.

Indirectly, even the theft of facially innocuous information, like a phone number or one's purchase history, may still violate privacy by enabling malicious actors to de-anonymize a person's data trail across the web (see Chapter 2). From this perspective, privacy is supplemental, an add-on to the baseline of computer security.

With these properties in mind, we can return to Narayanan's analysis. We can appreciate now that in describing cryptography as more successful in promoting security than privacy, he was saying something like the following: In purchasing or building the technological infrastructures through which they operate, corporations and governments generally aspire to be technically secure. They aspire to have effective access controls over customer data, to prevent data tampering, etc. They rarely, however, aspired to achieve the forms of unlinkability in terms of which computer scientists actually define privacy. And indeed, even in the hopeful post-Snowden moment of shifting world-making around privacy, technologists primarily pursued their mandate using forms of encryption directed towards confidentiality, integrity, and authentication, but not towards unlinkability, transparency, and intervenability. Even as technologists involved in the push for universal encryption held encryption out as the key to privacy's future, they promoted the kinds of "crypto-for-security" necessary but not sufficient for privacy, and not the kinds of highly specialized, domain-specific crypto-systems that someone like Narayanan would acknowledge as actually capable of producing privacy.

Technological Impunity

In February 2016, the FBI filed an application for assistance with a federal magistrate judge in California. The application related to its ongoing investigation of the December 2015 mass

shooting in San Bernardino. The shooting had resulted in fourteen deaths, and the FBI believed it to have been inspired by the Islamic State, a terrorist group.

With the application, the FBI sought an order compelling Apple to "assist" it in searching an iPhone used by one of the San Bernardino shooters. The government had seized the iPhone after the shooting, believing it to contain crucial evidence. Despite possessing lawful authority to do so, the government had been unable to complete its search due to encryption protecting the phone's data. The FBI argued that Apple retained the "exclusive technical means" necessary to bypass the phone's encryption, but had declined to voluntarily assist the government in so doing.

As civil liberties lawyers would soon note, the All Writs Act cited as the application's statutory basis was little known and infrequently used. Nonetheless, on its face it provides federal courts with the authority to issue "all writs necessary or appropriate in aid of their jurisdiction." Existing case law interpreted this language as including the authority to compel third parties to provide technical assistance in the execution of valid search warrants. Magistrate Sheri Pym granted the FBI's application on the day of its filing, ordering Apple to provide the FBI "reasonable technical assistance" in its search.

Despite this dry language, American computer scientists and security experts quickly mobilized alongside major technology companies and civil liberties advocates to defend Apple. They identified the FBI's filing as a radical and surreptitious assault on encryption in general. Magistrate Pym's order technically applied only to a single iPhone. FBI officials insisted their sole concern lay in identifying, via search of the phone, any accomplices to the San Bernardino shooting. As FBI Director James Comey soon insisted to Congress, "this case is about this case." Looking past the order's limited scope, however, technologists believed its reasonable assistance

standard required Apple to subvert security precautions it designed specifically to prevent iPhones from being decrypted using the "brute force" technique preferred by the FBI. Given the iPhone's security architecture, this would be no easy task, effectively requiring Apple to build a custom operating system. In heated editorials, interviews, tweets, and other public pronouncements, technologists attributed a boundless proliferability to the "hacking tool" that would result. If created, they warned, it would weaken the security of all Apple devices and threatened to fatally compromise the security assurances of encryption in general.[26]

In critiquing the FBI's All Writs application, technologists linked it to a prior period of social conflict over the nature of the encryption accessible to the American public. It marked, they said, a return of the "Crypto Wars." In 1993, the Clinton Administration attempted to induce technology companies to incorporate a tamper-proof "Clipper Chip" into their products. The government represented that the Clipper Chip would offer the public strong encryption,[27] but on condition of providing law enforcement and intelligence agencies with a "backdoor" through which to decrypt communications traffic upon court order. I return in the following sections to the cultural work that technologists performed in telling the story of the original Crypto Wars.

---

[26] In a written submission to Congress, Taylor Peake Wyatt, the founder of an app development company, identified the entwined technological and legal vectors by which technologists understood the San Bernardino order to threaten all of encryption. The "realities of digital security," Wyatt wrote, belied the "tempting" belief that the FBI's request would only "affect one phone used by a terrorist." "Once Apple makes the tools required to help the FBI hack that one phone they will work on EVERY iPhone and it will be a target for hackers around the world. And, it's clear law enforcement around the country will also start lining up to use these tools regardless of whether it's a terrorism case or a case of unpaid taxes. Putting the genie back in the proverbial is not a real option" (Committee on the Judiciary, The Encryption Tightrope, 183).

[27] According to computer scientist Susan Landau, "strong" cryptography is a sliding term for any type of cryptography that is hard to break given the current state of technology (Committee on the Judiciary, Going Dark).

For now, what I want to emphasize is that if by 2016 technologists were primed to interpret the FBI's All Writs filing as a threat to encryption, it was most proximately because the FBI had been arguing for years that it was "going dark."

Starting in 2010, top FBI officials claimed in a series of highly-publicized speeches, interviews, and court filings that there was a growing gap between law enforcement's legal authority and its practical ability to intercept communications. As a result, vital investigative capabilities—surveillance capabilities—were under profound threat. Officials insisted that fulfilling their public safety and counter-terror mandates would inevitably and imminently require technology developers to implement some form of what they now called "lawful access," rather than backdoors.[28]

In its early articulations, the FBI attributed going dark to increasing fragmentation in the communications market and the growing popularity of "exotic," new internet-based communications devices and services (Caproni testimony on Going Dark).[29] FBI officials described going dark as primarily a challenge to investigative timeliness—the ability to immediately, completely execute court-ordered wiretaps. In 2014, however, the FBI began to attribute what it continued to describe as the "imminent" going dark of its surveillance abilities primarily to encryption's proliferation. And as the computer scientists who self-identified as

---

[28] In the remainder of this chapter, I follow computer scientists and security experts in using the term "exceptional access" rather than lawful access or the more colloquial "backdoor" to indicate that such mechanisms were not contemplated in the original relationship between technology developers and their users.

[29] Technologists viewed the demise of the telephone's monopoly on American communications to be significant in this context because while telephone providers are required under the 1994 Communications Assistance for Law Enforcement Act (CALEA) to provide law enforcement officials with an easy means to wiretap customers, CALEA's access mandate does not extend to internet-based service-providers.

Crypto War partisans were aware, this shift dated precisely to Apple's introduction of the encryption used in the San Bernardino shooter's iPhone (see Zittrain et al. 2016).[30]

In September 2014, Apple announced that the next version of its mobile operating system would significantly expand the categories of iPhone data protected by encryption. It would also encrypt such data by default. Given consumers' well-known propensity to rely on default settings, this move promised in-and-of-itself to significantly expand encryption's use in American consumer technology. It was precisely the kind of unilateral corporate action on behalf of privacy, which the attendees of Computers, Freedom, and Privacy had hopefully anticipated. The new operating system, however, would also provide "device" rather than "service-provider" encryption. This meant that iPhone encryption keys would now be tied to a user's chosen passcode and stored locally on his or her phone. Apple, as the service provider, would no longer retain any keys. "Unlike our competitors," Apple's new privacy policy provided, even presented with a valid warrant or court order, we will no longer be technically capable of providing access to encrypted iPhone data.

Within days of Apple's announcement, encryption become the focal point of the FBI's going dark narrative. As it did, the narrative took on something of the totalizing, teleological

---

[30] Against this backdrop, the FBI's All Writs application "initiated" the second Crypto Wars by extending into the consequential domain of the common law a semiotic campaign already being waged in the public sphere. For technologists, the San Bernardino case, with its lurking terrorist accomplices, was too perfect, too compelling. Indeed, for this reason it was rumored in the technology press (and later partially confirmed by Justice Department investigation) that the FBI had intentionally avoided trying to hack the San Bernardino iPhone. The FBI's goal, as technologists characterized it, was not to conclude its investigation so much as to create an authoritative principle capable of circulating in American law independent of San Bernadino's exigent circumstances. This would constitute an end-run around Congressional and Presidential approval, but one which would empower law enforcement agencies to finally commandeer the technology industry's innovative capacities in the fight against going dark.

logic previously attributed to encryption by the Cypherpunks. Take for example the March 2016 testimony of FBI Director James Comey before the House Judiciary Committee. During the hearing, Comey prefaced his remarks by carefully acknowledging encryption's social and economic value. He then insisted, however, that encryption's defense of online security and personal privacy stood in necessary tension with the imperatives of public safety and national security. Even I, Comey granted, "think it sounds great when people say, 'Hey, you buy this device, no one will even be able to look at your stuff.'" But "there are times when law enforcement saves our lives, rescues our children, and rescues our neighborhoods by going to a judge and getting permission" to look at our stuff (Comey testimony on The Encryption Tightrope). Privacy, grounded in encryption, might sound "awesome." But "stopping this kind of savagery and murder and pedophilia, and all the other things that hide in the dark spaces of America life, is also incredibly important…" (Ibid.).

The going dark problem as Comey now expressed it was no longer that encryption would make investigative techniques less efficient or timely—a more realistic assessment of its likely impact, according to computer scientists. Rather, it was that in the "not-too-distant future," encryption would spread to the extent that "all of our conversations and all of our papers and effects are entirely private" (Ibid.). This fantasmic future of perfect, absolute privacy would preclude law enforcement from 'bringing us' public safety. It would thus subject Americans to all the yet-to-be-imagined threats of contemporary counter-terror logic (see Massumi 2015; Masco 2014; Amoore 2013). It would also, however, undermine the rule of law, and in so doing disrupt public order. On the one hand, it would create "warrant-proof" spaces in American life, technologically-enforced zones of legal impunity(Comey testimony on The Encryption

Tightrope). Historically, Comey testified, there had been "no closet in America, no safe in America, no garage in America, no basement in America that could not be entered with a judge's order."[31] Now, by contrast, encryption allowed people to place themselves beyond the law. It told drugs dealers and killers, "do what you want… because law enforcement can't get into your phone."

On the other hand, the very act of implementing device encryption independently threatened the rule of law by effectively usurping the government's legislative prerogative. This was the language, at least, in which Manhattan District Attorney Cyrus R. Vance, a key proponent of exceptional access, dismissed Apple's claim that the FBI's All Writs application threatened customer security. In written testimony, Vance countered that in implementing its new encryption system Apple had "effectively decided they know better than our elected representatives…how to keep Americans safe" (Vance testimony on The Encryption Tightrope). Private companies, he continued, should not be able to dictate through their technologies "who can access key evidence in criminal investigations." Just because the law had failed to "keep pace with technology" they should not be able to "write their own laws" and "upset the balance between privacy and public safety established by centuries of jurisprudence."

---

[31] As technologists noted throughout the Crypto Wars, Comey's claim is historically dubious. There has never been a guarantee under American law that a subpoena or search warrant "will result in the revelation of the contents of a private message" (Landau et al. 1994, 34). Legal scholar William Stuntz (1995) argues to the contrary that the Fourth and Fifth Amendments exist specifically to constrain the state's regulatory power by making the collection of evidence for certain kinds of crimes unavailable to law enforcement.

The Costs of Surveillance

In positing technological subversion of the rule of law alongside unstoppable crime and terror as inevitable effects of encryption's spread, Comey and Vance invoked both the anarchistic yearnings of the Cypherpunks and the utopian visions of early internet activists. These visions were most famously illustrated in "A Declaration of the Independence of Cyberspace," a manifesto published in 1996 by Electronic Frontier Foundation co-founder John Perry Barlow. In it, Barlow asserted that the internet, as a self-ordering planetary commons that disregards geographic and thus jurisdictional boundaries, transcends government regulation. Such arguments may have been as much normative pleas for nations to resist regulating the internet as they were empirical assessments of some unique ability of it to escape government control.[32] The notion of technology's capacity to frustrate or rival the law lingered, however, in the work of scholars like law professor Lawrence Lessig (1999), who famously analogized software code to cultural norms and market prices in possessing the law-like ability to regulate human behavior.[33] If by 2014 encryption was culturally legible as a potential rival to the law, it was further because American law was itself already widely understood to be a machine-like "tool." As Annelise Riles (2005) has shown, since the mid-20th century, American lawyers have widely approached the law as a technical practice judged in terms of its ability to achieve stated social ends. This

---

[32] See Lessig (1999); Wu (1997). For a discussion of the technological context of this "folkloric notion," see Kelty (2008, 51-7).

[33] See Reidenberg (1997) for a similar contemporary argument. See Ziewitz (2016) for a recent argument transposing the locus of technology's law-like capacities from code to "algorithms." It is worth noting that in declaring "code is law" Lessig acknowledged that code could be used to evade government regulation. His primary concern, however, was that the internet's transformation into a medium of commerce was creating a technological architecture that threatened to inadvertently perfect the government's ability to control the public.

instrumental modeling of the law upon technology has become so foundational to the distinctions between law, politics, and philosophy that the law has become hard to conceive otherwise.

In considering the FBI's claim that encryption subverts the rule of law, it is tempting to dismiss it as a mis-recognition of the "technical orientation" (Forsythe 2001, 44), which animated the post-Snowden campaign for universal encryption. This is to say that what law enforcement characterized as an illegitimate intervention in the legal regulation of human behavior might be more properly understood as a more limited effort to regulate the behavior of technical systems, one required by technologists' understanding of their professional obligations.

Recall in regard to this proposition that the rolling wave of Snowden-related reporting published between 2013 and 2014 revealed not only the existence of NSA mass surveillance, but also its means. In particular, a series of articles published in September 2013 disclosed that the NSA, in its quest to capture the entirety of global internet traffic, had actively subverted the cryptographic protocols widely used to secure everything from e-mail and web searches to the global banking system (see Perlroth, Larson, and Shane 2013; Ball et al. 2013). It had, moreover, actively compromised the technical infrastructure of America's most prominent technology corporations (see Gellman and Soltani 2013).

Many technologists responded to this news with a profound sense of shock and betrayal. The NSA, after all, is charged not just with monitoring foreign governments and terrorists, but also with protecting America's information systems from foreign spies and hackers. While much of its work to weaken and circumvent encryption was surreptitious, a portion of it relied on the

collaboration—willing or not—of private sector technologists.[34] A prominent line of discourse running through technologists' related public discussions held that mass surveillance of the kind practiced by the NSA is problematic in specifically technological ways. It was an engineering problem. Regardless of whether one generally approved of the NSA's mission, by exploiting vulnerabilities in commercial internet technologies and actively working to undermine encryption, the NSA was exacerbating the fundamental insecurity of the modern internet.

We can gain some purchase into the logic of this reaction via a remarkable document published by the Internet Engineering Task Force (IETF). The IETF is a non-governmental organization comprised of internet engineers, network operators, and computer scientists responsible for developing the protocols that enable the internet to function. In May 2014, based on extensive internal discussion, the IETF released RFC 7258, a new "best current practices" memo. It expressed the IETF's consensus "technical assessment" that the "pervasive monitoring" of internet communications constitutes an "attack on the privacy of internet users and organizations."[35] RFC 7258 defined pervasive monitoring as widespread, often covert

---

[34] A European privacy scholar told me in 2015 that Snowden was a "wake-up call" to many engineers that there were 'bad people in their profession.' Lawyers, she observed, are used to being portrayed as villains. Engineers, however, took from the news of the willing cooperation of some Silicon Valley companies in NSA programs the difficult lesson that their profession is subject to the "same corrupting influences as politicians."

[35] RFC 7258 specified that pervasive monitoring is an attack in the narrow "technical sense" that it subverts the intent of communicating parties without their knowledge or consent. It is in this respect functionally equivalent to other known attacks, like changing the content of a communication. The IETF argued that, as an attack, pervasive monitoring relies so heavily on known vulnerabilities that it does not constitute a new category of technical exploit. The IETF instead distinguished it on the basis of it "being indiscriminate and very large scale." Despite historical efforts to mitigate known vulnerabilities through protocol design, RFC 7258 acknowledged that the IETF had inadvertently facilitated pervasive by failing to seriously contemplate the possibility of an attack at the scale of the entire internet.

surveillance effected through intrusive gathering of internet content and metadata. Without mentioning the NSA, it cited as the constitutive components of pervasive monitoring the hacking techniques the NSA had now been revealed to employ.

As engineers, the IETF's members generally view themselves as obligated to build technologies for society that are at least theoretically capable of being secure. In addition to establishing pervasive monitoring's "technical nature," RFC 7258 thus also set forth the IETF's determination to "mitigate" its technical aspects wherever possible via protocol and architectural design. In recommending such mitigation, the authors identified the intentions behind pervasive monitoring as being necessarily out of scope. If they thus declined to account for the NSA's motivations, it was not because there was no clear organizational consensus on the NSA's actions, although this was also true.[36] Rather, it was because they believed there is no way to determine the morality of 'acts on the wire.' From the perspective of technique and effect pervasive monitoring conducted for purposes of nation-state surveillance could not be distinguished from pervasive monitoring conducted either for "legal but privacy-unfriendly" business purposes or as criminal acts. If the IETF were to "defend against the most nefarious actors" by mitigating pervasive monitoring—as it would any protocol vulnerability—it would have to defend against monitoring by other actors "no matter how benevolent some might

---

[36] Foreign members of the IETF quickly noted that despite its origins in the U.S., the IETF is a global organization, and that U.S. legal protections against surveillance apply only to U.S. citizens and legal residents. Regardless, therefore, of whether a citizen believed the NSA acceptably balanced national security and civil liberties, NSA mass surveillance was indefensible from a global perspective. It violated the laws and privacy interests of foreign nationals without rebounding to their benefit in the way it did, at least arguably, for Americans.

consider them to be."[37]

As reflected in RFC 7258, technologists registered the Snowden revelations through a self-consciously articulated engineering perspective. Under its logic, internet engineers were obligated to deploy their expertise to prevent a newly recognized form of harmful system behavior—attacks at the scale of the entire internet—regardless of the morality of the human behavior thereby indexed.

In making this observation, I do not mean to suggest that technologists were not also deeply engaged with the perceived social and political implications of mass surveillance. As suggested by the FBI, there was indeed a sense in which the post-Snowden campaign to encrypt all the things sought to perform the law-like function of "determining what people can and can't do" (Lessig 1999). This was implicit, for example, in the observation of Apple CEO, Tim Cook, in a March 2016 interview: "[A]t the end of the day, we're going to fight the good fight not only for our customers, but for the entire country. We're in this bizarre position where we're defending the civil liberties of the country against the government" (Grossman 2016).

Still, if as this suggests technologists sought through universal encryption to intervene in citizen-government relations, it was not by promulgating socially authoritative, law-like rules. And it was in no way reinforced, like the law, by the threat of state sanction. By contrast to the Cypherpunks, the computer scientists and internet engineers involved gave no indication they intended to technologically supplant the law. Technologists flirted with asserting their

---

[37] An IETF member argued in one of the post-Snowden email discussions, "The point is not that the NSA can surveil you, the point is that _anyone_ can. The NSA is just who publicly did it recently. The lesson is not 'Okay, so let's stop law enforcement from from eavesdropping.' It is 'holy shit, we are really vulnerable.'"

technological will in the service of fundamental rights "regardless of government policies," as Snowden put it at HOPE. But they still widely identified mass surveillance as being ultimately a political problem requiring political solutions. They consistently expressed the belief that computer security must be grounded in law, even if in the meantime this was no reason to neglect technical mitigation.[38]

Rather than an assault on the rule of law, the campaign to encrypt all the things expressed a critique of the law as a modality of regulation. We can see this critique most clearly in the sensibility of technical efficacy, which marked the post-Snowden claims to the stewardship of privacy. Recall that when I described the feeling of cautious optimism I encountered at Computers, Freedom, and Privacy, I observed that many attendees specifically identified engineers and not lawyers as the social agents best positioned to leverage it in the service of meaningful privacy reform. This conclusion was predicated in part on the perceived power of technology companies vis-a-vis the affordances of their mass market products. At a more intimate scale, though, it was predicated on a certain self-identified hacker sensibility, a belief in engineers' superior speed and efficacy vis-a-vis lawyers as agents of social change. Some conference-goers, for example, justified valorizing technologists in terms of the glacial pace and perceived dysfunction of legislative and judicial processes. In their introductory session on government surveillance reform, Mark and Sandy thus reminded the audience that Congress hadn't passed a federal privacy bill since 1986. A privacy scholar separately observed to me at lunch, relative to technology's appeal, that she had filed a human rights abuse lawsuit in 1996,

---

[38] As an IETF member noted in a related email discussion, the existence of laws against burglary is no reason not to develop better locks.

which had just gone to trial in 2015.

Over the broader arc of my research, I found that many technologists further questioned the law's efficacy as a vector of privacy reform based on the way the law constantly "lags" behind technology. Such lag manifested for them in the obvious disconnect between the law's presuppositions about technology and the actual privacy-invading capabilities of contemporary technology in operation.[39] Following the Snowden revelations, for example, many technologists faulted existing American surveillance laws for being tailored to analog technological systems no longer in use. The law, for example, protected communications content against surveillance under many circumstances, but offered no protection for communications metadata (see Chapter 2). It thus left entirely unregulated analytic techniques with vastly greater privacy-invading powers than those applicable to content like recorded phone calls.[40] Why trust laws, many asked, that discount the exceptional analytic potency of metadata when every technologist knows better?

---

[39] The related ideas that the law "lags" behind developments in technology and that judges and lawmakers lack the technical literacy to make good law regarding technology were, like the implementation gap, constant objects of concern for computer scientists and internet engineers. Lessig (1999 Ch. 9) identifies the "latent ambiguities" produced by the recurring gap between the presuppositions about technology embedded in the law and the privacy-invading capacities of technology's leading edge as a key driver of American privacy law's historical development.

[40] Metadata is data about data. In the context of the internet, it generally includes time, duration, and location information as well as identifying information for the sender and receiver of any communication. Metadata is often highly revealing in-and-of-itself. Technologists often offer, in this regard, the hypothetical example of a phone call to an addiction treatment center. One doesn't need to know the content of such a phone call to infer that the caller might prefer to keep the fact of the call, revealed through metadata, private. Technologists also consider metadata to be a powerfully invasive object of surveillance because it is generated in machine-readable format. It is thus easily subjected, by contrast to content like a recorded phone call, to machine-learning based forms of inferential analysis.

By contrast to the law, for those trained in its ways, technology held the appeal of immediacy and efficacy. As Snowden reminded his audience at HOPE in 2014, a dedicated hacker might fill the gaps left open by America's patchwork privacy laws by quickly building and publishing to the web new privacy-enhancing tools and widgets. Or she might implement already-existing but underused privacy-enhancing technologies. And she could do either without coordinating large numbers of other people—a drawback of the law from technologists' perspective. At the same conference, Peter Eckersley, lead technologist for the Electronic Frontier Foundation, suggested further that technologists' ability to build solutions to mass surveillance somehow inoculated them against the privacy malaise— the widespread resignation to privacy's ongoing erosion—which otherwise hampered efforts at reform. "If we go out and built tech structurally secure against dragnet surveillance," Eckersley explained, "then that's the world we're going to live in."

What I would argue further is that the "form and performative character" (Riles 2005, 22) through which technologists carried out the post-Snowden campaign to encrypt all the things was indeed modeled on another domain of social practice. This domain was not the law, however, but the market. What best characterized, in other words, the particular ways in which technologists conceived and carried out the campaign for universal encryption was a commitment to and appreciation for an aesthetics of market-like efficiency.

We can access the contours of this market aesthetics by reconsidering the terms in which technologists' characterized the nature of the problem that universal encryption was to solve. As illustrated in RFC 7258, from an engineering perspective, this problem was at its most basic the failure to secure much of internet traffic against surveillance at the scale of the internet. Outside

the purified realm of IETF documents, however, security experts frequently attributed the fact of mass surveillance not to the absence of security technologies but to having allowed internet surveillance to become "too cheap."

This understanding of mass surveillance as being a function of the costs of surveillance as much as of the poor state of internet security deeply conditioned the post-Snowden campaign for universal encryption. We can begin to see its outline in the November 2013 technical plenary delivered by security expert Bruce Schneier at the IETF's annual meeting. Schneier is among the nation's best known security experts, the author of numerous popular books on security, surveillance, and the data economy.[41] In his talk, he explained to his fellow engineers that there are "a lot of technical things we can do. The goal is to make eavesdropping expensive. That's the way to think about it." "The first thing encryption does," Schneier continued, "is it makes it harder. If it is easy for the NSA to grab all of Google's traffic as it goes between data centers, it will do it. If it's encrypted, it can't do that. Right?"

We can gain further insight into this market aesthetic by considering the inherent limitations that technologists ascribed to encryption as a solution for mass surveillance. Despite the self-identified proclivity of engineers to say, 'we can solve this problem completely,' with universal encryption, internet engineers were at pains to acknowledge that it was unlikely to fully resolve the problem of NSA surveillance. There were numerous reasons for this likelihood, as

---

[41] Months earlier in 2013, Schneier served as a technical consultant to the reporters entrusted with the NSA papers, tasked with helping them understand in practical terms the technical capabilities revealed therein. Based on this experience, upon publication of the reports detailing the NSA's subversion of encryption, Schneier accused the government of "betraying" the internet. He became a leading voice alongside Snowden calling for the engineering community to 'take the internet back' (see Schneier 2013b).

technologists understood it. These included the exploitable vulnerabilities that were sure to continue to be found in widely used encryption protocols, as well as the belief, given the technical difficulties involved, that any meaningful protection of metadata against surveillance would be best grounded in the law rather than in encryption. Moreover, technologists argued, whatever effects encryption might have on mass surveillance, a sufficiently motivated attacker with the resources afforded by the US "black budget" would remain capable of compromising any individual suspect targeted.[42]

Given such limitations, what precisely did technologists hope to achieve through universal encryption? How did they imagine that raising the costs of surveillance would serve to curtail mass surveillance? In practical terms, increasing the use of encryption would make surveillance "harder" by forcing the NSA to conduct "active" attacks rather than "passive" surveillance. Technologists anticipated this would increase the computational complexity of the NSA's work (see Barnes et al. 2015), requiring the NSA to deploy "racks and racks of servers, which will require significant amounts of cooling power." The direct monetary cost of such infrastructure might require the NSA, engineers surmised, to divert funds used to subvert encryption standards or conduct mass surveillance towards the targeted surveillance of specific, individual terrorist suspects. Indirectly, the expanded infrastructure's sheer physical presence

---

[42] Indeed, for many technologists the post-Snowden focus on preventing specifically mass surveillance served not just to mitigate technical vulnerabilities, but also to defend the rule of law. This understanding was predicated on Snowden-related reporting to the effect that the government had relied on secret and potentially unconstitutional legal authorities to justify its mass surveillance of American citizens. From this perspective, to force the NSA to shift resources from mass to targeted surveillance was equivalent to forcing the government to rely on legal authorities, especially Fourth Amendment probable cause requirements, which are clearly established in public law.

would also impose costs. It would undermine the NSA's ability to operate covertly, to hide the evidence of its activities in the utility closets of network access providers, for example[43] (see Farrell, Stephen, and Hannes Tschofenig. 2014.). Technologists hoped that this would be force the NSA to incorporate the reputational and other negative costs of "discovery and attribution" (i.e., loss of access, arrest, prosecution, etc.) in deciding whether and how to exploit any given vulnerability (see Schneier 2013c; Barnes et al. 2015). "You know oddly," Schneier observed to the IETF, "near as I can tell the NSA," under sway of the national security mandate to 'collect everything,' "doesn't weigh the costs and benefits of their programs." Still, he continued, "the NSA has a big budget, but they are not magical. They have the same limitations as everybody else."

What we see in such statements is that if technologists responded to Snowden by deploying their technical expertise, they did so guided by a model of the world derived from neoclassical economics. This model presumes that all social actors, including collective entities like the NSA, are—or should be—rational, self-interested 'satisfaction'-maximizers operating in a world of relative scarcity (see Tribe 1972). With ubiquitous encryption, then, technologists sought to change the NSA's behavior not by directly stymying its technical capabilities but by inducing a shift in how different forms of surveillance entered into the NSA's decision-making. Or more precisely, per Schneier, technologists sought to force the NSA, perhaps for the first time, to make decisions in service of its counter-terror mandate as a "calculative agency" (Callon

---

[43] In 2006, a long-time communications technician for AT&T revealed that the NSA had been diverting and siphoning the entirety of the internet traffic received by AT&T via an interception facility surreptitiously located in AT&T's Folsom Street, San Francisco office building (see Singel 2007)

1998a). The strategy of raising the costs of surveillance thus ultimately turned on inducing the NSA to adopt a subject position from which it would be obligated, like all market actors, to identify and rank its goals and efficiently allocate its finite resources to maximize top priorities. Only after pushing the NSA past the imagined threshold of impunity afforded by its nearly unlimited budget would the market intervention effected by encryption's spread reliably mitigate against indiscriminate mass surveillance and in favor of more efficient targeted surveillance.

Whether or not government officials would recognize any truth in this representation of the NSA as operating beyond the constraints of the market, the image is quite revealing. On the one hand, it suggests that the intermingled paranoia and awe that the agency elicits in American public imaginaries derives not just from its long designation as a state secret (see Masco 2014). It also derives from the NSA's perceived market exceptionalism, its status as the rare social institution apparently able to operate outside the value logics of the market[44] otherwise viewed today as the measure of all social worth (Povinelli 2011). On the other hand, it shows technologically-mediated market intervention to be the only imaginable avenue of social recourse available to reign in the anti-social tendencies of counter-terror institutions widely viewed to have escaped traditional forms of public accountability grounded in the law (see Masco 2017).

[44] The NSA is joined in this exceptional category by America's nuclear weapons complex, which as the third most costly post-WWII federal expenditure after social security and the non-nuclear military, has never conformed to market logics (Masco 2010).

Mandated Insecurity

In June 2015, I attended the annual Privacy Enhancing Technologies Symposium (PETS), a conference dedicated to innovations in privacy technology, held that year in Philadelphia. Matt Blaze, a well known computer scientist based at the University of Pennsylvania, delivered a keynote address on the conference's final day. Rather than discussing a new development in privacy technology, Blaze organized his talk around the apparent disconnect between the going dark narrative and a set of statistics on federal and state wiretap applications compiled by the Administrative Office for United States Courts. As Blaze recounted, the FBI characterized going dark not as a future problem, but a current one requiring attention from the highest levels of government. This suggested "there's a major epidemic of crimes not being solved." But the annual Wiretap Report showed that in 2014 the number of wiretaps in which investigators encountered encryption actually decreased, from 43 to 22. In only four cases did officials report being unable to decipher encrypted communications. Why would the FBI squander "political capital" on encryption, Blaze asked, given such numbers? What compelled them to rehearse the same "terrible" arguments the FBI had first made, unsuccessfully, 23 years prior?

Blaze was a particularly apt messenger for such skeptical appraisals of the factual basis for the resurgent Crypto Wars. In 1993, as a young cryptography researcher at AT&T's Bell Laboratories, Blaze made what technologists view as a key contribution to their victory in the initial Crypto Wars. That year, to Blaze's surprise, AT&T introduced a new secure telephone, the first (and as it would turn out, the last) commercial product to incorporate the NSA-designed Clipper Chip. At the time, there was very little public information regarding the design of the Clipper Chip or its key escrow system. Through a series of experiments on a prototype

encryption key generation device, Blaze quickly found a serious flaw in the Chip's "backdoor" escrow mechanism. The flaw enabled users to easily subvert its security features and disable the government's ability to decrypt user communications (see Blaze 2011). Blaze's discovery made national headlines and was widely cited in July 1994 when the Clinton Administration abandoned the Clipper Chip proposal.

If Blaze was a singular messenger, the mood of exasperation that suffused his address was characteristic of the way technologists talked about the Crypto Wars. Looking back to the original Crypto Wars, Blaze described it as a kind of "denial of service attack" on the computer science community. Rather than spending the 1990s actually securing the internet, the Crypto Wars forced computer scientists to spend their time explaining "all the obvious reasons" why encryption backdoors were a bad idea (Blaze 2011). The debate was "fun" to the extent computer scientists were "completely right" and thus did not need to compromise. "What they're saying doesn't justify what they're asking for," Blaze concluded.

It was not simply Blaze's mood but the form in which he expressed it, which I found to be characteristic of the Crypto Wars. While I occasionally encountered the Crypto Wars in other forms,[45] I primarily encountered it as a story, one told by computer scientists like Blaze as a techno-moral fable. These stories followed a recognizable arc. In them, computer scientists located the contemporary debate over encryption as a continuation of the government's dogged historical efforts to limits encryption's availability and strength. They described how American policymakers had faced "just this issue" in the 1990s (Bankston 2015), and how at that time,

---

[45] In February 2016, for example, I waited in the frigid evening air outside an Apple store on Chicago's Michigan Avenue for what would turn out to be a very sparsely attended demonstration in support of strong encryption, prompted by the Apple-FBI standoff.

computer scientists had actively collaborated with a coalition of allies in civil society and the nascent internet technology industry to explain to the government the technical risks entailed by encryption backdoors. The coalition eventually "won" the Crypto Wars. They convinced American policymakers that the political, economic, and security benefits of readily available strong encryption far outweighed the risks of going slightly dark. This victory was evidenced in the Clipper Chip's demise and the abandonment of the export controls on encryption. It was further evidenced in the findings of numerous government reports to the effect that encryption was vital to securing the internet and thus to the likelihood of the internet achieving its social and economic potential (see, e.g., National Research Council 1996). If the "lessons" of the original Crypto Wars needed to be shared, it was because the government had now commenced again upon the same fruitless pursuit, even though the arguments which had won the original Crypto Wars remained equally valid, if not more so, today.

I sketch out the common features of this story not to assert its historical accuracy, but to ask what was at stake for computer scientists in rendering the Crypto Wars in this way. Riles (2005, 11), drawing on the sociology of knowledge production, observes that the generally-accepted historical sequences important to expert communities should be understood as rhetorical achievements. Science and technology studies scholars have repeatedly found, for example, that even the apparently incontrovertible facts 'discovered' through empirical science gain their truth value only when subsequent "collections of statements" retrospectively attribute certainty to them (see, e.g., Latour 1987; Shapin and Schaffer 2011). From this perspective, in rendering the debate over the availability and strength of encryption as a historical sequence technologists figured it to be a kind of "black box," a normatively incontestable part of the

common sense of engineering practice. The question then becomes not whether the story's historical claims are right or wrong, but what they achieve or perform by being rendered in this way (Riles 2005).

On the one hand, in telling the story of the Crypto Wars, technologists provided a compelling vision of their project, a unifying fantasy about the practice of technical knowledge as a form of consequential political action (see Riles 2005; Biagioli 1990). From this perspective, the significance of the Crypto Wars lay in its demonstration of computer scientists' willingness to defend their vision of the internet's moral and technical order (Kelty 2008) by actively engaging in political processes.

As indicated in RFC 7258's refusal to address the intentions behind mass surveillance, American computer scientists and engineers have historically ascribed to a form of "political agnosticism" (Coleman 2013). Despite the American technology industry's fondness for claiming that its "innovations" will "change global culture" (Lanier 2013), individual technologists have explicitly resisted and disavowed the inclusion of politics, popularly understood, in technical work. In her study of Artificial Intelligence engineers, for example, Diane Forsythe (2001, 44) found that while such engineers recognized social matters as "troublesome," they tended not to find them to be 'interestingly' or importantly problematic in the way they did technical matters. Chris Kelty (2008) later found that free software programmers practiced a form of politics, but one limited to preserving the technical and legal means of their own organization (see Chapter 3). It is arguably only over the past half decade that American computer experts have exhibited a pronounced willingness to deploy their expertise on behalf of a robust if heterogeneous activist politics (see Coleman 2014). Given this

historical orientation, in telling the Crypto Wars' story as computer scientists' first major political engagement alongside the nascent internet community (Bankston 2015), they offered a vision of computer security expertise as properly directed towards and engaged in coordinated, collective political projects like public education and advocacy.

On the other hand, the significance of the Crypto Wars story lay perhaps more precisely in the perception that this entree into the fray of politics had been "a rousing success" (Bankston 2015). To appreciate this point, and the cultural work it allowed technologists to perform, we must first look at the nature of the core technical argument forwarded in the Crypto Wars. As Blaze put it at PETS, from a "policy perspective [exceptional access] sounds kind of perfect. It we could only develop this magic key, keeping bad guys out and letting us in." But from a "technical perspective," it was always "kind of obvious" why the Clipper Chip's key escrow system—and indeed any potential form of exceptional access—was a bad idea. Why? Because as a matter of implementation and operation, it was infeasible, it was a "technological impossibility." In a report on the going dark debate, Bruce Schneier elaborated: "As technologists, we can't build an access system [of the kind demanded by law enforcement] that only works for people of a certain citizenship, or with a particular morality, or only in the presence of a specified legal document" (Zittrain et al. 2016). This was so because at the level of technical operation, any exceptional access mechanism would be functionally indistinguishable from a "backdoor." By its existence it would introduce a new, vulnerable path to the unauthorized, surreptitious recovery of data. It would, by definition, constitute an "engineered vulnerability" (Bellovin et al. 2013).

If exceptional access thus constituted a security vulnerability based on its function within the schema of systems engineering, this conclusion was compounded and lent urgency by what computer scientists understood to be the security challenges common to the design and engineering of all large-scale technical systems. In particular, as I observed at Mozilla, it is a central axiom of not just security engineering but software development in general that software coding is an inherently insecure process, that "all code has bugs." Technologists attribute this tendency to a wide variety of factors, including poor design, human error, the propensity of certain programming languages to introduce security flaws, and difficulties involved in thoroughly testing and debugging software updates (see Tschoofenig 2014; Bellovin et al. 2013).

The practical consequence of code's inherent bugginess, though, is that complexity is "the enemy of security." The more code required for any piece of software, the greater the number of exploitable flaws it will likely contain.[46] Given the enormous inherent complexity both of encryption[47] and of the kind of large-scale re-engineering an exceptional access mandate would entail, computer scientists predicted that any exceptional access system would inevitably introduce flaws that would provide attackers with precisely the forms of system access they most desire. It would thus be functionally equivalent to weakening encryption, eliminating the inherent guarantee of security afforded by non-access crypto-systems (Bellovin et al. 2013; The

---

[46]At PETS, Matt Blaze posited the principle of simplicity, of reducing the number of system components to which engineers must attend, as "one of the few software engineering techniques that we actually know works." See Herley and Orrschot (2017) for an analysis of computer security's failings as a science.

[47] Computer scientists describe the implementation of even well-known free software crypto protocols as requiring hundreds of different functions to be selected, with even apparently simple details capable of producing serious security vulnerabilities.

Chertoff Group 2016). Most computer security configurations might operate on a sliding scale, but in this regard encryption does not. It is a binary with no middle ground.

This then had been the perceived significance of Blaze's Clipper Chip research, to practically demonstrate that unanticipated security vulnerabilities could easily be found, as predicted, even in a crypto-system designed and tested by the NSA's vaunted cryptographers. Whatever the ultimate cause for the original Crypto Wars' conclusion, technologists maintained, it had proven them right on the substantive merits.[48] By the time of the Crypto Wars' resurgence, computer scientists portrayed their dire assessment of exceptional access as having been further vindicated by history. It was a hard-learned lesson of security engineering by this point that "if a vulnerability exists in a security system, it is likely that someone will take advantage of it sooner or later" (Bellovin et al. 2013, 65). It was thus common in discussing the Crypto Wars for technologists to reference the endless series of corporate and government data breaches that punctuated the news in 2014 and 2015. They pointed, moreover, to the many access mechanisms, which since the original Crypto Wars had been surreptitiously and maliciously exploited to conduct illegal surveillance.[49]

In telling the Crypto Wars as a story of victory and vindication, computer scientists not only expressed their changing political orientation, but also enacted a form of social authority

---

[48] The non-technical claims which technologists described as having been historically vindicated included the arguments that strong encryption was essential to the internet's flourishing as a commercialized medium of mass communications and that the global proliferation of encryption had by the mid-1990s already exceeded the ability of any US-based law to control.

[49] They frequently cited in this regard a series of data breaches enabled by access mechanisms implemented by the telephone companies to comply with CALEA (see Bellovin et al. 2013, 65-6).

grounded in expertise.[50] As indicated by the emphasis on technical and operational feasibility in such stories, this was, in the first instance, an expertise of means. In telling such stories, technologists took no issue with the government's desire for access to encrypted data. But they insisted that science serve as the ultimate arbiter of that desire's feasibility (see Abelson et al. 2015, 20). Even so, technologists established the infeasibility of exceptional access not via claims of theoretical certainty, but rather through displays of professional humility. They asserted, based on their extensive collective experience designing and operating large-scale technical systems, that building a reliably secure exceptional access system exceeded the current abilities of computer science as a science.[51] They thus enacted themselves, by implicit contrast to law enforcement officials, as the kinds of persons willing to honestly and openly account for the limitations of their knowledge and skill (see Lempert 2005; Goffman 1967).

Closely entwined with technologists' enactment of a socially authoritative expertise of technical means was an expertise of technical and social risk. Connecting them was the claim that law enforcement's refusal to acknowledge the limits of cryptographic engineering necessarily implicated a failure to honestly account for the harms that would follow from engineering vulnerabilities into encryption. In venting their frustration with the Crypto Wars' resurgence, computer scientists and engineers broached this idea via the concept of "mandated insecurity."

---

[50] See Hilgartner (2000) and Carr (2010) for discussions of the performative techniques through which experts establish socially consequential forms of credibility and authority.

[51] As Matt Blaze explained at PETS, "computer security is a big, fat mess getting bigger, fatter, and messier by the day." It's the only branch of engineering where it's expected that products will fail after they are shipped to users. "Building engineers," by contrast, "don't expect buildings to fall down."

Specifically they claimed that while exceptional access might serve public safety in certain limited instances, it could only do so by mandating insecurity for all (Abelson et al. 2015). An access mechanism might ensure the FBI's ability to obtain a criminal suspect's or victim's encryption keys in some limited number of cases. But the exploitable flaws it would inevitably introduce materially increased the risk that the keys of all systems users would be stolen. The real tension, from this perspective, was not between security and privacy, as Comey had insisted, but per Schneier, between "less security and more security" (Zittrain et al. 2016). As Schneier testified to Congress, framing the Crypto Wars in terms of public safety threats enabled by encryption's spread myopically obscured the wider variety of threats that encryption secures against. Criminals might use encryption to "hide their planning from authorities." This was true, though, of many "aspects of society's capabilities and infrastructure," including cars and restaurants," which we generally accept "can be used by both honest and dishonest people." Weakening encryption, by contrast, was akin to secretly poisoning all the food in a restaurant. "Yes, we might get lucky and poison a terrorist before he strikes," Schneier wrote, "but we'll harm all the innocent customers in the process" by exposing them unnecessarily to the forms of theft, fraud, and harassment endemic to the internet (Ibid.).

Through stories of their vindication, computer scientists and engineers defined the policy question of encryption's relative strength and availability as a question of engineering feasibility and risk. Doing so created an opportunity for them to convert their technical authority into a kind of moral authority (see Hirschman and Berman 2014, 786; Nissenbaum 2005), to demonstrate the transferability of expertise derived from practical experience with large-scale technical systems to the domain of political relations and change (see Biagioli 1990). In rehearsing the

"lessons" of the original Crypto Wars, technologists delimited the potential scope of this transferability. At stake in the original Crypto Wars, as technologists described it, had been the internet's viability as a mass medium, and thus all of the cultural, economic, and civic benefits that the internet conferred on society. By 2014, the internet's flourishing was no longer in doubt. It was now inextricably intertwined with human sociality in all its dimensions. The societal stakes of the second Crypto Wars tracked encryption's extension alongside the internet into every aspect of contemporary life.

If technologists' expertise over technical means and risk was now theoretically transferable to any social process or domain mediated by the internet, they most noticeably sought to actualize it by asserting authority to speak to the necessary objects of national security. They did so by demonstrating the contemporary grounding of America's critical national security infrastructure in its commercial internet technologies. In 2016 testimony, security expert Susan Landau illustrated this relationship using the example of the smartphone (Landau testimony on The Encryption Tightrope). Smartphones, Landau testified, had in recent years become centralized repositories of personal information, the place where we keep all our photos, music, calendars, and contacts. They thus presented obvious computer security and privacy risks to their owners. But, Landau continued, they also now provide us, through login authentication, with access to proprietary corporate and government computer systems. Any flaws introduced into smartphone encryption thus implicate not just personal data, but also corporate intellectual property and critical national infrastructure. They risk, for example, exposing the login credentials of a low-level "HVAC employee who's going to service a power plant" to "organized crime or a nation-state." Smartphones, however, were only one example of a broad-scale

convergence of America's national security and defense infrastructure with its civilian and commercial information technology infrastructure.[52] For purposes of determining encryption's future, the salient consequence was that national security could no longer be considered apart from the security guarantees of private sector technologies and infrastructures.

According to Masco (2014), the traumatic failure of US intelligence agencies to predict the September 11, 2001 terrorist attacks continues to "haunt" US national security operatives, 'affectively recruiting' them to a "state of constant crisis." The resulting "politics of shock" creates a boundless drive for militarization. National security agencies fulfill this drive by increasing surveillance and by using spectacle, hypotheticals, and scenario planning to conjure images of danger from unrealized terrorist futures that may then be preempted as though they were real and imminent. The increasingly coercive and repressive uses of state power justified by such images of imminent danger have been shown to produce not assurance or comfort but widespread forms of insecurity (Masco 2014; Glück and Low 2017). These include both the subjective uncertainty, fear, and anxiety experienced by populations subject to state security operations (Goldstein 2010; Jusionyte and Goldstein 2016) and the emergence of retaliatory terrorist threats potentially implicating the entire nation (Masco 2014).

---

[52] According to Landau, this convergence originated in the post-Cold War decline in military spending and the related shift in government procurement towards commercial off-the-shelf technologies. For historical analysis of the shift in the government's approach to technology research and design, see Stowsky (2014). Whatever its origin, the fact and extent of the convergence between civilian and defense infrastructure was illustrated for technologists by the Snowden papers. Based on his experience as a technical consultant to the reporters entrusted with the papers, Bruce Schneier told the IETF in late 2013 that a primary lesson to be drawn from them was that NSA surveillance was parasitic of the surveillance capabilities of the internet economy: "The NSA didn't wake up and say, 'Let's just spy on everybody.' They looked up and said, 'Wow, corporations are spying on everybody. Let's get ourselves a copy'" (Schneier 2013c).

The "mandated insecurity" identified by computer scientists differs from other insecurities perpetuated by the counter-terror state in noteworthy ways. For example, the vulnerabilities to be introduced into consumer technology by exceptional access would not be limited, like many other harms flowing from the war on terror, to already marginalized populations (see Goldstein 2010). Indeed, while they would affect all technology users, given the premium commanded by strong security on the consumer technology market, they would disproportionately effect the global elite. Moreover, by introducing widespread vulnerabilities into America's critical national infrastructure, mandated insecurity would directly threaten the technical systems, which by some accounts (Masco 2014; Lakoff 2007) have supplanted the population as national security's primary object of care. Mandated insecurity thus foregrounds the self-devouring nature of the counter-terror state, its tendency in pursuit of preempting hypothetical future threats to spread certain harms through the very networked systems in whose name it purports to act[53] (see Suchman et al. 2019).

For purposes of understanding privacy's changing cultural coordinates, technologists' identification, via the concept of mandated insecurity, of commercial encryption's centrality to American national security suggests that technological stewardship has enrolled privacy in a relationship to 'security' this is far more complex and ambiguous than indicated in preceding sections. As we have seen, thanks to the historical debts of privacy engineering to security engineering, privacy has become linked in new conceptual, institutional, and material ways to the

---

[53] This tendency is all the more remarkable given that the federal government had by this point formally recognized the national security import of network computer systems for almost 20 years (see Department of Homeland security 2003). For an analysis of the historical failure to fully "securitize" so-called cyberspace, see Hansen and Nissenbaum (2009).

practices, tools, and concerns of computer security. Some technologists view this subordination as effectively handicapping privacy, limiting its capacity to protect the public from the threats to individual autonomy and flourishing posed by contemporary data practices.

At the same time, as reflected in the references to privacy's lack of charisma at Computers, Freedom, and Privacy, privacy has long figured in law and culture as the weak man of American civil liberties. This has been especially true of privacy as it has been continually 'rebalanced' against national security in the post-9/11 era (see de Goede 2014; Masco 2017). As demonstrated in the field of "securitization" studies, though, successfully defining a security matter as an existential threat to state or society effectively places it outside of ordinary politics, relocating it to the extrapolitical realm of emergency laws and exceptions (see Pedersen and Holbraad 2013). This is true of any securitized object, even ones like privacy, which the state has historically compromised during periods of perceived national emergence. With policymakers' uptake of the language and logic of mandated insecurity, computer scientists arguably securitized privacy, if indirectly, via its grounding in commercial encryption. In so doing, technologists effectively reconfigured the spatial and social scale and affective intensity of privacy harms. With regards to scale, for example, the legitimation of commercial encryption as a necessary object of national security transposed the risks posed by engineered vulnerabilities in commercial encryption from device to nation, from technology's individual users and collective corporate developers to the public and the nation's networked critical infrastructures, thus endowing them with potentially planetary scope. Affectively, identifying the technical security of America's commercial internet technologies as vital to national security elevated its perceived stakes from relatively localized harms to imminent, existential ones. Privacy, materially and conceptually

grounded like computer security in encryption, thus gains new affective intensity and resonance in the context of the Crypto Wars and with it new potential to resist the demands of the counter-terror state using the state's own techniques of affective calibration.

The Compromise Layer

In closing this chapter, I'd like to shift focus away from the substantive concerns that animated the Crypto Wars and attend instead to the Crypto Wars' character as a public event. In particular, I'm interested in the reflexive insistence on the part of both law enforcement and technologists that the Crypto Wars be conducted as an instance of public debate. At first blush, this framing is entirely unremarkable. The Crypto Wars may have involved clearly identified and antagonistic partisan 'sides.' And technologists may have mobilized historical memory of the original Crypto Wars to commemorate their collective political origins. But otherwise, the Crypto Wars were in no conventional sense a "war." What else then might they have been other than a public debate? As Director Comey observed in introductory remarks to his Congressional testimony, "democracies resolve these kinds of really hard questions through robust debate." Such belief in the power of a particular form of communication to sustain social order by resolving social conflict remains a powerful American inheritance of the Enlightenment (Habermas 1989; Povinelli 2001).

And yet, there are reasons to question the descriptive and normative characterization of the Crypto Wars in terms of public debate. As we have seen, technologists and law enforcement identified a wide array of occurrences as forming part of the Crypto Wars. These included, respectively, the filing of a court order and the implementation and marketing of a new security

architecture. Even if expressive content can be attributed to such disparate occurrences, they hardly register as offers or exchanges of truth claims among free and equal citizens, as called for under relevant theories of public reason and deliberative democracy (see Povinelli 2001, 326-27; Habermas 1996; Chambers 2003) Given technologists' exasperation with the resurgence of a policy issue viewed to have been long-resolved, why commit again to a public exchange of ideas? What work was achieved through the mutual determination to engage in the Crypto Wars as a public exchange of reasoned arguments? Let's look at each side in turn.

For computer scientists and security engineers, the insistence on public debate reflected in part encryption's historical designation as a state secret. Even after the development of public key cryptography and the blossoming of cryptographic expertise in civil society, national security and law enforcement officials continued to invoke unspecified classified information to bolster the demand for exceptional access. It remained an unanswered question in the early phases of the original Crypto Wars whether a debate over national cryptography policy could even be carried out on an unclassified basis (National Research Council 1996). The subsequent determination by the National Academy of Sciences (NAS) that classified material was not "essential to the big picture" of how US policy should evolve (Ibid.) subjected encryption policy for the first time to public, expert scrutiny. This newly endowed technologists with a semiotic power of dispute, which they were loath to cede.

Technologists' insistence on public debate further reflected the fact that the post-Snowden calls for technological stewardship of privacy never solely contemplated the development and implementation of privacy-enhancing technologies. At virtually every privacy conference I attended, for example, at least one session addressed technologists' civic duty to improve the

law's apprehension of technology. Generally conducted by technologists either with government regulatory experience or trained jointly in the law, these sessions offered basic primers on lawmaking and judicial processes. At PETS in 2015, for example, the lawyer and computer scientist Jonathan Mayer offered a tutorial on surveillance law. Mayer justified the tutorial as necessary to equip technologists with knowledge of the legal system: "We want judges and policymakers to take us seriously," Mayer said. We want surveillance programs to be built on good computer science and thus need to introduce some "actual computer science facts" into the post-Snowden public debate. "To be taken seriously, we have to take law and policy seriously." Begrudging as it may have been, technologists' engagement with the Crypto Wars as public debate extended the strategy expressed here. To begin to improve technology-policymaking not by 'educating' lawmakers but first by learning to 'speak the language' of policy and thereby gain strategic access to policymakers' point of view.[54]

What then are we to make of law enforcement? Why was it import for them to at least be seen as engaging in public debate, rather than shaping encryption policy through the courts or private lobbying of Congress and the President? What I would suggest here is that law enforcement invoked the necessity of public debate not to engage in an open and honest exchange of facts and viewpoints so much as to impose a semiotic structure of accountability on technologists. As Lezaun (2010) argues, ritualized institutional forms of social interaction may aid in the resolution of conflicting social claims by imposing a structure of accountability and moral order on claimants. Lezaun derives this conclusion from analysis of legal trials. He shows

---

[54] In this regard, see Biagioli's (1990, 205) argument that in the realm of the sciences, members of an emerging discipline may "use the language of the adversary" as a toehold for the 'invasion' of an entrenched disciplinary domain.

that the rules governing the admissibility of evidence and courtroom procedure require litigants to adhere to clearly articulated norms of mutual intelligibility. Courtroom examinations, for example, are generally structured around conversational "adjacency pairs." The questions posed by a judge or lawyer thus constrain subsequent speech, creating a structure of expectations that conditions what can be recognized as an appropriate, relevant response. As Lezaun shows using the example of a famous Scopes trial exchange between Clarence Darrow and William Jennings Bryan, a court may take deviations from such expectation, or failures to correct identified deviations, as admission of the paucity of a litigant's claim, elicited by the trial's revelatory machinery.

Outside of certain exceptional institutions settings, public debate is clearly not ritualized to the degree of courtroom procedure. It doesn't impose formally sanctionable speech rules and thus cannot produce a common linguistic grid with the disciplinary power exercised by that at play in trials. Still, public debate imposes its own structure of accountability and moral order. One might assume this structure to be modeled in some way on the Enlightenment ideal of public reason. As such, public debate would involve an exchange of truth claims and would obligate participants to evaluate and respond to opposing arguments on the bases of truth, sincerity, and legitimacy (Povinelli 2011). As actually mobilized in the Crypto Wars, however, public debate did not appear directed towards establishing truth. It was not invoked as a revelatory machine, like a trial, but rather as a commensurating one. In referencing public debate, lawmakers and law enforcement figured it as a mechanism by which collective group "interests" are voiced, compared, and ranked. Upon this basis, the Crypto Wars, as public debate, would ultimately produce a "tradeoff." It would submit the competing social "interests" understood to be

implicated by encryption backdoors—privacy, public safety, national security, innovation, etc.—to the commensurating, socially-necessary imperatives of "balancing."[55] In prefatory remarks before Congress, Director Comey made explicit the role of public debate in establishing the epistemic basis for interests to be balanced: "Thank you for hosting this conversation, and for helping us all talk about an issue that I believe is the hardest issue I've confronted in government, which is how to balance the privacy we so treasure, that comes to us through the technology that we love, and also achieve public safety, which we all very much treasure." Later, he linked the necessity of balancing privacy and national security to the very possibility of public order, noting "this is a nation founded on balance."

Figured in this way, we can think of the structure of accountability imposed by invoking public debate as requiring as a threshold matter that participants acknowledge the pluralist

---

[55] While the predominance of balancing as a form of reasoning within law and policymaking is now largely taken for granted, it is of relatively recent origin. As a form of social reasoning, balancing generally involves identifying competing social interests implicated by a social conflict and comparing and valuing them. Aleinikoff (1987) shows that as a form of judicial reasoning, balancing first appeared in Supreme Court case law in the late 1930s. Balancing subsequently rose to prominence in the 1950s and 1960s and remains a key mode of judicial decision-making (Aleinikoff 1987; Tribe 1985). Tribe sketches a similar trajectory for cost-benefit analysis—described by Aleinikoff as a stylized, putatively more "objective" variant of balancing— as a mode of decision-making in American policymaking. According to Tribe (1972), cost-benefit analysis entered the field of policy analysis in the 1950s with its 'takeover' by economists (see also Porter 1995). It became widely applied across the US government in the 1960s following apparent early success in resolving non-economic problems in terms of economic tradeoffs. For a sustained critique of the forms of ideological distortion imposed on legal and policy reasoning by balancing and cost-benefit analysis, see Tribe (1972, 1973, and 1985).

reality of other, competing interests.[56] Indeed, participants in the Crypto Wars performatively rehearsed the norm that predicates any successful resolution of public debate on mutually legitimating the interests perceived to be in conflict. The executive summary to an influential National Academy of Sciences report observed to this effect, "All of the various stakes involved are legitimate… Informed public discussion of the issues must begin by acknowledging the legitimacy both of information gathering for law enforcement and national security purposes and of information security for law-abiding citizens and businesses" (National Research Council 1996).

What I would suggest further here is that for law enforcement, foregrounding the necessity of public debate served simultaneously to impose a structure of mutual accountability on computer scientists and engineers and to accuse them of being already in violation of it. Consider in this regard the particular terms in which law enforcement justified public debate's necessity. As indicated in the introductory quote to this section, participants in the Crypto Wars pointed to general principles of democracy to justify resolution through public debate. They often pointed more specifically, however, to the idea that it was the incommensurability of positions at the heart of the Crypto Wars, which necessitated mediation through debate.

---

[56] In this regard, the Crypto Wars demonstrate the influence of pluralist theory on contemporary American enactments of the ideal of public reason. Pluralist theory holds that government policy is the product not of reasoned discourse but of agonistic competition among organized "interest groups." According to McFarland (2007), pluralist theory originated as a mid-20th century effort by political scientists to explain the generation and use of "power" in political processes. By the late 1960s, however, a "vulgarized" version of pluralist theory had become the predominant "public philosophy" operative in American politics (McFarland 2007, citing Lowi 1964). The idea that policy is the product of bargaining between government and interest groups had come to "permeate the values" of decision-makers at every level of American government (McFarland 2007, 54; see also Tribe 1973).

Participants characterized this incommensurability in different ways. Some described it as an inherent tension between privacy and national security or public safety, others as an inherent tension between technology and the rule of law. Law enforcement officials, however, frequently attributed it not to an inherent tension, but rather to incommensurable tensions generated by technologists' irresponsible privacy "absolutism."

During a 2015 interview with the technology reporter Kara Swisher, President Obama gave voice to this accusation. Obama had run for office in part on a platform of defending strong encryption. Asked to comment on the ongoing tension between the FBI and Apple, Obama reiterated this position: "I'm a strong believer in strong encryption….and lean more in favor of strong encryption than do some in law enforcement." Still, Obama acknowledged, from the presidency's "bird's-eye view…smack dab in the middle of these tensions that exist," he sympathized with law enforcement. "[T]here are times where folks who see this through a civil liberties or privacy lens reject that there's any tradeoffs involved. And, in fact, there are. And you've got to own the fact that it may be that we want to value privacy and civil liberties far more than we do the safety issues. But we can't pretend that there are no tradeoffs whatsoever."

To be fair, technologists' articulation of the security risks posed by exceptional access coincided in the Crypto Wars with the forwarding by civil liberties lawyers and activists of privacy-based arguments against weakening encryption. But as implied in Obama's quote, the accusation of absolutism as it applied to technologists communicated that their arguments against the feasibility of reliably secure exceptional access were in fact an ideological smokescreen. They obscured either absolute fealty to privacy's primacy as an American value or a dogmatic vision of technologically-enforced absolute privacy. From technologists' perspective, of course,

this charge was without merit. They opposed exceptional access not because they refused to submit privacy to balancing. They were not opposed to balancing per se. They had concluded, however, that the balance proffered by law enforcement—a technological mechanism that enabled search warrants to be executed without thereby exposing the public to new insecurities—was not supported by their collective experience.

It is still worth pausing here to consider the possibility that if law enforcement could not accept the infeasibility argument at face value it was because it was not culturally legible as credible. Consider, for example, the exceptional powers claimed by Silicon Valley technology companies in the period between the first and second Crypto War. By 2014, Silicon Valley entrepreneurs, venture capitalists, and engineers had spent years making fantastic promises to solve formerly intractable social problems using internet technology. They had consequently become associated in American public imaginaries with a quasi-mystical form of technological solutionism (see Morozov 2013; Lanier 2013). Law enforcement's skeptical appraisal of the technological impossibility argument, with its implicit endorsement of the belief that, in fact, there is no limit to innovation's power to resolve the unresolvable, partook of these imaginaries. In his 2015 House testimony, for example, Director Comey was asked whether, as technologists' claimed, requiring companies to build encryption backdoors into their products constituted a form of mandated insecurity. Comey responded that while the issue had to be resolved by talking to "true experts," he himself remained optimistic: "I actually don't think we've given this the shot that it deserves. I don't think the most creative and innovative people in our country have had an incentive to try and solve this problem." In 2015 testimony before the Senate, Comey observed similarly, "I think Silicon Valley is full of folks who when they stood in their garage

years ago and were told, 'Your dreams are too hard to achieve.' Thank goodness they didn't listen."

I think the more likely scenario, though, is that, by accusing technologists of privacy absolutism, law enforcement effectively marked them as in violation of the norms and moral order of public debate (see Biagioli 1990, 185) Such marking placed the onus on technologists to correct their infraction, to preserve the legitimacy of their political voice by yielding to public debate's disciplinary force. To do so, technologists would have to demonstrate their willingness to constructively engage in the pluralistic system of interest competition by recasting the conclusion of technological impossibility in American policymaking's preferred language of compromise and tradeoffs.

As the political scientists Landman and Lauth (2019, 3) note, the prominence of concepts of tradeoffs and balancing in politics reflects the American belief that it is generally impossible to simultaneously accomplish all legitimate political goals. Tradeoffs and balancing respond to this perception by creating a political presumption of gradual mediation and partial satisfaction. Such half-measures are probably achievable, they assert, on "most points of contention in political disputes, but "not cases that make normative claims to truth" (Ibid. at 239).

As illustrated by Matt Blaze, technologists were aware of the bind thus placed on them by the related ideologies of public debate and balancing. In 2015, at PETS, Blaze observed that technologists had largely engaged in the Crypto Wars at "the technical layer." By this he meant that they engaged in the Crypto Wars, under the influence of their training, by trying to "find the right answer." If technologists thus judged their participation in the Crypto Wars according to the engineering standard of correctness, the government, meanwhile was holding the debate at the

"policy layer." And this, Blaze acknowledged, is the "compromise layer." Technologists were disadvantaged in the Crypto Wars to the extent that "[w]e don't really know how to engage in these kinds of discussions where compromise is involved."

For technologists to engage in the Crypto Wars at the compromise layer, for them to repair the rupture in the structure of expectations imposed by public debate, would effectively require them to legitimate what they otherwise categorically rejected as magical thinking: the possibility of a reliable secure lawful access mechanism. In so doing, they would undermine the basis of their own authority in American policymaking by ceding their expertise over technical means and risk. In this regard, law enforcement's invocation of public debate, like the accusation of absolutism, functioned as what Mario Biagioli (1990, 186) calls a "strategy of non-dialogue." Biagioli, like Lezaun, rejects the idea inherited from analytic philosophy that incommensurability is a linguistic phenomena that results when theories or viewpoints are expressed via heterogeneous grammars across which translation is not possible (see also Povinelli 2001). Based on study of claims of incommensurability in the history of science, Biagioli anticipates Lezaun in arguing that incommensurability is a pragmatic achievement. For cultural objects like privacy and security to be perceived as pointing to incommensurable futures, some authority must establish this relationship between them.

Biagioli takes as his primary example the perceived incommensurability between the theories of buoyancy respectively forwarded in the early 1600s by Galileo and the philosophical school known as the Tuscan Aristoteleans. This debate was marked, like the Crypto Wars, by profound skepticism on both sides. Galileo, for example, repeatedly accused the Aristoteleans

both of refusing to constructively engage with his theories and of being unable to comprehend them because of their mathematical illiteracy.

Biagioli proposes to explain these features of the debate on buoyancy in terms of the trespassing of professional and disciplinary hierarchies. These included especially a long-standing disciplinary hierarchy subordinating math to philosophy. He argues that in dismissing Galileo's work as incomprehensible, the Aristoteleans invoked and sought to reinforce this social hierarchy. It was as much about rejecting the applicability of mathematics to what the philosophers viewed as their rightful domain—explanations of physical phenomena—as it was any particular failing of Galileo's mathematical treatment of buoyancy. To an Aristotelean, given their significantly more secure position in the rewards systems of the day, to accept the abstractions of math as a means of explaining physical phenomena would be to "learn the language of a previously subordinate 'other' now turned alien invader" (Biagioli 1990, 204).

Following Biagioli, we can think of law enforcement's invocation of public debate and accusation of privacy absolutism as strategies of non-dialogue. Here, non-dialogue serves not to delegitimate math vis-a-vis philosophy as an explanation of physical phenomenon. Instead, it expresses law enforcement's rejection of technological feasibility and risk as appropriate rubrics through which to determine the objects of public safety and national security. To identify a danger to the state or society and have it be socially legitimated is to claim and express a certain exclusive political power (see Buzan et al. 1998). While any socially authoritative actor can in theory define a risk to the nation, given the increasing centrality of national security to state-making and to justifications of state power, governments generally claim the exclusive authority to do so. From this perspective, the invocation of public debate and absolutism reflected not a

rejection of the conclusion of technical impossibility in-and-of-itself but of the general legitimacy of technologists-in-security-policymaking.

If law enforcement could not accept as legitimate the application of engineering expertise overs means to questions of public policy it was largely because in thus applying their expertise technologists had marked a potential limit to the pathologies of post-9/11 counter-terror logic. As explored by Masco (2014), the form of militarization sustained through preemptive counter-terror logic is so exceptionally productive because its imaginative conjuring of danger from alternative futures is delinked from evidence, facts, or the observable. Against this backdrop, Congress' ongoing refusal to impose an exceptional access mandate on technology companies is noteworthy. It reflects a rejection, if limited, of post-9/11 America's totalizing counter-terror logic, one precipitated in part by technologists' insistence that the material and operational realities of encryption not be black-boxed (Riles 2005; Latour 1987, 2-3) in any related policymaking. This insistence was why technologists found it so frustrating and irresponsible that the government declined in the second Crypto War to propose a concrete exceptional access mechanism. As Diffie, Landau, Schneier, and other prominent computer scientists wrote in July 2015 in "Keys Under Doormats," a highly influential report on the likely social and technical effects of exceptional access, any access policy would have to be accompanied by concrete, analyzable technical requirements before it could be "taken seriously." Because real risks always "lurk in the technical details," only with a "concrete technical proposal" could the government credibly claim to have accounted for the likely impact of exceptional access.

On what bases are we to conclude that it was technologists' insistence on grappling with the real that provoked law enforcement's strategies of non-dialogue? Take, for example,

President Obama's expression of sympathy. Later in the same interview, Obama elaborated on its nature: "I think the only concern is... our law enforcement is expected to stop every plot. Every attack. Any bomb on a plane. The first time that attack takes place, where it turns out we had a lead and couldn't follow up on it, the public's going to demand answers."

Here was an absolutism, but it wasn't that of privacy over all other American values. Instead, it was the absolute drive for anticipatory control over the future, as manifested in the panic-driven trauma of the officers tasked with carrying it out. Technologists' perceived ability to resist the counter-terror state's affective logic was further illustrated in the apoplectic response of lawmakers to technologists' resistance to the logic of exigent circumstances. Again and again during the Congressional hearings of 2015 and 2016, lawmakers challenged technologists to mark the limit of technological impossibility. During the 2016 House Judiciary hearing, for example, Congressman King proposed the following hypothetical: I think it's an accepted public fact, that the Islamic State terror group (ISIS) "is seeking a nuclear device." If "American consciousness" accepted that ISIS was on the cusp of nuclear capability, "do you think that would change this debate we're having here today?" Or consider the questions submitted to Apple's legal counsel by Congressman Bob Goodlatte following the hearing: "Is there any instance of national security, terrorist attack, or major gang-related violence affecting our communities where you would offer your assistance in unlocking a single iPhone? In other words, if you were presented with a "ticking time-bomb" scenario, would you offer to produce the necessary software work-around giving law enforcement access?" In these exchanges, lawmakers accustomed to the bending of democratic norms and the rule of law before images of danger conjured from hypothetical futures confront a style of reasoning that would not bend.

CHAPTER 2: CONTAINING CONTEXT

Consider two visual artifacts of the internet age, published 20 years apart and marking an arc in changing American perceptions of the internet. The first is a single-panel cartoon published by *The New Yorker* in July 1993, a year before Netscape launched the Navigator browser, inaugurating the web's popularization. The cartoon features two dogs, one perched atop a desk chair pawing the keyboard of a personal computer, the other seated before the desk on the floor. "On the Internet," the computer-savvy canine confides, "no one knows you're a dog."

Among *The New Yorker*'s most frequently reproduced images, the dog cartoon captures the playful if cautious spirit that accompanied public access to the internet. As the cartoon suggests, central to early public fascination was the perception that the internet augured not just a new medium of mass communication, but also a newly amplified form of privacy. This perception reflected the fact that nothing in the protocols that facilitate computer networking technically requires users to provide an identity. It was thus anticipated the internet would anonymize its users. On the internet, it seemed, one would never know what sort of persons—what kind of beings, even—might lie at the far end of one's connection.

Alongside such hints of ontological uncertainty, early public fascination contemplated the new possibilities for experimentation with personal identity that the internet seemed to promise. In contemporary analysis of the dog cartoon, sociologist Sherry Turkle captured the imaginary anchored by this perception: "You can be whoever you want to be. You can completely redefine yourself if you want. You don't have to worry about the slots other people put you in as much.

They don't look at your body and make assumptions. They don't hear your accent and make assumptions. All they see are your words" (Turkle 1995). Through the internet's mediation, the socially salient markers of personal identity would be stripped away. Freed of the cultural biases carried by material bodies (boyd 2014; see Barlow 1996), individuals would be known entirely through intentionally communicated speech and the beliefs and ideals thus expressed.[57]

Now consider a second artifact, the ten-part comic, *The Private Eye*, published digitally in 2013. A noir mystery modeled on cinema's technologically-saturated near futures, *The Private Eye* tells the story of Patrick Immelmann. A private investigator, Immelmann is professionally stymied by the universal use, circa 2076, of secret identities. Confounding Immelmann, however, are not the internet-enabled digital aliases, which Americans once imagined would enable a dog to pass as a human. Rather, they are physically-assumed identities involving elaborate, technologically-enhanced masks and costuming. Indeed, in *The Private Eye*'s future, the internet no longer exists at all. It was banned following a cataclysmic cyber-security breach, a great "bursting of the cloud," which publicly revealed all the private information accumulated on the internet for a biblical forty days, and with predictably disastrous consequences.

The Private Eye in certain respects presents the dystopian inversion of the emancipatory vision expressed by *The New Yorker*'s dogs. In depicting the messy aftermath of the internet's revelatory divulgence it literalizes then-nascent critiques of the modern internet. By 2013, far from enhancing privacy, the internet had rendered it effectively unattainable. Rather than

---

[57] In danah boyd's formulation, it was hoped that the internet would produce a "color-blind and disembodied social world," thus serving as a "great equalizer" of racial and class-based inequalities (2014 23).

ushering in a new era of mastery over self presentation, the internet seemed poised to eliminate the very possibility of selective revelation.

Still, *The Private Eye'*s critical thrust is better understood as deriving not from an inversion of the internet's early promise, but rather from carrying this promise to logical extremes. Recall in this regard that in developing the web's underlying technologies, Tim Berners-Lee was not motivated by concern for personal identity but rather for access to information. Berners-Lee sought to make the internet easily accessible to non-experts via the web specifically to make the world's knowledge universally available. Read through this history, *The New Yorker* dog cartoon illustrates not the internet's promise of anonymity, but rather the involutions it promised in the relationships between self-revelation, personal identity, and access to knowledge. Under this reading, a dog's ability to surf the web speaks to the emergence of a world in which access to knowledge is no longer limited only to those with elite credentials and immediate, physical access. Indeed, a world in which, one might obtain total knowledge without divulging any personal information at all. *The Private Eye*, from this perspective, speaks to the uncertainties introjected into social life when re-organized around a principle of unmediated access to total knowledge. As represented in *The Private Eye*, a world of total information access is necessarily one in which intimate information regarding the self can no longer be contained by the body, the home, or technology, with destabilizing consequences for the forms of everyday social reasoning reliant on so doing.

In this chapter, I explore the questions of privacy, identity, context, and meaning raised in these artifacts of the internet age through a study of Containers, a privacy-enhancing Firefox feature under development in 2016 and 2017. At the level of technical infrastructure, the

Containers project involved re-engineering Firefox's data isolation mechanism to introduce new limits on website access to user data. At the level of the browser interface, meanwhile, Containers required Mozilla's designers to surface Firefox's new data isolation capabilities in an aesthetic and experiential form that Firefox users could recognize as providing them with distinct, segregated browsing "contexts."[58] In this way, Mozilla's engineers sought to return to users the ability to inhabit unique, partial, and contextually-specific identities— a capacity they understood to have been effectively eliminated by the internet. As I will show, in their desire to redress something of the epistemic and social uncertainty unleashed by the modern internet's data practices, Mozilla's engineers participated, like the creators of *The Private Eye*, in a broad-scale reconceptualization of privacy. Under the increasingly prominent vision of privacy they helped enact, privacy operates primarily as a social and material technology for establishing and maintaining contextual boundaries.

In subsequent sections, I first describe the social and technological aspirations animating the team responsible for Containers. On the one hand, these aspirations drew on a minor tradition of American privacy concerned with individual control over representations of the self. On the other hand, Mozilla's engineers developed Containers in response to the perception that the commercial internet, as it had actually developed circa 2016, is an unparalleled source of "context collapse." As an attempt to limit the forms of epistemic uncertainty that flow from context collapse, I argue, Containers was designed to "socialize" the web. With Containers, Mozilla's engineers sought to refigure browser "use" along the lines of human social interaction,

---

[58] In Chapter 3, I consider the distinct material and semiotic aspects of browser features in greater detail, and analyze the significance of the feature form for privacy's future.

modeling it on an understanding of social life as necessarily involving different kinds of relationships maintained with different kinds of persons through practices of selective self-revelation.

In the second section, I perform my own contextualizing move. I locate Containers' concern with the boundedness of context within a broad-scale shift towards contextually-aware privacy approaches in both American technology and policymaking. Comparing the visions of context pursued at Mozilla and in policymaking and tech communities, I show that attending to the contextual factors understood to inform privacy preferences formally serves the interest of individual autonomy. As actually taken up by American policymakers and businesses, however, it effectively undermines the feasibility of any categorical approach to protecting privacy. In practice, technology companies use research into the contextual factors that influence privacy preferences to map the limits of the socially tolerable and thus overcome consumer resistance to corporate surveillance and behavioral targeting.

In the third section, I turn to the problems of usability that haunted Containers' drawn out development. I describe Mozilla's years-long effort to make Containers more effective and appealing as a privacy feature through user experience research and design. Mobilizing findings from the socio-linguistic study of context, I show that context should not be thought of as a pre-existing social fact, as engineers and computer scientists often do, but rather as a contingent social achievement established and transformed in interaction. Using the related semiotic literature on "contextualization cues," I argue that the usability challenges that Mozilla's professionals attributed to Containers expressed a mismatch between the feature's static

representation of browsing contexts and users' lived familiarity with context's dynamic, interactional construal.

In the final section, I shift focus away from Containers and examine the post-Snowden concern among technologists with the law's differential treatment of metadata and communications content. Unpacking the popular intuition that 'mere metadata' poses a less significant threat to privacy than does the content of conversations, I describe a genre of tutelary discourse employed by technologists to educate the public about metadata. Through this analysis, I identify metadata's ability to yield sensitive personal information as being specifically "indexical" in nature. The law's failure to account for metadata's privacy implications, I argue, is the product of powerful linguistic and semiotic ideologies, which privilege referential speech over the indexical modes of signification that actually make metadata so "personal" and "sensitive."

Privacy's Shifting Locus of Containment

In late 2016, when I first learned of it, Containers already existed as an experimental browser feature available in Firefox Nightly. Nightly is an early-release version of Firefox, the first of three "channels" or repositories through which all code moves in Firefox's release management process.[59] Mozilla's engineers use this multi-stage process to test and debug the quality, stability,

[59] At the time of my fieldwork, the release management process involved four channels, with code moving from one channel to the next every six weeks. In 2017, Firefox eliminated one of the release channels. As of 2021, code now moves through the remaining three channels every four weeks. For analysis of the privacy and security implications of the tech industry's post-2000 shift towards rapid, "agile" software development methodologies, see Gürses and van Hoboken (2017).

and security of new code before "shipping" it to Firefox's general-use population.[60] Containers' inclusion in Firefox Nightly signified that the project was, as an engineering challenge, largely complete. The engineers responsible for it had already identified a privacy threat, attributed it to identifiable aspects of browser operation, designed a technical solution, and implemented it in Firefox's "backend" technology platform.

Despite the completion of Containers' engineering work, in early meetings, Mozilla's privacy professionals identified it as a compelling object of study. A lawyer I'll call Peter told me that if, as I had described, I wished to gain insight into both the technological and institutional mediation of privacy, I should consider focusing in part on Containers. As Head of Trust and Privacy, Peter was one of the primary officers responsible for implementing Mozilla's privacy commitments in its products. According to him, Containers was of interest because of the significant organizational challenges it posed for Mozilla. Like many privacy-related engineering projects, he told me, Containers raised questions of how to balance "what we want to accomplish" as a browser developer with Mozilla's commitment to preserving privacy.[61] With Containers, this tension took the forms of ongoing problems of "usability." The Containers experience, Peter said, was simply "too foreign" to easily incorporate into the typical browsing habits of Firefox users. Given the tricky tradeoffs between "user experience, engineering, and

---

[60] Unlike the users of Firefox's official release version, users of Nightly tend to be what Mozilla's engineers call "power" users. Such users are generally well-versed in the technical operation of web technologies like the browser and make heavy use of browser features and capabilities. Often technology professionals themselves, they use Nightly to remain appraised of upcoming browser features or to contribute to Firefox's ongoing open source development.

[61] See Chapter 3 for analysis of the constraints that Mozilla's market-based strategy for realizing its techno-moral vision imposes on the form that privacy can take within Mozilla and its technologies.

product" involving in making Containers more usable, the feature remained subject to ongoing study and design. At this point, Mozilla had yet to determine whether Containers would ever be "surfaced" to Firefox's full user population as a privacy feature, and if so, in what form.

In later sections, I return to the challenges of meaning, comprehension, and usability that Mozilla's designers and engineers understood Containers to pose for users. In this section, I focus first on the nature of the privacy problem addressed by the feature. What does the nature of this privacy problem, I ask, tell us about changing American valorizations of privacy and the social interests it serves?

In a paper delivered to the 3rd International Conference on Information Systems Security and Privacy in December 2016 (Vyas, Marchesini, and Kerschbaumer 2017), three of Mozilla's senior privacy and security engineers described the engineering effort behind Containers. In so doing, they characterized Containers as a technical response to the problem of pervasive third party web tracking. Citing a series of influential academic studies, they observed that circa 2016 virtually every website included embedded content or resources used to track and correlate user behavior "whenever we surf the web" (Ibid. at 472). The business practices of advertisers and ad-tech platforms might bear ultimate responsibility for web tracking's 'omnipresence.'[62] Functionally, however, the behavior of local data storage mechanisms like browser cookies within browser developers' technically control made such tracking possible.

Prior proposals by internet engineers to address pervasive web tracking largely pursued a strategy of blocking. Such proposals called for browsers to unilaterally block third parties from

---

[62] See Chapter 4 for analysis of the increasing prominence of the web's surveillance-based business model in the logic by which technologists explain privacy's ongoing erosion.

setting cookies, thereby preventing data brokers and advertisers from collecting personal data.[63] By contrast, with Containers, Mozilla's engineers proposed not to block web tracking, but separate it.[64] By introducing new layers of separation into the browser's data isolation mechanism, they would limit the kinds of user data that any given website could access. This would not restrict the overall amount of personal data collected during browsing. It would, however, complicate the ability of individual advertisers to aggregate the entirety of a user's browsing data, ideally narrowing the scope and inferential reach of web-based tracking.[65] In this

---

[63] Such proposals were considered both by individual browser developers and by relevant standards bodies. The proposals were met with howls of protest by advertisers and digital advertising trade groups, who argued that the web economy would collapse if advertisers lost the ability to conduct targeted advertising. See Kristol (2001) for discussion of the fraught history of browser cookies. In Chapter 3, I briefly describe DoNotTrack, a failed effort to convince advertisers to respect user requests not to be tracked online.

[64] In addition to blocking and separation, prominent strategies employed by privacy-preserving technologies include *minimizing* the amount of personal data generated by a system and *obfuscating* personal data by mixing it within a system with plausible but fake data. On the latter, see Brunton and Nissenbaum (2016).

[65] Historically, Firefox, like other browsers, stored all cookies and other website-accessible user data together in a single, all-encompassing "Cookie Jar." Firefox isolated such data according to the so-called "Same Origin Policy," a security mechanism fundamental to all modern browsers. Same Origin Policy defines the "origin" of a web resource as the combination of its application layer protocol (e.g., TCP, http, https), its domain host name, and the port number used to reach it. Under Same Origin Policy, browsers condition a website's ability to access stored data on a comparison of the website's origin with that of the data. If different in any way, the browser would automatically prevent the site from accessing the requested data. At an engineering level, Mozilla's engineers enabled Containers, by building into Firefox a generalized ability to "extend" the Same Origin Policy by adding new attribute labels to the definition of an origin. Once enabled, Firefox can automatically enforce the new attributes alongside the original elements of the Same Origin Policy whenever a website attempts to store or retrieve data. Containers leverages this ability by adding to Cookie Jar data a new user-controlled attribute known as userContextId. When Containers is enabled, Firefox labels stored data with both its origin and the user-defined userContextId. Henceforth, Firefox prevents websites from accessing any data labelled with a userContextID that differs from that assigned to the website, even if the website and data otherwise share the same origin.

respect, Mozilla's aspirations for Containers were deeply pragmatic. On the one hand, Containers' development reflected an acknowledgement by Mozilla's engineers of the effective impossibility, for most web users, of fulfilling the popular injunction to "take control" of one's privacy online.[66] On the other hand, their proposed intervention was conditioned by advertisers' repeated defeat of more radical solutions. As an intentionally "lightweight" alternative to existing privacy tools, Containers' promise was simply to enable users to access the internet's riches without fully acquiescing to the web's surveillant tendencies.

As one of the paper's authors explained to me, its descriptions of Containers' technical underpinnings were ultimately directed to other browser vendors and web developers. Sharing innovative engineering solutions is one way Mozilla engineers seek to improve privacy on the web in general, and not just for Firefox users. However, he told me, while technically accurate, the paper's description of Containers in terms of third party tracking misrepresents something of the project's animating aspirations.

Consider, in this regard, the name "Containers." From the perspective of browser engineering, the name suggests a desire to contain either browsing data or the browser storage mechanisms surreptitiously used to collect it. For Mozilla's engineers, however, technically containing tracking vectors in fact served as a means to the more abstract goal of containing 'context.' Specifically, by providing Firefox users with newly segregated browsing contexts

---

[66] See Acar et al. (2014, 2) an influential, large-scale study of web-based user tracking mechanisms. In it, the authors conclude based on the wide-spread prevalence of web tracking and the practical difficulties in blocking it that "even sophisticated users may not be able to take control of their privacy online without loss in web content and functionality." Marwick and boyd (2014) separately conclude that on social media the "networked" nature of privacy—the fact that in the persistent, widely-shared ecosystem that constitutes social media decisions about what to post often implicate the privacy of others—effectively removes privacy from individual control.

within which to conduct different kinds of browsing activity, Containers aspired to provide users with limited control over the selective revelation of personal information. Containers thus purports to restore users' ability to inhabit contextually-specific identities online, in ways Mozilla's engineers understood to be demanded by privacy but increasingly precluded by online tracking.

In a June 2016 blog post announcing Containers' launch in Firefox Nightly, Tanvi Yvas, the project's lead engineer, explained this aspiration. With Containers, she wrote, Mozilla sought to "empower" Firefox users to segregate their online identities "in the same way [they] can segregate [their] real life identities." To illustrate this ability and its value, Vyas offered examples drawn from her own family, social, and consumer life:

> We all portray different characteristics of ourselves in different situations. The way I speak with my son is much different than the way I communicate with my coworkers. The things I tell my friends are different than what I tell my parents. I'm much more guarded when withdrawing money from the bank than I am when shopping at the grocery store. I have the ability to use multiple identities in multiple contexts. But when I use the web, I can't do that very well….The Containers feature was designed for a single user who has the need to portray themselves to the web in different ways depending on the context in which they are operating.

Note how embedded in the examples offered here is a theory of personhood, which contemplates the individual expression of multiple, contextually-specific identities. As Vyas portrays it, identity does not refer to an inherent or unchanging individual essence.[67] Rather, identity here is a function of interactional perception. It is the set of socially-legible characteristics attributed to individuals during interaction. From this perspective, underpinning Container's prospects as a

---

[67] In this respect, it diverges from the prominent contemporary vision, grounded in understandings of genetics and biology, of identity as an ineluctable essence (see Comaroff and Comaroff 2003, 455).

privacy technology is the fact that the characteristics attributed during interaction vary, and they does so in ways partially determined by practices of selective self-revelation. Per Vyas, in social life unmediated by the internet, individuals seek to achieve different social effects by calibrating what they say and how they say it to different intended audiences.

As Vyas depicts it here, the inhabitance of multiple identities is a prosaic, everyday occurrence, not a notable achievement or skill.[68] Vyas' examples indeed leave open the possibility that shifts between projected identities can occur beneath the threshold of individual intention or awareness. This is not to suggest, however, that Mozilla's engineers believe the ability to be without practical significance. Rather, as Vyas portrays it, the play of identities is part-and-parcel of social life. It is what enables individuals to be, for example, in one moment a responsible parent and in the next a productive colleague. As illustrated by the guarded demeanor Vyas describes adopting in banks as compared to grocery stores, shifting between partial identities facilitates the negotiation of interpersonal trust and the managing of perceived threats to personal safety and security, broadly conceived.

In designing Containers as a tool for controlling depictions of the self online, Mozilla's engineers grounded its legibility as a privacy-preserving technology in a minor tradition of American privacy. According to the legal historian James Q. Whitman (2004), when compared to

---

[68] In early internal presentations on Containers, Mozilla's engineers used cultural familiarity with acting as a form of identity-like role inhabitance to introduce contextual identities. To this end, for example, an August 2015 presentation offered in-character photos of the actor Christian Bale. Under the heading, "We have multiple identities in life," a series of photos presented Bale as Bruce Wayne in the *Dark Knight* films, as a 1970s con man in *American Hustle*, and as a deranged Wall Street banker in *American Psycho.* Accompanying text identified the photos as illustrating that we inhabit distinct identities when "At Work," "With Friends," and "Anonymous Online," respectively. In using such examples, Mozilla's engineers risked suggesting that assuming new identities is only possible under exceptional professional circumstances.

those of other nations, U.S. laws display distinctive intuitions and sensibilities regarding what properly falls within the scope of privacy's protection. What characterizes the unique "culture of privacy" expressed in American privacy laws, Whitman argues, is an intense concern with questions of liberty (see also Nelson 2002). American privacy laws primarily seek to protect individuals against unwanted intrusions by the state, or to a lesser extent, by other private citizens and entities. Continental European privacy laws, by contrast, emphasize not liberty, but the individual right to personal dignity and respect as against the potential humiliations of public exposure.[69]

As Whitman acknowledges, the distinct cultures of privacy that he identifies in American and European law are not mutually exclusive. For example, while it has enjoyed greater prominence in legal and philosophical theorizing than in legislation or common law, a tradition of American scholarship has long sought legal recognition of privacy's role in sustaining human dignity (see, e.g., Bloustein 1964). Indeed, the first systematic American analysis of the need for a legal right to privacy, Warren and Brandeis' 1890 law review article, "The Right To Privacy," can be understood as an effort to "graft" a European-style dignity standard onto American law (Whitman 2004, 1204-08). In formulating their famous call for a "right to be let alone," for instance, Warren and Brandeis contemplated a form of "retreat from the world" defined not with

---

[69] For Whitman (2004, 164), such privacy laws comprise one element of a broader class of continental legal protections for interpersonal respect. Whitman traces European law's concern with protecting individuals from public shame and humiliation to the rigid class and status hierarchies that historically characterized continental society. It is the product, he argues, of a process of historical "leveling-up" whereby privileges and norms of respect long-reserved for nobility and the wealthy were progressively extended to other social classes. The German practice of generalizing honors throughout society, Whitman shows, was not a negative reaction to fascism, as is often claimed, but was rather directly prefigured in the Nazi's regime practice of honoring low-status persons, so long as they were racially German.

respect to the government, but to "[r]ecent inventions and business methods"—instant photography; newspaper gossip—which enable the press to overstep "the obvious bounds of propriety and decency" (1890, 196). Only by recognizing a right to control how one's "thoughts, sentiments, and emotions" are shared with others, Warren and Brandeis argued, could the law secure for individuals exclusive possession of the self in its psychological and spiritual integrity (Ibid. at 206; see Glancy 1979).

In promising to grant users control over the data shared with websites during browsing, Mozilla's engineers squarely aligned Containers with this tradition of American privacy concerned with the self-determination of self-presentation. Rather than citing online police or state surveillance, for example, Mozilla's staff generally defined the need for Containers in terms of the practical risk that circulating representations of the self pose to individual social status and reputation.[70] Bram Pitoyo, a visual designer assigned to Containers, thus told me, in "real life you have different faces that you present to different people." As such, "your browser should allow you to do the same, because this is how you do it in real life." In identifying the inhabitance of contextually-specific identities as a fundamental social ability, which the internet has eliminated, Containers effectively critiques the modern, commercial internet in dignitary terms. In Whitman's language, Containers figures the internet as 'leveling users down.' Through

---

[70] In describing Containers as a technology designed to "empower" users, Mozilla's engineers arguably also aligned it with ideals of autonomy derived from America's predominant, liberty-oriented privacy tradition. The autonomy that Containers seeks to enable, however, is of a limited kind. It is not the autonomy characteristic of sovereign political actors when duly protected against state intrusion. Nor is it the autonomy sometimes theorized as necessary for democratic citizens to cultivate and exercise their civic capacities (Cohen 2012). Instead, it is the autonomy expressed when, through selective self-revelation, individuals define themselves for others in life's heterogeneous circumstances (see Nissenbaum 2010, 22; Goffman 1959).

pervasive tracking, it extends to the entire web-browsing population diminished norms of dignity (i.e., enhanced forms of surveillance) theoretically reserved for threats to America's social fabric.[71]

Recall, however, that in announcing Containers, the project's lead engineer described the human need to inhabit multiple identities in terms of the demands of social context. In describing the technical innovations enabling Containers, Mozilla's engineers predicated achievement of the feature's privacy potential specifically on 'containing context.' How should we understand the nature of the relationship presumed by such statements between context and privacy? What should we make of their implication that the internet has compromised the boundaries between social contexts, which structure 'real life?'

In predicating the form of privacy offered with Containers on reinforcing contextual boundedness, Mozilla's engineers drew together the minor tradition of dignitary privacy described above with a different American privacy tradition. This tradition, which Deborah Nelson (2002) traces to the Cold War, is organized around a generalized, mobile anxiety about zones of sovereignty. According to Nelson, while privacy has long appeared in American history as an object of social anxiety, it was only after the 1950's that it achieved visibility in its contemporary form as nostalgic "lost thing" (Nelson 2002, 37). This "sudden" emergence, she shows, was precipitated by a prominent strain of Cold War thought, which identified the sanctity of the private sphere as the key social feature distinguishing democracies from totalitarian

---

[71] Browne (2015) and Bridges (2017), for example, document the ways the U.S. government has historically subjected African-Americans and women, respectively, to deeply invasive forms of surveillance.

states.[72] Such thinking marked privacy, as symbolically grounded in the family home, as the font of democratic liberty. Thus invested with new political and psychological significance, to tolerate any penetration of the private sphere was to effectively invite totalitarian oppression.

According to Nelson, in heightening the significance of any perceived permeability between private and public life, the Cold War prefigured the periodic "crises" of privacy, which new surveillance and information technologies have repeatedly precipitated in the late 20th century. In Nelson's telling, the Cold War also provided the narrative grammar through which Americans understood these crises.[73] Of central importance to such "scripting" was containment ideology and the generalized metaphor of containment that it introduced into public imaginaries.

As elaborated in U.S. foreign policy circa 1948-1960, containment ideology initially referred to the goal of containing the spread of the Soviet army and ideology. Over time, Americans came to view such containment as requiring containment of the internal divisions of American society.[74] Containment ideology transformed into a distinctive cultural formation through which Americans perceived the events of late 20th century history. Americans came to

---

[72] Hannah Arendt, for example, famously argued in the *Origins of Totalitarianism* (1968) that while all tyrannies destroy public life, totalitarianism is distinguished as a mode of domination by its efforts to also destroy the private sphere.

[73] From this perspective, the distinct culture of American privacy described by Whitman can be understood as a byproduct of Cold War intellectual and political dynamics.

[74] In the Cold War's paranoid ferment, the specter of the "enemy within"—the possibility that subversive forces already lurked in the sanctuary of private spaces—perpetually threatened such containment. According to Nelson, there was an infinitely expandable logic to such fear, illustrated by the FBI's contemporary expansion of the category of so-called "subversives" to include not just avowed Communists, but also civil rights activists, feminists, and eventually any of their respective "fellow-travelers" (2002, 12-13).

imagine all kinds of entities as bounded 'spaces' haunted by a frightening permeability that must be sealed at all costs.[75] Privacy itself became subject to its own containment anxieties.

As Nelson shows, traces of this mobile, "topological" anxiety can be found in Constitutional jurisprudence of the Cold War era. Given the Cold War's deep psychological investment in the privacy of the home, for example, for Nelson it is no surprise that the first Supreme Court case to recognize a Constitutional right to privacy was decided in its name. The claims before the Court in 1965's Griswold v. Connecticut concerned not the home's legal status, but rather the state's ability to regulate personal uses of contraception. Nonetheless, in ruling that the government cannot "intrude" on a married's couple decision to buy and use contraceptives, the Court did so in terms of the home and the "zone of privacy" it defines. Writing for the majority, Justice Douglas asked, "Would we allow the police to search the sacred precincts of marital bedrooms for telltale signs of the use of contraception?" Answering in the negative, the Court mobilized the image of the home as the seat of democratic domesticity to locate its paradigmatic relationship—marriage—within the scope of Constitution's implicit privacy protections. In subsequent decades, U.S. courts applied such topological imagery analogically to other privacy cases. These cases expanded the range of entities and activities recognized as falling within privacy's protections[76] to the individual mind and nation in addition to the body

---

[75] For analysis of the semiotic logic that facilitates such analogical reasoning, see Gal (2005, 2002).

[76] In Katz v. United States, for example, also decided in 1965, the Court ruled that 4th Amendment protections against unreasonable search and seizure require police to obtain a warrant before wiretapping a public pay phone. This was so even absent a physical search, the Court wrote, because the Constitution ultimately "protects people, not places." Katz thus detached the locus of American privacy protections from the home and grounding them instead in the individual—analogically figured to be like the home a bounded space in need of continual shoring up against state intrusion (see Nelson 2002, 21, 112).

and home.

Mozilla's engineers never explicitly connected Containers to such Cold War-inflected concerns with individual sovereignty. By framing Containers' contribution to privacy in terms of context, however, they partook of the topological privacy logic that Nelson describes. Promising to improve privacy by containing context posits context as being like the home, nation, and body, a boundable entity definitive of privacy's scope. To appreciate this claim and how Mozilla's engineers understand the internet to threaten the boundedness of context, it is helpful to return to Vyas and the distinction she draws between context in "real life" and as it operates in internet-mediated life.

To draw out this distinction, imagine conducting some of the tasks that Vyas highlights, but in a world as you might remember it from before the internet's popularization. Imagine you've finished work for the day, but before returning home to your family you must deposit a check and buy groceries for dinner. You first drive to your local bank, a physical building staffed by clearly-designated employees and populated perhaps by a smattering of customers. As you enter, you join a line to wait for a teller. You might exchange pleasantries with a neighbor, but otherwise speak only to the teller. You hand over your check and deposit slip, wait for the receipt, and then make your way back to the car for your next stop—quickly grabbing ingredients for dinner.

In each of these two contexts, through your behavior, appearance, and speech—through your perceptible, socially-salient characteristics—you 'reveal' information about yourself. Some but not all of this you intentionally communicate. To deposit your check, for example, you must communicate to the teller, either through words or by providing appropriate institutions markers,

that you are a customer in good standing. Even as you do so, however, the style and condition of your shoes may lead the grocery bagger to surmise information about you that you do not explicitly communicate—that you are an office worker, for example. Similarly, a fellow shopper, overhearing a brief phone call with your daughter as you peruse the grocery's vegetable aisle, might conclude from your accent, that you were born in the American South.

Whether intentionally communicated or inferred, you can rest assured that any knowledge of your personal life gained in these two social contexts will remain more-or-less 'in place.' Such knowledge is likely to be locally drawn, based primarily on the limited information conveyable during fleeting interactions. Moreover, it is unlikely to circulate beyond the individuals and institutions that you directly encounter. The regulatory mechanisms that inhibit such circulation are numerous. They include the physical organization of space within the bank branch and grocery stores, the distance separating the buildings from one another and from the other places in which you conduct your affairs, the fallibilities of human memory, institutional restraints, including federal regulations regarding personal financial information, and especially social norms, including norms of privacy and civil inattention (Goffman 1963), which compel us in appropriate circumstances to avert our eyes, pretend we have not overheard, or hold our tongue. In this example, the broad categories of social activity around which you structure your everyday life comprise social contexts as Mozilla's engineers generally used the term. It is relative to the material, institutional, and psychological restraints described that Mozilla's engineers understood context to 'contain' or hold personal information it in its 'proper,' contextual place, thereby preserving privacy and enabling individuals to inhabit multiple identities.

Now consider the same scenario transposed to a world like our own in which virtually every aspect of social life is circuited through the internet. To carry out your errands you turn to your computer not your car. You open it and launch a web browser, navigating first to your bank and then to an online grocer. Using the relevant website features, you carry out your errands without traversing physical space, entering 'the public,' or directly interacting with other people. Certain experiential markers, however, still appear to demarcate your errands as occurring in something like the distinct, self-contained contexts familiar from 'real life.' Using the browser, for example, you call forth the respective webpages of your bank and online grocer using distinct URLs. Via design, branding, and login forms, the webpages further hold themselves out to you as corresponding and manifesting distinct institutional entities. As a Mozilla engineer noted on this point in the Containers project wiki, "In meatspace, there are more clues about how your information will be shared… and information leakage is easier to track and contain. By contrast, intuiting the total amount of personal information shared online that can be inferred via a vast array of technologies (cookies, web bugs, search engines, user-supplied application data, log data) is difficult." Other than the cookie popups that you likely click through unthinkingly, the web browsing experience itself provides little indication that the relationship between context and informational circulation with which you are familiar from the physical world operates any differently online.

And yet, despite the evidence of your screen, on the internet context is said to "collapse" (boyd 2014; Marwick and boyd 2014; Marwick and boyd 2011). As browser engineers describe it, this is so in two respects. First, businesses exploit the design and operation of modern web technologies, including browsers, to track individuals across the web. They thereby aggregate

data corresponding to the entirety of an individual's web browsing behavior, not just that portion of it that, in life lived off the internet, one might selectively reveal to specific individuals under contextually appropriate circumstances.[77]

Second, by applying powerful analytic techniques to the data they aggregate, businesses infer personal knowledge that individuals may never have 'revealed' to anyone at all. With respect to such powerfully invasive inferences, internet engineers and computer scientists now sometimes say that "all data is sensitive." That is, when aggregated and fed through the right algorithm, even apparently non-sensitive, non-identifying, or otherwise innocuous information may unlock deeply intimate, embarrassing, or otherwise personal information. The net result of these dynamics is that information circulates promiscuously online, heedless of contextual boundaries. Businesses are thereby able to effectively 'look past' individual's attempts to project partial identities by selectively revealing different sets of information in apparently discrete contexts.

While Containers' engineers approached context collapse in terms of its problematic effects on personal identity, scholars and technical experts have in fact analyzed it from a number of conceptual angles. In the general legal and philosophical literature, for example, when information "properly" belonging to one sphere migrates to another, this may in itself constitute a form of injustice (Nissenbaum 2010). Such literature approaches information-out-of-place, with its inherent capacity to take on new meaning in unanticipated contexts (Amoore 2013;

---

[77] Per Nissenbaum (2010, 117), on the internet one can no longer count on the disinterested observation that typifies most fleeting social interactions. One also cannot presume that personal information will remain distributed across disparate observers, each of whom "can only take so much information in."

Caduff 2012), as a potential source of financial, reputational, psychological, and even physical harm. For Viktor Mayer-Schönberger (20011), by contrast, context collapse is of interest primarily for its temporal and spatial effects. Focusing on the tendency of online information to persist across time, he characterizes the internet as a kind of technologically-enabled perfect memory auguring the end of both forgetting and selective disclosure. The practical consequence, for Mayer-Schönberger, is to draw future relations into the present, forcing individuals to constantly anticipate hypothetical future reactions to actual present behavior.[78] In this way, he writes, context collapse threatens to eliminate the ability to decide and the capacity for abstract thought (Ibid. at 13, 22, 113-18; see also Gerstein 1978).

For Marwick and boyd (2010; 2014), meanwhile, context collapse manifests in the challenge teens face in 'reading' and thus potentially managing or controlling the "social situations" in which they find themselves entangled on social media. As boyd (2014) elaborates, context collapse precludes teens from easily anticipating the likely audiences of their social media posts. Thus faced with the "impossible" task of accounting for the full range of plausible interpretations of their behavior, teens attempt to "stabilize" or contain context by employ ad hoc strategies like visualizing their intended audience.

What binds these analyses and connects them to the critique expressed with Containers is an account of the internet's profound destabilization of the reliability of everyday social reasoning. By undermining context's ability to keep personal data in its proper place, as defined by culturally-salient norm and practices, the internet effectively invests personal data with

---

[78] In this respect, Mayer-Schönberger echoes Enlightenment era concerns with the danger of living for the estimation of others, but extends the scope of the potentially scrutinizing other indefinitely into the future.

unbound epistemic potential. The issue is not only as Mayer-Schönberger and Marwick and boyd each suggest, that present actions become laden with the specter of uncontrollable future interpretations. Rather, given the unpredictable inferences that emerge when apparently innocuous, unrelated pieces of information are aggregated and analyzed, all personal information becomes laden with latent connections to all other information. The online behavior that generates such data becomes, in turn, similarly imbued with revelatory potential. The practical consequence is a perceived diminution of privacy's usefulness in managing and predicting the epistemic effects likely to follow from any given action online. In this respect, the anxiety expressed in the Containers project regarding context's containment foregrounds the epistemic functions performed by privacy. It shows much of everyday social reasoning to be undergirded by culturally-informed presuppositions regarding the likely circulation of different kinds of information in different kinds of spaces, settings, and relationships. In this respect, privacy is not just a normative guide to appropriate personal behavior, but also a tool to think with (see Gal 2005; 2002) and a guide to whether and how one's information is likely to become available to and be interpreted by others (Palen and Dourish 2003; see Grudin 2001).

To the forms of context collapse already identified, we might add the ontological confusion that haunts the modern internet. As I describe in greater detail below, early in my fieldwork a senior Mozilla privacy engineer told me that her motivations for developing Containers included a personal dislike of "the feeling" of precisely targeted advertising. Without naming it, she gestured here to "creepiness." Much commented upon (see, e.g., Tene and Polonetsky 2014), creepiness is arguably the defining feeling of the internet, that characteristic reaction that most prominently distinguishes the internet from other evocative technologies, like

the computer and nuclear bomb. In both popular and expert narratives, Americans widely use the term to index the uncanny, visceral sensations elicited by networked technologies, including surveillance-driven behavioral ads, that appear to personally "stalk" users across the web.

Such uses of creepiness express the perception that the internet is populated by uncanny, other-than-humans beings displaying human-like intelligence and intentionality. This perception, of course, is informed by the contemporary proliferation of automated bots intentionally designed to mimic human communication. Historically, it was prefigured by the anonymizing potential attributed to the web, as illustrated in *The New Yorker's* 1993 dog cartoon. It has been amplified, however, by the way that the internet frustrates the everyday ability to definitively attribute circulation personal information to the specific communications of particular actors.

Read against the internet's frustration of everyday social reasoning, Containers can thus be understood as not just 'mitigating third party tracking' or 'containing context,' but as an attempt to socialize the web. Containers' defense of the ability to inhabit multiple, partial identities, from this perspective, does not simply critique the dignitary insults of online tracking. Rather it reflects an effort to refigure browsing or internet "use" on relational terms, not as the asocial navigation of a seamless information space but as a series of encounters with discrete and identifiable entities, modeled on characteristic aspects of human social life. Key to such socialization is the understanding that Mozilla's engineers expressed of social life as constituted by heterogeneous relationships defined and distinguished by the varying degrees of interpersonal intimacy along which they are conducted. With Containers, Mozilla's engineers provided Firefox users with a technical means to hold themselves out to websites and the entities they represent in partial, selectively revealed form. This constituted both a direct technical intervention in the

circulation of personal data online and an interpellation, an invitation to respond in kind, to acquiesce to the natural limits on knowledge that define and sustain human social interactions.

Respect for Context

In an early interview, a Firefox engineer I'll call Dana described the impetus behind Containers to me in personal and affective terms. A senior privacy and security engineer, Dana served both as Containers' technical lead and as a vocal internal advocate on its behalf. As someone who, in her own words, takes personal privacy and security very seriously, Dana joined Mozilla years prior drawn by its official commitment to privacy. She had been gratified to find, she told me, that unlike other Silicon Valley tech companies with which she was familiar, at Mozilla product managers do not view privacy and security engineers as a nuisance to be avoided.

According to Dana, Containers developed in part out her own experience trying to frustrate web surveillance. Over lunch at Mozilla's Mountain View headquarters in October 2016, Dana described herself as someone who "doesn't like the feeling" of precisely targeted advertising. This sensibility, she suggested, lay at the root of her ongoing efforts to prevent advertisers from using her browsing data to infer knowledge of the kinds of major life

developments, which they are known to covet.[79] To limit the kinds and quantities of personal information that different websites and third parties could collect about her, she approached her use of internet technologies with a self-discipline informed by expertise. Given her familiarity with its deeply invasive data practices, for example, Dana refused to install the Facebook app on her phone. While she did use Facebook on her personal computer, to prevent Facebook from tracking her across the web, she signed out of her account after each session.

More notably with regards to Containers, Dana's personal privacy practices included an "ad hoc" strategy for segregating browsing data.[80] This strategy involved using multiple browsers, with specific browsers dedicated to different kinds of browsing activity. Like many privacy-minded tech professionals, she told me, she reserved one browser exclusively for online

---

[79] The contemporary practice of individually targeting consumers during key changes in the course of their lives was first brought to popular attention by Charles Duhigg in a 2012 New York Times article, "How Companies Know your Secrets." As Duhigg reported, retailers view life events, such as pregnancy, illness, marriage, and home buying as opportunities for the most effective forms of targeted behavioral intervention. During such events, consumers face new commercial needs. Their otherwise deeply-engrained shopping habits are thus believed to become unusually and temporarily flexible. Among other revelations, Duhigg showed that major retailers seek to intervene in shopping habits in such moments based on inferred knowledge, which individuals consumers have not yet shared with even family or close friends. As a former statistician for Target told Duhigg, "If we send someone a catalog and say, 'Congratulations on your first child!' and they've never told us they're pregnant, that's going to make some people uncomfortable. We are very conservative about compliance with all privacy laws. But even if you're following the law, you can do things where people get queasy." Target thus learned to mask its individualized consumer knowledge by hiding targeted ads among unrelated promotions. Duhigg's reporting helps clarify that the negative visceral reactions elicited by targeted advertising can be partially attributed to the experience of knowledge-out-of-place, knowledge that has escaped the container of the self and is known without ever having been expressly communicated.

[80] As a Mozilla user experience designer later explained to me, such strategies are considered "ad hoc" in that they involve the use of multiple internet technologies in combination, and often in ways that their developers did not intend or anticipate.

banking. She maintained a second browser solely for online shopping. By thus visiting different kinds of websites in different, dedicated browsers, Dana attempted to protect her sensitive personal information and prevent websites from aggregating the whole of her browsing data. Social scientific accounts of the context-bound nature of identity may have ultimately informed Mozilla's design for Containers.[81] In Dana's telling, however, the feature initially emerged from the desire to translate the burdensome segregation practices employed by professionals like herself into a simple, purpose-built tool for non-experts.

Despite such grounding in the habits and frustrations of specific engineers, Containers merits inquiry in part because its development participated in a significant shift in American approaches to privacy. In broad outline, under this shift, American scholars, technologists, and policymakers increasingly defined privacy in terms of individual, context-sensitive expectations, preferences, and choices. Indeed, early in my fieldwork, as I met with Mozilla's privacy managers to identify projects for study, I was initially drawn to Containers precisely because context and its relation to privacy loomed so prominently in preliminary research. Context-based approaches to defending privacy were, for example, a frequent feature of the privacy conferences I attended while developing my research. In these settings, technologists often linked the treatment of privacy as a problem of context to the influence of a single scholar. This individual, the analytic philosopher Helen Nissenbaum, was herself a familiar presence at privacy events. She spoke frequently on the need for values-based approaches to technology design and was well-known for collaborating with computer scientists, activists, and artists interested in privacy.

---

[81] The project wiki for Containers, for example, cites Erving Goffman's work on impression management and Carl Jung's persona theory as relevant prior art.

Given its prominence in relevant professional discourses, Nissenbaum's scholarship provides entree into the context-based approaches to privacy, which have increasingly circulated among American technologists and policymakers in recent decades.

In 1998, Nissenbaum published, first as a series of scholarly articles and later as a book, a theory of privacy as "contextual integrity," which would become influential among computer scientists and internet engineers. Diverging from predominant legal and philosophical theories, Nissenbaum (2010) characterized digital privacy not as a right to control or restrict access to personal information, but as a right to information's "appropriate" flow. The question of whether any given information exchange should be considered appropriate, she argued, is one that can only be answered relative to "context." More precisely, it is a question of contextually-specific informational norms. Such norms, according to Nissenbaum, prescribe the particular the terms under which people ought (or ought not) to transfer information in a given context.[82] They thereby structure and differentiate contexts from one another. In Nissenbaum's terminology, "contextual integrity," as a benchmark for privacy, is preserved when a technologically-induced alteration of informational flows "respects" the applicable contextually-specific informational norms. Conversely, when technology causes information to flow in ways that contravene such norms, contextual integrity is lost and what we think of as privacy is violated.

---

[82] The contextual factor that Nissenbaum identifies as informing perceptions of appropriateness include the type of information being shared, the social role of the sending, receiver, and information subject, and the mechanism of transmission. These "principles of transmission," include concepts familiar from American privacy law, including confidentiality, reciprocity, dessert, notice, consent. In proposing contextual integrity, Nissenbaum refigures informational control and access as simply two among many possible contextually-applicable transmission principles.

Drawing on a range social theories concerned with the ordering of social life,[83] Nissenbaum defined contexts, like "health care" and "school," as abstract representations of the structured social settings of everyday life. While the contextual integrity framework largely takes contexts' existence for granted,[84] it nonetheless posits them as defining and sustaining social life's essential activities, relationships, and interests. Efforts to preserve contextual integrity by preventing technologies from radically altering informational flows are thus morally justified in that they maintain the very "fabric of social life" (Ibid. at 3).

As suggested by this concern with justification, Nissenbaum developed contextual integrity as a normative framework for morally evaluating new technologies.[85] In explaining the need for such a tool, Nissenbaum pointed, like Dana, to affect. Nissenbaum, however, was concerned not with her own feelings but with the collective anxiety and alarm consistently elicited by new surveillance technologies. She argued that as expressions of long-evolved expectations regarding informational flows, contextually-specific informational norms condition

---

[83] Nissenbaum points in particular to Pierre Bourdieu's field theory and Michael Walzer's spheres of justice as inspiring her treatment of context.

[84] In addition to contextually-specific informational norms, the contextual integrity frameworks calls for contexts to be formally characterized in terms of key elements including paradigmatic roles (i.e., the typical capacities in which a person may act in context—as a client, teacher, congregant, etc.), activities (e.g., browsing goods in a shopping context, singing hymns in a religious context), orienting values or purposes (e.g., transmitting knowledge in the schooling context), and other behavior-guiding norms.

[85] The contextual integrity framework flags novel practices that depart from entrenched informational norms as *prima facie* evidence that contextual integrity has been violated. Analysis of moral legitimacy, however, does not stop with the identification of such violation. Nissenbaum instead calls for such practices to then be compared to entrenched norms on the basis of how effective each is in promoting relevant contextual values. Novel informational practices that violate informational norms may still be justified on moral grounds so long as they support the attainment of both context-specific values, ends, or purposes and general societal values, such as equality, justice, and fairness.

the embodied experience of information practices. The alarm elicited by novel information practices thus constitutes a reliable social signal that privacy expectations have been violated.[86] Given the regular contemporary recurrence of privacy panics, Nissenbaum concluded, prevailing approaches to evaluating new technologies must be too "heavy-handed" (Ibid. at 103-04). They fail to account for the full range of contextual factors that inform appropriateness. They thus lack the expressive grammar needed to guide fine-tuned responses to the alterations in informational flows wrought by ongoing technological development.

Following publication of Nissenbaum's framework, computer scientists working across a range of sub-disciplines began to integrate it into privacy engineering practice (see Benthall et al. 2017).[87] Computer scientists applied contextual integrity, for example, to projects like the design of technical agents capable of reasoning about context. Contextual integrity resonated with computer scientists' long-held desire to design systems capable of adapting to changing social

---

[86] Nissenbaum argues, for example, that approaches to privacy rooted in the public/private distinction frequently result in legitimate complaints being dismissed as irrelevant to privacy. Given its sensitivity to the complex contingencies that make a privacy claim morally and socially legitimate, contextual integrity is better equipped, in her view, to identify privacy threats obscured in such accounts, including ones dismissed out of hand for manifesting in putatively "public" spaces.

[87] Barth et al. (2006), which Nissenbaum co-authored, is generally cited as inaugurating computer science's engagement with the contextual integrity framework. A more recent survey (Benthall et al. 2017), which Nissenbaum also co-authored, found that computer scientists have not taken up contextual integrity in its full theoretical scope. In particular, computer scientists generally elide or bracket Nissenbum's accounts of privacy's ethical dimensions and contextual integrity's normative dimensions.

and environmental conditions.[88] It offered them a formal structure for expressing informational norms as technical constraints on information flows and for modeling privacy expectations in well-defined contextual settings (Ibid.).

Over the same historical period, a concern with the relationship of context to privacy emerged in parallel among American policymakers.[89] Circa 2010, policymakers began to call upon government and business entities to adopt context-aware approaches to data privacy. In 2012, for example, the Obama White House published a report (the "White House Report") listing "Respect for Context" among seven principles intended as a new baseline of consumer privacy protection.[90] Despite advancing communications technologies and business models, the White House Report asserted, "[c]onsumers have a right to expect that companies will collect,

---

[88] Since the 1990s, for example, the field of ubiquitous computing has aspired to replace the personal computing paradigm with information processing capabilities embedded throughout the physical environment (Barton and Kindberg 2001; see Weiser 1993). Its proponents have thus long grappled with how to design computing applications to automatically adapt to changes in the physical and electronic environment, or context, in which individuals interact with computational services (see Dey et al. 2001). Benthall et al. (2017) cite such context-aware computing research as a key influence on the definition of "context" that many computer scientists employ today.

[89] The turn towards context in computer science and policymaking appears to have been co-constitutive. Benthall et al. (2017), for example, hypothesizes that the contextually-sensitive proposals described below contributed to computer science's uptake of contextual integrity. The White House Report, meanwhile, cites Nissenbaum for the proposition that there are important connections between privacy and context.

[90] The same year, the Federal Trade Commission ("FTC"), acting as the primary U.S. regulator of consumer privacy, urged businesses to adopt a "context of interaction" standard. This standard requires companies to provide consumers with notice-and-choice before collecting and using their personal data, unless such data practices are "consistent with the context of the transaction or the company's relationship with the consumer." According to the FTC, given the dramatic contemporary increase in forms of corporate data collection, a context-sensitive standard was necessary to provide consumers with an "easy-to-use" choice mechanism without imposing an overly onerous "choice-related burden" on them.

use, and disclose personal data in ways that are consistent with the context in which consumers provide the data."[91] Businesses should thus limit uses of personal data to those purposes "consistent with both the relationship that they have with consumers and the context in which consumers originally disclosed the data…"

Policymakers framed their embrace of contextually-sensitive privacy approaches in terms of individual autonomy. In justifying the need to update American consumer privacy, for example, the White House Report observed that since the 1970s, the U.S. federal government had regulated privacy through the sector-specific application of the Fair Information Practice Principles ("FIPPS"). While this continued to provide a "solid foundation," privacy protections now confronted a radically altered business and technology landscape. Data processing was far more pervasive, varied, and dynamic; more companies processed more data for a widening array of purposes. As a result, business practice risked diverging irretrievably from consumer privacy "expectations." The logical implication was that such divergence would eventually undermine consumer privacy choices as legitimate expressions of individual autonomy.

Concern for autonomy was further evidenced in the particular ways policymakers proposed to mobilize context. To fulfill consumers' right to "control" the collection and use of personal data, for example, the White House Report proposed that businesses should not just provide easy-to-use choice mechanisms, but should do so under contextually-relevant conditions. Notice and choice should be surfaced to consumers "at times and in ways that enable consumers to make meaningful decisions."

---

[91] See Chapter 4 for a study of the role of the business model concept in deflecting responsibility and blame for the invasive data privacy practices of Silicon Valley technology companies.

As suggested, however, by their emphasis on the importance of the 'innovative' reuse of previously collected data, policymakers also embraced context to introject new "flexibility" into consumer privacy. The White House Report was explicit on this point: Any changes to America's privacy framework would have to both address the unique privacy challenges posed by networked technologies and preserve American businesses' ability to offer innovative products and services. As figured in related policy documents, in a networked world, America's technological and thus economic potential can only be realized if its business are empowered to use personal data in new and unanticipated ways. In thus grounding innovation and growth in the creative reuse of personal data, policymakers figured any categorical limits on possible uses to be socially and politically intolerable.

The history of the FTC's embrace of context is instructive on this point. Following a year-long consultation process, in December 2010, the FTC issued a preliminary privacy framework. In it, the FTC addressed how businesses should provide privacy-related choices to consumers. For certain categories of "commonly accepted" business practices, such as product fulfillment and internal operations, the FTC proposed that business not be required to provide consumers with any choice. In this way, the FTC purported to "simplify" choices for both consumers and industry. In related public comments, however, businesses warned that approaching data consent as a matter of categorical standards would stifle innovation, freezing business practice in place (Federal Trade Commission 2012, 36). Media and technology companies leveled a similar charge against the FTC's proposal to limit corporate data collection to uses accomplishing "a specific business purpose." To be clear, this language already represented a retrenchment from FIPPs' "purpose specification" requirement, which limits

collection to business purposes specifically articulated and disclosed to consumers. Nevertheless, in related comments, Facebook argued that if companies were precluded from collecting data for purposes other than existing business needs Netflix's video recommendation feature would never have been developed (see Ibid. at 26).[92]

As the FTC acknowledged in its 2012 final report, such critiques led it to abandon its initial proposal in favor of the more "flexible" context of interaction standard. This history suggests that despite the ethical aspirations attached to it by technologists, in its broader circulations, context has contributed to an ongoing relativization of privacy. The more privacy is figured in terms of context, the more privacy's value is understood to vary in contingent, localized ways, and the less thinkable becomes any substantive or categorical legal protection. As illustrated by the FTC's context of interaction standard, context mediates consumer-corporate relations in part by mobilizing fantasies of technologies never invented and innovations never pursued. Through such hypotheticals, context authorizes corporate uses of personal data that might otherwise be prohibited for contravening personal autonomy.

In the Containers project, context figured as a set of boundaries demarcating social life into domains conducive to practices of selective revelation and thus to the inhabitance of partial and shifting identities. Context performs a similar function for Nissenbaum, carving life into into a social topography in which the heterogeneous relationships and values that imbue society with moral heft can flourish. In seeking to preserve contextual integrity, however, Nissenbaum effectively operationalizes context, mobilizing it for insight into whether new information

---

[92] Netflix originally collected video preference information from subscribers in order to send them specifically requested videos. Only later did Netflix leverage this data to begin generating personalized video recommendations.

technologies are likely to disrupt entrenched norms and trigger social outrage. Nissenbaum purports to balance such analysis with consideration of the relative contribution of contextual informational norms to morally-legitimating social values. There is still a sense, however, in which the contextual integrity framework uses context to probe and map the limits of the socially tolerable.

Similarly, as it has been taken up in policymaking, and indeed by many technologists not involved with Containers, context is best understood a map of the perceptions and sensibilities that condition privacy preferences and inform privacy choices. In July 2015, for example, I traveled to Ottowa, Canada, host city that year for the Symposium on Usable Privacy and Security (SOUPS). Held annually since 2005, SOUPS attracts a mix of academic and corporate researchers interested in the contributions of user interface research and design to digital privacy and security. Having been directed to the conference by a technologist with a DC-based tech policy non-profit, I arrived at SOUPS with little knowledge of the vibrant research community that has formed around the usability of privacy technologies. On morning of the conference's first day, I found myself in the 2nd Annual Privacy Personas and Segmentation workshop. As the workshop prospectus explains, given increasing awareness of "the importance of context in privacy concerns, and the complex relationship between concerns and behaviors," further research is needed on the "complex interactions of factors that influence privacy attitudes and preferences."

In the first session I attended, a team from Carnegie Mellon University, home's of the nation's first master's degree in privacy engineering and a leading center of privacy usability research, presented findings from a recent study. The study explored the factors that contribute to

internet users' willingness to share personal data for use in targeted advertising. By presenting research subjects with a set of different browsing scenarios, the study examined whether such willingness changes based on (i) who is collecting data (e.g., was it being trust by Facebook or a relatively trusted actor like CNN?), (ii) the stated purpose of collection (e.g., would the data be used to improve a visited website? would it be shared with third parties?), or (iii) the type of data collected (e.g., gender, zip code, online behavior, IP address). For our purposes, the study's findings are less important than the observation that its mode of questioning was representative of the research presented throughout the workshop. Also representative was its animating presumption: if only the right combination of contextual factors could be identified, consumer resistance to sharing data could be overcome.

In the workshop's closing keynote address, Norman Sadeh, a Carnegie Mellon professor, acknowledged this ambiguity. The formal motivation for research into privacy preferences, he observed, is to support notice-and-choice. Understanding what about privacy people "really care about" theoretically enables businesses, for example, to highlight the most salient portions of their privacy policies. This in turn facilitates "more meaningful decisions." Understanding the nuances of privacy preferences might further help businesses expose "meaningful" privacy settings to users. Through such calibration, for example, businesses can offer location privacy settings at the scale or "granularity" that users actually care about. Still, Sadeh acknowledged, among the challenges facing research into privacy's contextual factors is the concern among prominent technologists that it "will be used to box people in, to take away the ability to have privacy, to follow people for the rest of their lives."

Lost in the Browser

In early October 2017, I met with Jonathan Kingston remotely via Skype from his home in the United Kingdom. A self-identified front end developer and Javascript programmer, Jonathan had not done much actual front end work since joining Mozilla's security team two years prior. Mozilla hired Jonathan to work on the security of Firefox's mobile operating system. When Mozilla cancelled that project shortly after his hiring, Jonathan bounced between different teams and projects. This was typical, according to Jonathan, of the "craziness" of working for an organization like Mozilla, which seemed to him to be in a near-constant state of internal reorganization.

I was meeting with Jonathan in his capacity as a developer of Containers, one of a series of interviews I conducted with the project team in 2016 and 2017. In them, I sought to understand both the form of privacy that Mozilla aspired to offer with the feature and the nature of challenges that privacy poses for technology developers like Mozilla when translated into a user-facing product.

Jonathan couldn't remember the precise conditions under which he had been drafted to work on Containers. By that point, however, the "principle…for how Containers actually works, how it isolates…and what the work was for the Origin Attributes side of things" was "well established." Indeed, following a three year effort, the "back end" work of implementing Origin Attributes in the Firefox—"a large, fundamental platform change"—was complete. The result, in Jonathan's estimation, was a "radical" web technology unique to Firefox. "The thing that I think never was potentially taken as a big problem," he continued, "is the user interface, which arguably is just as big of a problem….We're trying to address to the more general user. And that

has been the life-cycle of me on this problem. It was quickly realized that we couldn't release it to the general purposes users, because they wouldn't be able to understand it at all."

As Jonathan indicates here, his primary task when drafted to work on Containers was to help build its user interface. To not just build the interface, however, but make it intuitive and reliably usable by non-experts, Jonathan helped manage a long, iterative process of experimentation, testing, and redesign. To be fair, he told me, competing corporate priorities meant that Containers' initial interface was designed without input from Mozilla's user experience and visual design experts. Designers had, however, eventually been consulted. They participated in a "design sprint," working over five days with the full project team to brainstorm ideas for how to present the Origin Attributes technology from the user's perspective. Subsequently, they conducted multiple "polish" passes on the interface's visual elements. The interface was now "definitely much better," according to Jonathan, but still "potentially wrong." General purpose users still struggled to "understand how to use [Containers] in a way that is secure, and two, doesn't hinder them too much."

By the time I met Jonathan, Containers' usability issues were internally perceived to be so intractable that Firefox executives had determined not to ship it in the browser. More precisely, while Containers would not be released as a "native" or default browser feature, it would be offered as an optional browser "extension." Suitably motivated and knowledgeable users would thus be able to download and install it from the collection of simple, open source programs available to customize Firefox. A new application programming interface would also be released, which would enable web developers to build their own extensions using Origin Attributes.

In our discussions, the engineers, user researchers, and designers responsible for Containers were circumspect regarding its usability. "[I]f you just think about it," a user researcher told me, "if you're the engineering person,' We did this really hard thing. We're going to save everyone.' And people are like, 'Oh, I can color code it like my folders.' That's hard." In Jonathan's words, it was frustrating, both for the team and for Mozilla in general, not to be able to "solve these [privacy] problems that we have a solution for."

Despite such frustrations, the risks of shipping Containers in its current form were clear. On the one hand, shipping a confusing feature risked alienating Firefox users and driving them to a competitor. Luke Crouch, a young engineer who worked closely with Jonathan during Containers' user testing, described this possibility in the following terms:

> The product team, the reason they really want to do [user testing] is because it is such a complicated concept, the Containers concept is. And so, you can't just pref it on to a bunch of users. Because if you do and they start making a bunch of containers, and a bunch of things break, they are not going to blame Containers. They don't even know how to disable Containers. What they are going to do is blame Firefox and go use something else…. [I]f you try to put it on an unwitting population, they will just get angry at Firefox and leave Firefox.

As a user researcher who worked on Containers helped me appreciate, the risk identified by Luke here stems from the browser's ambivalent status for most users. Browsers might play a central role in mediating everyday access to and experiences of the internet. But, they have "become so commoditized that people don't even know what the browser is. I'll ask people what their browser is and they'll tell me, 'Facebook.' Or they'll tell me, 'Windows.' That's very common, whatever your background." The average user might perceive any number of different browsing issues as 'breakage'—webpages that load slowly, only partially, or not at all; a webpage's failure to remember account login credentials. The ultimate cause might stem from network issues,

website programming error, the intervention of a third party, such as an advertiser, or browser behavior itself. The easy availability of multiple free, mostly interchangeable browsers means, however, that many users are all-too-willing to simply abandon Firefox, regardless of technical cause.

On the other hand, shipping a confusing feature simply risked confusing users. Even when a feature technically operates as intended, Luke told me, confusion might still arise if its operation diverges from user expectations. The need for features to work in accordance with user expectations posed particular challenges for privacy features. "Quite often," according to Luke, "users go read up on an article that says, 'If you set this setting in Firefox, you will be more private.' And that setting is, 'disallow all third party cookies everywhere,' [or] 'never send an origin referrer,' blah, blah, blah. So, I'll just click that button. They forget they do that and now Firefox is breaking everything on the web for them." As implied by this reference to 'just clicking the button,' Luke's experience was that non-expert users often approached privacy tools as quick fix, all-encompassing solutions, as though privacy were a binary proposition, or could simply be turned on like a light switch. Clicking such settings would indeed make users 'more private.' Users failed to appreciate, however, that they could only do so in partial ways, and only by intervening in the accumulated historical entanglements of the web's surveillance capabilities with technologies like cookies that enable much of the modern web's core functionality.

Regardless of competitive concerns, confusing users was further problematic to the extent it threatened to derogate Firefox's obligations as a self-identified "user agent." Consider in this regard the following scenario offered by Luke when asked to explain Containers' usability issues: "So with Containers," Luke responded, "it would be very easy to kind of onboard

somebody on to it." Thereafter, the user might open a container, navigate to Facebook, and log in, thereby instructing Firefox to set a cookie in the corresponding Cookie Jar. This would enable the user to remain permanently signed in *so long as* she remembered to only open Facebook in the Facebook container. If instead the user navigated to Facebook in a different container, Facebook wouldn't be able to access its login cookie. "And then they go, 'Every time I open a tab on Facebook I'm not signed in. What happened? Facebook never remembers that I'm signed in.'… I mean, technically, Containers are doing the right thing. They are protecting your default container from Facebook. But the user doesn't remember that. The way you explain it to a user is very important or they will get lost in their own browser."

In sharing this scenario, Luke was primarily concerned to convey that the "breakage" haunting Containers was not objective, technical breakage, but breakage from the perspective of user expectation. The story further illustrates, however, that in the gap between technical operation and user expectation lies the risk that Firefox might transform from a facilitator of user desire into a hindrance. The key to appreciating this risk is Luke's image of becoming lost in the browser. Beyond the broader, techno-moral aspirations for the web embodied in the Mozilla Manifesto (see Chapter 3), Firefox maintains a basic technical commitment to serving as the user's "agent" during web browsing. Most proximately, this requires Firefox to act on behalf of and represent users in the communications with website servers that facilitate web access. More abstractly, it entails helping users enrich their lives by providing open, easy access to the internet's riches. As against such ideally seamless mediation, users led astray by a confusing feature might instead experience Firefox as distracting or frustrating. In Luke's example, rather than help a user catch up with old friends on Facebook, Firefox mires the user in uncertainties

stemming from its own design and operation. Users thus lost can no longer use Firefox, as intended, to easily access web content in service of personal interests and goals, rendering Firefox an obstacle to its own basic purpose. Even worse, to 'lose' the user of a privacy feature like Containers risks fatally compromising the feature's privacy guarantees, in this case its promise to restrict the "slices" of user data made accessible to different websites.

If these were the perceived risks, in the remainder of this section I focus on the nature of the usability challenge posed by Containers. I propose to add specificity to the generalized assertion that Containers was "too difficult to explain" by introducing semiotically-grounded understandings of context and browser "use." Thus informed, I argue that Containers' usability challenges were the function of a mismatch between its representation of context as a fixed, intuitively apprehensible social fact and users' lived experience with context as a contingent, dynamic interactional achievement.

To gain purchase on this argument, consider the insights of a discipline, which shares privacy engineering's interest in context as both a tool of knowledge-making and a challenge of comprehension in its own right. The discipline is anthropology, which as Marilyn Strathern (1988) observes, historically derives its disciplinary identity from its contextualizing claims and practices. Specifically, Strathern shows, anthropology has long premised its contributions to human understanding on making sense of the apparently alien or irrational beliefs and behaviors of others by putting them in proper social context (see also Dilley (1999)). As an early, influential illustration of anthropology's vision of what constitutes 'proper context,' Strathern cites Malinowski's insistence that the ritual exchange of Trobriand wedding valuables be made intelligible relative to Trobriand inheritance and land tenure rules, and not to the wedding

practices of other peoples.[93] Anthropology's analytic commitment to contextualization roughly overlaps with popular understandings of context's clarifying powers. Both presume the meaning of social events to lie beyond events themselves in the other phenomena in which they are embedded (Goodwin and Duranti 1992). Only relative to this encompassing surround can one find the cultural resources needed to successfully interpret events from the perspective of their actual participants.

As Strathern notes, anthropology's contextualizing practices implicitly figure societies as self-contained wholes characterized by internal integrity. Over its modern history anthropology may have resoundingly rejected this image of cultures as self-contained entities. Its ideals of coherence and wholeness, however, have persisted, lingering implicitly in the discipline's contemporary analyses of "structures" and "systems." To the extent understandings of "culture" have gradually converged with those of context,[94] the intelligibility of something like Helen Nissenbaum's depiction of context as a self-evident, self-contained social 'sphere' (Marwick and boyd 2014, 14) thus bears the mark of anthropology's influence.

---

[93] Indeed, Malinowski's mythical status as founding father of modern anthropology rests precisely on the attribution to him of the theory that a people's practices must be understood in relation to its other practices. Anthropology's adoption of Malinowskian long-term, first person observation as its characteristic methodology is similarly understood to have followed from the conclusion that it was the only reliable means of recording human practice in the immediate social context that makes it intelligible (Strathern 1988, 254). Strathern, for her part, rejects the historical attribution of ethnographic fieldwork to Malinsowki, arguing that Malinowski's true innovation was literary not methodological. Specifically, Strathern credits Malinowski with first organizing an ethnographic monograph around a reconstruction of the fieldwork experience, thus using the figure of the fieldworker "entering a culture" to convey for Euro-American readers the distinct experience of foreign cultures.

[94] Per Dilley (1999, 4-5) where culture once referred to what a people say, do, and think, it is now often treated as the general interpretative context for such behavior.

In addition to examining context's role in the interpretation of social events, anthropology has also analyzed how what we understand to be context comes to be meaningful in people's lives. Here we turn from socio-cultural to linguistic anthropology and semiotics, i.e., the study of signs and meaning making, including but not limited to linguistics (Gal and Irvine 2019, 14-15). The semiotic literature generally concurs with cultural anthropology that context is fundamental to the interpretation of broad-scale social phenomena. It extends this role, however, to every act of semiosis, holding that interpreting any utterance requires interpreting both the meaning it communicates as an example of its linguistic form and the meaning communicated by its occurrence as a socio-historically located, or contextual, signal (Gal and Irvine 2019, 103; Hanks 1992).[95] Crucially, the semiotic literature insists that no matter how solid and separate the socio-historical world seems from human action, it is in fact created in and through semiotic processes (Gal and Irvine 2019, 110). Context, in other words, does not precede or exist outside of semiosis. It is not a set of material and social facts given as such (Silverstein 2003, 201-02), but rather a social achievement. In the here-and-now of actual historical interactions, people use a repertoire of semiotic means, described below, to jointly construe or create the "context" by which they understand both individual utterances and overall events of interaction (Auer 1991, 3-4).

It is important to appreciate the expansive scope of this theory. It contemplates that contextualizing activities guide participants' understandings of a vast array of potentially relevant

---

[95] This role of context in guiding the interpretation of utterances is illustrated by the ability to effectively communicate different meanings using the same words on different occasions (Culler 1981). Indeed, the semiotic understanding of context originates in the study of "shifters," a category of linguistic sign defined by the fact that their meaning always relies on some contingent, contextual aspect of their occurrence of use (Silverstein 1976; see Nakassis 2018).

aspects of every interaction's context. These range from the "genre" of an interaction, i.e., the larger activity-type in which participants understand themselves to be engaging (see Briggs and Bauman 1992), to the interaction's topic. They also extend, however, to the social roles of participants and the nature of the relationships between them.[96] Semiotics rejects, in other words, the social scientific presumption that members of a given society automatically recognize the social scenes they confront (see Suchman 2007, 63-65; Goodwin and Duranti 1992, 28). The relative success or intelligibility of a typical visit to a doctor's office, for example, requires doctor and patient to establish themselves as inhabiting such roles. Cultural schema exist, which define, institutionalize, and regulate social access to these roles. From the patient's perspective, however, a doctor does not inhabit her role by virtue of a diploma, but rather by actualizing her potential inhabitance of it in a given interaction. The necessarily provisional nature of all such achievements is illustrated by the ease with which doctor and patient can, through relevant interactional moves, renegotiate context, shifting temporarily into the roles of "old friends" and the activity of "catching up" (Auer 1991, 22, 27).

Not even the material, spatial, or temporal environment of an interaction can be taken for granted as immutably fixed or 'out there' (Goodwin and Duranti 1992). Interactionally negotiated contexts may include facts of the material or social surround as observable by an outsider. These facts, however, must still be made 'visible,' brought to consciousness as the common grounds of an interaction (Auer 1991, 22-25). While some aspects of context are thus

---

[96] Silverstein (2003) extends this argument from interactional roles to the perduring categories of personal identity, including age, gender, class, and profession. The apparently intuitive ability of such categories to partition social space, he shows, is constituted in real-time interaction by mobilizing ideologically laden schemes of categorization.

"brought along" from the observable surround, others are instead "brought about" or entailed. These aspects of context cannot be predicted from the material or social environment prior to interaction.[97] They rather emerge from shared facets of knowledge as they are made relevant to participants' cognitive schema and frames.

As elaborated by linguistic anthropology, the repertoire of means used to realize potential aspects of context as relevant to an interaction includes verbal signals, but also extensive non-verbal, non-referential ones. These include gesture, mimicry, gaze, and posture. Such "contextualization cues" draw on socio-cultural knowledge, but do not necessarily carry information. Instead they suggest lines of inference or reasoning through relational signaling, by establishing contrast and oppositions (Gumperz 1991, 50), and by evoking structures of symbolic value (Silverstein 1992, 73-4). As reflected in the breadth of such contextualizing cues, human interaction is generally 'multi-channel.' Verbal and non-verbal context cues frequently co-occur in time, and their relative synchronization, or the lack thereof, is relevant in guiding participants' interpretations (Auer 1991, 13). Relatively dense bundling, for example—the simultaneously pointing of multiple cues all in the same interpretive direction—can enable complex inferences unachievable by single cues (Ibid. at 29).

The semiotic study of context invites us to analyze web browsing, and interactive media in general, as forms of interaction shot through with semiosis. It cautions against taking "browsing" or other forms of "using" networked technology as self-evident (see Nakassis 2020). Rather, it calls for attending to how Mozilla, through the browser, deploys contextualizing cues

---

[97] Conversely, contextualizing signals may render perceptible features of a situational context interactionally irrelevant even if they cannot entirely substitute them (Auer 1991, 28; Goodwin and Duranti 1992).

to project and manage the interactive relational schema that we understand in terms of browsing. It further compels us to consider how Mozilla uses such signaling to characterize the nature of browsing, marking it with salient social characteristics, including presuppositions about the nature of the relationship of users to the browser's interface and developer, and to website servers (see Gumperz 1991, 41-3). With respect to Containers, the question becomes, what contextualizing strategies did Mozilla employ in its efforts to guide users towards safe, reliable use of the feature?

As a threshold matter, observe that the repertoire of communicative channels and contextualization cues available to mediate human-browser interaction is significantly less robust than in human-human interaction (see Suchman 2007, 125-75). Many non-verbal signals, for example, largely escape the available channels of human-computer communication, thus limiting the resources available to Firefox and users to socially frame and organize the space of human-browser interaction. There is nothing in the available channels of interaction, for example, to accommodate the form of intentional, mutual gaze, which in human-human interaction both indicates "intense focus" and enables full monitoring of a partner's non-verbal activity (see Auer 1991, 34). Despite the introduction of haptics in contemporary smartphones, the use of the body and gesture as framing devices remains limited to the hands via keyboard, mouse, and trackpad. Similarly, despite advances in voice recognition technology, one cannot directly converse with a browser, thus limiting the use of linguistic and non-linguistic verbal signals to cue or mutually align around different interpretive frames (see Goodwin and Duranti 1992, 7).

Nonetheless, in seeking to cue, manage, and characterize interaction with Containers' users, Mozilla's engineers were not at total disadvantage. As Webb Keane (2003) argues, the

capacity of material objects to communicate meaning is predicated on the existence of an applicable "semiotic ideology." These are sets of background assumption in a given socio-historical setting about the nature of signs, how they work in the world, and the kinds of entities that are capable of signification. Following Keane, for Containers to successfully signal to a Firefox user that she was, with the feature, 'browsing in a secure context,' users would first need to be able to recognize pixels-on-screen as vehicles of interpretable meaning. In this regard, the Containers team thus benefitted from existing cultural familiarity with the possibility of "reading" software interfaces as text-like repositories of meanings in need of interpretation. To confirm the existence of such a semiotic ideology, one need look no further than the humble hyperlink. With its characteristic blue coloring and underlining, the hyperlink deploys a minimal aesthetic grammar to set itself off from surrounding text, indexing the presence of a navigable URL, and effectively inviting users to "click here."[98] Relative to the hyperlink, the Containers team might take comfort that Firefox users would recognize the color, shape, placement, and movement of Containers's visual elements as intentional, interpretable signals regarding the feature's operation.

To appreciate the relational schema that Mozilla needed to enact with Containers, consider the interface itself. For Containers' experimental launch in June 2016, Mozilla's designers and engineers implemented what they described as a basic, minimal interface. In this initial design, Containers consisted of little more than a drop down menu and a limited set of new

---

[98] Beyond the relational schema thereby invoked, the existence of a hyperlink further speaks to the historical presence of the website developer. Within web communities, it arguably also invokes structures of symbolic value that would mark a website developer as the kind of modal web citizen who takes appropriate advantage of the web's novel capabilities and affordances.

visual design elements. Once installed, the feature resided in a new menu option accessed via the browser's File Menu. Upon clicking the New Container Tab option, as it was called, users could to select from one of four default "container" options. These included: Personal (for websites, including email and social networks "you frequently use at home"), Work ("to use at the office"), Banking, and Shopping. By selecting one of these options, users instructed Firefox to open a new browser tab. Mozilla's engineers described such tabs as being "in" the selected container in the sense that they corresponded at the level of interface to one of the newly isolated Cookie Jars defined in the Firefox platform by the relevant userContextId. Opening a website in a container tab thus effectively instructed Firefox to isolate all cookies set by that site in the corresponding Cookie Jar.

In publicly announcing Containers, Mozilla's engineers were at pains to communicate that the feature would not change Firefox's "normal browsing experience" (Vyas 2016). They specified, for example, that normal tabs and windows would still "look and act as you'd expect." The visual appearance of non-container tabs would not change and they would continue to have access to all previously stored browsing data, such as browsing history. The offer of this assurance, and the related incorporation of Containers into Firefox's existing visual and organizational grammar, supplied an inferential basis to surmise that installing Containers would not change the genre of activity mediated or offered by Firefox (see Auer 1991, 4).

Containers, from this perspective, did not need to summon its interactionally relevant context whole cloth. It only had to cue an alternation, a shift from 'normal' browsing into 'contextually-contained' browsing. We might think of this shift in terms of the inauguration of a new browsing sub-genre, one distinguished at a first approximation by the attribution of new

qualities to browser tabs. In contextually-contained browsing, for example, users can no longer presume that all browser tabs represent undifferentiated, interchangeable instances of Firefox. Some now instead represent distinct, isolated arenas of self revelation. In this respect, Containers requires users to apply a category concept to browser tabs, with container tabs distinguished from default tabs by their connection to and visual manifestation of Firefox's newly multiple Cookie Jars. Further, in contextually-contained browsing, opening and managing browser tabs no longer implicates only browsing's relative ease and convenience. It instead requires users to weigh such considerations alongside whether or not to project a contextually-specific identity, and if so, which one. Only by consistently using Firefox tabs mindful of this new possibility could users reliably activate the contextual-containment enabled by Containers. In addition to such shifts in the presumptive characteristics of browser tabs, as a new sub-genre, contextually-contained browsing also entails subtle shifts in the presumptive social characteristics of Mozilla, Firefox, and its users. Browsing with Containers, for example, shifts Mozilla from its general role as a browser vendor into its more specific role as a user agent or "data steward," i.e., as the trustworthy caretaker of user data and privacy. Similarly, to browse with container tabs is to fold into the figure of the user a presumptive willingness to actively, intentionally manage one's personal privacy while online.

To cue and sustain these shifts in interactional context, Firefox's designers relied primarily on Containers' new visual elements. They assigned a unique icon and color, for example, to each of the four default container options. These included a stylized blue fingerprint for the Personal container and a pink shopping cart for Shopping. The relevant name and icon now appeared in the URL bar of every open Containers tab, and each container's assigned color

also formed a slight border across the top of each tab. According to the project team, the new tab colors served to visually distinguish container tabs from default tabs and from other container tabs of different types.

With container names, by contrast, Mozilla sought not just to differentiate container tabs but also to characterize this difference. The categories were originally recommended by Mozilla user researchers based on users' identification of them as the most common or useful contexts given their general browsing habits. The Banking and Shopping categories, however, also overlapped with browsing activities, and thus the types of personal data, that Mozilla's engineers believed technical experts generally sought to protect when employing ad hoc data segregation. As a set then, the container categories and icons provided inferential clues to the nature of the contextual identity services provided by the feature. They did so by linking the taken-for-granted organizing contexts of everyday browsing to categorical distinctions between activity and data types, which already invoked privacy-related sensitivities.

In selecting and designing Containers' defaults and icons, the Containers team worried about successfully "representing" the container types for Firefox users. Such concerns reflected in part the global nature of the Firefox user base. In comments posted to Mozilla's software development platforms, for example, they questioned the likely reception in Muslim countries of a piggy bank as the proposed icon for the Banking container. They similarly wondered whether the dollar sign, given its status as the currency of a particular, if hegemonic nation, could iconically represent "Banking" to a global audience.

In posing such questions, however, and indeed in presenting the concept of distinct browsing contexts to users in the form of pre-set defaults, Mozilla approached as a problem of

representation what linguistic anthropology (see also Dourish 2004) identifies as one of interactional positioning. Mozilla effectively figured context as a pre-existing, reified and self-contained fact of the world, an image inherently in tension with users' lived experience of context as the object of rapid, fluid, and mutual negotiation. To a certain degree, the Containers team recognized this tension and the challenges of comprehension that it posed for Containers. Based on feedback to Containers' initial design, for example, Mozilla eventually enabled Firefox users to create and name their own custom containers. They thereby introducing a degree mutuality to Containers' definition of context, effectively recognizing the futility of attempting to unilaterally impose its own vision of interactional context.

The team also recognized context's dynamic nature when discussing the relationship between the Work and Home containers. Despite the conceptual distinction identified by user researchers, the Containers team found that many users did not rigidly separate Work and Home. As Bram Pitoyo told me, "sometimes, you know, browsing is not that straight forward. It's not that. It's very permeable. Sometimes you do work stuff at home and home stuff at work, and stuff like that." Recognition of such malleability may not have motivated the ongoing redesign of Containers, but it did justify related decisions. Bram, for example, cited both the sharing browser history across all container types and the segregation of container types by tab, rather than by browser window or profile, as design decisions that practically aligned with and accommodated the flexibility of the work/home distinction.

It's Just Metadata

In October 2016, Harlo Holmes, Director of Digital Security at the Freedom of the Press Foundation, gave a talk at Mozilla's Toronto office, which was live-streamed to Mozilla's global staff and volunteers. Holmes had been invited to speak as part of a monthly series in which technologists, academics, and civil society experts lectured on issues relevant to Mozilla's non-profit mission and corporate objectives. As Holmes explained, at the Freedom of the Press Foundation, her work consisted primarily of training journalists and filmmakers to securely use digital technology. The Mozilla Foundation had, moreover, once awarded Holmes a fellowship to help the New York Times develop computer-assisted reporting capabilities. Despite these connections, Holmes had come to Mozilla not to discuss digital security or investigative journalism but rather metadata.

After introducing her work with the Freedom of the Press Foundation, Holmes moved on to the topic at hand. She began, as technologists frequently do when publicly discussing metadata, by defining the term. Starting from the standard shorthand definition of metadata as "data about data," Holmes walked the audience through various approaches to categorizing metadata, enumerating specific types of it along the way. When looking at commonalities across these categories, she asserted, metadata reveals itself to be at its core "data about events." Whether one is dealing with local computer files or data packets in transit across the internet, metadata refers not to a file's semantic content but to its machine-readable characteristics and event structure. Metadata describes things like the size of a file, for example, and the date of its creation. For files such as photographs, metadata identifies things like who took the photo and, when it had been accessed.

With origins partially in library science's classificatory impulses, Holmes acknowledged, metadata seems "kind of boring." In a separate talk delivered in December 2014, however, Holmes succinctly identified its social significance and relevance to Mozilla's mission. In this earlier speech, Holmes began as she would at Mozilla by defining metadata. She did so in this instance, however, by conducting a Google search. "By way of illustrating what metadata is," Holmes said, "I decided to actually go to the image search, because that right there, *this* is what metadata is. You will get the answer to your question…just by being *around* the question, and seeing what kind of data pops up to the fore." Showing a screen capture of her search results, Holmes directed the audience to an image category prominently suggested by Google: "You'll notice that one of the most popular associations with the term metadata…is, of course, 'NSA.' Living in the post-Snowden world, you guys all understand how this word got associated with the NSA in such a way that it was pushed to the top of Google image search. That's an interesting way to think about what metadata is."

Per Holmes' allusion here, if metadata and its definition had become objects of significant concern to civic-minded technologists by 2014, it was because metadata had been spectacularly thrust into public imaginaries by the Snowden revelations. On the one hand, the NSA documents released by Edward Snowden revealed communications metadata to be the central object of the government's secret mass surveillance programs. The initial Snowden-related news article, for example, published by Glenn Greenwald in The Guardian on June 6, 2013, disclosed the existence of a secret court order. This order required the telecom provider Verizon to give the NSA certain information regarding all telephone calls in its system, whether placed between the US and other countries or entirely within the US. The court-mandated

information did not extend to the content of calls, but it did include telephony metadata, such as the phone numbers of the parties to a call, location data, unique identifiers, and call time and duration. Subsequent reporting revealed the existence of additional NSA programs dedicated to the bulk surveillance of e-mail and internet metadata, including browsing history and map searches.

On the other hand, in the immediate aftermath of the Snowden revelations, federal officials repeatedly mobilized the metadata/content distinction to legally and morally justify the NSA's programs. Regardless of their surprising scale and ambition, public officials argued, the NSA's programs did not actually infringe the privacy of American citizens. This was so because the programs collected and analyzed only communications metadata and not communications content. On June 7, 2013, for example, the day following The Guardian's initial reporting, President Obama was asked at a press conference to react to the news of secret phone and internet surveillance. Obama began by questioning the characterization of the NSA's surveillance as "secret." The programs in question might formally be classified, he noted, but bipartisan congressional majorities had nonetheless repeatedly authorized them. Turning to program specifics, Obama continued, "When it comes to telephone calls, nobody is listening to your telephone calls. That's not what this program is about. As was indicated, what the intelligence community is doing is looking at phone numbers, and durations of calls. They are not looking at people's names and they're not looking at content. …So, I want to be very clear, some of the hype that we've been hearing over the last day or so—nobody is listening to the content of people's phone call."

Senator Diane Feinstein of California, then chair of the Senate Select Committee on Intelligence, similarly foregrounded the exclusion of the human voice from the NSA's programs. At a June 6 news conference, Feinstein told reporters, "As you know, this is just metadata. There is no content involved. In other words, no content of a communication." Feinstein repeated the point in an October 2013 op-ed, adding, "The NSA only collects the type of information found on a telephone bill: phone numbers of calls placed and received, the time of the calls and duration." Through such statements, government officials analogized the object of NSA surveillance to the mundane if pervasive forms of information that populate the bureaucratic peripheries of contemporary life. This was the stuff of billing, not of the intimate, private self, vulnerable to capture when exposed as speech. Public officials who mobilized the content/metadata distinction thus dismissed Snowden-related public alarm as unwarranted 'hype,' undermining privacy's legibility as a legitimate basis for moral critique of the NSA.

It is easy enough to understand why officials like Obama and Feinstein might turn to the content/metadata distinction to defend government surveillance. Surveys of American attitudes consistently demonstrate that Americans worry more about the kinds of interpersonal privacy issues likely to involve close family or friends overhearing a conversation than with the corporate and government surveillance effected using metadata. A 2013 survey conducted by the Pew Research Center, for example, showed Americans to be most concerned about controlling access to the content of their email. Conversely, Americans were least concerned about controlling access to the times of day they use the internet, a form of metadata (Raine et al. 2013). Similarly, in American popular culture, depictions of privacy-violating government overreach frequently revolve around scenarios involving police wiretaps, bugs, and other

surreptitious means of capturing the human voice or written word. Public officials who defended the NSA on the basis that it does not 'listen to what you say or read what you write' invoked such popular, concerns and depictions, implicitly framing the written and spoken word as the seat of legitimate privacy interests.

Whatever assurance the metadata/content distinction offered the general public, for many American technologists it was cold comfort. As Edward Snowden himself reminded the engineers of the IETF in July 2015, undeniable practical harms follow from metadata-based surveillance. For proof, one need look no further than the ongoing US drone war in Afghanistan. Even if "[w]e don't know who's holding it at the time we launch the missile," Snowden said, military operators target suspected Taliban fighters based on unique cell phone identifiers and other metadata. The former NSA director Michael Hayden admitted outright in a 2014 public debate, "'We kill people based on metadata.' And he's not lying," Snowden added.

Beyond such troubling practical effects, technologists shared a widespread belief that metadata does directly implicate personal privacy. Indeed, they warned, metadata is often significantly more revealing than the content of conversations. In a June 2013 op-ed addressing NSA surveillance, computer scientist Matt Blaze of Clipper Chip fame wrote to this effect, "[T]here's more to privacy than just the sound of our voices: Content may be what we say, but metadata is about what we actually do. And unlike our words, metadata doesn't lie… Metadata is our context. And that can reveal far more about us—both individually and as groups—than the words we speak" (Blaze 2013). The government's mis-recognition of this inherent expressivity especially troubled technologists because the metadata/content distinction is not only operative in the murky recesses of national security law. Rather, the distinction's legal significance

originated in and remains foundational to the criminal law doctrines that govern the Fourth Amendment's application to electronic communications.

Legal scholars generally trace electronic surveillance law's incorporation of a content/metadata distinction to the 1967 Supreme Court case, *Katz v. United States* (see Kerr 2010, 1023). As noted, *Katz* addressed the question of whether the Fourth Amendment prohibition of unreasonable search and seizure applies to the police when attaching a recording device to a telephone booth's exterior. In ruling that it does, Justice Stewart wrote for the Court, "One who occupies [a telephone booth], shuts the door behind him, and pays the toll that permits him to place a call is surely entitled to assume that the words he utters into the mouthpiece will not be broadcast to the world." (389 U.S. at 352). *Katz* thus established that the Fourth Amendment protects the content of telephone calls and sketched the contours of what would become the Fourth Amendment "reasonable expectations" test. Under it, to prove that police surveillance constitutes a protected "search," individuals must both demonstrate "an actual (subjective) expectation of privacy" and show that "society is prepared to recognize [that expectation] as 'reasonable'" (*Katz,* 389 US at 361 (Harlan concurring)).

Because *Katz* exclusively concerned the protections that apply to the contents of telephone conversations, the Supreme Court did not consider in 1967 whether any protection extends to call-related metadata. It finally addressed the Fourth Amendment status of such non-content in 1979's *Smith v. Maryland.* In this case, the Court considered whether the warrantless use of a pen register device to record the numbers dialed from a phone violates the reasonable expectation of privacy. Ruling that this would not constitute a Constitutionally-protected search, the Court distinguished use of pen registers from the content-acquiring listening devices at issue

in *Katz*: As a device that "do[es] not hear sound," a pen register does not disclose "the purport of any communications between the caller and the recipient of the call" or "their identities."[99]

With *Katz* and *Smith*, the Supreme Court formalized a legal distinction between communications content and non-content. Implicit to the distinction was the presumption that metadata poses a less significant threat to privacy than does oral or written content. It thus does not merit the same degree of Fourth Amendment protection.[100] The net effect of the distinction as incorporated in the law has been to absolve federal agencies engaged in metadata surveillance from following the extensive substantive and procedural safeguards that govern efforts to compel disclosures of content.

Given their sense of metadata's expressivity, news of the NSA's bulk metadata collection reinvigorated technologists' concern with the law's differential treatment of content and metadata. Technologists were left to contemplate why it was that metadata collection did not register either for lawmakers or in popular intuition as the "obvious" problem of privacy that it appeared to them to be. As they did when critiquing the government's "impossible" proposals for an encryption backdoor, they generally answered this question by pointing to policymakers' relative lack of technical expertise. Here, as in the resurgent Crypto Wars, technologists figured

---

[99] The Court's decision further turned on the conclusion that the dialing information had been "voluntarily conveyed" to a third party, i.e., the telephone company, for the purpose of connecting the call. Because "a person has no legitimate expectation of privacy in information he voluntarily turns over to third parties," the Court determined, society was not prepared to recognize the existence of a reasonable expectation of privacy in dialed telephone numbers.

[100] Relatedly but separately, under the so-called third party doctrine announced in 1976's *United States v. Miller* and affirmed in *Smith*, individuals presumptively waive any Fourth Amendment protections that might apply to communications metadata "voluntarily" disclosed to a third party service provider, such as a telecom, ISP, or email provider, regardless of its relative sensitivity.

technology and technological development as moving at speeds beyond the the law's institutional grasp.

We can see a version of such logic at work in "It's Too Complicated: How the Internet Upends *Katz*, *Smith* and Surveillance Law," a 2016 law review article authored by a group of prominent computer scientists, including Matt Blaze and Susan Landau (Bellovin et al. 2016). In the article, Blaze, Landau and their colleagues largely avoid the question of metadata's inherent sensitivity. They reject the viability of the content/non-content distinction as a basis for meaningfully regulating law enforcement access to communications data. They do so, however, by showing the doctrine to be built around outdated assumptions regarding the ontology of data and the operation of communications technology.

Self-consciously adopting a "a technological vantage point," the computer scientists open "It's Too Complicated" by analyzing the architecture of the mid-twentieth century wireline telephone system, the predominant communications technology at the time of *Katz* and *Smith*. The distinguishing characteristic of this system, they posit, was its localization of intelligence in telephone companies' internal infrastructure. This confinement of intelligence to centralized telephone switches, they argue, both technically and economically prohibited phone companies from offering services beyond call dialing and answering. The network's design further required phone companies to directly provide all services themselves. As a result of these features, the phone network operated according to a structurally simple division. Under it, information was either a dialed number or a conversation, either metadata or content, with no intermediate categories. Given the telephone system's treatment of content and metadata as distinct technical

concepts, the computer scientists conclude, it was "quite plausible for the courts to draw a bright line between content and metadata" (Ibid at 35).

The authors next turn their attention to the internet, the architecture and operation of which, they observe, is significantly more complex than that of the phone network. In circuit-switched networks like the wireline telephone system, for example, a dedicated circuit is built for every individual call and used exclusively for the call's duration. By contrast, as a packet-switching network, the internet breaks all communications into multiple small packets. Each packet may travel a distinct route through the network before being reassembled at the endpoint. Among its other practical effects, the computer scientists argue, the internet's open and dynamic architecture creates a communication environment in which the status of an individual unit of data may change as it travels from sender to recipient. What functions as content at one layer of the network may operate instead as non-content or metadata at another, and vice-versa. Whether something is "content" requiring police to obtain a probable cause warrant thus becomes a function of the precise location in the system at which the question is analyzed. By rendering the physical and legal distinction between content and non-content too difficult for courts to consistently apply, the authors argue, the internet undermines the distinction's utility as a means of identifying the aspects of communications that require Fourth Amendment protection.[101]

---

[101] The computer scientists take a similarly dim view of the internet's effects on the third-party doctrine. The complexity of IP-based communications and services, they argue, is such that even the most technically sophisticated user may be unable to discover or comprehend what information she communicates to third parties. Unlike the dialed digits shared with a telephone company, these interactions—and indeed the presence third parties themselves—may be completely invisible to the user of internet services. If *Miller* and *Smith* predicate the voluntariness of conveyed information on a showing that consumers know they are disclosing the information at issue to third parties, the authors conclude, third-party doctrine can no longer provide meaningful legal guidance for regulating government access to internet-based data.

In its focus on technical comprehension and the out-of-synch temporalities of technology and the law, "It's Too Complicated" can be considered representative of much of technologists' post-Snowden discourse regarding the law's failure to protect metadata. For Blaze, Landau and their co-authors, the law's problematic treatment of metadata reflects an understanding of the operation of communications technology keyed to technical infrastructures no longer in widespread use. Other computer scientists and engineers concerned with the absence of legal protections for metadata looked instead at intervening developments unrelated to the internet's complexities. These included, most prominently, the contemporary explosion in the kinds and the overall amount of metadata generated about individuals, and the proliferation of increasingly powerful and thus revelatory data analytic techniques.

Whatever access such considerations provide into the law's treatment of metadata, they seem not to tell us much about related popular intuitions. The common sense nature of the public attitudes invoked by Obama and Feinstein suggests not simply a lack of familiarity, but some affirmative reason to perceive metadata as being inherently less threatening to personal privacy than communications content. In closing this chapter, I propose to tease out the logic of such public intuitions not by reference to technical literacy, but rather to a different analytic, one which, I will show, technologists themselves employed even if they did not reflexively name it. Specifically, I will argue that if not just the law but popular intuition failed to recognize metadata's inherent sensitivity, this was as much a function of metadata's mode of signification —i.e., of the semiotic "grounds" through which metadata communicates meaning—as it was of public's relative capacity to correctly apprehend the workings of technology.

To appreciate this argument, it is helpful to examine the genre of tutelary discourse that technologists employed after Snowden to familiarize the public with the pervasive contemporary fact of metadata and to demonstrate its sensitive nature. Consider as illustration the anecdote offered by Harlo Holmes during her talk at Mozilla. By this point in her talk, Holmes had developed the argument that the widespread aggregation, analysis, and monetization of metadata implicates the forms of freedom possible in contemporary life. In extreme cases, such as the surveillance of cellphones in refugee camps, it may even facilitate militarized forms of behavioral modification, or "digital herding." If metadata can have such social effects, Holmes suggested, it is in part because when it is operationalized at mass scale, it effectively becomes "part and parcel of the user themselves in a way that it not separable as they traverse the services to do their thing and live their best life." Even so, the capacity of metadata to "reify" the self, to serve as the basis of a process "by which a person comes to know themselves from the representation they get back from the services they use," can be exploited in the service of security and freedom.

To illustrate this point, Holmes shared a technique developed in her work educating journalists and activists. "In my digital security training," she said, "I often do this one thing first. It's like a magic trick. I tell them, 'Go to Account on Google, and look at their Web and App Activity.' People usually gasp. It's like one of those commercials, like, 'You can save on car insurance!' People are literally blown out of their minds by what they never knew was being collected and then being presented back to them.' Having introducing the trick, Holmes simulated it for us, displaying screenshots of her Web and App activity data: "Here's my October 9th and 10th history. It tells me what I did in search, what I've done in developer land, what I've

done in the ad space, etc. What's most interesting is the info on apps and when I use them. … [T]he clock app… shows when I wake up and press snooze. I can query for patterns in my day-to-day: When I'm most likely during the day to be doing research for work. When I'm involved in a Twitter thing. When I do my banking (usually [it's] when I get in a cab). Even though info about [the] context of conversations and parties I'm conversing with is not available, it does say a lot about 'when.' That gives me a lot of pause when I engage in services, such as [Google's] Android."

Despite Holmes' characterization of it as a trick, the form of tutelary discourse illustrated here is better understood as a regular feature of technologists' post-Snowden discourse regarding the content/metadata distinction. Such metadata demonstrations served a dual purpose for technologists, both informing non-expert audiences about the pervasive reality of contemporary metadata surveillance, and ideally, motivating them to adopt practices of responsibilized care towards their metadata. As illustrated by Holmes, such demonstrations unfold across two analytically distinct phases. In the first, surprise phase, individuals are simply exposed to their own data. Such exposure ideally 'blows people out of their mind,' reproducing in personalized miniature the national shock of the Snowden revelations. If, per Holmes, there is a certain magic to the technique, the magic resides in the provision of experiential access to a domain of contemporary life that is at once remote and intimate, removed from everyday perception, but unnervingly suffused with personal information.

It is the second phase of such demonstrations that is of primary interest, however. In this phase, technologists move beyond revealing the 'hidden' reality of metadata to training individuals to recognize and appreciate the kinds of information that metadata can communicate

and the semiotic mechanism by which it does so. As Holmes noted in displaying her data for us, metadata doesn't capture the content of conversations, but it does capture is 'what I did.' It captures technologically-mediated events and interactions. Analyzed over time, such individual events are said to correspond to and reveal patterns of personal behavior, as in the example of Holmes' daily sleeping habits. As Holmes showed in reviewing her smartphone clock, metadata not only captures the actual occurrence of events—sending a tweet; logging in to a banking app—but also characterizes it. In Holmes' examples, it does so by establishing relations of co-occurrence between such events and the measures of clock time. In this respect, metadata describes not an intrinsic property of certain kinds of data, but rather a relationship between two pieces of data, with one understood as being "about" and thus qualifying the other in some way (see Kift and Nissenbaum 2017)

In her talk at Mozilla, Holmes quickly moved through the metadata demonstration. Holmes' examples showed that metadata can capture behavior and locate it in time, but Holmes offered little in the way of a reflexive theory of metadata's expressivity. To gain further insight into the kinds of information that metadata can communicate and into the logic of the revelatory power attributed to it, consider another example of this genre of tutelary discourse. Here we turn to the *amicus curiae* or 'friend of the court' brief submitted by a group of prominent computer and data scientists in *ACLU v. Clapper.* A 2015 district court case, *Clapper* involved an ACLU challenge to the legality of the NSA's bulk phone metadata program. Writing in their capacity as technical experts, the brief's authors, including Matt Blaze and Bruce Schneier, sought to refute the government claim that metadata is inherently non-sensitive. In introductory remarks, the computer scientists wrote that when aggregated metadata generates comprehensive records of

people's habits. Collected in bulk, call records are thus not "just" metadata as Senator Feinstein put it, but "intimate portraits of the lives of millions of Americans" (ACLU v. Clapper Brief Amicus Curiae, 4) The computer scientists insisted, however, that aggregation only "intensifies" the sensitivity of the information yielded by metadata. To illustrate that metadata yields sensitive information "even at the level of individual calls," they proffered the example of single-purpose telephone lines. While the telephone metadata generated during individual calls may "on its face appear innocuous," they wrote, "it is anything but." "A call to a hotline or another dedicated, single-purpose phone line provides perhaps the starkest demonstration of the power of metadata to reveal deeply private and sensitive information about a single call or caller. An hour-long-call at 3 A.M. to a suicide prevention hotline; a thirty-minute call to an alcohol addiction hotline on New Year's Eve; or a fifteen-minute call to a phone sex service—the 'metadata' from these calls, even in the absence of the 'content' of the conversation, still reveals information that virtually anyone would consider exceptionally private" (Ibid. at 8-9). Beyond the limited case of single-purpose hotlines, the computer scientists argued, individual calls to certain kinds of organizations may still reveal an individual's religious or political affiliation, their sexual orientation, and even aspects of their career and social life.

As such examples demonstrate, if metadata can communicate information by locating events in time, it also does so by locating them in social space. They further show that while technologists often describe metadata as being "inherently" expressive, it is more precise to speak of metadata's inferential expressivity. Metadata does not communicate personal information in-and-of-itself, in other words. Its capacity to reveal sensitive personal information requires that the events it captures be linked to and supplemented by shared cultural knowledge

and presumptions. Metadata might establish the fact that an individual placed a call to a phone number registered to mental crisis health-line, for example. But to move from this documented fact to some "likely narrative," as the computer scientists put it, which locates the individual in social space as a certain kind of person—as someone who 'is suicidal,' for example—requires a series of intervening inferential moves. Such a narrative would seem to require, for instance, a cultural understanding of telephone calls as involved a form of motivated behavior, behavior that responds to and thus expresses individual agency and intention. It might further turn on cultural norms regarding the appropriate time for telephone calls. Only relative to such a norm might one reasonably presume, for example, that someone placed a phone call at 3am because they were experiencing some form of personal crisis.

What is thus ultimately important to observe in such demonstrations is the way they turn on training audiences to recognize in metadata a basis for inferential reasoning and narrative elaboration. Metadata may appear in a physical form popularly associated with opaque bureaucratic and machinic processes. It nonetheless yields meaningful personal information by instantiating partitions of social space that appear to point to data subjects as being a certain kind of person, with certain kinds of characteristics. From this perspective, we can specify that if metadata reveals personal information, it does so "indexically."

In characterizing metadata in this way, I draw upon Peircean semiotics, the same theoretical corpus from which linguistic anthropology derives the theories of context and contextualization discussed in preceding sections. In his influential theory of signification, the philosopher and linguist Charles S. Peirce identified three fundamental sign types. Each sign type is defined in Peirce's system by its distinct semiotic "ground" or mode of generating

meaning by linking signs to the objects for which they stand. The simplest sign type described by Peirce is the icon. As illustrated by the example of a portrait, icons signify through relations of perceived similarity or likeness between a sign and its object. By contrast, symbols, the second sign type, are typified or arbitrary signs. They signify by virtue of some specific delimited system. While the rules of grammar are the most familiar of such systems, symbols are not limited to linguistic forms. They may be deciphered relative to any code with conventionalized rules, including the color-based code used in traffic lights (Gal and Irvine 2019, 101). The index is the third and final sign type, and it is said to signify through an "existential" relationship of some kind between sign and object. Among the classic examples offered to illustrate indexicality is the relationship of smoke to fire. As a sign, smoke does not "stand for" fire representationally or symbolically, the way the word "fire" does (Duranti 1997, 17). Rather, smoke acquires its meaning from its relationship of physical and temporal contiguity with fire. By virtue of such co-presence an index like smoke is said to semiotically 'point' to its object (fire) as its understood cause (Gal and Irvine 2019, 96-97). In this respect, indexical signs intrinsically relate to and can only be fully specified by appeal to to the contextual conditions of their event of use (Nakassis 2018, 286, 289).

For the purpose of understanding American popular intuitions regarding metadata, identifying metadata's indexical nature is significant. To appreciate why this is so, one must know something of the challenge that indexicality's theoretical elaboration has posed for certain deeply-lodged Western assumptions about the purpose of language and the nature of meaning. Post-Enlightenment Western thinking about language is largely preoccupied with how words and expressions represent the objects of the world for which they stand (Gal and Irvine 2019).

Western accounts of language accordingly emphasize speech's "referential" function, the way through its form language appears to label a pre-existing world and convey true statements about it (Silverstein 1976; Agha 2007; Hill 2008). Such thought treats language as an autonomous, purely symbolic system that is divorced from both the world it describes and from events of language use (Nakassis 2018, 284-85). Under its logic, language exists to serve as an "instrument" for informing or describing the world (Lempert 2005; see Duranti 1997; Culler 1981). Under the sway of this so-called referentialist ideology (Hill 2008; Andrus 2015), Western thought treats symbols as the fundamental linguistic form and reduces meaning to referential meaning (Silverstein 1976).

Since the 1970s, linguistic anthropology has called the tenets of referentialist ideology into question. It has done so by examining language not as an abstract, autonomous system but as a form of meaningful social behavior, a task accomplished in significant part by elaborating the nature and function of indexes (Nakassis 2018). In an influential 1976 paper, Michael Silverstein inaugurated this tradition, building from the observations that most signs employed in actual events of speech use are not referential and that most utterances in sequence are in fact multifunctional. Whether or not they contribute to reference, utterances independently accomplish other important socially-constituted, or pragmatic, ends. Among the communicative functions Silverstein explored in support of this claim is the "privacy function" of the language. By this he meant the individual communicative ability, using phraseology that only some people will understand, for example, to set social boundaries on an interaction. As Silverstein observed, such a pragmatic social effect is otherwise generally achieved through non-linguistic behavior, such as by erecting a physical barrier. Other identified pragmatic functions of language include

marking and changing the social status of participants to an interaction, and entailing particular kinds of relationships between interactants (Silverstein 1976; see Duranti 1997). Indeed, we can take the list of language's non-referential functions to be co-extensive with the set of social effects achievable through the semiotic constitution of context, as described in preceding sections. As Silverstein established (1976, 54), it is because indexes are existentially grounded in aspects of their situations of use that what we understand to be context can be established and transformed.

Linguistic anthropology has enjoyed considerable success in demonstrating language's power to accomplish deeds beyond mere description and symbolic classification. Referentialist ideology nonetheless remains deeply entrenched in American popular imaginaries where it appears as a form of common sense (Andrus 2016; Hill 2008). As a linguistic ideology, referentialism shapes and constrains American ideas and discourse regarding the functions of language and the nature of meaning (see Gal and Woolard 2001; Gal and Irvine 2019). By glossing over the action of language, it complicates Americans' ability to recognize that language functions interactionally and pragmatically, contributing to meaning by establishing indexical connections to contextual factors of use such as speakers, settings, topics, and institutions (Ibid.; Hill 2008, 42). In this respect, it can be understood to undergird the popular intuition that metadata does not implicate individual privacy. Because metadata signifies indexically rather than symbolically, it is ideologically unrecognizable as the kind of sign capable of yielding the identifying, 'sensitive' information that is privacy's object.

Referentialism thus constitutes part of the ideological backdrop against which NSA surveillance can be dismissed as involving 'just metadata.' Recall, however, that in mobilizing

the content/metadata distinction public officials placed particular emphasis on the spoken, oral quality of communications content. Similarly, in *Smith v. Maryland*, it was because pen registers do not capture sound that the Supreme Court dismissed the possibility that the metadata they collect can "disclose the purport of any communication." Such references suggest that if the surveillance of content intuitively appears to be more problematic from a privacy perspective than does surveillance of metadata, it has something to do not just with modes of signification but also with the medium of communication, i.e., with the human voice itself. The relative social illegibility of metadata's indexical mode of signification would not seem to account for this concern with the special, privacy-related sensitivity of spoken. We can nonetheless make sense of this feature of privacy-related discourse by turning to a different language ideology, one which Hill (2008) and Andrus (2015) describe as being closely aligned with referentialism.

Under "personalism," as this ideology is known, individuals are presumed to have an invisible, interior self, which is continuous across time and is the seat of beliefs and intuitions (Rosaldo 1982, 218). Speech, in turn, is presumed to be the product of intentional efforts to communicate one's beliefs by choosing words that best match them (Hill 2008, 64). Personalism thus holds that meaning is privately 'owned,' originating in the beliefs and intentions of individual speakers (Hill 2008; Duranti 1993; Rosaldo 1981). It grants primacy to such mental states in determining what an utterance is understood to constitute (see Goodwin and Duranti 1992, 18). From this perspective, personalism is what allows us to presume that when an individuals speaks he or she "means something by it" (Hill 2008, 89). From the perspective of personalism, intimate knowledge of the self, of a persons's beliefs and intuitions, can only be known as the individual explicitly communicates them in speech.

CHAPTER 3: FEATURING PRIVACY

In the preceding chapter, I focused on the usability challenges that Mozilla's designers and engineers described as complicating the effort to release Containers to Firefox's general use population. I showed how, during the feature's long development, the Containers team struggled to help users understand and reliably, safely use the feature as a privacy tool. Analyzing Containers in this way allowed me to situate its animating vision in relation to a widespread contemporary shift among American technologists and policymakers towards context-based approaches to privacy. It also allowed me to draw out some of the epistemic and political issues raised by related efforts in industry and academia to map the contextual factors understood to inform consumer privacy attitudes. Restricting my analysis to the issue of how to make Containers comprehensible and usable, however, required me to elide another significant tension, which also marked the feature's development. This tension, over precisely what the feature should be and whom it should serve, raises important questions about privacy's form and cultural image under conditions of technological stewardship.

To gain access to these questions, it is helpful to return to Containers and examine the nature of the initial user response to the feature. Consider, for example, the June 2016 blog post, which announced Containers' release in Firefox Nightly. As I described in Chapter 2, the post identified as the motivation for the feature's development the difficulty of maintaining contextually-specific identities online. It further explained how Containers purports to technically address this issue, and using screenshots of Containers' interface, illustrated how to

properly use the feature. What I wish to emphasize here however, is that the post was as much a solicitation as an offering of information. Explicitly framing Containers as an experiment, it attempted to enlist Firefox users in its testing and refinement. "We don't have the answers…yet" to the question of "the right User Experience" for Containers, the post offered, "but hope to uncover them with user research and feedback." The post concluded by urging interested readers to try Containers and then share with Mozilla how it might "iterate" on its design to make it more convenient and usable.

The "Contextual Identities" post identified multiple channels through which users could provide feedback on Containers, including a dedicated email address, a survey, and a link to Bugzilla, Mozilla's bug-tracking system. Many readers chose, however, to respond directly, submitting comments either to the blog itself or to its re-posting on Hackers News, the popular Silicon Valley news aggregator. Perhaps surprisingly, according to Bram Pitoyo, one of Containers' user experience designers, based on the comments, "we found that they get it. They're really excited. And they instantly know what to use it for." Indeed, scanning through the comments, what is immediately notable is that many take the form of "feature requests." Rather than asking for clarification regarding the feature's purpose or operation, they identify specific ways in which Mozilla might modify or extend its functionality. "It would be nice to customize the names associated with the identities," one suggests. "That looks neat," another writes, "and if it were possible to cheaply create and delete contextual identities on the fly it would even fix an issue I had today."

Such responses appear to belie the concern described in Chapter 2 that Containers was 'too complicated' to explain to users. As Pitoyo clarified by interview, however, it was not the

casual or general Firefox user who instantly knew how to use Containers. Rather it was Firefox's "elite technical audience," the kind of expert user likely to frequent Mozilla blogs and the Hacker News forums. Moreover, while such users correctly identified uses to which Containers could be applied, they directed their enthusiasm less to the feature's intended privacy benefits and more to its perceived usefulness as a tool of organization and task management.

As such expert users immediately appreciated, and as the Containers team were aware, the segregated cookie jars implemented to address the problem of contextual identities lent itself to multiple uses. Cookies, after all, are themselves a multi-purpose technology. In popular imaginaries, cookies primarily figure, if at all, as a vector of online surveillance. Cookies, however, were initially built to introduce into the nascent web new forms of "statefulness" and interactivity. As recounted by Lou Montulli (2013), the inventor of cookies, in the web's early days, websites were little more than static documents, functionally equivalent to printed pages of paper. As such, they lacked the ability to identify individual users. In 1994, when Montulli developed the cookie for Netscape, he did so as part of the early rush towards online commerce. Cookies enabled developers to create website "shopping carts" capable of remembering user selections. Only years later did advertisers co-opt the cookie to surveil users as the basis for personalized ad targeting. Websites today continue to widely use cookies for such purposes, as well as to authenticate users and enable account logins.

For Firefox's expert users, it was in relation to this latter function that Containers appeared to lend itself to task management. As Bram Pitoyo explained, Firefox's "pro users" tend to have multiple accounts with major internet services. They might have multiple Gmail accounts for example, or maintain both personal and professional Twitter accounts. For such users,

Containers' isolation mechanism offered a new means for simultaneously logging in to multiple accounts with the same service within one browser. In thus facilitating easy switching, for example, between 'personal' and 'professional' tasks, Containers presented an alternative to the burdensome ad hoc account management practices expert users otherwise employ.

Within the web community, expert users are also well known for being "heavy" users of browser tabs. Such users might regularly maintain more than 100 open tabs, for example, each a reminder or pathway to some professional resource or unconsummated task. Whatever benefit expert users derive from this practice, heavy tab usage strains the technical capacities of modern browsers, and presents challenges of organization and self-administration. For a vocal segment of Firefox's expert users, Containers' appeal derived not from its cookie segregation at all, but rather from the "layer of visual organization" that it added to Firefox's interface (Crouch 2017), and the new opportunities this afforded for grouping and managing open browser tabs.

As such responses illustrate, Mozilla's engineers may have built Containers' data isolation technology as the underpinnings for a new privacy-preserving browser feature. Firefox's technically-savvy users, however, instead identified in its implemented form a potential new tool of productivity. In Chapter 2, I argued that we might think of the usability issues ascribed to Containers in terms of the challenge of developing a semiotic grammar capable of reliably cuing and sustaining for users a new sub-genre of "contextually-contained" browsing. If usability thus entails interactionally guiding users towards the desired interpretation of a web technology's interface, the expert embrace of Containers as a productivity tool shows it also requires non-desired interpretations to be effectively dissuaded. Containers, from this perspective, suffered from an over-abundance of the stories that could be told about it as a

product. The multiple "use cases" that Containers fulfilled—the different kinds of browsing-related problems for which Containers appeared to offer a solution—meant that there were too many ways to frame how its technical architecture could look and behave for Firefox users. The organizational challenge represented by Containers was thus not limited to how it could be made easily usable, but also encompassed a struggle over what Bram described as the "shape" it would take as a feature, the function it would actually serve. Attending to the struggle of the Containers team to, on the one hand, preserve Containers' potential to protect user privacy, and on the other hand, usher it to widespread release, throws into relief the tension between Mozilla's stewardship of privacy and the other elements of its techno-moral vision for the web, in particular its commitments to "openness" and user empowerment.

In this chapter, I use the tension between privacy and productivity evidenced in the Containers project to analyze the feature form as an object and site of contemporary world-making around privacy. First, I describe the role of the techno-moral ideal of "openness" in Mozilla's historical emergence as a non-profit and open source project. I show that for Mozilla openness represents both an essential characteristic of the web in its idealized form as a globally accessible, empowering commons, and a model for its own corporate structure and operations. Exploring the value attributed to openness in browser engineering helps throw into relief that, as an institutional value, privacy historically served for Mozilla as a means to the end of the open web. From this perspective, Mozilla valorized privacy-related user control over data as an instance of the more general kinds of control over the user experience of the web required by openness. Thus pursued, privacy in part names as a tool for cultivating the collective sociality and affective ties that empower not users but rather Mozilla itself as it pursues its vision for the

web. Next, I describe how during my fieldwork in 2016 and 2017, Mozilla's privacy and security engineers faced a new institutional mandate to more ship privacy-enhancing browser features. As part of its revised competitive strategy, Mozilla determined that privacy represented a product domain in which it enjoyed structural and reputational advantages over its competitors. The mandate to clear the backlog of privacy features under development lent new institutional urgency to Mozilla's privacy commitments but also subjected the features under development to the economy of quality and attention by which Mozilla determines the browser's composition. Using Containers as a case study, I show that Mozilla's decision to materialize privacy in Firefox features constrains the form and political capacities that privacy may wield under conditions of technological stewardship. The gap between a feature's material and semiotic manifestations exposes privacy to the competing needs and aspirations of Firefox's technically-savvy expert users.

The Conspiracy of Open

In 2014 and 2015, as I began to explore the proliferation of tech-based efforts to defend privacy, I found my attempts to secure a long-term fieldsite repeatedly frustrated by privacy folding in on itself. As I traveled to privacy-related conferences and conducted preliminary meetings, the growing legion of privacy lawyers, activists, and technologists I met expressed support for an anthropological study of privacy. When it came to institutionally hosting such a study, however, encouragement quickly turned to reluctance.

In March 2014, for example, I spoke by phone with Jeff, the chief technology officer of a prominent Washington, D.C.-based technology policy non-profit. In his professional capacity,

Jeff works with lawmakers, the executive, and standards organizations to shape technology policy and infrastructure in the service of individual rights. Though he declined my request to conduct fieldwork with his organization Jeff agreed to help me understand why he thought such an arrangement would be unworkable. By phone, Jeff identified as the primary obstacle to fieldwork, the fact that his organization "plays the inside game." It was "a very D.C. mentality," he acknowledged, but working with policymakers meant that much of his work could never become public. His ability to extract concessions from lawmakers while negotiating proposed changes to surveillance law, for example, might depend on such negotiations remaining confidential. Jeff expressed similar concerns regarding the effect of a researcher's presence on his ability to "work the innovation angle," i.e., to exert influence through private negotiations and collaborations with corporations on the privacy-related properties of new consumer technologies.

The following month, I met with Will, an investigative researcher with the Electronic Frontier Foundation (EFF), the digital rights non-profit. At the time, the EFF's staff technologists were deeply involved in the post-Snowden effort to shift the web towards the default use of encryption. I had expressed interest specifically in conducting fieldwork with the EFF's engineers and computer scientists. Will told me, however, that the organization's sensitive legal work made such a proposal impossible. Here it was not privacy's role in cultivating and maintaining trusted professional relationships that was at issue, but rather a kind of spatialization of privacy. Within the EFF's cramped San Francisco offices, Will told me, certain rooms housed files related to the organization's ongoing national security litigation and were subject to court-ordered seal. Even the physical presence of an independent researcher, Will suggested, would

thus impose unmanageable supervisory burdens on the organization. Regardless, the staff's technologists had bristled at the possibility that the fieldwork I proposed would entail physically sitting in their offices, looking over their shoulders as they typed.

During our call in March, Jeff suggested that rather than embedding myself in a particular institution, I might try organizing my fieldwork around a more broadly defined privacy community. As possible entry points for such an approach he identified the regular meetings of the relevant standards organizations, as well as a weekly technology policy meetup and a book club hosted by prominent local technologists. A year later, in an exploratory trip to San Francisco, I tried something like Jeff's approach. My primary goal in planning the trip had been to attend the Yahoo Trust UnConference, a one day gathering of experts in cryptography, browser security, and anonymity, organized in the wake of the Snowden revelations as an exercise in community-building. While in town, however, I also learned of and attended the monthly meeting of PrivacyLab, a recently formed meetup for Bay Area engineers and others interested in using technology to enhance user privacy. At the meeting that night, hosted in the community space of Mozilla's San Francisco office, overlooking the Bay Bridge, I met Stacy Martin, one of PrivacyLab's organizers.

As Mozilla's Senior Data Privacy Manager, Stacy described her work as being largely behind-the-scenes but focusing on "external-facing" privacy issues. As part of the Mozilla Foundation's digital literacy efforts, for example, Stacy created privacy-related teaching kits. She also coordinated the volunteers of a task force dedicated to educating the Indian public about privacy. Stacy's work spoke to the breadth of Mozilla's privacy-related activities, its use of policy advocacy and education in addition to product engineering to fulfill its commitments to

privacy. Her career also tracked in certain ways privacy's emergence as a professionalized concern of American technology corporations. In 2000, the year she remembered as the beginning of widespread corporate engagement with privacy,[102] Stacy found herself by happenstance writing the first privacy policy issued by Hewlett-Packard, her former employer. The same year, she attended the first meeting of the International Association of Privacy Professionals, an organization whose membership grew dramatically in her subsequent years of participation. Though not single-mindedly dedicated to privacy, she told me, she believed generally in its value. Professionally, she enjoyed the challenge of applying privacy principles to a constantly changing legal and technological environment. Working at Mozilla combined for Stacy privacy a job "with the ability to do good in the world."

Over the following months, as we corresponded by phone and email, Stacy introduced me to friends and colleagues working at the intersection of privacy and technology. Eventually, we discussed the possibility of conducting fieldwork at Mozilla. As part of its formal commitment to privacy, Stacy told me, Mozilla has an institutional interest in facilitating such research. As part of its effort to build community and spread privacy-related knowledge, the Mozilla Foundation in fact routinely sponsors privacy research. Either way, Stacy intimated, given Mozilla's formal commitment to conducting its business in the transparent spirit of open source programming, it had little basis in its institutional norms upon which to deny my request.

In his 2008 ethnography of free software and open source programming, Chris Kelty describes the valorization of openness within tech communities as an element of and response to

---

[102] Across my fieldwork more broadly, my interlocutors offered competing periodizations for privacy's contemporary history.

the reorientation of knowledge and power effected by the internet. As figured in the practices of free software programming, openness is both technical and moral, means and end. According to Kelty, as "geeks"[103] imagine it, it is through the open, decentralized structure and operation of its protocols and architecture that the internet has introduced to the world an empowering new form of global, universal publicness. Preserving the internet's open, distributed interconnectivity is thus a key concern of geeks in their operation as a "recursive public," i.e., as a self-determining, politically independent collective committed to maintaining the technical, legal, and conceptual means of its own existence.

According to Kelty, this concept of openness has roots in post-Enlightenment ideals of freely circulating goods and knowledge.[104] It took its contemporary form, however, in the early 1980s through struggles over the definition of "open" computer systems (2008, 142-78). These struggles resulted in the creation of a partially articulated "infrastructure" of technical components, including the TCP/IP protocols, and moral components, including new demands for fair and open competition, open markets, and open standards processes in software production. In the late 1990s, the existence of this infrastructure helped facilitate the emergence of free and open source software, an emergence historically marked for Kelty by Netscape's 1998 release of the source code to the Navigator browser. As Kelty observes, in popular culture, Netscape is best

---

[103] As Kelty uses the term, Geek refers not to a self-ascribed identity but a mode of thinking and working and form of elite affinity (2008, 35).

[104] See Caduff (2012, 351) on the liberal imaginary of unrestricted sharing and free exchange. In the Habermasian public sphere, matters of common concern are decided on the basis of reason, not status. Public opinion thus wields political authority to the extent it is understood to be grounded in ideals of rational open dialogue (see Gal and Woolard 1995, 5). Warner (1990) similarly attributes political authority to openness. Openness for Warner, however, is achieved not through active participation but rather as an abstraction of individual characteristics enabled by the de-contextualizing anonymity of print.

remembered for its sensational 1995 IPO, which inaugurated the venture capital-driven mania of the dotcom boom. For hackers and programmers, however, equally significant was Netscape's release of what would become by 2004 the Firefox source code. For many years, the Firefox browser was the most visible and widely used consumer product produced through free software processes. Its popularization, according to Kelty, precipitated a wave of critical self-reflection among geeks, through which free software formed as an independent, recursive public capable of checking the spread of proprietary technical systems and monopoly power within the tech industry (Ibid. at 107-15, 143-45).

For some at Netscape, including CEO Jim Barksdale, the decision to take the browser open source was primarily a question of economic value, productivity, and efficiency.[105] When it was initially released in 1995, Navigator had been wildly popular as the first feature-rich browser. By 1998, however, Navigator had largely been eclipsed by Microsoft's Internet Explorer.[106] In Barksdale's justification, by releasing Navigator's source code, Netscape would extend its developer community beyond its small staff. Barksdale hoped to attract code from a number of outside contributors sufficiently large that "no one company could afford to have working on any one project" (Winton 2000). For Netscape's engineering staff, meanwhile, many with prior experience in open source, creating the Mozilla organization and making Navigator open represented the best chance "for the code to actually prosper" (Kelty 2008, 107), to "build a

---

[105] Kelty (2017) traces these concerns within open source programming and other contemporary development methodologies to a history of technocratic industrial research into workforce participation.

[106] Microsoft's incorporation of Internet Explorer into the Windows operating system, with its monopolistic market position, accounted in large part for this reversal of affairs. In a widely publicized trial brought by the U.S. Department of Justice, the practice was in fact found to violate anti-trust law.

shared asset…that lots of people cared about succeeding in the market" (Baker Remembering Mozilla's History), and thus to "validate the idea of open or sharing" as the basis of a great consumer product (Baker 2014).

In 2004, when Mozilla released the first version of the Firefox browser as an open source project, Internet Explorer controlled 90% of the browser market. Within Mozilla, Firefox's success between 2004 and 2009 is remembered as ushering in an age of renewed openness on the web. In such recollections, Firefox marked the return of choice to the browser market, and demonstrated the necessity of fair, open competition to browser quality and thus user sovereignty on the web. Through its adherence to and support for open web standards, Firefox further spurred development of the interoperable technologies necessary to preserve the openness of web access.

To appreciate the contours of openness as taken up Mozilla's mission, consider the remarks offered at the Decentralizing the Web Summit, held across two days in June 2016. Hosted by the Internet Archive in its San Francisco headquarters, a former church building described by executives as a "temple to knowledge," the Summit brought together internet and web engineers, activists, and scholars, and featured talks by both Vint Cerf, co-developer of the TCP/IP protocol, and Sir Tim Berners-Lee, the web's inventor.

In a keynote address, Berners-Lee described the logic behind the gathering. While it had never been clear that the web would "take off," Berners-Lee noted, when it did it was accompanied by a certain euphoria. Key to this sentiment was the web's openness, the fact that "anyone could make a website. Anyone could publish." It had been imagined, Berners-Lee said, that by enabling anyone to "put their ideas out there," the web "would completely change society. It would be a big leveler…[a] web of really intelligent discussion that could resolve the

problems of the day." In recent years, however, following the Snowden revelations and given the contemporary web's "siloization" by internet giants like Facebook, Berners-Lee said, "[t]here's an unease out there. Thinking, 'wait a minute, that utopian leveling of society and reinvention of systems of government and debate? What happened to that?' The people we hoped would be making their own websites? No, they're just on one big website."

If the gradual closure of the web described by Berners-Lee constituted the impetus for the Summit, in an opening address, Wendy Hanamura, an Internet Archive official, sketched the collective vision of openness shared by the Internet Archive and its cohosts, including Mozilla, the EFF, and the Wikimedia Foundation, developer of Wikipedia. This "conspiracy of open," as Hanamura deemed it, believes that "the web belongs to everyone and we have to keep it that way." If the web had become less and less open over the past 20 years, she explained, "locking the web open for good" would place it again "beyond the control of any entity." It would secure the web, rid it of censorship, and fulfill its promise of providing universal access to all knowledge.

Joining Hanamura on stage, Mitchell Baker, Mozilla's Chairman, described the technical and experiential traits that define the web in its ideal, open form. The web, she said, should be immediate and safe, accessible with nothing more than a URL. It should be open in the sense that anyone with a network connection should be able to publish content without permission or third party intervention. It should be universally available on any device, operating on any computing platform. Finally, something in the system should represent the end user, mediating the web experience on users' behalf such that content providers "don't have 100% of the power." Citing the way browsers enable users to "chang[e] font size, chang[e] the color of links, all the way up

to protecting from you from malware and DoNotTrack," Baker concluded that it is the browser that continues to fulfill this role on the web.

According to Kelty, the "revelatory" experience of a public emerging from the "arcane" practices of coding, reusing, and modifying software, is such that open source programmers tend to attempt to "port" it to other aspects of life (2008, 7-12). Indeed, while Mozilla continues to produce Firefox according to open source methods to which anyone can contribute,[107] Mozilla has also applied openness as a model to its general corporate operations. It is not only the Firefox code, which is open, in other words, but also the previously internal-only design decisions and the knowledge produced by its various professionals. While there are limits to this openness—employee personnel records, for example, remain confidential—it extends to Mozilla's administrative functions, including marketing. As a Firefox director described it in 2013 talk on Mozilla's engineering culture, "Because we have a mission we try to be very open….We try not to keep secrets, unless we have to. So, if you're having discussions on a mailing list, someone will often ask, 'Does this really need to be private? Can we put it in public so anyone interested can see why we're doing what we're doing?' It's important to get people outside Mozilla to participate, to contribute to our discussions and work."

As suggested here, Mozilla's staff traces the normative openness of their work to Firefox's origins as open source project but also to Mozilla's institutional commitment to an open web. Whatever its institutional value, however, Mozilla staff also valorize working in the open as

---

[107] As of 2017, Mozilla claimed to have approximately 10,000 global volunteers. While some portion of this volunteer base contribute code, others contribute to the mission by translating webpages, or by demoing Mozilla technologies and "evangelizing" on Mozilla's behalf at technology conferences.

a refreshing change from the general secretiveness of the tech industry and as a source of professional fulfillment. As a user researcher told me, in her prior jobs in the tech industry and financial services, she had been frustrated because her work "was very closed. Meaning my work was confidential. You kind of like "white label" everything to an extreme… I feel like it put a cap on the knowledge sharing I could do with other practitioners." Working in the open at Mozilla was thus appealing because it empowered her to actively participate with her peers "in the bigger UX community."

As illustrated by the concern expressed above about work being conducted via private email exchange, the relative openness of Mozilla's processes requires active maintenance, and is an object of frequent discussion. One problem, as the engineers describe it, is that Mozilla simply uses too many different communication tools. Paid staff themselves occasionally complain of the difficulty of identifying where precisely in the organization's communications systems "the conversation [is] happening." The expectation that team conversations be publicly memorialized and that their documents be publicly posted imposes on staff an ongoing obligation to determine how, relative to the ever-changing technologies and practices of the modern web, their work will be most accessible and inviting to non-employees.

I focus on the contours of openness at Mozilla here in part simply to explain its role in helping me sidestep the privacy-related norms and practices, which otherwise precluded access to the institutional sites of contemporary privacy world-making. Understanding the techno-moral value of openness in browser engineering, however, is also important for appreciating the nature of Mozilla's commitment to privacy. Attending to Mozilla's valorization of openness throws into

relief its overarching value structure, revealing its privacy commitments to be only one among multiple potentially competing components of Mozilla's techno-moral vision for the web.

Consider again in this regard Mitchell Baker's reference to DoNotTrack at the Decentralizing the Web Summit. DoNotTrack is a user-configured browser setting, which Firefox implemented in 2011. It was developed as part of an industry-wide effort to provide web users with a simple, enforceable means of opting out of online data collection.[108] Note how in Baker's treatment, however, DoNotTrack figures as one among many capabilities enabled by Firefox in its guise as a user agent. From the perspective of user sovereignty, privacy figures primarily as a configurable aspect of the user experience of the web, one Baker implicitly equates to other browser-mediated abilities, like controlling a website's font size or the color of its hyperlinks.

While I do not mean to suggest that someone like Mitchell Baker believes privacy to be morally equivalent to control over Firefox's visual interface, Mozilla's approach to privacy does define it and link it to organization's overarching techno-moral vision. Whatever the personal and professional commitments of Mozilla's individual staff to privacy, as an institutional value Mozilla has historically approached privacy as a means to the end of the open web. To appreciate this claim, let's turn to a privacy training session offered by Stacy Martin and Allison

---

[108] As a privacy tool, DoNotTrack (DNT) is both a technology and a policy. When enabled, it sends a machine-readable header, or signal, indicating to websites that the user does not wish to be tracked online. DNT was developed in part by a working group organized by the World Wide Web Foundation (the W3C) in 2011, which included both browser vendors and advertising companies. At the time, the Federal Trade Commission offered public support for DNT as a potential alternative to regulating online tracking. Following an acrimonious public process, in 2012, the W3C group failed to settle on the policy that websites would apply upon receiving the DNT signal. DNT is generally perceived within the web community to have largely failed to alter website tracking behavior.

Naaktgeboren, a Firefox Engineer, in early 2013. Held on the occasion of International Privacy Day, the training session represented an effort to provide Firefox engineers with a pragmatic, usable understanding of Mozilla's privacy principles. In introducing the session, Stacy and Ally observed that Firefox engineers occasionally think that because Mozilla employs privacy professionals, "I don't have to think about it." In truth, however, as a function of the generalized responsibility to uphold the Mozilla mission, privacy is everyone's responsibility. "We build it, we ship it," they said, "so we're responsible for making sure we're responsible to our principles."

As Stacy and Ally explained, while there are many ways to understand privacy, for Mozilla, privacy is defined by user control. Under the Manifesto, Firefox users must have the ability to "shape their own experiences on the internet. And nothing shapes user experience more today than data by and about you." The "good news," they offered, is that a matter of day-to-day work, Mozilla's privacy principles mostly require engineers to "write more stuff down and … make more proactive planning [decisions] during the [product] development lifecycle." To respect the principle of user control, for example, engineers must plan in advance of writing code whether any given new feature will generate data and if so what kinds. They need to decide who within Mozilla can access the data, where it will be stored, and who will be responsible for its safety. Additional planning is required to determine how users will interact with any data generated, whether they can export it or delete it, for example, and how they can otherwise control it.

To ensure that such planning also honors Mozilla's commitment to openness, Stacy and Ally explained, engineers should discuss the risks and benefits of important data related decisions in public forums like an email list serve. Openness further requires, however, that

engineers practice "active transparency," documenting tradeoffs and design decisions made with respect to user data. By contrast to the "passive transparency" engineers tend to fall into when "up to [our] elbows developing features," active transparency requires design documentation to be easily searchable and locatable. Active transparency may be harder to accomplish than passive, but is necessary. The teams who will eventually inherit any given project "can't read your mind or go back in time to figure out what you were thinking." Moreover, "[t]his is how Mozilla builds community. If you can't find the source code or documentation, it's not really open."

As articulated here, privacy for Mozilla has historically been, like the browser itself and many of its component technologies, dual use. On the one hand, as a means of providing individuals with culturally meaningful forms of control over personal data, privacy in theory directly empowers users. It aligns in this form with predominant legal and philosophical understandings of information privacy. On the other hand, Mozilla valorizes user control over data because it understands it to be an instance of the kinds of the generalized control over the user experience of the web required by its commitments to openness and user sovereignty. Here, privacy serves not so much to directly facilitate individual self-flourishing, as in liberal political philosophy, as to facilitate forms of freedom grounded in the techno-moral properties of the internet. Finally, in Stacy and Ally's concluding statement linking openness to community-building, privacy figures as a tool for cultivating the collective sociality and affective ties that empower not users but rather Mozilla itself as it pursues its mission to preserve the openness of the web.

This conceptualization of privacy as a tool of affective management and recruitment was relatively implicit in Ally and Stacy's training. On other occasions, however, Mozilla's privacy professionals explicitly embraced it, articulating it in the idiom of trust. On Data Privacy Day in 2017, for example, Peter, Mozilla's head of Trust and Security, published an article commemorating the occasion by distinguishing Mozilla's approach to trust from that generally employed in the tech industry. Because data is essential to most of the internet-based products and services we use, Peter wrote, trust is critical to the health of the modern internet. Trust, however, is hard to discern on the internet, and the entreaties of most internet companies for user trust are accompanied by the presumption that users 'not ask too many questions.' Mozilla, by contrast, seeks to cultivate trusting relations with Firefox users through privacy practices rooted in its open source culture. This approach, he explained, was why Mozilla not only provides settings and tools to directly control data collection, but also encourages users to ask questions, "[and] give[s] users tools to answer those questions." When developing new browser features, for example, just as Mozilla makes its code available, it also makes available the code for any analytics that it applies to collected data. It also provides clear notices and documentation of the practices internally governing the data's use. In this way, users are theoretically empowered to evaluate Mozilla's data practices and hold them accountable for any mistakes. From this perspective, privacy names a strategy for affectively calibrating users relations with Mozilla and its technologies, cultivating the trust understood to be necessary for users to rely on Firefox as pursues its vision for the open web.

Shipping Privacy

In 2016, in the early months of my research in San Francisco, as I explored the Bay Area's vibrant, privacy-oriented public life, I met with Mozilla's privacy managers to determine the ultimate shape of my fieldwork. While I had previously been granted permission to conduct research at Mozilla, when I arrived in the fall I was greeted with confusion. The other members of Stacy Martin's public policy team, for example, wondered whether I was coming in as a new hire, and what precisely I would be doing. A manager on the privacy and security engineering team suggested I would produce "less friction" navigating the organization if I secured an institutional designation—as an intern, for example, or a contractor—that would make me recognizable as "an insider."

To resolve the impasse over my status, I met with Firefox's lead product counsel. Mozilla was happy accommodate my research, she told me, but wanted to minimize any potential disruption caused by my presence. Rather than having unfettered access to Mozilla's offices, therefore, I would be granted status equivalent to one of Firefox's trusted open source contributors. Such "vouched Mozillians" gain privileged access to Mozilla meetings and technical systems in order to support and facilitate their valued contributions. After selecting projects around which to focus my research, I would be assigned a contact for each who would serve in effect as my project manager, helping me establish relationships and resolve any further obstacles I encountered.

To help me identify the best projects on which to focus, I was directed to Peter, the author of the article on trust described in the preceding section. A lawyer by training, Peter, I was told, was the officer with primary responsibility for implementing Mozilla's privacy commitments in

its products. Given his role in managing any tensions between Mozilla's product development and its privacy principles, Peter held a privileged perspective on privacy's varied appearance as challenge and opportunity across Mozilla's work.

When I met with Peter, he explained to me that Mozilla's work generally implicates privacy issues in one of two ways. On the one hand, Mozilla develops features, which directly empower Firefox users to protect their privacy online. On the other hand, Mozilla develops features that do not offer privacy-related benefits, but may nonetheless implicate Mozilla's privacy commitments. This happens, for example, when engineers conclude that offering the feature will require Firefox to collect either new forms of personal data or already-collected forms of data, but on some newly expanded scale.

Both scenarios, Peter explained, raise unique organizations issues for Mozilla. Historically, he continued, Mozilla interpreted its commitment to privacy to entirely preclude Firefox from collecting user data. For much of its history, therefore, Mozilla had no ability to analyze how people actually use the browser, no insight even into which features they use most frequently. Mozilla had eventually determined that some data collection was necessary to produce competitive, empowering web technologies. Nonetheless, Peter said, collecting user data necessarily increases the risks facing Mozilla. Any time an engineer proposes collecting data, therefore, Mozilla must work to ensure that only the absolute minimum necessary is collected. It must also effectively communicate to users why Firefox seeks to increase its data collection, and what benefit or value will accrue to them as a result. For privacy features like Containers, meanwhile, Mozilla faces the challenge of "balancing tricky tradeoffs between user experience, engineering, and product." In determining how to make privacy technologically, and in what

form to “surface” it to users, Mozilla must balance what it “wants to accomplish” with its commitment to improving privacy.

Though I did not appreciate it at the time, Mozilla in 2016 and 2017 was undergoing a period of significant organizational change, a shift not in its mission but in the strategy for achieving it. As I would learn, this reorganization implicated the nature of Mozilla’s institutional approach to privacy, a reconceptualization of how in protecting privacy Mozilla contributes to the open web.

I first became aware of these changes in part through conversations with Karen, Mozilla’s Senior Privacy Engineering Manager. These meetings were, like the majority of my interactions at Mozilla, conducted remotely via Mozilla’s video conferencing software. Though Karen lived in nearby Mountain View, CA, where Mozilla’s headquarters were then located, she was used to working remotely. At the time, approximately half of the Mozilla Corporation’s 1200 employees did so. Hiring talented engineers uninterested in moving to the Bay Area, I was told, was another way that Mozilla seeks to compete “asymmetrically” with its better funded rivals. With primary responsibility for managing the engineers working on privacy and security-related features, Karen’s supervisees were scattered not just between San Francisco and Mountain View, but also in Tulsa, Portland, Oregon, and Taipei, and across Canada and Germany.

When we first met in March 2017, Karen was only six weeks into her job at Mozilla. Previously, she had worked as a project manager in various divisions of another major Silicon Valley corporation. With no significant professional background in either privacy or open source software, Karen was during this period both acclimating to and attempting to shift Mozilla’s non-hierarchical corporate structure.

Like Stacy Martin, Karen described being drawn to privacy in terms of the intellectual and professional challenges it offered. “Personally,” she told me, “I want to be in control of my information, what I want to share. With people I know, I’m very open. I share pretty much all the information. But I don’t think anyone should have the right to see what I’m doing, how my point of view is, without my permission.” Still, in taking her job with Mozilla, Karen had primarily been drawn to its non-profit status and commitment to the open web. Having recently completed a major project for her previous employer, Karen had been interested in new professional challenges when Mozilla reached out to gauge interest in a job. “I had worked on portions of encryption and decryption systems for a satellite,” Karen said, “but it’s not my core competence. But to me it’s interesting, especially with the new [Trump] administration, and with so much data in the cloud space. And we use the browser every day. So much of our personal information is being delivered and transmitted through that. I find that very interesting.”

According to Karen, it was precisely the ways in which her background diverged from the average Mozilla hire, which attracted the company to her. In particular, Karen pointed to her prior success developing and shipping new features for both startups and mature companies. She had, she observed, demonstrated the ability to work with engineers and researchers in developing new features from scratch. “So, I think a lot of it,” she continued, “is due to the fact that I demonstrated the ability to work in an environment that maybe is a little chaotic, and am able to lead.” Mozilla’s engineers were, from her perspective, like those of her old employer in being “very technically driven. They look at it, ‘I’m going to build that.’” To ship features in such an environment required, in Karen’s experience, an ability “to harness that,” to manage and support project engineers while persuading them to work towards the goal of shipping.

If Mozilla was drawn to such skills at the time, it was in part because of the backlog of privacy features residing in its code repositories. There were, per Karen, a "good number" of privacy features that had been developed "many cycles ago" but never "turned on." Such features had been implemented in Firefox's code base but remained "behind preferences," accessible only to the technically savvy expert users motivated and capable of finding and configuring the browser's advanced settings.

For Karen, the existence of this backlog followed in part from the historical absence of management process around privacy. During her initial months at Mozilla, as Karen sought to build relationships with her engineers, she encountered a plethora of privacy projects but an absence of clear guidance on how to push them through to completion. "When I first came on," she told me, "it seemed very chaotic. … There were so many projects. Everyone seemed free to do what they want to do." What was missing, she concluded, were guidelines regarding how Mozilla's privacy engineers should, individually and collectively, prioritize projects to make effective use of their limited resources. Karen devoted her first six months at Mozilla to identifying top privacy and security priorities and to implementing a methodology under which such projects would "succeed fast or fail fast."

Still, if such managerial concerns might apply to any team at Mozilla, privacy features, Karen acknowledged, face unique hurdles. "It's historically extremely difficult to ship privacy and security features," Karen told me, "because first, users don't understand what we are trying to protect them from. They don't know the data is being collected." Mozilla's "high motivation," expert users might "understand everything about privacy and security" and be willing and able to configure the browsers' privacy settings accordingly. The majority of Firefox users, however,

don't understand privacy features, they aren't "going to understand the privacy significance or security significance of the features they use." As for "low motivation" users, they might "care about privacy, sort of, but they are not going to spend a second or two to use it." They might, for example, use the browser in "unsafe ways." If a privacy-preserving feature were then to prevent or otherwise interrupt such use, "to a lot of users its added inconvenience, so they would avoid using it." Given such obstacles, to clear Mozilla's backlog of privacy features, Karen concluded, required investing additional resources to "reach" users, to make privacy "very layman-term easy to understand," not necessarily in terms of a feature's technical details, but certainly in terms of ease of use and "at least understanding why we're doing it. How they can benefit."

If this first category of privacy feature, including Containers, was haunted by the question, "do people really want to use it," for a second category of feature there was clear consumer demand, but other complicating risks. The desire among browser users for an ad blocking feature, for example, had long been demonstrated by the overwhelming popularity of ad blockers as third party browser add-ons. Some years prior, Karen told me, Mozilla had in fact developed an ad blocker. Like other privacy features, however, Mozilla had never turned it on by default. Doing so would undoubtedly provide a measure of enhanced privacy to users. Mozilla executives worried, however, that it would alienate website authors, many of whom rely on digital ads for funding. One might wonder why an organization committed to user privacy and sovereignty would concern itself with the viability of ad funded businesses. But from Mozilla's perspective the products and services funded through ad revenue are a key element of what makes the web so rewarding and enriching for users. Principle 9 of the Manifesto thus states,

"Commercial involvement in the development of the internet brings many benefits; a balance between commercial profit and public benefit is critical."

In our conversations, Karen and I primarily discussed the changes she was trying to institute in the process by which potential privacy features moved from experiment to product. Her very hiring, however, with its mandate to ship privacy features, reflected broader ongoing changes at Mozilla. In recent years, Mozilla had grown significantly in size, more than doubling its number of employees. For some long term engineers, this growth coincided with Mozilla becoming "less hackery" and open and more bureaucratic and closed. Following the recent failure of a high-profile effort to push into new product categories, Mozilla was, moreover, redefining its corporate priorities. Between 2011 and 2016, Mozilla devoted significant resources to building an open source alternative to Google's Android, the dominant mobile operating system. If successful, executives argued, Mozilla would ensure the viability of the open web values embedded in Firefox, even as global internet traffic continued its rapid shift away from personal computers to mobile smartphones.[109] When Mozilla abandoned the project in 2016, the web community widely viewed Mozilla as having allowed its core product, Firefox, to fall into technical disrepair, and Firefox's market share had fallen into the single digits. Between 2016 and 2018, Mozilla thus refocused its resources around Firefox, working first to improve and

---

[109] The general dynamic at play here, the desire on Mozilla's part not to 'miss' the next wave of commercial computing, repeated itself throughout my fieldwork. For Mozilla, such shifts represent both threats to the browser's role in mediating the user experience of the internet and opportunities to embed Mozilla's techno-moral principles in new networked technologies, preserving Mozilla's values if not the open web itself. Prior to my fieldwork, Mozilla sought to influence the shift from personal to mobile computing with its mobile operating system project. During my fieldwork, Mozilla initiated efforts to influence the spread of networked computing in everyday consumer devices (i.e., the so-called Internet of Things), and, following the launch of Amazon Alexa and Google Home, in the shift towards voice command as a computing platform.

modernize its technical underpinnings (see Chapter 5) and then to use new browser features to differentiate Firefox from its competitors. As Karen explained to me, Mozilla leadership had identified privacy as a product domain in which it held a potentially defensible advantage over rival browser vendors, a capacity to offer desirable features that its competitors would be unwilling or unable to. The mandate to ship new privacy features thus constituted an effort to recapture market share by "building on [Firefox's] brand as the browser that supports privacy and security."

While there is a certain counterintuitive quality to an organization controlled by a non-profit thus finding itself compelled to cater to market forces, in fact, it is in terms of Firefox's market share that executives articulate Mozilla's ability to have "impact" in the world, to actualize the Manifesto. Consider in this regard, a speech delivered by Mitchell Baker at a 2014 All Hands meeting. In her talk, Baker addressed the general topic of decision-making at Mozilla. She questioned how well the decision-making framework implemented during Mozilla's early days as an open source project continued to serve its goals as Mozilla had grown, changed, and become more organized.

As Baker observed, one area in which Mozilla's decision-making frequently came into question was around the issue of free and open source software. Mozilla had an important role in the free and open source ecosystem, she continued, but unlike some organizations, its goal is not to create as much open source software as possible. "The goal of Mozilla is to build an Internet that is a global public resource, open and accessible to all…. [to] change the nature of the internet." Given this goal and Mozilla's comparatively limited resources, Baker argued, Mozilla could not afford to make decisions "based solely on open web choices and how the web should

be." Instead, Mozilla has to make decisions "based on where we can change things," on which options will have the most "impact" on "build[ing] the traits of the web into the future." In the area of product, Mozilla often faced painful choices "where we… can imagine something that is perfect for" the open web. Such products represent a dangerous temptation, however, because they align with "how the web should be," regardless of whether or not they hold any consumer appeal. "But Mozilla has impact by being in the market, by being successful. That means this axis of consumer appeal is critical. …Consumers using it, consumers wanting it, gives us the leverage to build the internet we think the world should have."

In a 2014 town hall address on Firefox's annual goals, Jonathan Nightingale, Vice President of Firefox, elaborated on the ways that market position is understood to contribute to Mozilla's ability to effectuate its vision. In the talk, Nightingale discussed the recent trajectory of the top line user metrics tracked by Mozilla, including the total number of daily Firefox users and the total number of monthly hours of Firefox usage. While the precise numbers are not important for our purposes, what is illuminating is the justification Nightingale provided for tracking these metrics:

> Usage is important. If you launch the browser and don't do much with it, it's not clear how much we can influence your experience of the web, or the view the billions of websites and developers out there see in the mix of browsers. The more we are used, the more we can take care of you, the more influence we can exert on your experience, and the more we show up in the broader web ecosystem.

As Nightingale indicates here, usage empowers Mozilla to advance its mission at multiple social scales simultaneously. At the level of individual users, usage matters because Mozilla's engineers understanding themselves to be building the techno-moral values of the Mozilla mission into Firefox technologies. Only when individuals use Firefox, therefore, can they directly reap the

benefits to their experience of the web that flow from such mission-oriented design. At the same time, only by demonstrating to peers, competitors, and partners the existence of meaningfully large numbers of Firefox users can Mozilla shape the broader trajectory of the web. Mozilla's executives describe market share as necessary, for example, to induce website authors to build websites that are compatible with Firefox. In turn, only if a preponderance of popular websites work as well on Firefox as they do on Chrome or Safari can Mozilla continue to introject choice into the browser market. As I was told by privacy experts on Mozilla's public policy team, usage also grounds Mozilla's effectiveness in advocating for open web principles in the web's standards settings organizations and with policymakers. As a Firefox director put it, to "take a stand on some principle that causes users to abandon Firefox and our other products" means we "won't be able to get anything good done ever again."

Recognizing that Mozilla pursues its mission through competition in the browser market is important for understanding the state of privacy under conditions of technological stewardship. This is so, in part, because the nature of browser competition, as understood at Mozilla, imposes a political economy of attention on the browser. In a 2014 town hall on the state of Mozilla's market share, Mozilla executives described browser competition as generally unfolding along dimensions of price, placement, promotion, and product. No vendor, however, has charged for a browser since the 1990s, making all equivalent on price. Meanwhile placement, or distribution channel, has always presented problems for Mozilla. "It used to be," an executive explained, "that Internet Explorer had distribution through Windows, Safari through Apple, and we had none at all." "Now," given the connections of Chrome, the market leading browser, due to its to

Google's suite of popular products, "it's worse." Mozilla faced similar asymmetries in the domain of promotion, or marketing.

This left product as the sole dimension of competition practically available to Mozilla. And here, Mozilla faces the challenge that the average user views browsers as being effectively interchangeable. "So, to compete on product," the executive continued, "we have to find ways to be different that competitors can't chase us on." By 2016, Mozilla's product strategy explicitly contemplated attracting new users by prioritizing the development of such "Uniquely Firefox," features, features that help users "shape and control" the "deeply personal" experience of browsing the web and thus reflect "the reasons users chose us in the first place." As an existing strength of the Firefox brand, and an area where Mozilla, as a not-for-profit, was perceived to have room to offer products that its for-profit competitors would not, privacy was foremost among these prioritized feature domains.

As illustrated by Karen's hiring, and her mandate to "unblock" the development of long-gestating privacy features, the identification of privacy as a promising basis of differentiation on the browser market lent new institutional energy to Mozilla's long-standing formal commitment to privacy. It also shifted, however, the ways in which Mozilla effectively articulates privacy's value. Historically, Mozilla's privacy-enhancing tools directly empowered users in the sense of providing them with means to exercise control over the personal data generated online. They also, however, served as a tool of affective calibration, one deployed in service of preserving the internet's viability as a medium of human flourishing. As it emerged from and was articulated with respect to Mozilla's open source processes, privacy served to sooth the fears and anxieties associated with using internet devices and services, and by developing trust, alleviate users of

some of the cognitive burdens associated with privacy and security choices. As an object of prioritized feature development, privacy was now imbued with the perceived potential to generate not just trust and comfort, but desire, the kind of affective pull that might distinguish Mozilla from other browser developers and attract new users.

Even so, the aspiration to ship more privacy's features subjected Mozilla's privacy technologies to the aesthetic and market standards according to which which it allocates precious feature space in the browser. Alongside its "Uniquely Firefox" strategy, Mozilla now aspired to "Uncompromised Quality" in its features, applying a form of market-adjacent institutional self-discipline to the determination of the changing suite of features that constitute the browser. As Firefox's Director of Engineering explained in an email to web developers, any feature in the browser should either be "great or dead." Every feature, in other words, should be "polished, functional, and a joy to use." If after "spending time to make [a feature] great," Mozilla engineers find they "can't get it to that state, we shouldn't do it at all." Any features already included in the browser and surfaced to users, but which failed to meet these standards, should be removed. As an engineer elaborated in an onboarding session for new hires, exercising such economy is necessary because it "improves the quality of the rest of the product by having more focus on that."

CHAPTER 4: SIN, INTIMACY AND DISAVOWAL IN THE INTERNET'S BUSINESS MODEL

In August 2014, Ethan Zuckerman, a well-known media scholar, marked the occasion of the web's 25th anniversary by inaugurating a new genre of public confession (Zuckerman 2014). Such *mea culpas* for internet technologies gone wrong would soon become a regular feature of public discourse. As the Director of MIT's Center for Civic Media, Zuckerman had long passionately promoted internet technologies as tools of civic empowerment. Given his interests, Zuckerman's decision to explore the rising public disenchantment with the internet was hardly surprising. What was surprising—and briefly made Zuckerman the target of both light-hearted public ribbing and anonymous death threats—was the revelation around which Zuckerman organized his essay: In 1994, Zuckerman had helped initiate the web's transformation into a hellscape of distraction, harassment, and misinformation by creating, in his own words, "one of the most hated tools in the advertiser's toolkit," the pop-up ad.

At the time, Zuckerman was a recent college grad working as the webmaster for Tripod, Inc., a Massachusetts-based startup. Tripod had been founded to produce an online magazine for recent college grads, which offered practical advice on the challenges of transitioning to adulthood. The magazine, however, was losing money. Between 1994 and 1999, like other startups rushing to unlock the web's commercial potential, Tripod thus cycled through dozens of new potential products and revenue models. It finally found itself with a hit when, during a late-night work session, a colleague of Zuckerman's hacked together an easy-to-use webpage builder. At the time, most people still weren't on the internet. Many who were hadn't yet learned to

navigate to websites beyond the "walled gardens" of access providers like AOL. Tripod's new tool allowed even people with no technical expertise to build and publish their own webpages. For a certain generation, Tripod soon came to define the experience, aesthetics, and ethos of the early web alongside other early hosting providers, like Geocities and Lycos.

Like many of the corporations responsible for commercializing and popularizing the web, Tripod adopted surveillance-based advertising as its business model for a combination of ideological and pragmatic reasons. Ideologically, Tripod believed that offering its web hosting service at no monetary cost would provide the greatest number of people with the opportunity to "express themselves and be heard" on a global scale. In so doing, Tripod would thus help fulfill the original techno-moral vision of the web as a universal information commons, free and open to all (see Berners-Lee et al. 1999). Pragmatically, public anxiety about the safety of using credit cards online, and a lack of necessary enabling technologies, limited Tripod's ability to actually charge users for its service. Given such cultural, technological, and economic realities, Tripod concluded that advertising was its best bet: "At the end of the day," Zuckerman writes, "the business model that got us funded was advertising. The model that got us acquired[110] was analyzing users' personal homepages so we could better target ads to them."

According to Zuckerman, Tripod's adoption of an advertising-based business model quickly unleashed a corrosive influence, which continues to eat at the web today. The immediate vector of this corruption was the few lines of code that Zuckerman wrote to square his techno-moral aspirations with the promissory visions Tripod had conjured to secure investor financing. Functionally, this code relocated Tripod's advertising space to a new window, which 'popped-up'

---

[110] In 1998, Lycos, an internet search provider, purchased Tripod for $58 million.

in front of the active browser window whenever a user visited a Tripod webpage. Effectively, the pop-up provided wary corporate advertisers with a form of visual deniability. It assured that they would not be seen as endorsing the content of Tripod webpages, without requiring Tripod to censor the kinds of content that users could post.

If Zuckerman's code was the vector, the corruption it carried initially expressed itself as an escalating assault on the senses. At first pleased with the elegance of his solution, Zuckerman soon saw the pop-up mutate. Web companies rushed to adopt the pop-up, drawn by the way—unlike the web's then-dominant banner ad format—it claimed the user's field of vision. As users learned to ignore pop-ups, versions quickly appeared, which actively resisted any attempt to minimize or close them. Pop-unders appeared, for example, lurking in wait behind active browser windows to grab users by surprise. With their implicit imprimatur as tools of legitimate business, pop-ups became a favored tool of cybercriminals, an easy means of surreptitiously capturing credit card and other personal information. Whatever techno-moral aspirations inspired their creation, under the competitive pressures of the nascent online advertising market, pop-ups became an undeniable, relentless source of frustration, grief, and fraud on the web.

According to Zuckerman, by the time the major web browser vendors began automatically blocking pop-ups in early 2000s, it was too late. The pop-up had already helped prove the viability of the free-to-use, surveillance-driven, advertising business model, which became the default means of supporting the most popular websites and web-based applications. It was the business model, in turn, which inured the American public to the idea that highly detailed, deeply intimate data—privacy—was a reasonable price to pay for "free" services like email and social media. On this basis, Zuckerman argued in 2014 that the business model was

"the Internet's original sin," the root cause of all internet evil. This evil continues to register most directly for web users through the visceral frustrations and disturbances imposed by ever-mutating ad formats. For Zuckerman, however, such disturbances also index the broader set of individual and societal harms, which by 2014 were increasingly recognized to flow from the ever-intensifying surveillance and targeting that form the foundation of the internet economy.

A New Explanatory Logic for Privacy's Decline

Popular culture has long understood privacy's fate to be intertwined with technological progress. Historically, when privacy has appeared in the national public imaginary, it has most often been as an endangered object of care and concern that is threatened by the impact of new technologies on culture and law (see Nissenbaum 2010). With predictable regularity, Americans have responded to the introduction of new technological systems, from photography, to the telegraph and telephone, as alarming and potentially fatal threats to privacy (Nelson 2001).

The ongoing revolution in computer-based information technology has been no exception. Since the 1960s, public discourse has anxiously catalogued the ways the personal computer, massive databases, and the internet greatly expand and intensify the ability of private and public entities to track, monitor, aggregate, analyze and disseminate personal information (see Brin 1998). Given the popular understanding of privacy as a perpetual casualty of technological progress, Zuckerman's diagnosis of the contemporary threat to privacy—his deferral of its locus from technology and technological design to their enabling business model —is noteworthy.

Certainly, technology has never been alone in the rogue's gallery of American privacy. A full accounting of privacy's foes would have to include the institutions of American policing and national security, for example (see Browne 2015; Masco 2017). The supposed moral failings of the technology-consuming public—the public's supposed refusal or inability to take personal responsibility for its own privacy—remain a constant source of vexation for privacy professionals. Yet in diagnosing the basis of the entwined decline of privacy and the web, Zuckerman rejects these and other culturally-legible candidates. Where Zuckerman might have directly indicted technology corporations or their famed inventor-entrepreneurs, he instead insists that "[t]he Internet doesn't spy on us because Zuckerberg, Brin and Page are scheming, sinister masterminds." His diagnosis gestures toward an indictment of capitalism, but in truncated form, limited to a particular strategy adopted by a particular class of corporation to ensure its social reproduction and flourishing in the face of changing conditions of competition and commercial possibility.

Zuckerman's diagnosis stands out not only in relation to cultural understandings of privacy, but also to the American history of corporate malfeasance. For decades, cigarette makers suppressed public knowledge of the link between cigarettes and cancer, but no one blamed the cigarette industry's "business model" for the millions of lung cancer deaths annually caused by smoking. Nor have people generally rationalized the contributions of the fossil fuel industry to global climate change in terms of Big Oil's business model. In 2015, when Volkswagen was caught using software tricks to mislead regulators about the true emission levels of its diesel vehicles, public blame settled on VW's autocratic corporate culture, not its business model.

Nonetheless, between 2014 and 2017, I frequently encountered privacy-minded technologists, activists, and academics in the San Francisco Bay Area and beyond who mobilized logic like Zuckerman's to make sense of privacy's precarious state in the internet age. In a wide variety of contexts, ranging from blogs, to Tweets, everyday conversations, and TED talks, technologists and other Silicon Valley insiders argued that the contemporary state of privacy had to be understood in terms of the fundamental opposition of the internet's business model to privacy (see e.g., Schneier 2013).

As Zuckerman's article illustrates, the technical community's long-simmering concerns with the internet's business model are not limited to its effects on privacy. They rather encompass a panoply of individual and societal harms, from online misinformation, to harassment, distraction and attentional hijacking, and other forms of emotional and behavioral manipulation. Following the 2016 election, such concerns began to seep from the expert communities in which they had gestated into the general public sphere, prompting the technical community and policymakers to grapple in newly public ways with the unintended social effects of social media and other internet services. New and surprising actors including prominent Silicon Valley investors and entrepreneurs followed Zuckerman's precedent, issuing mea culpas for their own unintentional roles in spreading surveillance-based advertising (see e.g., McNamee 2019). They amplified the technical community's warning that tech industry apologies and promises to correct past mistakes would come to naught unless a new business model was found for the internet.

As evidenced by its increasingly common appearance everywhere from the national media to congressional hearings, the business model has become part of the general social

coding of the internet and internet corporations. It now provides a readily available grammar for identifying the origins of the social ills increasingly recognized to propagate through the internet and for critically parsing Silicon Valley's products and public statements. The cultural proliferation of the business model concept as a necessary rubric through which to understand privacy, the internet, corporations, and the relationships between them bears explanation. What exactly is the internet's business model? How do technologists understand it to explain the diminished state of privacy in the internet age? How should we make sense of its growing salience and explanatory power in expert and popular culture?

"Surveillance is The Internet's Business Model"[111]

The internet's business model, and its connection to privacy, are easy enough to understand in the abstract. Generally, when technologists like Zuckerman talk about the internet's business model they have in mind a simplified scenario like the following: Business A provides content (e.g., news, information, videos) or services (e.g., search, storage, email, social networking) to internet users at no monetary cost. Simultaneously, Business A charges Business B to serve advertisements to the users attracted to Business A's content or services, at least in part by their gift-like appearance.

In certain respects, the arrangement is not dissimilar to the historical practice of newspaper and magazine publishing. Publishers have long used advertising revenue to supplement subscription and sales revenue. Where technologists understand the internet's

---

[111] This heading is taken from a conference speech delivered in 2014 by the influential computer security expert Bruce Schneier.

business model to differ is in its third term. To sustain its relationship with advertisers and compete against traditional print and broadcast media, Business A surveils its users. More precisely, it tracks user browsing habits (what sites users visit, what they look at, what they post and purchase), aggregates it with data purchased from third party data brokers, and analyzes the aggregated data for insights into users' preferences, habits, relationships, and emotional states. It is on the basis of such surveillance that Business A is able to offer Business B the promise of highly precise, measurably effective consumer targeting.

As one prominent critic of the internet's business model puts it, before the internet advertisers "lived in the Dark Ages " (Ceglowski 2014). They might try to direct their ads to particular categories of consumers based on a publication's subscriber demographics. Such efforts, however, were inevitably both over- and under-inclusive. They also provided no means for determining whether an ad was ever actually seen. The historical arrival of the internet was thus "Christmas for advertisers. Suddenly you could know exactly who was looking at your ads, and you could target them by age, sex, income, location, almost any criterion you wanted" (Ibid.).

In the early 2000s, the promise of reliably reaching the consumers most likely to purchase one's goods, just when they were most likely to do so, convinced advertisers to begin shifting their spending from traditional media to the internet. Following the 2000 crash in the public market for technology stocks, this promise effectively brought Silicon Valley back from the dead, propelling it to new heights of wealth and cultural influence. By 2018, Facebook and Google had leveraged the promise of targeted advertising to build the internet's dominant advertising businesses, together accounting for roughly 60% of total U.S. digital ad revenue. In

one form or another, however, the model accounted for the vast majority of the annual revenue generated by both Silicon Valley's iconic internet startups —the Twitters, the LinkedIns— and the sprawling ecosystem of individually-maintained blogs and websites that give the web its breadth and character. Both defenders and critics insisted that the web could no more exist as we know it without the business model than it could without the physical infrastructure that comprises the internet.

When called to account for their privacy practices over the years, tech companies have been at pains to insist that they don't actually "sell" user data. Be that as it may, Silicon Valley professionals generally define privacy in terms of individual control over the dissemination of personal information—control over who can know what about us, under what circumstances. By their own definition then, to the extent internet companies sustain themselves by selling access to deeply invasive personal inferences derived by surveilling users, their survival places them in ongoing tension with privacy.

Laying out the internet's business model in the abstract helps us to appreciate the fundamental nature of the antagonism that technologists argue it directs towards privacy. In this regard, the explanatory mobilization of the business model serves as a corrective to years of public discourse, which at least implied that unauthorized data access—hacks, database breaches, intercepted internet traffic—constituted the internet's primary threat to privacy. In the years leading up to the 2013-2014 Snowden revelations, for example, breathless news reports about the growing scale of data breaches punctuated the public sphere with increasing regularity: 22 security clearance forms stolen! 110 millions credit cards exposed! Such reporting introduced the public to the fact that a bewildering array of corporations and government agencies were quietly

accumulating stockpiles of personal information. But it also lent itself to the misleading impression that criminals and other indisputably bad actors represented the primary threat to privacy introduced by the internet. It further lent credence to the idea that the key to preserving privacy lay in simply doing computer security better.

By contrast, when technologists blame the business model for the internet's threat to privacy, they communicate that this threat is, in the engineering parlance, a feature, not a bug. That is, it is the web's normal functioning as a socio-technical system, the use of the corporate content and services that dominate the popular experience and understanding of the web as they were designed to be used, which imperils privacy. This re-articulation of the nature of the threat has potentially important implications for the defense of privacy, suggesting that technological interventions, such as the universal adoption of encryption on the web, which don't in some way address the stranglehold of the internet's business model, are ultimately insufficient to the task. And indeed, over the course of my fieldwork, I repeatedly found that the engineers, lawyers, and activists behind the privacy-enhancing technologies that piqued my curiosity actually designed them as technology-based efforts to shift the internet's economic foundations away from a reflexive reliance on surveillance and targeted advertising (see Chapter 1).

Believing in Free Lunch

Whatever insight we may glean from examining the internet's business model in the abstract, it doesn't take us very far in understanding the concept's growing cultural salience and explanatory power. What does the business model logic add to other compatible explanations of the internet's antagonism towards privacy? What kind of cultural work do technologists and others accomplish

when they mobilize it? What does it tell us that we need a special term to make culturally legible the facts that internet companies make money and that they do so using surveillance?

In regards to these questions, consider the extraordinary exchange which took place between Senator Orrin Hatch and Facebook CEO Mark Zuckerberg during the 2018 Senate hearings on social media, privacy, and data abuse. Congress convened the hearings in the wake of reporting that in 2016 Cambridge Analytica, a political consulting firm, used data surreptitiously harvested from more than 87 million Facebook users to influence voters on behalf of the Trump election and Brexit Vote Leave campaigns. Cambridge Analytica's CEO claimed after the election that his firm's work had been instrumental to Trump's surprise victory. The claim was never corroborated, but it resonated with outstanding fears that Russian agents were exploiting internet technologies to compromise America's electoral processes.

The news that Facebook had for years allowed third party developers to collect data on anyone in a user's friend network exacerbated outstanding critiques of Facebook's cavalier approach to privacy. Along with the immediate aftermath of the 2016 election, the 2018 Cambridge Analytica scandal marked a major moment in the penetration of the business model's explanatory logic into public discourse. The technologist and sociologist Zeynep Tufecki (2018), for example, took the opportunity to argue that it was Facebook's business model and not its technical systems that served as the common denominator uniting Cambridge Analytica and Russian manipulation with Facebook's data practices. From the perspective of the business model, she observed, there was no meaningful difference between a corporate ad buyer, a legitimate political campaign, and a malicious foreign agent. Each was free to use the largely

unsupervised tracking and advertising systems demanded by the model's global aspirations to 'sell' whatever it saw fit, whether shoes, a healthcare reform proposal, or social discord.

During the hearing, Zuckerberg testified that Facebook's mission—to connect everyone in the world and bring the world closer together—required that it offer a service that "anyone could afford." "Well, if so," Senator Hatch asked, "how do you sustain a business model in which users don't pay for your service?" "Senator," Zuckerberg replied following a brief pause, "we run ads."

Following the hearing, clips of this exchange circulated widely on social media and in the news, eliciting mockery of Senator Hatch and quizzical appraisal of Zuckerberg. The poor grasp of social media's technical operation displayed by many senators during the hearing had been expected. Not long had passed, after all, since senators regularly boasted about never having used email. But here was apparent evidence that the lawmakers charged with regulating one of the nation's most powerful industries didn't even understand its basic economic premise.

In fact, as the broader context of the hearing makes clear, Hatch did understand. In opening comments, Hatch dismissed the very possibility that users could fail to appreciate how they 'paid' for free services like Facebook:

> The recent stories about Cambridge Analytica and data mining on social media … touch on the very foundation of the internet economy and the way the websites that drive our economy make money. Some have professed themselves shocked—shocked! that companies like Facebook and Google share data with advertisers. Did any of these individuals ever stop to ask themselves why Facebook and Google…don't charge for access? Nothing in life is free. Everything involves tradeoffs. If you want something without having to pay money for it, you're going to have to pay for it in some other way, it seems to me.… And these great websites that don't charge for access, they extract value in some other way. And there's nothing wrong with that as long as they're you're upfront about what they're doing.

In a culture suffused and structured by the norms and assumptions of the market, the metaphysics of value Hatch employs here are so widely held as to generally defy scrutiny. In its extreme form, this metaphysics reduces all interaction to exchange and all exchange to market exchange, i.e. to the agonistic transactional pursuit of self-interest under conditions of scarcity (Davis 1996).[112] Despite plentiful evidence in the social scientific archive contravening Hatch's totalizing, zero-sum vision (see Graeber 2001; Zelizer 2005; Mauss 1990), it continues to wield the authority of the common sense.

Nonetheless, we should not dismiss the plausibility of the charge leveled against Hatch. We should rather take it as an index of the sociologically-noteworthy fog of uncertainty that surrounds the economic bargain entailed by internet services. Doing so reveals this uncertainty it to be in part a product of historical contingency. That is, it reveals the internet's iconic websites and services to have developed under conditions, which lent cultural plausibility to the existence of a class of corporation characterized precisely by the delivery of something-for-nothing. These conditions included, for example, the fact that both the internet and the world wide web were originally designed for use not by private individuals or businesses but by scientists, universities, and other research institutions. Until the early 1990s, businesses were in fact legally prohibited from offering commercial services on the internet. Even following legalization, whether and how companies would actually make money online remained open questions subject to significant public debate (see e.g., O'Reilly 1996).

---

[112] See Zelizer (2010) describing the economic ideology which reduces caring, friendship, family, etc., to "nothing but" advantage-seeking individual choice.

Moreover, during the period when the internet and web entered public imaginaries, they did so invested by early evangelizers with countercultural, utopian, and libertarian commitments (see e.g., Barlow 1996). These included the notion, inherited from mid-20th century debates about science and society, and amplified by the free software movement, that digital information should be freely accessible (Johns 2009). In line with this ethos, in the early years of the commercial web the public showed itself willing to pay for internet access, but generally resisted paying for website content outside of certain specialty categories, including porn.

The web's subsequent transformation from an arcane tool of scientific research into the mediating substrate of everyday life was marked by the appearance of a class of startup, which appeared to flout the basic physics of the market. Such startups offered content and services over the web at no monetary cost to consumers, and often with no proven plan for generating revenue whatsoever. Based on the speculative promise of the global audiences accessible through the web, they nonetheless enjoyed spectacular success on the venture capital and public markets. Prominent examples of this class of startup included Mozilla's predecessor, Netscape. As success stories like the Netscape IPO accumulated, economists, journalists, and corporate consultants concluded that the so-called "new economy" inaugurated by globalization and the internet had not just changed how commercial goods could be purchased and delivered, but had entirely upended the rules of corporate competition and success (Boyle 2019; Feng et al. 2001).[113] In the new economy, the idea of a corporation that offered consumers products gratis, in the apparent form of a gift, yet remained economically viable, became not just thinkable, but actual and

[113] See Barry and Slater (2002) and Thrift (2001) for accounts of the full set of changes to economic processes that were popularly ascribed to the new economy.

commonplace. The internet startups of this era, it seemed, had realized the capitalist fantasy of abundance without scarcity (see Slater 2000).

With the bursting of the dot-com bubble in March 2000, the capital markets showed renewed if ultimately temporary interest in startup earnings records. Within years, a new generation of startups, including Google and Facebook, rode the free-to-use model to new heights of success. Startups soon converged en masse on Google and Facebook's practice of using surveillance-based targeted advertising to "monetize" the global audiences attracted to their services. The widespread adoption of this business model changed the economic logic of the relationship running between internet startups and their users. It retrospectively recast the character of the content and services that had been offered to the public in the guise of gifts. It revealed them to, in fact, have always been the object of a commodity-like exchange, if one whose return was unusually deferred in time and obscured by the substitution of user data for currency as the medium of payment. Did users know that the nature of the bargain being offered to them had changed? Certainly they complained when ads began to interrupt their Tumblr and Instagram feeds. But did they appreciate that they 'themselves'—their data and what it revealed about them to advertisers—were the source of the economic and imaginative premium that tech companies attached to such advertising? As late as 2009, Chris Anderson, the editor-in-chief of Wired magazine, published an entire book arguing that the internet's iconic websites were "[r]eally free—no trick." Notably, Anderson asserted this claim despite describing within the book the surveillance-based advertising that companies like Google widely used by 2009 to subsidize their operations (Anderson 2009, 119). In his 2018 testimony, Zuckerberg himself

continued to hold Facebook out as free, assuring Senator Hatch, "There will always be a version of Facebook that [like today's] is free."

As Zuckerberg's testimony suggests, in 2018 Silicon Valley's startups continued to benefit from the public impression that they offered something-for-nothing. We can attribute this fact in part to the way the constant churn in website aesthetics, features, and preference options obscured the changes startups instituted to operationalize the internet's business model. Consider in this regard the historical experience of using Facebook. Despite Mark Zuckerberg's famously dismissive attitude towards privacy, when it was incorporated in 2004, Facebook explicitly promoted privacy as a selling point of the service. Unlike other social networks of the time, Facebook limited registration to individuals who could prove affiliation with one of a handful of elite universities. Facebook's 2004 customer support page described such segregation-by-school as a design feature intended to "make sure that your information is seen by people you want share it with, and not seen by folks you don't." Early users in fact identified it as a big part of Facebook's appeal. The privacy afforded by restricting the set of people who might view a user's page made Facebook's early users feel as though the service had been made exclusively for them (Graham 2013b).

Then, one morning in late 2006, Facebook's 12 million users logged in to find the social network radically changed. Overnight, Facebook had launched a new feature, which fundamentally reorganized how the site displayed user information. Previously, to see what a friend posted to Facebook required users to actively navigate to the friend's profile. Now, with the launch of "News Feed," Facebook instead broadcast all friend activity, no matter how personal, to a rolling queue on the user homepage. With the passage of time it's easy to forget,

but Facebook's introduction of the News Fee prompted strong negative reactions from many users. At the time, Facebook argued that users retained the ability to limit the sharing of their activity to friends and that no new categories of information were being shared. Despite such assurances, reducing the effort required to access such information undeniably changed the perception and experience of privacy on the site, leading some users to organize a short-lived boycott.

This experience, of waking up to find a beloved internet service suddenly reconfigured, is a routine feature of life on the web. Sites constantly add new features and deprecate existing ones. They overhaul their styling and organizational flow, reorganize and reset user preferences and settings. Such changes can be subtle and incremental, but can also be drastic, upending user expectations of how a site operates or what purpose it serves. This functional and aesthetic churn reflects the fact that, despite the intimacy and familiarity we may establish with them, websites are ephemeral manifestations of constantly changing code bases. Internet-based distribution may conspire with software's relative immateriality to enable websites to be quickly, remotely, and cheaply reconfigured, but it also compels internet companies to constantly offer new and improved functionality in pursuit of rapidly changing markets.[114]

While not all such changes implicate a company's data practices or business logic, it can be hard to tell when they do. The intensifying surveillance that has accompanied such changes over time takes advantage of inherited design features in the web's hardware, software, and communications protocols. It thus generally operates invisibly, perceptible only indirectly via the affective disturbances users attribute to ads that appears to stalk them across the web. Internet

---

[114] See Slater (2002a) describing the new economy's "routinization" of innovation.

companies consistently describe any changes to their services as necessary to better fulfill their corporate mission or serve user needs. By design or omission, in so doing they often conflate the data used to maintain and improve their products with the data collected to more precisely target users. Even where a change obviously implicates the scope of a corporation's data practices and thus the form of privacy available in relation to it, tech companies are loathe to indicate that it contributes to their ability to generate ad revenue. In this destabilizing phenomenal landscape, it's little wonder that most casual users accept the web's magical conveniences at face value, muddling along passively with little concern for overarching competitive or economic logics.

The role of the web's constant churn in obscuring the changing nature of the corporate-user relationship from even expert observers is vividly illustrated by a series of infographics published in May 2010. Matt McKeon, a software engineer then working at IBM Research, posted the graphics to his personal blog under the title, "The Evolution of Privacy on Facebook." As he explained in related commentary, McKeon admired Facebook as a feat of technical engineering and as a means of keeping in touch with old friends, but believed it "hadn't always managed its users data well." To mitigate public confusion generated by Facebook's ongoing changes to its privacy policies, McKeon created a simple, multi-level pie chart for each year between 2005 and 2010. Using shading and structure, the charts illustrate the default privacy settings for each data category that Facebook recognized in a given year. For example, they show whether, in a given year, gender was viewable by default only by a user's confirmed friends, by a user's friends and their friends, by all Facebook users, or by the entire internet. Together, the five graphs reveal a clear progression. In 2005, none of information users posted to Facebook was available to the general internet. Only a few categories could be viewed by the general Facebook

population beyond a user's direct and indirect friend networks. By 2010, by contrast, every category of information other than contact info and birthdate had been made available to the entire internet by default.

In related publicity materials, Facebook stated that it had redesigned its privacy policies in order to clarify them and to give users greater "control" over their information. But for McKeon, tracking the changes to Facebook's default settings in indisputable visual form pointed to an alternative explanation: They reflected an effort to increase the burden on users to preserve some form of privacy on the site. Despite Facebook's invocation of transparency and control, McKeon conclude that the changes in fact served Facebook's then little-known effort to "correlate, publish and monetize" the massive database of personal information that it had quietly amassed since its inception.

At a time when the economic foundations of internet businesses remained shrouded in mystery, McKeon's graphics circulated widely in tech and civil liberties circles. Both privacy engineers at Mozilla and lawyers with the Electronic Frontier Foundation, for example, cited the graphics as striking confirmation that the ongoing changes to Facebook's features, policies, and settings were directly related to its growing ad business. They revealed Facebook to have transformed from a private space into a public platform by incrementally eroding user privacy, and to have done so largely unnoticed. It may appear perfectly evident in retrospect that the internet's promise of something-for-nothing was always marked by a significant qualifier. McKeon's graphics, however, foreground the significant cultural investment that was required for even those with privileged positions within the tech industry to appreciate that the value logic of the internet's free services changed as corporations adopted the internet's business model.

Innovations in Disavowal

The internet's business model proliferated under cover of the techno-moral aspirations projected onto the early web and of the destabilizing churn in the features, looks, and legal terms of the services that define it for most Americans. But the changes the business model inaugurated in corporate-user relations further escape appreciation because they reflect a novel form of exchange. When we use 'free' internet content and services, we are legally and materially enrolled in a form of exchange with aesthetic, social, and temporal properties that diverge significantly from those typical of commodity exchange in the consumer era. This form of exchange remains largely illegible to users as such because internet companies deploy novel techniques to actively disavowal it. By foregrounding for users the intimate magic of their technologies, the value they confer, and the unobjectionable social missions they forward, they obscure the particular mechanisms through which they extract value from users and convert it into revenue. From this perspective, when civically-minded technologists like Zuckerman mobilize the business model as an explanatory heuristic, they are attempting to make internet use legible as exchange by naming its underlying value logic. Doing so theoretically serves to empower users to bring to bear on internet use the kinds of instrumental reasoning and strategy, which economic theory understands individual to deploy to pursuing personal interests during exchange. To draw out the distinctive properties of the form of exchange implicated by free internet services, let's compare it to an idealized example of market exchange.

Imagine entering a 1950s-era mom-and-pop store to purchase a tangible good such as a stapler. In so doing, your intentions and expectations will likely appear perfectly self-evident. A mass of unruly papers has overrun your office. You wish to tame it, so off to the store you go.

On further reflection, however, one can appreciate that such self-evidence is possible only by reference to a set of common rules governing exchange. These rules are themselves a significant social achievement. They are the product of a vast array of material and institutional devices, procedures, and arrangements, which create the conceptual and experiential "frame" within which, of all the possible forms of interaction, of all the possible courses of action open to individuals, the one with the characteristics that we think of in terms of exchange is contextually enabled (Callon 1998b; Goffman 1971).

For example, if pressed to explain why you sought out a stapler at mom-and-pop's store in particular you might refer to your prior experience shopping there. Framing cannot rest on subjective commitments alone, however, but must also be buttressed by appropriate physical markers (Ibid.). Thus, if upon your initial visit to mom-and-pop's you arrived to find signage reading, "Day Care," or a building with the architectural features characteristic of a church, you might well have never entered.

As is, you arrive at mom-and-pop's to all the appropriate material and aesthetic markers of a store. Upon entering, the general layout, the clear, bright lighting, and the careful, intentional display of goods further indicate that you have entered one of the kinds of spaces culturally set aside for exchange. As such, in entering the store you effectively initiate a genre of social event governed by well-known, culturally- and historically-specific rules of interaction and interpretation. In various ways, this framing alleviates the uncertainties and questions about

identity and intention that inflect less ritualized social interactions. It allows you to take for granted, for example, that the kindly-looking gentleman standing behind the cash register is acting in the role of the store's agent—the seller. As such, you face no question of Pop's desire to see a sale initiated and concluded. By reference to these rules, you naturally slip into the other social role assumable in such an event—the buyer—and adjust your expectations, behavior, and comportment accordingly.

The intersubjective understandings and material arrangements that frame your shopping trip as such further condition both the appropriate timing and sequencing of the actions that constitute and eventually draw the event to a close and the moral and social significance assignable to them (Ibid.; Davis 1996). The law may provide for an exchange to be reopened upon certain unlikely future contingencies. Generally speaking, though, the event of exchange and the circumscribed social relationship it temporarily enters you into are formally drawn to a close once the culturally proscribed sequence of actions—the handing over of money, the bagging and ritual transfer of the stapler—has been performed (Callon 1998a).[115] You may thus proceed with the rest of your day confident that your new stapler will soon assume a place among the largely invisible, unsurprising background artifacts that populate life's periphery. In the words of Maciej Ceglowski, a well-known web developer and critic of the internet's business model, in commodity exchange "[i]t's easy to understand what's happening. You give money, they give you [the goods]. You don't need to say more about it (Ceglowski 2013)."

---

[115] Buck-Morss (1995) describes social indifference as one of the characteristic features of capitalist market exchange.

In the context of exchange, framing has a further significance. It is what not only allows us to recognize when it is culturally appropriate to engage in exchange, but also what endows us with instrumental rationality, the capacity and inclination to engage in the agonistic calculative action upon which market exchange is at least formally predicated.[116] To unpack this claim, think of market exchange as involving in the abstract distinct agents who pursue and seek to maximize divergent interests by opposing one another in transactions resolvable through the compromise of price (Callon 1998a). So defined, market exchange requires participants to identify possible goals, rank them, and conceive of the actions required to attain them. The capacities required to do so are only achievable through the extensive collective work and investments involved in framing. In the case of our stapler, these include not just the material arrangements and intersubjective understandings already described, but also Pop's careful display of prices, the weekly local newspaper insert that allows you to compare and contextualize the relative bargain such prices represent, and institutional devices, including especially the property rights without which you could not be certain that the stapler can be conclusively alienated to you.

Unlike our stapler, internet applications confound the norms of consumer exchange at every turn. And they do so in ways that undermine the individual ability to recognize their use as exchange and thus to bring one's instrumental reason to bear on them (see Callon, ed. 1998). As noted, for example, as an object of exchange, the software application is a trickster of a good. A purchased stapler might chip or crack as it suffers the normal wear and tear of use. But for the most part it will simply exist, collecting dust on a shelf until enlisted for its modest contributions

---

[116] Miller (2002) and Slater (2002b) argue that gifting, like market exchange, often involves complex modes of calculation, and that market exchange itself is not distinguishable in terms of a purified mode of calculation.

to the battle against entropy. As a purchaser, you need not worry that one morning you will reach for the stapler only to find a can opener, or that the stapler now sends its manufacturer a log of every document you bind, or automatically charges your credit card a fee after every 100th use.

By virtue of their origin in code and remote delivery and maintenance, internet applications contain just such a kernel of contingency. This kernel is sustained and actualized via legal and material technologies, which grant the software developer the authority and means to unilaterally enact its technical and competitive decisions. From the user's perspective, the practical effect of this contingency is the possibility, indeed the likelihood, that no matter how much time and energy one devotes to mastering an application, no matter how intensely one incorporates it in one's projects of self-fashioning (Foster 2007; Miller 1987), it may be suddenly transformed from afar according to inscrutable logics. The socially significant characteristics attributed to goods as apparently self-evident as a stapler may always be the result of social and symbolic processes and thus subject to ongoing contestation (Callon et al. 2002; Appadurai 1986). But the contingency that haunts internet services arguably differs in its immediacy—it can be effectuated at the speed of a software update—and in its scope, which extends from the service's perceived characteristics to the nature of the economic bargain by which it is made available to users.

Today, of course, much of everyday commerce in tangible goods occurs over the internet, and relevant norms and expectations have adapted accordingly. Despite the absence of many of the traditional cultural markers of exchange, there is little doubt, for example, what kind of event one is engaged in when scrolling through Amazon's endless listings, adding items to a "shopping cart" and then "checking out." With free websites and services like Facebook and Gmail, things

are different. Like most software today, such services are no longer sold as a discrete, tangible good at a definite point in time. Software delivered-as-a-service extends the duration of the exchange event, drawing it out over an often indeterminate timespan. Through their persistent delivery of content or service and ongoing surveillance of users, internet companies similarly transform the exchange relationship from a discrete, bounded engagement into an ongoing affair.

The form of exchange entailed by the use of internet applications further diverges from the consumer ideal in that it necessarily involves not just a buyer and seller but also the third party advertisers who convert collected user data into revenue. As Ceglowski (2013) notes, "As soon as you start talking about business models in computer land, it gets really complicated somehow…The users don't get to pay directly for services. …There's all these bizarre flows and things. The company is trying to demonstrate to its investors it should be funded further. When advertising is added it gets even more complicated. The advertisers are producing something that you have to then sort of shovel to the user. The user is generating data in your system that gets passed back up to the advertisers."

Advertisers don't actively contribute to the development of internet applications, and aren't party to the terms of service that formally determine the rights and obligations running between internet companies and users. They nonetheless haunt the use relationship. The circuit of value opened when Business A offers its services to the public at no monetary cost cannot be consummated unless advertisers enter into their own relationship with it. As discussed in greater detail below, technologists have long feared that accommodating advertisers in the design and engineering of new consumer internet technologies requires them to compromise their professional commitments to quality engineering and to properly serving end users.

The uncertainties surrounding the role third parties play in the bargain between internet companies and users are reflected in a maxim often used as a short-hand critique of the internet's business model: "If you're not paying for it, you're the product." As used in this context, the maxim stands for the proposition that users who think of themselves as the customers of companies offering free but ad-supported products suffer from a form of category mis-recognition. Under the business model's value logic, the structural position that users actually occupy is not that of customer but product, one secretly cultivated and sold to the 'real customer'— advertisers—through the power that personal data is understood to grant to influence user behavior. From this perspective, the continual failures of internet companies to properly steward user data—their refusal to implement adequate data security safeguards; their use of personal data in ways that plainly harm users interests—merely evidences the truth that users aren't their real customers and thus exercise no moral claim over them.

Ultimately, if users fail to recognize the bargain of data-for-service represented by the internet's business model, it is because internet companies offer their products up to the world in a manner designed to obscure any culturally legible indicators of an exchange event. There is, for example, no evidence of the specifically economic nature of the user relationship on the face of the services themselves. No price is ever presented to or negotiated with the user. Users can freely navigate between websites that do and don't employ commercial tracking technologies with no perceptible means of discerning the difference. The absence of recognizable signifiers extends to the consummation of the exchange, which involves no transfer of money or property, no receipts. Internet corporations instead devote considerable engineering and design resources to ensuring that "on-boarding" new users is as smooth and friction-less as possible, involving the

fewest possible clicks, the fewest cognitively off-putting decisions or choices. The industry's very use of the term "user" rather than customer or consumer, while understandable given Silicon Valley's organization around computer engineering, serves to foreclose recognition of internet technology 'use' as involving an ongoing economic relationship.

The fact of an exchange relationship is instead pragmatically asserted, legally construed to adhere upon the presentation of the developer's terms of service. In the U.S., the presentation of the terms upon which access to free online services is formally conditioned is effectively the only moment in the ongoing relationship between user and corporation when its reciprocal nature is directly communicated to users. Even then, despite years of effort by privacy activists, lawyers, and technologists to make terms of service comprehensible to normal users, it's well known in Silicon Valley that consumers simply do not read or understand them. And while a corporation's terms of service generally describe how it collects, uses, and shares user data, they do not describe how in doing so corporations capture value for themselves and convert it into revenue.

Perhaps most significantly, Silicon Valley's internet companies hold themselves out to the world in a manner that confounds their recognition as sellers or exchange partners. Over the past 25 years, Silicon Valley's corporations have developed an outsized presence in public imaginaries thanks to the magical conveniences they have brought to global audiences, their disregard for business convention and authority in general, and their promises to remake the world using technology. Even so, despite the long cultural shadow they cast, the general public continues to encounter and know them primarily in their technological and social guises. On the

evidence of the senses, there exists little to indicate to users how they make money from the services they provide, and thus more generally, to mark them as economic actors.

In this respect, unlike the store clerk in our stapler purchase, the motivations of technology companies remain culturally obscure, clouded by the shadow third parties cast on the user relationship and by the strategies internet companies deploy to promote their public perception as other-then-businesses.[117] We might think of the resulting ambiguity regarding their motivations in the economic language of information asymmetries and disproportionate bargaining power. But it is telling that technologists and other critics more frequently treat it as a kind of ontological confusion, another category error, in this instance with respect to the nature of internet companies as social actors. This is the thrust of the claim present in many critical articulations of the internet's business model that Silicon Valley's companies aren't "really" technology companies at all, but instead advertising or even surveillance companies.[118]

This is not to say of course that technology companies are literally unknowable in their economic guise. Part of the cultural fascination they command derives precisely from their perceived exceptionalism as economic actors, from the fabulous personal fortunes they generate and the stratospheric valuations investors bestow on them. Rather, their limitless economic ambitions and the mechanisms they use to make money are an open secret. They are knowable with effort, but otherwise obscured by their circulation in semiotic networks distinct from the

---

[117] Thrift (2001) identifies the idea of business as a creative outlet rather than a financial endeavor as a specifically romantic notion and traces it to the new economy-era effort by consultants and business academics to institute a spiritual renewal of market culture.

[118] As PayPal founder Peter Thiel (2014) acknowledges, Silicon Valley startups themselves seek to manipulate the industry or market with which they are associated in order avoid regulation, attract investment, or achieve other strategic goals.

ones internet companies activate to define their services in relation to their social missions. The temporary disturbance of this settled state of affairs by a major corporate hacking or other privacy-related scandal may confront the general public with the value logic through which internet companies actually sustain themselves. Otherwise, entrepreneurs generally limit communication of it to the closed-door venture capitalist meetings that mark the early stages of startup development or in the dense financial and legal disclosures required of public companies by the SEC.

This separate circulation should be considered an inherent feature of the internet's business model, one not evident when the model is analyzed in the abstract. It constitutes an innovative addition to the repertoire of techniques of disavowal practiced by American corporations. In recent decades, anthropologists and other social scientists have catalogued a number of such techniques. As a genre of corporate practice, they seek to warp the political environment in which corporations operate to facilitate the expansion and intensification of corporate capital (Benson and Kirsch 2010). They variously foster doubt about the existence or causality of corporate harms (Oreskes and Conway 2010), co-opt the language of critics, and adopt the rituals of audit culture (Power 1994) to manage or neutralize public critique and the threat of regulation. Examples of corporate disavowal include the adoption of "corporate oxymorons" (e.g., "sustainable mining," "safe cigarette"), which pair a positive cover term with an original literal term to obscure the harms attributed to the latter and deflect consumer concern (Benson and Kirsch 2010; Benson 2010).

Internet companies partake of many of the techniques of disavowal described by social scientists. Notably, they attempt to shift responsibility for protecting privacy onto individual

users through invocation of their systems of notice and consent. They hide their extensive use of surveillance and when challenged pretend it has no privacy implications. Where they add to the existing repertoire is in deploying semiotic idioms and strategies of circulation alongside the material arrangement and presentation of their services to preclude users from identifying their services as transactable objects and themselves as calculating economic agents (see Callon, ed. 1998). In other words, they not only refuse to acknowledge the existence and extent of the social harms inflicted through their business practices, but also hold themselves out in a fashion designed to preclude users from recognizing them as the kind of social actor that engages in the practice of business at all. In this respect, internet companies have adapted for the corporate form a tactic of manipulatively framing exchange long practiced by individuals. In describing the many ways that individuals manipulate the perception of exchange to their advantage, for example, Davis (1996) cites the record executive who foregrounds his friendship and monitorship of an aspiring musician in order to mask his commercial interest in securing favorable contractual terms. The techniques employed by internet companies, of course, cannot fully preclude the possibility users will recognize them as for-profit actors, but they need not in order to serve their purpose. It is enough for them to provide those who may so perceive internet companies with a basis for rejecting or refusing to confront the implications thereof.[119]

The net effect of this separate circulation is that internet companies wear a double face in society, one techno-social, the other economic, the gaps between which constitute a space of opportunity for the expansion of their particular form of corporate capital. When civically-

---

[119] In this regard, Fortun (2010) highlights the Freudian interpretation of disavowal as operating through a disjunction or refusal to connect.

minded technologists mobilize the concept of the business model to explain the diminished state of contemporary privacy, they seek to frustrate the efficacy of this technique by deploying a semiotic technology that captures these two faces and simultaneously holds them up to the world, mapping the interrelations between them. It effectively posits the existence of a hidden logic animating the behavior and pronouncements of internet companies in their techno-social guise, a behind-the-scenes that must be brought to the fore to understand their past, anticipate their future, and engage with them as social actors on terms appropriate to their otherwise obscured instrumental reason. Part of the explanatory power of the business model thus derives from its invocations of the deep-seated Euro-American distinction between appearance and essence and surface and depth (as illustrated in the Christian distinction between body and soul and by the Freudian idea of the unconscious), as well as of the metaphysical presumption of the fundamental self-interest of all social actors.

When civically-minded technologists like Zuckerman identify the business model as the basis of the internet's threat to privacy, they are to a certain extent simply participating in the broad-based effort of privacy professionals of various types to defend privacy by educating consumers. As a defensive strategy and form of consumer politics (Nakassis 2013), foregrounding the business model works not by providing consumers with new information so much as with a new cognitive resource through which to expand the scope of the market frame and thus subject internet companies and their products to a fresh evaluation (Callon et al. 2002). By showing their techno-social and economic faces to be one-and-the-same, it seeks to enable the public to cut the cultural associations with technological and social progress, convenience, and delight, so carefully cultivated by internet companies and recognize and narratively elaborate

them as agonistic economic actors. Without such recognition, in theory, the public cannot bring its instrumental reason to bear, cannot apply the appropriate balance of value, the appropriate expectations of return for any given offer, and is forced to fall back on standards, such as those appropriate to gifting, that are insufficient to the task of voicing and pursuing personal interest (see Fourcade and Healy 2007).

In and of itself, the explanatory mobilization of the business model is unlikely to mitigate the individual or social harms emanating from the internet. Doing so would require not just making internet use culturally legible as exchange, but also making the harms themselves visible and calculable. As other examples of such corporate externalities indicate, to open the framing of exchange to enable such harms to be factored in specific economic transactions, they must be first made measurable. In this respect, the business model must be understood as complementary of a broader set of heterogeneous projects originating in civil society and academia to provide the public with the cognitive, material, and institutional resources to engage in meaningful negotiations with internet companies. Important examples of such efforts include browser extensions that visualize for users the extent of online tracking, and services that allow internet users to control and generate income from their own personal data.

Whither the Business Model in Silicon Valley?

During my research, as the cultural salience of the business model came into focus, I wondered how else and for whom business models mattered. Clearly the idea of the business model had come to provide technologists, and increasingly the non-expert public, with a useful grammar for making sense of internet companies and the unintended social harms they propagate. But did the

concept have any purchase for business practitioners themselves? If so, what precisely did they understand a business model to model? How if at all did a business model guide the choices and actions of the employees tasked with formulating them and carrying them out?

To pursue these questions, I looked for actual examples of startup business models and for opportunities to observe Silicon Valley entrepreneurs and investors discussing their formulation and execution. I focused on the pitch decks that entrepreneurs prepare to prime and structure investor meetings in the early rounds of fundraising. Without easy access to pitch meetings themselves, I surmised that the pitch decks might provide insight in what business models look like in practice and how entrepreneurs employ them in the spectacular promissory visions—"Tens of millions of customers. Billions in revenue"—they are trained to conjure for investors.

I found an early clue in a template publicly released by Sequoia Capital, a famed venture capital firm known for investing in Apple, Yahoo, Google, and LinkedIn. Along with other topics, including the startup's purpose, the template directs prospective entrepreneurs to include a business model in their investor presentations. In bullet point format, the template defines a business model in terms of a set of necessary elements that collectively answer the question, "How do you plan to thrive?" The elements include a revenue model, pricing information, a description of average account size and/or lifetime value, a sales and distribution model, and a list of actual and potential customers. Despite the template's directive, however, I found in reviewing publicly available examples of pitch decks that while many addressed some of Sequoia's elements, very few addressed them all. Even fewer did so under the rubric of an explicitly labeled "business model."

As I continued my research, I found the ambiguous presence of the business model in investor presentations to be matched by the business model's absence from the stories that Silicon Valley investors and entrepreneurs tell about achieving startup success.[120] For example, business models were virtually absent from the advice offered by Y Combinator, the prominent startup "incubator," to budding entrepreneurs in a 2014 course taught to computer science undergrads at Stanford. Titled How to Start a Startup, the course consisted of twenty lectures on a wide range of topics including product development, hiring, company culture, strategy, competition, management, and sales and marketing, none of which directly addressed business model development or execution. In an introductory lecture, Y Combinator's then CEO Sam Altman (2014) identified a great idea, product, executive team, and execution as the key building blocks of startup success. He defined these elements broadly—his definition of a startup's "idea," for example, encompassed some sense of how the startup would defend against competitors—but none seemed to account for the elements identified by Sequoia Capital. Through the remainder of the course, the Y Combinator partners and alumni who guest lectured more or less faithfully emphasized the same elements as Altman. In the few instances in which

---

[120] I draw for the observations in this section on a rich set of materials, including tweets, blog posts, essays, podcasts, interviews, and presentations produced and disseminated between 2014 and 2018 by investors and entrepreneurs associated with Y Combinator, a startup "incubator." Y Combinator is best known for the three-month long intensive bootcamp that it runs twice a year for cohorts of competitively-selected startup founders. The program consists of a series of discussions with successful entrepreneurs, one-on-one and group advising sessions with Y Combinator partners, and presentations to peers and prominent outside investors. Through such activities, Y Combinator seeks to teach entrepreneurs how to improve their products, while training them in how to best portray a startup's 'story' to secure large-scale venture capital funding. Having provided initial funding for startups, including AirBnb, DropBox, Stripe, and Instacart, valued at more than $100 billion as of 2018, Y Combinator exerts an outsized influence on Silicon Valley's trajectory and reflexive understanding of startup culture.

lecturers did mention business models, it was in passing, with a hint of embarrassment and occasionally something verging on disdain.

Take for example Lecture 19, in which Michael Seibel (2014), an entrepreneur and Y Combinator partner, taught the essentials of pitching startups to investors. According to Seibel, fundraising success requires clearly and concisely communicating to investors that a startup has unique insight into a massive market opportunity and is moving quickly to take advantage of it. To this end, he recommended that entrepreneurs start every pitch by describing in the simplest, least ambiguous terms possible what the startup does and the size of its market such that investors understand, "If we're big, if we really blow this company up, it could be worth billions of dollars." An entrepreneur should then demonstrate the startup's traction in penetrating its target market by citing either growth metrics or speed towards public launch. Only after presenting this core pitch should entrepreneurs describe the company's secret insight into the market that existing corporations have failed to appreciate, and answer in no more than one sentence the question, "How does your company make money?"

> You know your business model. I see so many founders run away from this question because they think things like, if I say advertising people are going to be like, "Oh, that's stupid." Just say it! Don't run away. If it's advertising, say advertising. Facebook's a massive advertising business. So is Google. If it's direct sales, it's direct sales. If it's, you know, a game and you're selling in-app add ups, like, that's fine. Just say it. Don't run away from the sentence. … Where founders get tricked on how you will make money is they say, "Well, we're going to run advertising. Maybe some virtual goods. We're going to figure out how to 'this,' and maybe 'this,' and maybe this." Well, now you're saying nothing. Now you've told me you have no idea how you monetize this. This was a check mark that I just wanted to write: 'And then I am going to monetize it.' Instead I am writing a big question mark. So do the thing that everyone else in your industry does to monetize 95 percent of the time, say it and move on.

Seibel's commentary here seems to undercut the significance assigned to business models by critics like Ethan Zuckerman. It confirms that the business model exists for entrepreneurs as a required element of the stories they are expected to tell to demonstrate potential and thus secure their startups' futures. But it suggests that neither investors nor entrepreneurs view business models as a particularly compelling element of such stories, and it gives no indication that a startup's business model actually holds any internal relevance a to its competitive strategy. Unlike a startup's target market or its market insight, Seibel indicates, no particular burden falls on the business model to produce an "aha!" moment for investors. Business models here appear to exist for entrepreneurs primarily as a set of templates available to be copied from successful predecessors and applied wholesale to new ventures with overlapping profiles or characteristics.

The creative and competitive potential of the business model was, from Seibel's perspective, largely exhausted by 2014 and its pursuit therefore misleading. As he elaborated in a 2016 blog post, a small percentage of early startups continued to propose new business models. Unless a startup's product somehow promised to reconfigure its market, however, founders should just be "honest with themselves" (Seibel 2016). "By and large," startups that haven't "figured it out" during their initial rounds of investment are going to "make money by growing big and turning on advertising." They shouldn't be embarrassed to admit that a startup would monetize with advertising "when clearly that was the only answer."

Seibel's characterization is all the more noteworthy because business models have become the object of intense interest, formalization, and elaboration in other contexts. The internet is rife with blogs, videos, podcasts, and presentations offering advice on the nature and function of business models. Business schools and consultancies teach business model design

and innovation, and business models are the object of an extensive body of academic management literature.

As explored in this literature, elaboration of the business model concept extends began in the mid-1970s, approximately (Ghaziani and Ventresca 2005). Engineers and economists initially used the term to describe the use of electronic spreadsheets to model the likely financial consequences of various changes to a business' operations (Magretta 2002).

When the business model eventually entered the public lexicon in the mid-1990s, it did so alongside the popularization of the web and the growth in the digital economy. Between 1995 and 2000, entrepreneurs, investors, and journalists used the term business model either to refer to the new forms of transaction and pricing associated with online commerce[121] or as a rough stand-in for a company's revenue model (Porter 2001). In his 2000 account of wealth creation in Silicon Valley, the journalist Michael Lewis thus dismissed the business model as a pointlessly obscure buzzword of the dot-com bubble: "it glorified all manner of half-baked plans. All it really meant was how you planned to make money" (Lewis 2000).

Following the 2000 crash, as the initial period of internet-related euphoria came to an end, the term business model took on the more precise meaning that remains in use today. Consultants and scholars of corporate management converged on the idea of a business model as a simplified, conceptual representation of the value logic that sustains a corporation (Fielt 2013). Business models, in other words, model how a business first creates and delivers value to customers and then captures part of it in the form of revenue and profit.

---

[121] Among these were the customer-to-customer auctions introduced by eBay and Priceline's reverse auction, under which buyers specify the price they are willing to pay to prospective sellers.

Whatever ambiguities the academic literature betrays regarding the precise definition or necessary elements of a business model, it is entirely unambiguous in insisting that business models matter. It argues that business models determine the relative economic success of technological innovations (Ibid.). A business school professor whom I observed teaching a class on the topic further insisted that business models are at least as if not more important than any technological innovation to startup success. By shifting what the management literature calls the basic unit of business—the items that actually show up on an invoice—away from traditional "goods" and "services," the business model provides management with a conceptual tool to decompose corporate offerings into constituent elements, and recompose and present them in novel ways, thereby opening up a new domain of strategic creativity and choice (McGrath 2010; Osterwalder and Pigneur 2003). While the business model concept was historically elaborated in relation to Silicon Valley's internet startups, the management literature insists that in a market economy characterized by consumer choice and rapidly changing competitive environments, 'mature' corporations must also constantly attend to their business models or risk being outcompeted.

Despite such claims, I found that when Silicon Valley entrepreneurs and investors discussed their startups they were generally reluctant to address the question of the business model. When confronted with the internet's advertising-based business model in particular, they expressed something closer to repulsion or anxious avoidance. In Seibel's description, entrepreneurs "run away" from advertising; it embarrasses them and causes them to dissemble. It's a curious reaction. In Silicon Valley's reflexive mythology, entrepreneurs only really need a business model early in startup development as a kind of performative prop. Identifying the

appropriate business model plausibly demonstrates to investors that a startup's founders have 'figured out' the monetization problem. It thus confirms the capabilities of the startup team, authorizing investors to consider the related due diligence obligation formally satisfied. The overarching goal of pitching one's startup is to tell a story that convinces investors the startup has a plausible path towards controlling a market that is or will be very large. And as Seibel observes (see also Graham 2013a), advertising is the business model that Silicon Valley's two most successful startups used to conquer their markets. Given such facts, it would seem easy enough for entrepreneurs to simply specify the advertising model and move on.

In a podcast he hosts on startup success, the serial entrepreneur and investor Reid Hoffman suggests that it is this very sense of inevitability surrounding the business model that makes entrepreneurs' reluctant to discuss them. According to Hoffman (2017a), the kind of entrepreneurship that Silicon Valley valorizes "is something new. … It's almost always a new game." Whether because of changes in technology platforms, the competitive landscape, or consumer demand, "almost always you really have no clue. You're throwing darts at a dartboard about how it plays out. The jump into the unknown where you're like, 'Who knows,' is at least a certain adrenaline rush that perhaps I'm addicted to." For an entrepreneurial class that defines its social value in terms of the ability to navigate the head-on rush into the unknown, the sense that Google and Facebook long ago "solved" the problem of making money on the internet contributes to entrepreneurs' reluctance to engage with the business model concept.

In How to Start a Startup, Y Combinator's founder, Paul Graham (2014), suggests in the alternative that the distaste with which many Silicon Valley entrepreneurs approach not business

models, but revenue and business in general, derives from their background in engineering.[122] In his experience advising startups, he reports, he found that entrepreneurs trained as software developers generally wished they could just write an elegant program, upload it to a server, and get paid lots of money. “They'd prefer not to deal with tedious problems or get involved in messy ways with the real world,” a preference Graham deems reasonable, “because such things slow you down.”

Growth’s Imagination

Whatever commitments engineers may bring to entrepreneurship, we cannot understand the business model’s ambiguous presence in Silicon Valley without understanding the cultural, institutional, and economic conditions that enable internet startups to defer revenue generation. Doing so requires in turn that we explore the vehicle through which entrepreneurs seek to realize their aspirations, i.e., the venture capital-backed startup, and the ideological vision of success to which is is tethered.

Recall that in advising entrepreneurs to be ‘honest’ about their business model, Michael Seibel identified “growing big and then turning on advertising” as the obvious solution for most startups. The reference here to growing big is not incidental. It rather indexes Silicon Valley’s

---

[122] Chris Kelty (2008) argues that computer and software engineers share a techno-moral social imaginary under which information infrastructures, as commons, must be built in accordance with principles of justice and fairness. They thus reject corporate practices expressive of a “lust for money” and of the pursuit of benefit for the few over the greatest good for the many. For Kelty, this moral judgment indexes an inherent tension engineers face. Their sophisticated, technocratic imagination of moral and technical order can only be viably realized through a social technology—the corporate form—which demands the maximization of individual gain through winner-take-all competition and the constant exploitation of instability.

significant cultural investment in a particular vision of rapid, exponential growth as a both normative ideal and pragmatics of startup success. Silicon Valley's organization around this ideal is evidenced everywhere: Startups aspire to achieve rapid, exponential growth but must also demonstrate its unfolding to the venture capitalists who gauge startup potential on its basis. Under breathless headlines about discovering the next "unicorn," business and technology reporters tell endless stories about which companies have growth and which will have it next.

The elaboration and valorization of the cultural ideal of growth can generally be traced to Y Combinator's Paul Graham. Graham has long argued that startups constitute their own category of corporation, one distinguishable from others not because they primarily produce technology, are newly incorporated, or are funded with venture capital, but because they were intentionally designed from the beginning for rapid, exponential growth. As Graham wrote in an influential 2012 essay, "The only essential thing is growth. Everything else we associate with startups follows from growth" (Graham 2012). It is rapid growth, for example, which sustains the stratospheric valuations that investors bestow upon startups with no record of revenue or profit. Growth further explains how startups appear to suddenly emerge from nowhere as already unavoidable intermediaries of our social, professional, and commercial lives.

As mobilized by entrepreneurs and investors in Graham's orbit, the ideal of growth has unmistakable heroic and romantic undertones and intertwined social and aspects. Per Graham's typology, to start a startup is to declare the ambition to build a very large company. Venture capitalists, of course, support such ambition in order to make lots of money. The entrepreneur and investor Peter Thiel argued outright in his *How to Start a Startup* lecture that every startup's goal should be to establish a monopoly. Silicon Valley's entrepreneurs, by contrast, are generally

more circumspect about the allure of riches. Within Y Combinator's mythology, a startup is more likely to achieve success when founded as the ulterior motive to an entrepreneur's curiosity than when money is the ulterior motive to the startup's founding. Entrepreneurs consistently articulate the desire to grow their businesses as large as possible in terms of having the greatest possible "impact" on the world. From growth's perspective, startups are both a vehicle of wealth creation and a technology—indeed, the optimal technology—for solving hereto-for unresolved societal problems.

Whatever the precise mix of its economic and social aspects, the ideal of growth that Y Combinator espouses is expansionary and totalizing. Investors and entrepreneurs in Silicon Valley recognize no legitimate end to a startup's growth. For startups, there is always "another domain or idea that can be devoured." To the degree it indexes inevitable technological and economic progress and intensification, growth derives its social legibility and force from Moore's Law (see Ceglowski 2014b). As an ideal, however, growth aspires to the speed of the exponential, which is to say, to speeds beyond human experience and perception and thus beyond the regulatory grasp of human institutions.[123]

Growth then is a form of techno-social manifest destiny (Ceglowski 2014b), one that unfurls along material, economic, institutional, and symbolic dimensions. A startup may measure and demonstrate growth through a core metric like monthly users, or less frequently, revenue.[124]

---

[123] For a highly evocative treatment of the phenomenal experience of the exponential, see the 1977 short film, *Powers of 10*, by Charles and Ray Eames.

[124] Y Combinator partners tells participants that the best metric to use in measuring growth is "good old revenue." But they recognize that many if not most participants don't charge for their products and services in their earliest stage of growth. In such a case, the partners suggest that the startup track growth in monthly users as a "reasonable proxy" for the revenue it will generate "whenever [it] does start trying to make money."

But growth as thus measured necessarily leads to growth of other kinds—the number of employees, bureaucratic structure and processes, technological infrastructure, valuation, and ideally, in the ambition and abstraction of its core idea. At a 2017 conference for female entrepreneurs, a Y Combinator partner observed in this regard, you start a site for college students "and pretty soon you realize you could expand to sign up the whole world if you wanted to" (Livingston 2017).

As Silicon Valley-based critics of the internet's business model note, growth's totalizing drive is compulsory and moralistic. Investors and other entrepreneurs dismiss as "dying or presumed dead" any startups that doesn't demonstrate rapid, exponential growth even if successful as measured by other standards, including that of profitability, but (Ceglowski 2013). Daniel Heinemeier Hansson (2017), founder of the project management company, Basecamp, and a prominent critic of venture capital financing, reports on his personal blog that when he decided not to pursue the path of growth and trimmed his startup's product portfolio, he was met with incredulity and anger. Peers within Silicon Valley communicated to him that if the eliminated businesses had financial promise, he was crazy to 'turn down' growth. Entrepreneurs who don't "keep milking and pumping" a promising enterprise—who leave 'value on the table'—fail in their moral obligations to the market, the startup community, and the future.

Within Silicon Valley, growth's compulsory nature is widely recognized to follow from the imperatives of venture capital. Because most startups fail, and because even experienced investors struggle to identify important startup ideas until they're proven by the world, venture capitalists employ their own hit-based business model. They depend on any given investment fund having one success significant enough to compensate for the remaining failures while still

delivering a windfall. To ensure they achieve these rare hits, venture capitalists pressure entrepreneurs in their portfolio companies to single-mindedly pursue growth regardless of its toll on self-respect, solidarity, or any other such social motivation for collective enterprise.

The moral imperative toward growth partakes of the long intellectual tradition of representing economic growth—the growth of corporations—as a necessary condition for human progress (Fourcade and Healy 2007).[125] It is locally grounded, however, in particular aspects of growth's pragmatics. In Graham's (2012) articulation, for a newly formed corporation to grow rapidly it must (a) make something lots of people want, and (b) reach and serve all those people. Reaching a mass audience in the age of software and networked computing presents little challenge. Making something lots of people want is another matter. As Sam Altman explained in the introductory session of *How to Start a Startup*, "The best thing of all worlds" is to build a product that people love from the get-go. But practically speaking, the technologies and infrastructures that enable startups to aspire toward growth also expose them to intense global competition (see Teece 2009; Porter 2001). The tech industry is too competitive and efficient for a startup to build and defend a business based on an obvious idea. When an entrepreneur identifies a potential product for which there is obvious need, she should therefore assume that "Google or Facebook will do it."

Y Combinator thus teaches entrepreneurs to expect their initial startup idea to be unlikely to address the kind of non-obvious but mass consumer need that can support defensible, rapid, exponential growth. Even so, an entrepreneur's initial idea about "which hill to climb," if a fairly

---

[125] Growth, in this respect, partakes of the post-industrial conflation of civilization with the boundless desire for and proliferation of goods Buck-Morss (1995).

good idea, is usually "better than they realize" and "adjacent to even better ones." Entrepreneurs are urged to use growth as a tool to locate these adjacent ideas. Y Combinator, for example, famously requires the startups in its incubator program to measure their growth rate every week of the 10-week term, and expects them to show between 5% and 10% growth per week. Producing such consistent growth requires startups to adopt growth as the rubric through which they make all decisions: "Should you spend two days at a conference?…Should you add x feature? Whatever gets you your targeted growth rate." The pragmatics of growth in this way reduce the "bewilderingly multifarious problem of starting a startup …to a single problem."

More significantly, constantly measuring and reorienting a startup to produce consistent growth—following "growth's imagination"— is understood to draw startups closer and closer to the kind of elusive idea that can propel explosive growth: "The fascinating thing about optimizing for growth is that it can actually discover startup ideas. You can use the need for growth as a form of evolutionary pressure. If you start out with some initial plan and modify it as necessary to keep hitting, say, 10% weekly growth, you may end up with a quite different company than you meant to start. But anything that grows consistently at 10% a week is almost certainly a better idea than you started with." While Silicon Valley distinguishes between organic, word-of-mouth growth and paid, "inauthentic" growth, growth in general is still often taken as a proxy for consumer desire, a measurable sign that a startup has identified, and is on its way to solving, a hereto-for unresolved user problem. It is in this respect that growth is always a good in and of itself. To reject growth is to refuse the market and social injunction to follow the

instructions of consumer need.[126]

As illustrated by growth's simplification of startup decision-making, the pragmatics of growth impose a political economy of attention and effort on entrepreneurs. This political economy carries forward the new economy demand that business managers passionately, romantically pursue their vision and goals on a 24/7 basis (Thrift 2001). But it also determines the appropriate objects of entrepreneurial devotion during each of the roughly defined growth stages that comprise the idealized startup trajectory. Over the course of a startup's life-cycle, successful entrepreneurs will be expected to progressively take on more and more taxing functions— hiring, managing, modeling corporate culture, etc. In the initial stages of development, however, the political economy of care and attention specifically calls for them to focus their energies exclusively on producing the best product possible as measured by growth. Y Combinator partner Jessica Livingston (2017) notes, for example, that in her 11 years advising startups, the most successful were always those run by entrepreneurs who focused "fanatically" on product and who resisted "distractions" like going to conferences, recruiting advisors, and arguing on social media.

In *How to Start a Startup,* Altman elaborated on the optimal sequencing and timing of early entrepreneurial effort. Young entrepreneurs, he argued, frequently make the mistake of "Playing House." That is, they waste effort going through the motions of setting up a startup.

---

[126] David Heinemeier Hansson (2017) argues that the obsessive interest in growth, and the way growth discounts the present in favor of the lure of the future, helps explain the "pass" American society gives internet corporations to inflict social harms on the public without repercussion. As long as such companies continue to demonstrate growth, everything they do is *ipso facto* right. "Mistakes may have been made, but tomorrow is an entirely new day, divorced from any of the days that went before it" (Ibid.).

They may, for example, focus unduly on networking or on learning the supposed secrets of fundraising. They'll "rent a nice office in SoMa and hire a bunch of their friends, until they gradually realize how completely fucked they are because while imitating all the outward forms of starting a startup, they have neglected the one thing that is actually essential, which is to make something people want."

Circa 2014, the external trappings of the corporate form that investors like Altman viewed as potentially fatal distractions of entrepreneurial effort included both actual revenue and a plausible plan for how it might eventually be generated. As one critic of the internet's business model observed, based on his experiences with venture capital, revenue had come to be perhaps not detrimental to getting funded, but at least "a little gauche" (Ceglowski 2013). Founders who demonstrated revenue risked being accused of wasting effort that should have been spent gobbling up market share.

That entrepreneurs are not just expected but able to defer expending effort on the question of a business model is a function of the various ways that growth's pragmatics discount the present in pursuit of the speculative future. This discounting is evident, for example, in the presumption of investors that businesses founded to address obvious present needs will be indefensible given the competitive environment in which startups operate. It is similarly evident in the strategies investors encourage entrepreneurs to adopt, such as developing applied technical expertise, to prepare themselves to receive inspiration for startup ideas from the future.

Growth further discounts the present in the sense that venture capitalists understand the vast majority of the value they attribute to startups to reside in the speculative future. Traditionally, depending on the industry, investors valued a business on the basis of some

prudent mix of current revenue and future prospects. Business growth under such conditions is limited by the portion of current revenue that a business can set aside to pursue perceived opportunities. With startups, by contrast, the attitude is, in the words of Heinemeier Hansson (2017), "all potential, all the time." To achieve exponential growth, startups bypass the constraints inherent to relying on present revenue. Outside investors instead fund growth in a startup's early stages in return for the speculative promise of their future potential, assuming as LinkedIn co-founder Reid Hoffman observed on his podcast (2017b) that, "if they have a billion unique visitors a month…they have a property that is going to be worth a ton of money in some way eventually."

Together, growth's political economy of attention and effort and its discounting of the present structure the temporal logic of the internet's business model. They determine when in the life-cycle of the modal internet startup the business model will generally make its initial appearance, and when the startup will initiate the series of changes to its systems, processes and modes of user engagement necessary to finally 'turn on' revenue. The historical and imaginative conditions that originally made plausible the idea of a class of corporation that offers free lunch have arguably dissipated, and the general public has become increasingly aware of the nature of the exchange represented by internet services. Silicon Valley's valorization and practices of growth nonetheless contribute to startups' ongoing ability to bear two faces to the world.

Note, however, that to the extent surveillance-based targeted advertising remains the default business model of venture capital-backed internet startups, this can only be so because it is exceptionally compatibility with growth's limitless ambition. Growth, from this perspective, is an inherent part of the internet's business model, another part not evident when analyzed in terms

of abstract entities and flows. As it applies to the internet's business model, growth determines both the way the model unfolds over time and the totalizing accumulation toward which it aspires. As web developer and Silicon Valley critic Maciej Ceglowski (2014b) points out, the lamination of growth's temporal logic and ambition onto the internet's business model renders the theory of the world at its core effectively irrefutable and drives the ongoing intensification in the visceral disruptions and privacy harms that the model has inflicted on users.

In a series of talks delivered at web developer conferences since 2013, Ceglowski called Silicon Valley to task for the ongoing failure of the internet's business model to deliver on the promise startups attribute to it. In a 2014 presentation, for example, Ceglowski walked his audience through the "shocking uselessness" of the targeted advertising that "all this surveillance is buying us." By this point, Ceglowski estimated, the major internet companies had already accumulated a decade's worth of his viewing, search, and email habits. Even armed with such highly revealing data, Ceglowski still found himself regularly confronted with ads for products he already owned or for chain restaurants with no local outposts.

Moreover, Ceglowski argued, the claim that internet users love targeted advertising is belied by advertising's constant proliferation of newly invasive, demanding forms of online ad. Each new form, from the banner ad, to pop-ups, to auto-playing video ads, "turn[s] out to be like poison ivy" (Ibid.). People "clicked them once" and quickly learn never to touch them again. Maintaining the future promise of advertising thus forces it to act "like the flu," constantly changing to avoid triggering an immune response.

Rather than undermining the internet's business model, Cegloski argues, the proliferation of increasingly insistent forms of online advertising and the ongoing "crappiness" of targeted ads

in fact sustain it. This is because what actually powers internet startups is not advertising per se, but rather a promissory form that Ceglowski (Ibid.) calls Investor Storytime. For Ceglowski, advertising consists of paying someone to convince a product's users that "they'll be happy" if they buy it. Investor Storytime by contrast exists when someone pays you to convince them how rich they'll be when your startup finally starts selling ads. It works by convincing investors "that advertising in the future is going to be lucrative in ways it just isn't today."[127]

Under its logic, any failure of advertising is simply grist for a more convincing story. "It means there's vast room for improvement. So many stories to tell the investors." Consumers may hate ads now, but just wait til we finally have enough data to offer perfect targeting. If the targeting algorithms don't work, this can only mean more data will improve them. Even if they do work, imagine how much more valuable their ads could be if we only had better data. Silicon Valley chases personal data "not because it's effective now, but because we need it to tell better stories."

Together with advertising's 'poison ivy'-like quality, Investor Storytime drives advertising's self-propagating character online. Its "edifice of promises" can only be sustained if

---

[127] LinkedIn's 2004 pitch deck is one of the few I reviewed that contains a clearly identified business model, and it nicely illustrates Ceglowski's theory of Investor Storytime. Citing eBay, PayPal, and Google as examples of startups whose success followed from attracting a massive user base, the pitch deck identifies growing a user base as LinkedIn's first priority. The same slide notes that PayPal, with more than $400 million in annual revenue as of 2004, waited until it had 4 million registered users before "turning on revenue." A later slide showing rapid growth in the search advertising market between 2000 and 2004 describes how LinkedIn's advertising will be better than existing models, specifically that offered by Google's AdWords service. In three brief bullet points, the slide argues that LinkedIn's advertising will be more lucrative because of the quality of the information the site collects on each searcher, because it focuses on white-collar workers whose cost-per-click is significantly more lucrative than that of other demographics, and because of LinkedIn's ability to monetize user impressions.

companies constantly find new ways to make advertising more invasive and ubiquitous. This explains the expanding and intensifying frontier of online surveillance, the push to link online behavior to offline, and the enlistment of data brokers in assembling ever more granular and comprehensive data profiles. It explains the replacement of the popup with auto-playing video ads as the propagated effects of the internet's original sin.

The Cruel Intimacy of Unscalable Things

Having reviewed the institutional and cultural conditions that enable internet companies to practice their techniques of disavowal, we still have not accounted for entrepreneurs' anxious avoidance of the internet's business model. To do so, we must consider the role of hand-forged, deeply attentive, intimate user relationships in Silicon Valley's understanding of the necessary preconditions for startup growth. That is, we must explore the way growth's imagination compels entrepreneurs to immerse themselves in the real worlds of their users in precisely the messy, socially dense ways that Graham warned them to avoid.

It is a curious fact that despite the popular association of Silicon Valley with efficiency, rationalization, and automation, when prominent investors and entrepreneurs describe the work required to propel startups to success, they use the language not of engineering but of artisanal craft and courtship. Two key convictions appear to animate this use. The first is that in the quest for a defensible product that lots of people will want it is better to start with something a small number of people love passionately than with something that inspires only moderate enthusiasm from a larger audience. The love of initial users is a precondition for growth to the extent it both indexes the depth of the consumer need addressed by a startup's product and helps enlist a

startup's initial users as its champions.[128] The second is that to cultivate passionate love among initial users entrepreneurs must, in the words of a famous Y Combinator mantra, "do things that don't scale." In Silicon Valley, scaling generally refers to the process by which entrepreneurs build out a startup's institutional and technical capabilities as its users rapidly explode in number. It is sometimes described in shorthand as 'building the company that builds the product.' More precisely, scaling involves a dual movement of corporate expansion—hiring technical, support, and managerial staff; securing more server capacity—followed by efficiency-achieving rationalization and automation.

Doing things that don't scale means, by contrast, doing things 'by hand,' person-to-person. When an entrepreneur stands on the streets of Palo Alto, or goes door-to-door among its stores, personally flagging down and convincing passers-by to download and try her new app, she is doing something that doesn't scale. When, rather than simply emailing a download link, she says to a passerby who agrees to try the app, "Give me your phone and I'll set it up right now," she is again doing things that don't scale. Similarly, when an entrepreneur manually performs a service, which her new product purports to provide through the magic of yet-to-be-developed software, she is doing something that doesn't scale.

Doing things that don't scale may sound simply like running a business, but Silicon Valley abhors perceived inefficiencies as "unscalable" blockages standing between entrepreneurs and the mass market. Individually persuading strangers may be an effective means of recruiting

---

[128] Silicon Valley views such social endorsement as especially important because, in its own reflexive mythology, its best products will be initially novel, strange, or even trivial according to the sensibilities of the present. See also Moore (2003) for a brief history of "viral marketing" in the new economy.

initial users, but it is too time-consuming and labor intensive to drive exponential growth. Manually setting up webpages may be a great way to cultivate trust with new customers, but it's no way to tap the kind of mass market demanded by growth's imagination.

Nonetheless, Y Combinator teaches that to make the kind of product that users love passionately, founders must at least for a time seek out initial users and get extremely close to them. To the extent possible, entrepreneurs should work in the offices or homes of initial users. Barring such access, they should talk to their initial users as much as possible, send countless emails, and give highly personalized help in making the best use of the startup's product. Such labor-intensive personal care provides entrepreneurs with the opportunity to find the product roadmap in their users' mind—to observe the surprising ways they use a nascent product, probe them with questions, and ideally, take on and empathetically experience the problems of their initial users as their own. Through such intimate, manual work entrepreneurs test the hypothesis their startup idea represents about the world and refine the product into something that a small number of users will passionately love, while instituting the feedback mechanism necessary to continually identify and transform new and changing user needs into corporate decisions.

To a certain extent, the injunction to do things that don't scale simply reflects technological and economic trends, which since the 1990s have led to rising consumer demand for interaction with businesses and forced businesses to reconfigure their operations around

provision of newly customer-centric interaction and service (Wirtz et al. 2010; Teece 2009).[129] Aligning the qualities consumers attribute to a good with the qualities they desire and expect from it has always been an object of corporate practice. Under the competitive conditions introduced by the internet and globalization, however, the reflexive alteration of such qualities has intensified and moved to the center of business strategy (Nakassis 2013; Callon et al. 2002; Slater 2002a). For internet companies in particular, value creation increasingly hinges on establishing and managing intense emotional connections between consumers and corporate brands (Foster 2007).

Even against this backdrop, the emphasis in the discourse on doing things that don't scale on cultivating love is noteworthy. The language entrepreneurs and investors use to describe the process of enrolling, learning from, and catering to initial users consistently draws on and mobilizes ideals of romantic courtship, love, intimacy, and—in the insistence that entrepreneurs enter into mutual being with their users—even kinship (see Sahlins 2013). These ideals are evidenced, for example, in entrepreneurs' descriptions of the visceral emotions elicited by negative customer service feedback. As the entrepreneur Stanley Tang observed in How to Start a Startup, "it's painful" when a user reports having a bad experience with your product: "[T]hat something you love and put so much effort into, to know you got it wrong or somebody didn't

---

[129] Management scholars argue specifically that the global connectivity provided by the internet, the growth of internet-based commerce, and the international outsourcing of many business activities, reduced information asymmetries between buyers and sellers and brought more companies into competition with one another by expanding geographic markets (Porter 2001). The internet allowed businesses to offer—and in the case of software, distribute—their goods directly to consumers, eliminating the need for costly sales forces and reducing barriers to entry by new competitors (Ibid.). Together these developments served to offer consumers more choices in the market while demonstrating the value of user contributions to the new online platforms (Teece 2009).

treat them right." Just as in a romantic relationship, however, the trick is to be open to hearing what the customer is telling you, and to put in the effort to fix what needs to be fixed. "Problems are inevitable," according to Tang. "You're not going to have the perfect product; things are going to break; things are going to go wrong." What's important is to "always make it right, to always go the extra mile and make that customer happy." Graham (2013b) similarly urges entrepreneurs to bring a lover's passion to early user interactions. According to Graham, they should strive to make not just their products great, but the entire experience of being a startup's user, from discovery, to use, to customer service. Even if a startup's initial product is incomplete and buggy, by giving users an "insanely great experience…you make up the difference with attentiveness." This dictum applies not just to new and ongoing customers, but to "churn customers." When a customer leaves, according to Tang, entrepreneurs should reach out and find out why. That personal outreach "can make the difference between leaving and staying; sometimes people just need to know that you care and it's going to get better."

Entrepreneurs thus seek to cultivate love by showing love, love above-and-beyond what we generally think of as likely or even being possible from a corporation. In thus modeling their interactions with users on courtship, intimacy, and kinship, entrepreneurs enact a vision of startups as the kind of social actor capable of sympathetic fellow-feeling, of exhibiting the moral concern and self-restraint otherwise understood to be limited to individual persons. We can attribute the relative success of such efforts in part to the simultaneous refusal of many internet startups to attach a monetary price to their services, a refusal which in mimicking gift-giving reinforces the startup's enacted disavowal of self-interest. As Graham notes, however, startups' performative cultivation of love works primarily by implicit contrast to the dismal, highly

opaque interactions that continue to characterize the customer service offered by most large American corporations. According to Graham (2013b), Y Combinator has to teach entrepreneurs the importance of personally cultivating intimate delight because "they've never experienced such attention themselves." Their understanding of customer experience is informed by the misery that cable companies, airlines, insurance companies, for example, routinely inflict on the public. Against this backdrop, the simple ability of a startup to speak and respond to users with a real, live human voice carries a kind of magic that Y Combinator insists confers a competitive advantage on startups.

It is in the context of the intimate, loving relationships that entrepreneurs cultivate in pursuit of growth that the business model ultimately intervenes. For startups that grow, the period of intimate courtship inevitably gives way to the demands of scale. As sincere as the disinterest of some entrepreneurs in revenue may be, if they have availed themselves of venture capital financing to develop their technology and amplify its potential impact on the world, investors will inevitably demand that they 'turn on' the revenue streams specified in their business model. This fact of revenue, the fact that the potential represented by a rapidly growing user base accrued under the auspices of venture capital must inevitably be economically realized, is another open secret of Silicon Valley that is obscured and disavowed by the mythological foregrounding of entrepreneur's curiosity and techno-moral drive.

The rationalization and automation of person-to-person relations that follows when internet startups scale their operations and turn on their business model of course cannot fully

purge user-corporate relations of their intimacy.[130] Forms of intimacy no doubt persist in the care and concern users exhibit towards internet services, and are powerfully present in the loyalties mediated by internet brands (Nakassis 2013; Foster 2007). The open secret of revenue, however, reveals the form of intimacy enacted by doing things that don't scale to be a cruel one. It is cruel in the sense that it enacts an optimistic vision of the world in which corporations treats customers as ends in themselves, but does so under structural conditions that corporate agents know can never persist (see Berlant 2011). It is cruel when cultivated in the service of the internet's business model because unscalable interactions binds users to a socio-technical system, which in the course of its normal operation will necessarily undermine their privacy and exposing them to visceral disturbances, misinformation, and systematic attempts at manipulation. The cruelty effected by the internet's business model is particularly perverse, involving as it does the introjection of often anonymous, manipulative third parties into a relationship established under the pretense of a startup's other-than-business interest in its users' needs and being.

Advertising-fueled startups have convinced themselves over time that reconfiguring their services to accommodate advertisers in fact serves user interests by helping them to discover the goods necessary to fulfill their desires. The comfort they thereby derive of course rests on a wildly impoverished vision of human flourishing, one entirely circumscribed by the presumption that humans are defined by the curse of infinite, individualistic wants (see Graeber 2011). And despite the claim, Silicon Valley engineers have a long history of opposing advertising as unethical, antithetical to engineering ideals, and deeply compromising of the user experience.

---

[130] Zelizer (2010) rejects the historical tradition within social thought, which presumes that the introduction of instrumental means such as commodification and cost-accounting into intimate social relations necessarily depletes them of their intimacy.

A notable early example of this techno-moral antipathy is captured in one of the ur-documents of Silicon Valley, "The Anatomy of a Large-Scale Hypertextual Web Search Engine," published by Google founders Sergey Brin and Lawrence Page in 1998. Written when Brin and Page were still computer science Ph.D. students, the paper details the concepts behind what would become Google search. In the careful style of academic research, it lays out the features, technical architecture, and applications of the Google search system. As a historical document, the paper presciently anticipates both the web's rapid growth and the engineering and organizational challenges such growth would pose for web-based applications. It is further prescient, however, in anticipating the corrosive influence that technologists like Ethan Zuckerman would later assign to the internet's advertising-based business model. In an appendix that notably deviates in tone from the rest of the paper, Brin and Page argue that the tension between the goals of advertising and those of providing high quality search results simply cannot be reconciled. Advertising is so likely to introduce insidious, effectively invisible bias into search results that Brin and Page conclude in a historical irony that search engine technology should never be allowed to become beholden to the corporate form.

What is interesting for purposes of understanding the business model's ambiguous presence in Silicon Valley is that targeted advertising reigns as the presumptive business model for websites and internet startups, entrepreneurs continue to found startups predicated in part on the rejection of surveillance and advertising. This was true, for example, of Tumblr, WhatsApp, and Instagram, founded in 2007, 2009, and 2010, respectively. WhatsApp founder Jan Khoum attributed his support of user privacy to the fear of monitoring he experienced as a child in the Soviet Union. As a reminder not to compromise the messaging experience provided by

WhatsApp, Khoum famously kept a handwritten note taped to his desk reading, "No Ads! No Games! No Gimmicks!" By 2018, however, under pressure from their new corporate owners, each of these companies had implemented the internet's business model.[131] In September 2018, the news that Khoum and his co-founder had quit WhatsApp and were urging users to delete their Facebook accounts in protest of its broken privacy promises played out as an undercurrent to the Cambridge Analytica story. We can see in such events that guided by growth's imagination and facilitated by doing things that don't scale, venture capital-backed startups have channeled the internet's potential into a recurring cycle of cruelly intimate solicitation, seduction, and betrayal in which the public either colludes (see Schüll 2012) or to which, as Zuckerman suggested, it has simply become inured. If the anxious avoidance elicited by the internet's business model reflects the structural challenges it poses to entrepreneurs' techno-moral commitments to quality engineering and the user experience, it suggests that the techniques of disavowal that internet companies practice on their users first operate internally on founders themselves.

[131] Yahoo! purchase Tumblr in 2013. Facebook purchased Instagram and WhatsApp in 2012 and 2014, respectively.

CHAPTER 5: BUTTERY SMOOTH

In December 2017, from my home in New Orleans, I opened my laptop and launched Firefox, the web browser. Scanning through my hoard of open browser tabs, I found and played the recording of Mozilla's "All Hands" meeting.

Once a year, Mozilla gathers its globally dispersed staff for a week of face-to-face work, demonstrations, and planning. Opening the 2017 meeting in a spirit of reflection, Mozilla's CEO, Mitchell Baker, described the waning year as one of stabilization, renewal, and transformation—as "the year we relaunched Mozilla." Central to this optimism was the successful recent conclusion of Project Quantum.

A year-long, company-wide undertaking, Project Quantum marked Mozilla's effort to "modernize" Firefox and recapture its reputation for technological innovation. Its goal was to build a "parallelized" engine for Firefox, thus radically improving Firefox's "performance."[132] In web engineering, performance refers to a quality of browser operation encompassing both verifiable speeds and subjective perceptions of accessing, loading, and rendering webpages. As the meeting progressed, Mozilla executives summarized performance's stakes: "Speed," one observed, "it doesn't make, but it can kill a product." "Performance alone," another said, "isn't

---

[132] A browser's engine is its technical core, responsible for analyzing webpage files and determining what to display. Project Quantum's engine is "parallelized" in that it assigns key tasks required to convert web files into onscreen pixels for independent, simultaneously processing by any idle CPU core. It thus enabled Firefox to finally join its competitors in making efficient, performance-enhancing use of the increased processing power introduced into computers with the mid-2000s shift to multi-core CPUs.

enough to keep users." For Mozilla to build "the best browser in the world," however, it had to "take back the performance crown."

Despite Project Quantum's prominence, I was reviewing the All Hands' recording to observe discussions not of performance but rather privacy. I was curious what had become of certain privacy-related projects, which I tracked during my fieldwork at Mozilla between 2016 and 2017. Like all Mozilla events, the All Hands was peppered with references to the challenges facing privacy online. Particularly interesting, however, was a presentation by Luke Crouch, a young, Tulsa-based privacy engineer whom I had previously interviewed.

Luke's talk described a recent study on the website "breakage" caused by privacy-preserving browser features. The study responded, as Luke put it, to an emerging "desire to turn up privacy" in Firefox, to not just offer but compete on privacy as a potential market differentiator. It examined the extent to which eight existing privacy features, which remained effectively hidden from most non-expert users, actually disrupt the browsing experience.

Chief among the study's findings, Luke highlighted results relating to the Tracking Protection feature.[133] To his surprise, the study's Tracking Protection users reported fewer average breakage problems than did users in a control group:

> What's going on with this? In the notes we get from the control group, over and over, 'Something on the page is slowing the loading speed significantly.' Those of us familiar with the nature of the web and tracking? Yeah, it's the trackers. You can see how much they slow down the web…. So we can claim, maybe, Tracking Protection actually fixes websites? By blocking tracking events that slow them down.

---

[133] Tracking Protection protects privacy by temporarily suspending all network requests while comparing their destinations to a regularly-updated list of known tracking domains. Upon identifying a match, Firefox cancels the request, preventing the domain from transmitting trackers to the browser.

By trackers, Luke referred here to the ever-changing suite of technologies—cookies, beacons, scripts—which third parties embed in website elements to record and track users across the web. Per his allusion, trackers are notorious among web engineers for being "heavy" and slow loading. Despite their costs to performance and privacy, trackers, along with the ads, analytics, and share buttons that host them, enable much of modern websites' functionality—hence Mozilla's concerns with breakage.

Released in 2015, Tracking Protection had effectively languished in Firefox's expert settings. Now, Luke said, alongside Firefox Quantum, Mozilla had 'surfaced' it to Firefox's general use population. The finding that Tracking Protection in certain ways improves the browsing experience factored into this decision. As Luke anticipated, soon thereafter, a blog post highlighting Tracking Protection's increased availability gave its privacy and performance benefits equal billing: "We just can't stop making Firefox faster and with our most recent release, we also made it easier for you to control how much you're tracked" (Novak 2018). Nick Nguyen, Firefox's Vice President of Product, reiterated the connection, explicitly coding privacy in performance's sensuous image. With Tracking Protection, he wrote, "in addition to protecting their privacy, users actually have a better, faster experience with the web" (Nguyen 2018). Users enjoy "faster, always on privacy."

Performance's role in justifying and promoting Tracking Protection piqued my interest because it resonated with a certain slippage that I encountered frequently in California. When I arrived, years of interaction with privacy professionals had trained me to think of the internet as a surveillance technology as much as a communications or information technology. Relative to its boundless appetite for personal data (Masco 2017), the internet, whatever its value, is an

undeniable source of harm to privacy and to the values like liberty, dignity, self-cultivation, and self-expression, which privacy theoretically serves (see Nissenbaum 2010). Steeped in such perspectives, I was later struck to find that the web community's reflexive analysis often dwelled squarely on online surveillance's threats not to privacy but to the sensory experience of browsing and to user attention. If I thus encountered attention where political philosophy and American history instead predict privacy, Tracking Protection, it seemed, further predicated privacy on a certain sensuous alignment with or support for attention.

What does browser performance have to do with attention though? And what does privacy's coding in performance's sensuous image tell us about privacy's changing value and relevance under conditions of technological stewardship? Herein, drawing on ethnographic observations and archival analysis, I explore privacy's entanglement with attention in browser engineering. First, I describe performance's significance to Mozilla as a software "table stake," a browser quality, which like a poker table's minimum buy-in, represents a perpetually escalating threshold of competitive viability. Next, I detail the sensuous contours of the performance ideal pursued in Project Quantum. Firefox engineers, I show, understood such "smooth," "snappy" performance as necessary to facilitate the form of sustained user attention that is central to their techno-moral aspirations.[134] Drawing on the semiotic study of "qualia," i.e., experiences of abstract qualities as the raw material of the sensuous present (Harkness 2015; Chumley and Harkness 2013), I argue the ideal is instead the emergent product of Mozilla's efforts to

---

[134] I cite here Kelty's (2008) characterization of the intertwined moral and technical ideals of order shared by free software programmers. Mozilla's engineers similarly exhibit a dual movement between the technical and moral, drawing upon the form of web technologies to articulate organizational ideals for social, political, and economic life, and analyzing technology in explicitly moral terms.

pragmatically[135] guide users towards experiencing browsing's qualia as indexes of both Firefox's "innovative" technology and Mozilla's paternalistic stewardship of user interests. Turning to Tracking Protection and related examples of privacy's convergence with attention, I argue that performance engineering is a key site of privacy's material and semiotic "bundling" (Keane 2003) with attention, with the consequence that the experiential state of "having privacy" online increasingly stands in an iconic relationship to that of "paying attention." I conclude by exploring the likely changes to privacy's utility and value as its practically available forms increasingly exhibit qualities of speed and smoothness, authorized as "qualisigns"[136] of networked life by the fantasies of global capitalism.

Online Life is Real Life

Conducting fieldwork with an organization like Mozilla, with globally-dispersed employees and volunteers, entails various methodological challenges. One advantage, however, is that in bridging their asynchronous schedules and honoring open-source programming's transparent spirit, Mozilla's engineers effectively produce a continuous archive of Firefox's development. Now, turning to the blogs, videos, IRC channels, wikis, listservs, and development platforms through which "Mozillians" conduct and document their work, I reconstructed something of Project Quantum's unfolding and identified the contours of its animating performance ideal.

---

[135] I use pragmatic in its Piercean semiotic sense to refer to the indexical processes involved in typifying, situating, and making action meaningful.

[136] Defined as embodied qualities acting as signs, with privileged roles in social value systems (Munn 1986; Chu 2010).

Reviewing this archive makes apparent that Mozilla initiated Project Quantum at a time of perceived peril. Upon debut in 2004, Firefox was beloved for speed and customizability. In 2010, it held 25% of the US desktop browser market. By 2017, however, this had fallen to 12%. Worse, Firefox held less than 1% of the increasingly important mobile market (Shankland 2017). Competition from deep-pocketed rivals like Apple and Google accounted for much of this decline. Citing the comparative difficulty of loading interactive webpages and multimedia files in Firefox, however, web developers and users accused Firefox of having fallen into disrepair. Mozilla's decision to address its waning market share via performance responded directly to such grievances. It further reflected the best practices promoted in its own developer training materials. "The longer it takes for a site to respond," they provide—i.e., the worse its performance— "the more users will abandon [a] site" (MDN Web Docs 2020a).

If poor performance threatened Firefox's user base, it also threatened Mozilla's mission, and at a moment perceived to be crucial to the web's future. To appreciate this, it is helpful to know something about both browser operation and Mozilla's role on the web. After spreadsheet and word processing programs, browsers rank among software's most widely used applications. They nonetheless figure peripherally in popular imaginaries. Often confused with search engines and social media, browsers are the technology through which people generally access the websites that reside on the internet and collectively constitute the web. Browsers load webpage files from remote servers, render them on screen, and facilitate user interaction.

As references to its mission suggest, Mozilla approaches Firefox as both an engineering challenge and an ethical undertaking. Mozillians express the ethical aspects of their work in part by insisting that Firefox serve as a "user agent," rather than simply providing web access. As

Baker explained at the All Hands, by "brokering experience on [users'] behalf" relative to developers, advertisers, and other web stakeholders, Firefox "deliver[s] more user agency to people." It "put[s] people in control of their online life and shap[es] the internet for the good of people."

Per Baker, users needed such stewardship now more than ever because in recent years "we learned how much tech in our own world is being corrupted and used for purposes antithetical to our mission." The web's embrace by "bullies, trolls, stalkers, fascism, and violence" made clear "the internet at scale is not working so well for human beings." Worse, the proliferation of such "hostile elements" coincided with the web's social extension and intensification. Citing the sub-title of a Mozilla podcast on the social stakes of the modern web, Baker summarized this change: "Online life," she said, "is real life." The web emerged in the mid-1990s as a technical curiosity and hobbyist endeavor; by 2017, it "increasingly permeates everyday life." It had become an effectively unavoidable mediator of work, romance, leisure, and finance. Given the increasing correspondence between online and real life, the web's "user experience" now implicated, to paraphrase Baker, people's very ability to live healthy, productive lives. In such conditions, Firefox's performance woes hampered the effective representation of user interests, interests that now extended alongside the web into the central concerns of everyday life.

Silky, Buttery Smooth

In its 2016 announcement, a Mozilla executive wrote that Project Quantum would produce performance so improved, "so noticeable that your entire web experience will feel different"

(Bryant 2016). In pinning Quantum's promise to the changed feel of browsing with Firefox, the announcement indexed a historical shift in theories of performance. Historically, web engineers approached performance as easily quantified and defined by browser-internal speeds. Such "objective" performance (MDN Web Docs 2020b) was machinic and temporal. Objective metrics like "load," for example, measured the milliseconds between initiation and conclusion of key browser tasks, like downloading website resources.

Such metrics originated when the web consisted of little more than an interlinked set of static documents. In the mid-2010s, however, the web's increasing interactivity and new development technologies introduced new complexities into browsing. Web engineers and developers argued the user experience was progressively decoupling from performance's objective measures (see, e.g., Souders 2013).

Consider in this regard the talk delivered by performance expert Estelle Weyl in September 2017. Mozilla invited Weyl in connection with Project Quantum to discuss ongoing research into the metrics that best capture user-perceived performance. In her talk, Weyl challenged the assembled Mozillians to define performance. One might equate it with speed, she proposed. Because many contextual factors influence speed, however, no universal metric for it exists. Users with weak connections or old devices, for example, often experience website performance as slow despite its measurable speed in controlled programming environments. Targeting multiple metrics to such cumulative factors would still not address the fact that "fast means different things to different people." Objective metrics, Weyl concluded, usefully reflect "what the computer tells you, but they don't reflect what the user is seeing." They "are in a computer." By contrast, "it's what the human sees, the user experiences, that really matters."

In positing subjective perception as a necessary corrective to objective metrics, Weyl expressed the now widespread conclusion that whether browsing holistically feels fast matters more than browser-internal speeds. Even had browsing not become more delay-prone, users only experience its perceptible characteristics. The resulting need to cater to user perceptions folded biological and social concerns into performance's mechanic dimensions. Henceforth, for example, engineers would have to account for users' ambiguous descriptions of performance. Mozilla researchers thus observed (Strohmeier and Kirschner 2017) that for engineers "performance is about numbers and processes" measured relative to individual tasks. Firefox users, meanwhile, often characterize performance using "unspecific" terms like "'this is amazingly fast,' 'SLOW!,' 'WOW!,' and NO!'"

To account for subjective perceptions, engineers and researchers began, circa 2015, to develop new user-centric performance frameworks and metrics. Like Google Chrome's influential RAIL model, the new approaches generally proposed to reorient performance engineering around the question, "What does the user feel?" (Irish and Lewis 2015). RAIL, for example, urged developers to "decompose" websites into browsing's key interactions (scrolling, dragging, clicking, and animations), optimizing each against a response-time "budget." Like RAIL, Project Quantum's performance metrics used targets adopted from industrial research into the perception of computer application delays, originally conducted in the 1960s, but later popularized by Jakob Nielsen, a prominent user researcher (see Nielsen 1993).

According to Nielsen, the ability to perceive response delays to user inputs adheres to a series of natural thresholds. If a website doesn't respond to user input within 100ms, for example, users will perceive a slight delay. Delays longer than 1000ms, however, will cause

users to lose focus on their task. After 10000ms, they will perceive the task to be "broken" and abandon it in frustration. Based on these thresholds, Nielsen recommended (2009) that website response times be limited to 100ms to preserve the illusion of directly manipulating an interface, to 1000ms to maintain the flow of task engagement, and to 10000ms to sustain user attention. The metrics incorporating these thresholds purported to map the interactional primitives that compose the browsing experience to innate facts of human perception, cognition, and attention (see Grigorik 2015).

In developing the new user-centric metrics, browser engineers and developers articulated a performance ideal with notable aesthetic and affective contours. Related presentations and documentation consistently characterized high-performance web technologies as "light," "natural," "crisp," and "snappy." Browsers should always be fast, they argued, but scrolling and animations should be additionally "buttery" and "smooth." As these mixed qualities suggest, engineers now approached performance as registering not visually or haptically but cross-modally (Ballestero 2019; Csordas 1993). Engineers argued, for example, that users experience "smoothness" when browsers consistently render animations at 60 frames-per-second. But smoothness also refers to a tactile quality of scrolling, tapping, and dragging. Engineers meanwhile promoted browser "responsiveness" as necessary to induce the satisfying illusion of directly manipulating website elements, and thus of being "in control." So articulated, performance runs along an intertwined spectrum from slow, painful and frustrating to fast, pleasurable, and efficacious.

With its references to the changed feeling of using Firefox, Project Quantum's announcement invoked this ideal, indexing "silky smooth" performance's centrality to Mozilla's

techno-moral aspirations. The "quantum leap" in performance would be so pronounced, the announcement specified, that users wouldn't be able to avoid noticing the difference. Reflected in this latter claim was the anxiety that given Firefox's complexity, the ambiguities of perceived performance, and browsers' marginal status, Project Quantum's accomplishments might pass unnoticed. To ward against this, Mozilla employed an array of techniques to stage Firefox's improved performance so as to draw reflexive user attention. Project Quantum's closing stages, for example, included "Quantum Flow," an unprecedented mobilization of staff and volunteers into a performance-debugging "strike force." Organizers warned that bugs lurking in the codebase might prevent Firefox's performance gains from surfacing to user perception, thus bringing Quantum's heroic engineering to naught.

"Photon," another sub-project, used research on user expectations to redesign Firefox's interface. Having found slow-closing tabs to frustrate users, for example, Mozilla re-sequenced the tasks involved. All "clean-up" tasks now happened "off camera," i.e., after tabs visually appear to be closed. Thus poetically aligning the browser's interface and task execution logic with user expectations directly improved perceived performance. Photon, however, also redesigned Firefox's icons, tabs, and toolbars. At the All Hands, an executive observed of this 'modernization,' "We knew we'd need a visual refresh to see this performance. If it looked like the browser of two years ago, even if faster, no one would notice. They wouldn't perceive the performance." Reshaping Firefox's tabs didn't improve performance per se. Firefox designers hoped, however, that perceptible changes in Firefox's appearance would orient users to the additional changes in the overall feel of browsing. Ideally, their coincidence across sense

modalities would induce users to analogically project (Harkness 2015) the modernized quality of the redesigned interface onto Firefox's faster, smoother feel.

Through such staging, Mozilla hoped to guide the public to experience Quantum's performance improvements as an irruptive, sensuous break from ordinary browsing. The desire to produce such attention-demanding performance was temporary but is nonetheless noteworthy. For Neilsen, perceptual thresholds mattered precisely because perceived delays 'steal focus,' drawing attention from and disrupting user tasks. As reflected in its adoption of Nielsen's thresholds, it was to create the conditions for user attention to reliably emerge and adhere online that Mozilla marshaled its resources to improve Firefox's performance. By producing performance in a form understood to allow the browser's mediating presence to recede into the experiential background, Mozilla hoped to facilitate agentive, empowering attention-to-content, partially fulfilling Firefox's obligations as a user agent.

Thus mobilized, performance operates through content-neutral processes. For browser engineers, high performance web technologies sustain attention not by algorithmically targeting users with outrageous content, as popular critiques of the "attention economy" suggest (see Williams 2016). They rather do so through perceptual and affective subtraction and deferral. Fast, smooth performance theoretically holds unpleasant perceptions, feelings, and thoughts in abeyance. It softens and soothes the pain and annoyance of slow-loading, "janky," and jittery webpages. Good browser performance thereby relaxes users, priming them to further personal goals using web-based resources. Only through such self-effacement of the browser's mediating presence, performance engineers presume, can the perceptual and affective space be created for attention to adhere, as users intend, in website content.

The value Project Quantum attributed to the qualities of smoothness and speed thus reflected their perceived ability to forward Mozilla's techno-moral goals. As the semiotic study of qualia shows, however, sensuous qualities are not simply harnessed to political projects but also emerge from them (see Gal 2013; Chumley 2017). Though reflexively taken to be given properties of things (Gal 2017), qualia—embodied, experienced instances of abstract qualities and feelings—are in fact cultural emergents (Harkness 2015). Sensations only become palpable and persuasive as sensations of abstract qualities like smoothness and snappiness through institutionalized practice (Chumley and Harkness 2013). Social actors engage in such conventionalization, expending effort, for example, such that browsing's sensations might be experienced as resembling the smooth textures of silk cloth and warm butter, because qualia have pragmatic consequences (Harkness 2015). Attributing qualities like "hardness" and "fanciness" to varieties of linguistic forms, for example, may signal person-types as social identities, anchoring reflexive, group-defining conduct (Gal 2013). Qualia further signal the ontological categories that apply to material objects, grounding forms of practical engagement with them (see Chumley 2017). Finally, qualia help stipulate the kinds of social relations understood to adhere between people (Lemon 2013).

Here, in developing and disseminating perceived performance's new metrics and practices, browser engineers provided the systematic instruction necessary to guide users to a subset of browsing's innumerable potentially-noticeable properties and to experience them in terms of abstract qualities like speed, smoothness, and snappiness. Mozilla's promotion of Project Quantum, with its explicit directives ("The first thing you'll notice is the speed. Go on, open some tabs…") and invocations of engineering prowess ("brand new tech stolen from our

advanced research group") and user empowerment ("designed to get out of the way and let you do what you do best") built on such work. They established the pragmatic links necessary for users to interpret performance's qualia as meaningful signs indexically establishing the "innovative" quality of Firefox's technology and the empowering care with which Mozilla represents user interests.

No Surprises

In October 2016, I listened remotely as Mozilla's privacy professionals conducted a privacy and security training for new engineering hires. Stacy Martin, a privacy expert on Mozilla's public policy team, opened the training by reading from Mozilla's Manifesto. Originally drafted in 2007, the Manifesto summarizes Mozilla's techno-moral vision for the web. Principle Four provides, "Individuals' security and privacy on the internet are fundamental and must not be treated as optional." It attests, Stacy explained, to privacy's critical importance, with privacy defined as users' ability to "control" Firefox's collection and use of personal data.

Stacy next presented Mozilla's Data Privacy Principles, developed in 2010 as a practical guide to preserving privacy during product development. The Data Privacy Principles direct engineers, for example, to minimize user data collection and to maintain security defenses in multiple layers. More notable, perhaps, is the Principles' "No Surprises" pledge. In privacy and security engineering, data minimization and defense-in-depth both represent widely recognized best practices. No Surprises, by contrast, is less an implementable engineering strategy than a shorthand gauge of a company's effectiveness in mitigating privacy risk. Despite No Surprises' prominence as an informal guiding mantra of corporate privacy projects throughout the Bay

Area, Stacy did not dwell on it. She described it as expressing Mozilla's commitment to transparent, value-conferring data practices, and then moved on.

As Stacy's reference to user knowledge suggests, if Mozilla, like most Bay Area tech companies, defines privacy as data control, it practically interprets control in terms of responsibilized choice (see Hull 2015). Privacy obtains under this interpretation when a company's legal disclosures and settings options conspire to provide users with "meaningful" choices regarding the collection, use, and sharing of personal data. Because corporations must not just offer but honor such choices, privacy further requires aligning the content of their disclosures and the operation of their data systems.

For Mozilla, with its sprawling technological and organizational infrastructures, aligning user understanding, legal disclosure, and system operation is no easy feat, hence No Surprises' heuristic appeal. To surprise a user might indicate, for example, that Firefox's terms of service fail to comprehensively communicate some data practice. Alternatively, user surprise might reveal that in actual operation Mozilla's data operations contravene its policies. Mozilla's software development system, Bugzilla, is indeed replete with bugs citing surprise as evidence of potential privacy problems. Engineers' personal experiences motivate some such bugs ("I found this behavior to be surprising. I would have expected…"). In other instances, bugs cite user reports ("Users clicking links are sometimes shocked when…"), and even imagined possibilities of surprise ("I am concerned that users will not expect this behavior…).

Like the embodied perceptions of pulse employed diagnostically in Chinese medicine (Farquhar 2013) and the sounds oceanographers use to coordinate deep-sea exploration (Helmreich 2007), the feeling of surprise functions for browser engineers as a form of embodied

practical knowledge (see Harkness 2015). Despite the political and economic valences Stacy attributed to No Surprises, as a practical tool, it thus operates experientially. Surprise might lead engineers through a chain of inferential reasoning to identify unanticipated privacy risks or violations. To do so, however, it must register phenomenally, drawing attention to itself and away from user tasks. Despite its practical utility then, surprise also constitutes the kind of sensory diversion or intrusion that performance engineers aspire to eliminate in attention's service. No Surprises thus effectively grounds the presumptive sufficiency of corporate privacy practices in evaluations of their effects on user attention. It authorizes as socially reasonable the presumption that, so long as a corporation's use of personal data have not interrupted user attention, it is effectively providing users privacy.

At the All Hands, during Luke Crouch's Tracking Protection presentation, it was just such a grounding of privacy in attention, which piqued my interest. Such grounding, however, was only one of a number of unusual ways in which technologists characterized privacy in terms of attention. In perhaps the simplest example, technologists invoked attention where one might instead expect to find privacy, effectively treating attention as privacy's substitute. Consider as illustration a pair of structurally analogous explanatory logics, which circulate in parallel among Bay Area technologists. Circa 2012, technology professionals began to mobilize these logics to grapple with rising public disenchantment with the internet. In the abstract, they explain certain increasingly prominent social harms to be unintended but unavoidable side effects of the modern, commercial internet. Characteristically, they identify as the ultimate driver of such harms not internet corporations, executives, or technologies, but rather the "internet's business model." By

this they mean the surveillance-driven targeted advertising that internet startups generally adopt to generate revenue from their "free" products and services.

When I first encountered this logic, I encountered it in the form of an explanation of privacy's contemporary precarity (see e.g., Zuckerman 2014). As I designed my study in 2014, in articles, speeches, and everyday conversation, engineers and computer scientists argued that privacy's ongoing erosion follows not from the internet's surveillance technologies but from the value logic that animates them (see e.g., Schneier 2013). Such explanations figure the business model as a corrosive force that irresistibly compels internet companies to deploy ever-more invasive surveillance, often despite their stated social commitments.

Given surveillance's strong association with privacy in American popular and legal imaginaries, I was later surprised to encounter a second, equally prominent version of this explanatory logic. This alternative formulation similarly traces the internet's harms to its business model. In so doing, it cites the same business and technology developments (the invention of cookies; the launch of Google's AdWords business) identified as signposts of the business model's encroachment upon privacy. This second version, however, elides privacy, characterizing the internet's business model not in terms of the sale of personal data, but rather the resale and manipulation of human attention (see, e.g., Lewis 2017; Kulwin 2018). When naming the business model's harms, it similarly substitutes user time and attention for user privacy.

The parallel circulation of these explanatory logics among technologists suggests a structural interchangeability of privacy and attention on the modern web. It implicitly posits freedom from distraction as a national civic value on par with the sanctity of privacy. No

Surprises and Tracking Protection similarly figure privacy in relation to attention, but as a dependency rather than a substitute. With No Surprises, the occurrence or non-occurrence of a particular form of attentional intrusion, surprise, presumptively indexes the relative sufficiency of a corporation's privacy efforts. Tracking Protection similarly predicates privacy on a certain relationship to attention. Here, privacy's availability rather than its sufficiency is at stake. By 2017, Tracking Protection's ability to effectively mitigate a pervasive privacy threat was well-established. It was only when its protections were shown to improve rather than impede attention-sustaining performance, though, that Mozilla released it widely. As figured in Mozilla's promotion of Tracking Protection, privacy's apparent dependence upon attention appears more precisely as an aesthetic and affective resemblance. Nick Nguyen's identification of Project Quantum's "faster, always on privacy," for example, implicitly instructs Firefox users to interpret performance's qualia as meaningful signs of privacy's presence. It thus presupposes privacy to be the kind of thing, like browser performance, which can and should 'have' the sensuous qualities valued as conditions of sustained attention. The issue is not so much that privacy cannot exist on the web without attention as that privacy must appear, if at all, in an aesthetic and affective form attuned to attention's demands.

These convergences and slippages suggest that privacy's increasing salience in American imaginaries as a property of internet technologies coincides with its recoding, through performance, in attention's sensuous image. Technological stewardship has enrolled privacy in corporate efforts to calibrate users' embodied engagement with networked products and services. With this enrollment, the nature of interventions on privacy's behalf theoretically shifts. Efforts to protect privacy must still address the will and intellect—to provide meaningful data control,

privacy choices embedded in technology settings must be "informed" and "intentional"—but also reach 'beyond' will and intellect to the nervous system. To facilitate the emergence of qualia experienceable as 'having privacy' or 'being private,' technological interventions on privacy's behalf must be calibrated against the supposed facts of human perception and attention.

Consider further the insights of Webb Keane on the semiotics of material objects. Building like scholars of qualia on Charles S. Peirce's semiotics, Keane (2003) insists upon signification's materiality. Rather than being merely coded messages, he argues, signs are necessarily located in concrete circumstances. As signs and the sensuous qualities that carry them cross contexts, their salience, value, and utility are therefore inherently vulnerable to reconfiguration. In the process Keane calls "bundling," the fact of material co-presence opens qualities embodied in objects to pragmatic linkages with other material qualities in ways enmeshed in social systems of value and authority. Bundling is thus not determinative but contingent, a possibility that "remains available, ready to emerge as real factors" through socially-realized convention (Keane 2003, 188-9; see Gal 2013).

Taking signification's materiality seriously suggests that understanding privacy's changing cultural contours requires consideration of its new material manifestations. Technologists generally analyze the form that privacy might take in a networked world through the rubrics of professionalized engineering. How, they ask, can privacy's legal definitions be translated into precise technological requirements and identifiable properties of large-scale technical systems (see, e.g., Gürses 2014)? From Keane's perspective, a more proximate question is, how does privacy's embodiment in internet technologies—and the new forms of practical engagement with privacy this enables (Gal 2017)—open it to sensuous qualities and

interpretative possibilities different from those available via privacy's iconic historical embodiments, such as a house's walls?

Through material co-presence, then, tech-based efforts to embed culturally recognizable forms of privacy in web technologies open privacy's qualia to bundling with other qualities embodied therein. If institutional guidance is necessary for sensations to be experienced as sensations of qualities, it is also required to establish the qualic frame of reference linking social categories in relations of formal similarity (Gal 2013; Harkness 2013). Whatever else may be aligning privacy with attention, performance engineering provides a material and semiotic bridge between them. Through it, the sensuous qualities of speed and smoothness, which already operate as meaningful signs indexically pointing to and standing in iconic resemblance of the experiential state of 'paying attention,' are bundled with qualities signaling 'having privacy' or 'being private.' It thus enables performance's qualia to function as a "semiotic relay" (Harkness 2013) linking privacy to attention in a potential if not always actualized relationship of aesthetic and affective iconicity.

Attention's Insults, Privacy's Injuries

> *"Online, we often visit sites that track us, but it isn't clear when this is happening or how the information is being used. Adding insult to injury, this often invisible tracking actually slows down web pages." —Nick Nguyen (2018)*

As the All Hands progressed, Mozilla's executives shifted focus to how having taken "away the barriers to giving Firefox a try," Mozilla would next "make it hurt to put [Firefox] down." To loud applause, a Firefox Vice President, announced that Mozilla's 2018 strategy contemplated ad blocking:

> It's something we've cared about for a long time. It all fits in the category of annoyance mitigation…. There's a lot of stupid shit on the web that makes it annoying to use. Ads, videos, Javascript being injected by your broadband or Starbucks somewhere. The list goes on. We have to look at this space as an existential threat to the web and browsers entirely….[A] year from now, if we can't show how we made the web dramatically better for the end user by getting rid of…these stupid fucking annoyances, we'll have failed.

As Mayo and his audience well knew, the website elements identified here as existential annoyances host the trackers that make the web so inimical to privacy. Across its whole, the All Hands clearly acknowledged surveillance-driven advertising's privacy-undermining effects. And yet it was not unusual for web professionals, in analyzing 'what has gone wrong' with the modern web, to dwell like Mayo on the sensuous intrusions of lag, jank, pop-ups, and auto-playing ads, and elide the privacy invasions also thereby indexed. For many web developers, it seemed, the web's privacy harms mostly merit mitigation when they coincide with attentional insults.

I share this observation not to uncritically favor privacy over attention as a social value but to begin to sketch how privacy's social image shifts as it is bundled with attention. For Keane (2003), when abstract social concepts like privacy becomes embodied, co-presence necessarily subjects them to bundling, but only contingently so. The potential for resemblance inherent to all sensuous qualities might never be realized as an iconic sign. Even when so, resemblance can only be with respect to certain qualities, selected through social struggle. Having identified privacy's bundling with attention, the task thus becomes to (i) analyze the historical circumstances and regimes of value (Gal & Irvine 2019) that privilege browsing's speed and smoothness as qualisigns (Munn 1986) of contemporary networked technologies and their users (see Chu 2010), and (ii) identify any resulting changes in privacy's relevance, value, or utility.

To these ends, it is helpful to begin from something of privacy's historical and conceptual relationship with attention. Throughout American history, privacy has repeatedly emerged as an object of care and concern apparently endangered by new technologies (Nissenbaum 2010; Nelson 2001). Political philosophy and jurisprudence generally justify privacy as a good in itself and as necessary to cultivate dignity (Bloustein 1964), intimacy (Cohen 2002), personhood (Allen 1988), and moral autonomy (Habermas 1991). American law and culture, however, most prominently mobilize privacy in its guise as a bulwark of liberty to designate vital areas of home and family life as beyond government intrusion (Whitman 2004).

Scholars occasionally include attention among the social interests served by privacy. Bok (1982) and Gavinson (1980), for example, argue privacy protects individuals against unmediated external access, whether physical, informational, or attentional (see also Nippert-Eng 2010). Such works, however, primarily treat attention as a privacy problem when individuals become subject to undue attention, not when their own attention is interrupted or diverted (but see Gavinson 1980, 446). Nor is the form of attention invoked in them entirely consonant with the kind of cognitive affordance implicated by performance engineering and critiques of the attention economy (see, e.g., Carr 2010; Hayles 2007).

Indeed, what is most noteworthy in privacy's comparative history with attention is the parallel paths the two concepts travel, with each operating as a register of technologically-induced social anxieties. As Crary shows (1999), like privacy, attention historically emerged as a social problematic in relation to new technologies. However, where privacy's history is keyed to that of surveillance technologies, including photography, the telegraph, and telephone, attention's

instead tracks technologies of display, projection, attraction, and recording (see also Turner 2013).

Browser engineers undoubtedly desire high-performance web technologies for their own purposes and satisfaction. They express sensitivity, for example, towards infringements of the uninterrupted programming time, which programmers in general prize as necessary to "translate between the chaos of human life and the line-by-line world of computer language" (Ullman 2017, 4). As Crary (1999) argues in his history of vision, however, both the idealization of sustained attentiveness as a condition of creativity, and the conceptualization of attention as manifesting the autonomous will, are deeply historical. Individuals only came to be understood in terms of a capacity for "paying attention" as a result of 19th century modes of industrial production (Crary 1999; 1992). By newly requiring workers to rapidly shift between an array of productive and spectacular tasks, 19th century industrial capitalism imposed the imperatives of concentrated attentiveness to which we remain beholden. As products of efforts within capitalist modernity to make the perceiving body productive and manageable, our norms and ideals of attention constitute a disciplinary regime. Attention persists as a problematic as capitalism relentlessly challenges us to adapt to new structures of perception, forms of stimulation, and flows of information (Crary 1999; 2014).

For Crary (2014), ongoing acceleration in product innovation and consumption has only intensified contemporary capitalism's demands for constant engagement, project-oriented life management, and social capital cultivation (see Boltanski and Chiapello 2005). Capitalism's most recent mutation, he argues, has gone "24/7," normalizing an ideal of absolute availability. In consequence, social identity now normatively conforms to a form of continuous operation

modeled on market temporalities. It is thus more than incidental that browser engineers not only build and passionately use the web, but like many professionals today, also rely on it as a medium of entrepreneurial selfhood (Rose 1998; English-Lueck 2002). Performance's speed and smoothness may most directly cite the perpetually-circulating flows of global information networks. Performance's tethering to attention, however, and attention's social origins in the rationalization of laboring bodies, suggest it is the capitalist fantasy of ceaseless, frictionless exchange, which privileges speed and smoothness as qualisigns of networked life. Through browser engineering, performant technologies are made, which in turn produce performant persons. Just as Project Quantum's parallelized engine enabled Firefox to efficiently use the full computing power of multi-core CPUs, so too does Firefox Quantum's improved performance empower users to efficiently use web resources in service of dexterous, continual self-administration. The prized qualities of speed and smoothness extend in this way from browser operation to users, aligning both with the market's rhythms on the promise of enhanced functionality (see Crary 2014, 45, 74).

As observed, the projection of performance's qualia onto privacy implicates the nature of privacy's sensuous presence in the world. Indeed, one implication of privacy's bundling with attention is the return to individuals of an embodied capacity to register privacy. In pre-internet life, perceptible signs—the opening of a blind; the pointing of a stranger's camera lens—provided direct experiential access to unfolding privacy violations. By contrast, privacy's re-circuiting through the properties of technical systems like the internet frustrates this ability, obscuring the visible-yet-unverifiable Panoptic gaze (Chun 2006, 9). Still, conventionalizing browsing speed and smoothness as signs of privacy primarily serves to register privacy's

presence—to register the experiential state of 'being private'—and even then only presumptively so. Meanwhile, as illustrated by engineers' practical use of surprise, it is through privacy's violation, and the forms of embodied anxiety and alarm thereby elicited, that privacy's political capacity to motivate social action becomes potent (see Nissenbaum 2010). A world in which the practically-available forms of privacy are valorized, like networked technologies and users, for 'performing' smoothly and speedily is one in which privacy is compatible with the adaptable self-administration and efficient technological use required via attention. But it is also one in which privacy-based surprise and alarm is ideally suppressed, and in which privacy thus loses the ability to mediate power relations and shape moral imaginaries (see Gal 2005). For individuals, privacy's bundling with attention represents a kind of injunction against semiosis, a suppression of the perceptible contrasts, which animate the ideological conjectures of social struggle (Gal & Irvine 2019, 19). Collectively, by precluding attentional events that might 'spill over' into the formation of a public, it inhibits the formation of "new political occasions" (Braun & Whatmore 2010, xxivi) and "political awakenings" (Crary 2014, 24).

BIBLIOGRAPHY

Abbate, Janet. 1999. *Inventing the Internet*. Inside Technology. Cambridge, Mass: MIT Press.

Abelson, Harold, Ross Anderson, Steven M. Bellovin, Josh Benaloh, Matt Blaze, Whitfield Diffie, John Gilmore, et al. 2015. "Keys under Doormats: Mandating Insecurity by Requiring Government Access to All Data and Communications." *Journal of Cybersecurity* 1(1): 69–79. https://doi.org/10.1093/cybsec/tyv009.

Acar, Gunes, Christian Eubank, Steven Englehardt, Marc Juarez, Arvind Narayanan, and Claudia Diaz. 2014. "The Web Never Forgets: Persistent Tracking Mechanisms in the Wild." In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, 674–89. New York: Association for Computing Machinery. https://doi.org/10.1145/2660267.2660347.

Andrus, Jennifer. 2015. *Entextualizing Domestic Violence: Language Ideology and Violence against Women in the Anglo-American Hearsay Principle.* Oxford Studies in Language and Law. New York, NY: Oxford University Press.

Agha, Asif. 2007. *Language and Social Relations*. Cambridge: Cambridge University Press.

Agrama, Hussein Ali. 2012. *Questioning Secularism: Islam, Sovereignty, and the Rule of Law in Modern Egypt*. Chicago; London: University Of Chicago Press.

Agre, Philip E. 1995. "Conceptions of the user in computer systems design." In *The social and interactional dimensions of human-computer interfaces*. Peter J. Thomas, ed. Cambridge Series on Human-Computer Interaction. Cambridge, UK: Cambridge University Press.

ACLU v. Clapper. Brief Amicus Curiae of Experts in Computer and Data Science. March 13, 2014.

Aleinikoff, T. Alexander. 1987. "Constitutional Law in the Age of Balancing." *The Yale Law Journal* 96(5): 943–1005. https://doi.org/10.2307/796529.

Allen, Anita. 1988. *Uneasy Access: Privacy for Women in a Free Society.* Totowa, NJ: Rowman and Littlefield.

Altman, Sam. 2014. "Ideas, Products, Teams and Execution Part I." CS183B: How to Start a Startup. Class lecture at Stanford University, Palo Alto, C.A., Fall Quarter. https://www.youtube.com/watch?v=CBYhVcO4WgI.

Amoore, Louise. 2013. *The Politics of Possibility: Risk and Security Beyond Probability*. Durham: Duke University Press Books.

Anderson, Chris. 2009. *Free: The Future of a Radical Price*. New York: Hyperion.

Appadurai, Arjun. 1988. *The Social Life of Things: Commodities in Cultural Perspective*. Cambridge, England: Cambridge University Press.

Auer, Peter. 1992. "Introduction: John Gumperz' Approach to Contextualization." In *The Contextualization of Language*. Peter Auer and Aldo Di Luzio, eds. Amsterdam; Philadelphia: J. Benjamins.

Baker, Winifred Mitchell. 2014. Interview with Marc Weber. Computer History Museum. https://www.computerhistory.org/collections/catalog/102740019.

Ballestero, Andrea. 2019. "Touching with Light, or, How Texture Recasts the Sensing of Underground Water." *Science, Technology, & Human Values* 44(5): 762-85.

Ball, James, Julian Borger, and Glenn Greenwald. 2013. "Revealed: How US and UK Spy Agencies Defeat Internet Privacy and Security." *The Guardian*, September 6, 2013, sec. US news. https://www.theguardian.com/world/2013/sep/05/nsa-gchq-encryption-codes-security.

Bamberger, Kenneth A., and Deirdre K. Mulligan. 2015. *Privacy on the Ground: Driving Corporate Behavior in the United States and Europe*. Information Policy Series. Cambridge, MA; London: The MIT Press.

Bamford, James. 1982. *The Puzzle Palace: A Report on America's Most Secret Agency*. Boston: Houghton Mifflin.

Bankston, Kevin S. 2015. "Written Testimony." Hearing on Encryption Technology and Possible U.S. Policy Responses: Before the U.S. House of Representatives Subcommittee on Information Technology of the Committee on Oversight and Government Reform, April 29, 2015.

Barlow, John Perry. 1996. "A Declaration of the Independence of Cyberspace." Electronic Frontier Foundation. https://www.eff.org/cyberspace-independence.

Barnes, Richard, Bruce Schneier, Cullen Jennings, Ted Hardie, Brian Trammell, Christian Huitema, and Daniel Borkmann. 2015. "Confidentiality in the Face of Pervasive Surveillance: A Threat Model and Problem Statement." Request for Comments RFC 7624. Internet Engineering Task Force. https://doi.org/10.17487/RFC7624.

Barry, Andrew, and Don Slater. 2002. "Introduction: The Technological Economy." *Economy and Society* 31(2): 175–93.

Barth, Adam, Anupam Datta, John C. Mitchell, and Helen Nissenbaum. 2006. “Privacy and contextual integrity: Framework and applications.” *Proceedings - 2006 IEEE Symposium on Security and Privacy* 184-198. https://doi.org/10.1109/SP.2006.32

Barton, John and Tim Kindberg. 2001. “The Challenges and Opportunities of Integrating the Physical World and Networked Systems.” HPL Technical report HPL-2001-18.

Bellovin, Steven M., Matt Blaze, Sandy Clark, and Susan Landau. 2013. “Going Bright: Wiretapping without Weakening Communications Infrastructure.” *IEEE Security & Privacy* 11(1): 62–72. https://doi.org/10.1109/MSP.2012.138.

Bellovin, Steven M., Matt Blaze, Susan Landau, and Stephanie K. Pell. 2016. “It’s Too Complicated: How the Internet Upends Katz, Smith, and Electronic Surveillance Law.” *Harvard Journal of Law and Technology* 30(1): 1-101.

Benson, Peter. 2010. “Safe Cigarettes.” *Dialectical Anthropology* 34(1): 49–56.

Benson, Peter, and Stuart Kirsch. 2010. “Corporate Oxymorons.” *Dialectical Anthropology* 34 (1): 45–48.

Benthall, Sebastian, Seda Gurses and Helen Nissenbaum. 2017. “Contextual Integrity through the Lens of Computer Science.” *Foundations and Trends in Privacy and Security* 2: 1-69. 10.1561/3300000016.

Berners-Lee, Tim, Michael Dertouzos and Mark Fischetti. 1999. *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web By Its Inventor*. San Francisco: HarperSanFrancisco.

Berlant, Lauren. 2011. *Cruel Optimism*. Durham [N.C.]: Duke University Press.

Bernal, Victoria. 2005. “Eritrea on-line: Diaspora, cyberspace, and the public sphere.” *American Ethnologist* 32: 660-75. https://doi.org/10.1525/ae.2005.32.4.660.

Biagioli, Mario. 1990. “The Anthropology of Incommensurability.” *Studies in History and Philosophy of Science* Part A 21(2): 183–209.

Blaze, Matt. 2013. “Phew, NSA Is Just Collecting Metadata. (You Should Still Worry).” *Wired*. Accessed October 13, 2022. https://www.wired.com/2013/06/phew-it-was-just-metadata-not-think-again/.

2011. “Key Escrow from a Safe Distance: Looking Back at the Clipper Chip.” In *Proceedings of the 27th Annual Computer Security Applications Conference*, 317–21. ACSAC ’11. New York, NY: Association for Computing Machinery. https://doi.org/10.1145/2076732.2076777.

Bloustein, Edward. 1964. “Privacy as an Aspect of Human Dignity: An Answer to Dean Prosser.” *New York University Law Review* 39: 962–1007.

Bok, Sissela. 1982. *Secrets: On the Ethics of Concealment and Revelation*. New York: Pantheon.

Boltanski, Luc and Eve Chiapello. 2005. *The New Spirits of Capitalism*. London: Verso.

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

Boyle, James. 2019. “Is the Internet Over?! (Again?).” *Duke Law & Technology Review* 18: 32-60.

Braun, Bruce and Sarah J. Whatmore. 2010. *Political Matter Technoscience, Democracy, and Public Life*. Minneapolis: University of Minnesota Press.

Bridges, Khiara M. 2017. *The Poverty of Privacy Rights*. Stanford, CA: Stanford Law Books.

Briggs, Charles. L. 1986. *Learning how to ask: A sociolinguistic appraisal of the role of the interview in social science research*. Cambridge University Press.

Briggs, Charles L., and Richard Bauman. 1992. “Genre, Intertextuality, and Social Power.” *Journal of Linguistic Anthropology* 2(2): 131–72.

Brin, David. 1998. *The Transparent Society: Will Technology Force us to Choose Between Privacy and Freedom?* New York: Basic Books.

Browne Simone. 2015. *Dark Matters: On the Surveillance of Blackness*. By. Durham, N.C.: Duke University Press.

Brunton, Finn, and Helen Nissenbaum. 2011. “Vernacular Resistance to Data Collection and Analysis: A Political Theory of Obfuscation.” *First Monday*. https://doi.org/10.5210/fm.v16i5.3493.

Bryant, David. 2016. “A Quantum Leap for the Web.” Mozilla Tech (blog). October 27, 2016. https://medium.com/mozilla-tech/a-quantum-leap-for-the-web-a3b7174b3c12.

Buck-Morss, Susan. 1995. “Envisioning Capital: Political Economy on Display.” *Critical Inquiry* 21(2): 434–67.

Buzan, Barry, Ole Wæver, Ole Waever, and Jaap de Wilde. 1998. *Security: A New Framework for Analysis*. Lynne Rienner Publishers.

Caduff, Carlo. 2012. “The Semiotics of Security: Infectious Disease Research and the Biopolitics of Informational Bodies in the United States.” *Cultural Anthropology* 27(2): 333-57. https://doi.org/10.1111/j.1548-1360.2012.01146.x.

Callon, Michel. 1998a. “Introduction: The Embeddedness of Economic Markets in Economics.” In *The Laws of the Markets*. Oxford; Malden, MA: Wiley-Blackwell.

———. 1998b. “An Essay on Framing and Overflowing: Economic Revisited by Sociology.” In *The Laws of the Markets*. Oxford; Malden, MA: Wiley-Blackwell.Michel

Callon, Michel, ed. 1998. *The Laws of the Markets*. Oxford; Malden, MA: Wiley-Blackwell.

Callon, Michel, Cécile Méadel, and Vololona Rabeharisoa. 2002. “The economy of qualities.” *Economy and Society* 31(2): 194-217.

Carr, E. Summerson. 2010. “Enactments of Expertise.” *Annual Review of Anthropology* 39: 17-32.

Carr, Nicholas. 2010. *The Shallows: What The Internet Is Doing To Our Brains*. New York: W.W. Norton.

Carsten, Janet. 1997. *The Heat of the Hearth: The Process of Kinship in a Malay Fishing Community*. Oxford Studies in Social and Cultural Anthropology. Oxford: New York: Clarendon Press; Oxford University Press.

Ceglowski, Maciej. 2014a. “The Internet With a Human Face.” Lecture presented at Beyond Tellerrand 2014, Düsseldorf. https://www.youtube.com/watch?v=fWFo1VaQNmU

———. 2014b. “Web Design - The First 100 Years.” Lecture presented at the HOW Interactive Design Conference, Washington, D.C., September. https://idlewords.com/talks/web_design_first_100_years.htm.

———. 2013. “Barely Succeed! It’s Easier!” Lecture presented at Web Directions South, Sydney, Australia. https://www.youtube.com/watch?v=5Vt8zqhHe_c.

Ceruzzi, Paul E. 2012. *Computing: A Concise History*. The MIT Press Essential Knowledge Series. Cambridge, Massachusetts: The MIT Press.

Chambers, Simone. 2003. “Deliberative Democratic Theory.” *Annual Review of Political Science* 6: 307-326. https://doi.org/10.1146/annurev.polisci.6.121901.085538.

Chu, Julie Y. 2010. *Cosmologies of Credit: Transnational Mobility and the Politics of Destination in China*. Durham, NC: Duke University Press.

Chumley, Lily. 2017. “Qualia and Ontology: Language, Semiotics, and Materiality; an Introduction.” *Signs and Society* 5(S1): 1–20.

———. 2013. “Evaluation Regimes and the Qualia of Quality.” *Anthropological Theory* 13 (1/2): 169-83.

Chumley, Lily, and Nicholas Harkness. 2013. “Introduction: QUALIA.” *Anthropological Theory* 13(1/2): 3-11.

Chun, Wendy Hui Kyong. Chun, Wendy Hui Kyong. 2011. Programmed Visions: Software and Memory. Software Studies. Cambridge, Mass: MIT Press.

———. 2005. *Control and Freedom: Power and Paranoia in the Age of Fiber Optics*. Cambridge, MA: The MIT Press.

Clark, David D. 2016. “The Contingent Internet.” *Daedalus* 145 (1): 9–17. https://doi.org/10.1162/DAED_a_00361.

Cody, Francis. 2011. “Publics and Politics.” *Annual Review of Anthropology* 40(1): 37–52.

Cohen, Julie E. 2012. *Configuring the Networked Self: Law, Code, and the Play of Everyday Practice*. New Haven: Yale University Press.

———. 2002. *Regulating Intimacy: A New Legal Paradigm*. Princeton: Princeton University Press.

Coleman, E. Gabriella. 2014. Hacker, Hoaxer, Whistleblower, Spy: The Many Faces of Anonymous. 1 edition. London; New York: Verso.

———. 2012. *Coding Freedom the Ethics and Aesthetics of Hacking*. Princeton: Princeton University Press.

Coleman, E. Gabriella and Alex Golub. 2008. “Hacker practice: Moral genres and the cultural articulation of liberalism.” *Anthropological Theory* 8(3): 255-77. DOI: 10.1177/1463499608093814

Comaroff, Jean, and John Comaroff. 2003. “Reflections on Liberalism, Policulturalism, and ID-Ology: Citizenship and Difference in South Africa.” *Social Identities* 9(4): 445–73. https://doi.org/10.1080/1350463032000174632.

Committee on the Judiciary, *Going Dark: Lawful Electronic Surveillance in the Face of New Technologies: Hearing before the Subcommittee on Crime, Terrorism, and Homeland Security of the Committee on the Judiciary*, House of Representatives, 112th Congress, 1st Session. 2011. Testimony of Valerie Caproni, General Counsel, Federal Bureau of Investigation.

Committee on the Judiciary, *The Encryption Tightrope: Balancing Americans' Security and Privacy: Hearing before the Committee on the Judiciary*, House of Representatives, 114th Congress, 2d Session. 2016.

Committee on the Judiciary, *The Encryption Tightrope: Balancing Americans' Security and Privacy: Hearing before the Committee on the Judiciary*, House of Representatives, 114th Congress, 2d Session. 2016. Testimony of Cyrus B. Vance, District Attorney, New York County.

Committee on the Judiciary, *The Encryption Tightrope: Balancing Americans' Security and Privacy: Hearing before the Committee on the Judiciary*, House of Representatives, 114th Congress, 2d Session. 2016. Testimony of James B. Comey, Director, Federal Bureau of Investigation.

Committee on the Judiciary, *The Encryption Tightrope: Balancing Americans' Security and Privacy: Hearing before the Committee on the Judiciary*, House of Representatives, 114th Congress, 2d Session. 2016. Testimony of Susan Landau, Professor of Cybersecurity Policy, Worcester Poly-technic Institute.

Crary, Jonathan. 2014. *24/7: Late Capitalism and the Ends of Sleep.* New York: Verso.

———. 1999. *Suspension of Perception: Attention, Spectacle, and Modern Culture*. Cambridge, Ma.: The MIT Press.

———. 1992. *Techniques of the Observer: On Vision and Modernity in the 19th Century*. Cambridge, MA: The MIT Press.

Crouch, Luke. 2017. "Firefox Containers Are Go!" *Firefox Test Pilot* (blog). October 10, 2017. https://medium.com/firefox-test-pilot/firefox-containers-are-go-ed2e3533b6e3.

Csordas, Thomas. 1993. "Somatic Modes of Attention." *Cultural Anthropology* 8(2): 135-56.

Culler, Jonathan. 1981. "Convention and Meaning: Derrida and Austin." *New Literary History* 13: 15-30.

Davis, John. 1996. "An anthropologist's view of exchange." *Social Anthropology* 4: 213-226.

Dean, Jodi. 2002. *Publicity's Secret: How Technoculture Capitalizes on Democracy*. Ithaca: Cornell University Press.

de Goede, Marieke. 2014. "The Politics of Privacy in the Age of Preemptive Security: Introduction." *International Political Sociology* 8(1): 100–04. https://doi.org/10.1111/ips.12042.

Davidoff, Leonore, and Catherine Hall. 1987. *Family Fortunes: Men and Women of the English Middle Class, 1780-1850*. Women in Culture and Society. Chicago: University of Chicago Press.

Department of Homeland Security. 2003. National Strategy to Secure Cyberspace.

Dilley, Roy. 1999. "Introduction: The Problem of Context." *The Problem of Context*. Methodology and History in Anthropology, v. 4. Roy Dilley, ed. New York: Berghahn Books.

Dourish, Paul. 2004. "What we talk about when we talk about context." *Personal Ubiquitous Computing* 8:19–30. https://doi.org/10.1007/s00779-003-0253-8.

Dwyer, Jim. 2014. *More Awesome than Money: Four Boys and Their Quest to Save the World from Facebook*. New York, New York: Viking.

Duranti, Allesando. 1993. "Intentions, Self, and Responsibility." In *Responsibility and Evidence in Oral Discourse*. Jane H. Hill and Judith T. Irvine, eds. Cambridge; New York: Cambridge University Press.

Dumit, Joseph. 2012. *Drugs for Life: How Pharmaceutical Companies Define Our Health.* Durham: Duke University Press Books.

Ellis, J. H. 1999. "The History of Non-Secret Encryption." *Cryptologia* 23 (3): 267–73.

English-Lueck, J.A. 2002. *Cultures@SiliconValley*. Stanford, Ca.: Stanford University Press.

Farrell, Stephen, and Hannes Tschofenig. 2014. "Pervasive Monitoring Is an Attack." Request for Comments RFC 7258. Internet Engineering Task Force. https://doi.org/10.17487/RFC7258.

Farquhar, Judith. "Reading Hands: Pulse Qualities and the Specificity of the Clinical." *East Asian Science and Technology Studies* (EASTS) 8 (1): 1-16.

Federal Trade Commission. 2012. "Protecting Consumer Privacy in an Era of Rapid Change: Recommendations for Businesses and Policymakers." FTC Report. https://www.ftc.gov/sites/default/files/documents/reports/federal-trade-commission-report-protecting-consumer-privacy-era-rapid-change-recommendations/120326privacyreport.pdf.

Feng, Hengyi, Julie Froud, Sukhdev Johal, Colin Haslam & Karel Williams. 2001. "A New Business Model? The Capital Market and the New Economy." *Economy and Society* 30 (4): 467–503.

Fielt, Erwin. 2013. "Conceptualising Business Models: Definitions, Frameworks and Classifications." *Journal of Business Models* 1: 85-105.

Forsythe, Diana. 2001. *Studying Those Who Study Us: An Anthropologist in the World of Artificial Intelligence*. Stanford, Calif: Stanford University Press. https:/catalog.lib.uchicago.edu/vufind/Record/4500601.

Fortun, Kim. 2010. "Essential2life." *Dialectical Anthropology* 34 (1): 77–86.

Foster, Robert J. 2007. "The Work of the New Economy: Consumers, Brands, and Value Creation." *Cultural Anthropology* 22 (4): 707–31.

Fourcade, Marion & Kieran Healy. 2007. "Moral Views of Market Society." *Annual Review of Sociology* 33(1): 285–311.

Landes, Joan B. 1988. *Women and the Public Sphere in the Age of the French Revolution*. Ithaca: Cornell University Press.

Fried, Charles. 1970. *An Anatomy of Values*. Cambridge, MA: Harvard University Press.

Gal, Susan. 2017. "Qualia as Value and Knowledge: Histories of European Porcelain." *Signs and Society* 5 (S1): 128-53.

———. 2013. "Tastes of Talk: Qualia and the moral flavor of signs." Anthropological Theory 13 (1/2): 31-48.

———. 2005. "Language Ideologies Compared: Metaphors of Public/Private." *Journal of Linguistic Anthropology* 15(1): 23–37.

———. 2002. "A Semiotics of the Public/Private Distinction." *Differences* 13(1): 77–95. https://doi.org/10.1215/10407391-13-1-77.

Gal, Susan, and Gail Kligman. 2000. The Politics of Gender after Socialism: A Comparative-Historical Essay. Princeton, NJ: Princeton University Press. https://catalog.lib.uchicago.edu/vufind/Record/11185341.

Gal, Susan, and Judith T. Irvine. 2019. *Signs of Difference: Language and Ideology in Social Life*. Cambridge, UK: Cambridge University Press.

Gal, Susan, and Kathryn Ann Woolard, eds. 2001. *Languages and Publics: The Making of Authority*. Encounters, v. 2. Manchester: St. Jerome Publishing.

Gavison, Ruth. 1980. "Privacy and the Limits of Law." *Yale Law Journal* 89: 421–71.

Gellman, Barton, and Ashkan Soltani. 2013. "NSA Infiltrates Links to Yahoo, Google Data Centers Worldwide, Snowden Documents Say." Washington Post, October 30, 2013, sec. National Security. https://www.washingtonpost.com/world/national-security/nsa-infiltrates-links-to-yahoo-google-data-centers-worldwide-snowden-documents-say/2013/10/30/e51d661e-4166-11e3-8b74-d89d714ca4dd_story.html.

Gerstein, Robert S. 1978. “Intimacy and Privacy.” *Ethics* 89(1): 76–81.

Ghaziani, A. & M.J. Ventresca. 2005. “Keywords and Cultural Change: Frame Analysis of Business Model Public Talk 1975-2000.” Sociological Forum: 523-559.

Gitelman, Lisa. 2013. *Raw Data Is an Oxymoron*. Infrastructures Series. Cambridge, Massachusetts: The MIT Press.

Glancy, Dorothy J. 1979. “The Invention of the Right to Privacy.” *Arizona Law Review* 21(1): 1-39.

Glück, Zoltán, and Low, Setha. 2017. “A sociospatial framework for the anthropology of security.” *Anthropological Theory* 17(3): 281–96. https://doi.org/10.1177/1463499617729229

Goffman, Erving. 1986. *Frame Analysis: An Essay on the Organization of Experience*. Boston: Northeastern University Press.

———. 1963. *Behavior in public places: Notes on the social organization of gatherings*. New York: Free Press.

———. 1967. “The Nature of Deference and Demeanor.” *Interaction Ritual: Essays on Face-to-Face Behavior*. New York: Anchor Books.

———. 1959. *The Presentation of Self in Everyday Life*. Anchor books edition. Doubleday Anchor Books; A174. Garden City, N.Y: Doubleday & Company.

Goldstein, Daniel M. 2010. “Toward a Critical Anthropology of Security.” *Current Anthropology* 51(4): 487–517. https://doi.org/10.1086/655393.

Goodwin, Charles and Alessandro Duranti. 1992. “Rethinking Context: An Introduction.” In *Rethinking Context: Language as an Interactive Phenomenon*. Studies in the Social and Cultural Foundations of Language, no. 11. Alessandro Duranti and Charles Goodwin, eds. Cambridge; New York: Cambridge University Press.

Gottleib, Calvin. 1996. “Privacy: A Concept Whose Time Has Come and Gone.” *Computers, Surveillance, and Privacy*. David Lyon and Elia Zureik, eds. Minneapolis: University of Minnesota Press.

Graeber, David. 2011. Debt: The First 5,000 Years. Brooklyn, NY: Melville House.

———. 2001. Toward an Anthropological Theory of Value: The False Coin of Our Own Dreams. New York: Palgrave.

Graham, Paul. 2014. “Before the Startup.” CS183B: How to Start a Startup. Class lecture at Stanford University, Palo Alto, C.A., Fall Quarter. https://www.youtube.com/watch?v=ii1jcLg-eIQ&t=53s.

———. 2013a. “How to Convince Investors.” *PaulGraham.Com* (blog). August 2013. http://www.paulgraham.com/convince.html.

———. 2013b. “Do Things That Don’t Scale.” *PaulGraham.Com* (blog). July 2013. http://paulgraham.com/ds.html.

———. “Startup = Growth.” *PaulGraham.Com* (blog). September 2012. http://www.paulgraham.com/growth.html.

Grigorik, Ilya. 2015. “Performance RAILS: The art & science of optimizing for silicon & wetware.” Accessed March 16, 2022. https://docs.google.com/presentation/d/13AJe2Ip4etqA8qylrva7bEgu1_hAvsq_VQiVOAxwdcI/htmlpresent.

Grudin, Jonathan. 2001. “Desituating Action: Digital Representation of Context.” *Human–Computer Interaction*, 16(2): 269-86. DOI: 10.1207/S15327051HCI16234_10

Gumperz, John J. 1992. “Contextualization Revisited.” *The Contextualization of Language*. Peter Auer and Aldo Di Luzio, eds. Amsterdam; Philadelphia: J. Benjamins Pub. Co.

Gürses, Seda. 2014. “Can You Engineer Privacy?” *Communications of the ACM* 57(8), 20-3.

Gürses, Seda, and Joris V. J. van Hoboken. 2017. “Privacy After the Agile Turn.” *Cambridge Handbook of Consumer Privacy*. Jules Polonetsky, Omar Tene, and Evan Selling, eds. Cambridge, UK: Cambridge University Press.

Habermas, Jürgen. 1989. *The Structural Transformation of the Public Sphere: An Inquiry into a Category of Bourgeois Society*. Cambridge, MA: The MIT Press.

———. 1996. *Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy*. Cambridge, MA: MIT Press.

Hafner, Katie and Matthew Lyon. 1996. Where Wizards Stay up Late: The Origins of the Internet. New York: Simon & Schuster.

Hanks, William. 1992. “The indexical ground of deictic reference.” *Rethinking Context: Language as an Interactive Phenomenon*. Studies in the Social and Cultural Foundations of Language, no. 11. Alessandro Duranti and Charles Goodwin, eds. Cambridge; New York: Cambridge University Press.

Hansen, Lene and Nissenbaum, Helen. 2009. “Digital Disaster, Cyber Security and the Copenhagen School1.” *International Studies Quarterly* 53. 10.1111/j.1468-2478.2009.00572.x.

Harkness, Nicholas. 2017. “The Open Throat: Deceptive Sounds, Facts of Firstness, and the Interactional Emergence of Voice.” *Signs and Society* 5 (S1): 21-52.

———. 2015. “The Pragmatics of Qualia in Practice.” *Annual Review of Anthropology* 44: 573-89.

———. 2013. “Softer soju in South Korea.” *Anthropological Theory* 13(1/2): 12-30.

Hayles, N. Katherine. 2007. “Hyper and Deep Attention: The Generational Divide in Cognitive Modes.” *Profession*: 187-99.

Heinemeier Hansson, Daniel. 2017. “Exponential Growth Devours and Corrupts.” *Signal v. Noise* (blog). February 27, 2017. https://m.signalvnoise.com/exponential-growth-devours-and-corrupts/.

Helmreich, Stefan. 2007. “An anthropologist underwater: Immersive soundscapes, submarine cyborgs, and transductive ethnography.” *American Ethnologist* 34 (4): 621–41.

Herley, Cormac and P. C. Van Oorschot. 2017. “SoK: Science, Security and the Elusive Goal of Security as a Scientific Pursuit,” *IEEE Symposium on Security and Privacy*, 2017, pp. 99-120, doi: 10.1109/SP.2017.38.

Herzfeld, Michael. 2009. “The Performance of Secrecy: Domesticity and Privacy in Public Spaces.” *Semiotica* 2009 (175). https://doi.org/10.1515/semi.2009.044.

Hilgartner, Stephen. 2000. *Science on Stage: Expert Advice as Public Drama*. Stanford: Stanford University Press.

Hill, Jane H. 2008. *The Everyday Language of White Racism*. Blackwell Studies in Discourse and Culture 3. Chichester, U.K.; Malden, MA: Wiley-Blackwell.

Hirschkind, Charles, Maria José A. de Abreu, and Carlo Caduff. 2017. “New Media, New Publics?: An Introduction to Supplement 15.” Current Anthropology 58(S15): S3–12. https://doi.org/10.1086/688903.

Hirschman, Daniel and Elizabeth Popp Berman. 2014. “Do economists make policies? On the political effects of economics.” *Socio-Economic Review* 12(4): 779–811. https://doi.org/10.1093/ser/mwu017.

Hoffman, Reid. 2017a. “Minted’s Mariam Naficy in The Money Episode.” Masters of Scale. May 10, 2017. Accessed November 11, 2021. https://mastersofscale.com/mariam-naficy-the-money-episode/.

———. 2017b. “Handcrafted with Brian Chesky, Co-Founder & CEO of Airbnb.” Masters of Scale. May 3, 2017. Accessed November 11, 2021. https://mastersofscale.com/brian-chesky-handcrafted/.

Hull, Gordon. 2015. “Successful failure: What Foucault can teach us about privacy self-management in a world of Facebook and big data.” *Ethics and Information Technology* 17: 89-101.

Igo, Sarah E. 2018. *The Known Citizen: A History of Privacy in Modern America*. Cambridge, MA: Harvard University Press.

Innes, Julie C. 1996. *Privacy, Intimacy, and Isolation*. New York: Oxford University Press

Irish, Paul and Paul Lewis. 2015. “Introducing Rail: A User-Centric Model for Performance. Accessed March 22, 2022. https://www.smashingmagazine.com/2015/10/rail-user-centric-model-performance.

Johns, Adrian. 2009. *Piracy: The Intellectual Property Wars from Gutenberg to Gate*s. Chicago: The University of Chicago Press.

Jusionyte, Ieva, and Daniel M. Goldstein. 2016. “In/Visible—In/Secure: Optics of Regulation and Control.” *Focaal—Journal of Global and Historical Anthropology* (75): 3–13. https://doi.org/10.3167/fcl.2016.750101.

Kant, Immanuel. 1996. “An Answer to the Question: What Is Enlightenment?” In What Is Enlightenment?: Eighteenth-Century Answers and Twentieth-Century Questions, edited by James Schmidt. University of California Press.

Keane, Webb. 2003. “Semiotics and the Social Analysis of Material Things.” *Language & Communication* 23 (3–4): 409–25.

Kehl, Danielle, Andi Wilson, and Kevin Bankston. 2015. “Doomed to Repeat History? Lessons from the Crypto Wars of the 1990s.” New America. https://www.newamerica.org/cybersecurity-initiative/policy-papers/doomed-to-repeat-history-lessons-from-the-crypto-wars-of-the-1990s/.

Kelty, Christopher M. 2017. “Too Much Democracy in All the Wrong Places: Toward a Grammar of Participation.” *Current Anthropology* 58: S77-S90

———. 2008a. *Two Bits: The Cultural Significance of Free Software*. Durham NC: Duke University Press.

———. 2008b. “Don’t be evil.” Paper Presented at the 107th Annual Meeting of the American Anthropological Association, San Francisco, CA. Nov 19-22.

Kerr, Orin S. 2010. “Applying the Fourth Amendment to the Internet: A General Approach.” *Stanford Law Review* 62 (4): 1005–49.

Kift, Paula and Helen Nissenbaum. 2016. “Metadata in Context - An Ontological and Normative Analysis of the NSA's Bulk Telephony Metadata Collection Program.” *I/S: A Journal of Law and Policy for the Information Society* 13(2): 333-72.

Kristol, David M. 2001. “HTTP Cookies: Standards, Privacy, and Politics.” *ACM Transactions on Internet Technology* 1(2): 151–98. https://doi.org/10.1145/502152.502153.

Kulwin, Noah. 2018. “An Apology for the Internet from the Architects who Built It.” *New York Magazine*. Accessed March 20, 2022. https://nymag.com/intelligencer/2018/04/an-apology-for-the-internet-from-the-people-who-built-it.html.

Lakoff, Andrew. 2007. “Preparing for the Next Emergency.” *Public Culture* 19(2): 247–71. https://doi.org/10.1215/08992363-2006-035.

Lakoff, Andrew, and Stephen J. Collier. 2004. “Ethics and the anthropology of modern reason.” *Anthropological Theory* 4(4): 419–34. https://doi.org/10.1177/1463499604047919

Landau, Susan, Stephen Kent, Clint Brooks, Scott Charney, Dorothy Denning, Whitfield Diffie, Anthony Lauck, Doug Miller, Peter Neumann, and David Sobel. 1994. *Codes, Keys, and Conflicts: Issues in U.S. Crypto Policy*: Report of a Special Panel of the ACM U.S. Public Policy Committee (USACM). New York, N.Y.: Association for Computing Machinery.

Landes, Joan B. 1988. *Women and the Public Sphere in the Age of the French Revolution*. Ithaca: Cornell University Press.

Landes, Joan B., ed. 1998. *Feminism, the Public and the Private*. Oxford Readings in Feminism. New York: Oxford University Press.

Landman, Todd, and Hans-Joachim Lauth. 2019. “Political Trade-Offs: Democracy and Governance in a Changing World.” *Politics and Governance* 7(4): 237–42. https://doi.org/10.17645/pag.v7i4.2642.

Lanier, Jaron. 2013. *Who Owns the Future?* New York: Simon & Schuster.

Lashinsky, Adam. 2005. “Remembering Netscape: The Birth Of The Web - July 25, 2005.” *Fortune Magazine*, July 25, 2005. https://money.cnn.com/magazines/fortune/fortune_archive/2005/07/25/8266639/.

Latour, Bruno. 1987. *Science in Action: How to Follow Scientists and Engineers through Society*. Cambridge, Mass: Harvard University Press.

Lempert, Michael P. "Denotational Textuality and Demeanor Indexicality in Tibetan Buddhist Debate." *Journal of Linguistic Anthropology* 15, no. 2 (2005): 171–93. http://www.jstor.org/stable/43104048.

Levine, Yasha. 2018. *Surveillance Valley: The Secret Military History of the Internet*. New York: PublicAffairs.

Levy, Steven. 2001. Crypto: How the Code Rebels Beat the Government-- Saving Privacy in the Digital Age. New York: Viking. https://catalog.lib.uchicago.edu/vufind/Record/4369067.

Lessig, Lawrence. 1999. *Code and Other Laws of Cyberspace*. New York: Basic Books.

Lewis, Michael. 2000. *The New New Thing: A Silicon Valley Story*. New York: W. W. Norton.

Lewis, Paul. 2017. "Our minds can be hijacked: the tech insiders who fear a smartphone dystopia." *The Guardian*. Accessed March 20, 2022. https://www.theguardian.com/technology/2017/oct/05/smartphone-addiction-silicon-valley-dystopia.

Lezaun, Javier. 2010. "Eloquence and incommensurability: An investigation into the grammar of irreconcilable differences." *Social Studies of Science*, 40(3): 349–75. https://doi.org/10.1177/0306312709358119.

Livingston, Jessica. 2017. "What's Different about 'Unicorns.'" Lecture presented to Female Founders Conference 2017, San Francisco, C.A. https://www.youtube.com/watch?v=Ygr3rx4hSsc.

Lowe, Lisa. 2015. *The Intimacies of Four Continents*. Durham: Duke University Press Books.

Lowi, Theodore J. 1964. "American business, public policy, case studies and political theory." *World Politics* 16:677–715.

Magretta, Joan. 2002. "Why Business Models Matter." *Harvard Business Review* 80: 86-92.

Marres, Noortje and Javier Lezaun. 2011. Materials and devices of the public: an introduction, *Economy and Society* 40(4): 489-509. DOI: 10.1080/03085147.2011.602293.

Marwick, Alice E., & boyd, danah. 2014. “Networked privacy: How teenagers negotiate context in social media.” *New Media & Society*, 16(7): 1051–67. https://doi.org/10.1177/1461444814543995.

Marwick, A. E., & boyd, danah. 2011. “I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience.” *New Media & Society* 13(1): 114–33. https://doi.org/10.1177/1461444810365313

Marx, Gary T. 2016. *Windows into the Soul: Surveillance and Society in an Age of High Technology*. Chicago: The University of Chicago Press.

Masco, Joseph. 2017. “‘Boundless informant’: Insecurity in the age of ubiquitous surveillance.” *Anthropological Theory* 17 (3): 382-403.

———. 2014. *The Theater of Operations: National Security Affect from the Cold War to the War on Terror*. Durham: Duke University Press Books.

———. 2010. “‘Sensitive but Unclassified’: Secrecy and the Counterterrorist State.” *Public Culture* 22 (3): 433–63.

Massumi, Brian. 2015. *Ontopower: War, Powers, and the State of Perception*. Durham: Duke University Press.

Mauss, Marcel. 1990. The Gift: The Form and Reason for Exchange in Archaic Societies. New York: W. W. Norton.

May, Timothy C. 1994. *THE CYPHERNOMICON: Cypherpunks FAQ and More*. Version 0.666. https://hackmd.io/@jmsjsph/TheCyphernomicon.

Mazzarella, William. 2019. “The Anthropology of Populism: Beyond the Liberal Settlement. *Annual Review of Anthropology* 48(1): 45–60. https://doi.org/10.1146/annurev-anthro-102218-011412.

———. 2009. “Affect: What Is It Good For?” In *Enchantments of Modernity*. Routledge India.

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

MacKinnon, Catherine. 1989. *Toward a Feminist Theory of the State.* Cambridge, MA: Harvard University Press.

McDonald, Aleecia M., and Lorrie Faith Cranor. 2008. “The Cost of Reading Privacy Policies.” *I/S: A Journal of Law and Policy for the Information Society* 4(3): 543-68.

McFarland, Andrew S. 2007. “Neopluralism.” *Annual Review of Political Science* 10 (1): 45–66. https://doi.org/10.1146/annurev.polisci.10.072005.152119.

McGrath, Rita. 2010. “Business Models: A Discovery Driven Approach.” *Long Range Planning* 43 (2–3): 247-261.

McNamee, Roger. 2019. *Zucked: Waking up to the Facebook Catastrophe*. New York: Penguin Press.

MDN Web Docs. 2020a. “Web Performance.” Accessed April 15, 2022. https:/developer.mozilla.org/en-US/docs/Web/Performance.

———. 2020b. “What is performance.” Accessed March 15, 2022. https://developer.mozilla.org/en-US/docs/Learn/Performance/What_is_web_performance.

Miller, Daniel. 2002. “Turning Callon the Right Way Up.” *Economy and Society* 31(2): 218–33.

———. 1987. *Material Culture and Mass Consumption*. Oxford and New York: Blackwell.

Miller, Peter, and Nikolas Rose. 1997. “Mobilizing the Consumer.” *Theory, Culture and Society* 14(1): 1–36. https://doi.org/10.1177/026327697014001001.

Montulli, Lou. 2013. “The Reasoning behind Web Cookies.” The Irregular Musings of Lou Montulli (blog). May 14, 2013. https://montulli.blogspot.com/2013/05/the-reasoning-behind-web-cookies.html.

Moore, Adam D. 2003. “Privacy: Its Meaning and Value.” *American Philosophical Quarterly* 40 (3): 215–27.

Moore, Robert E. 2003. “From Genericide to Viral Marketing: On ‘Brand.’” *Language & Communication* 23(3–4): 331–57.

Morozov, Evgeny. 2013. *To Save Everything, Click Here: The Folly of Technological Solutionism*. First edition. New York: PublicAffairs.

MozillaWiki. n.d. “Security/Contextual Identity Project/Containers.” Accessed October 12, 2022. https://wiki.mozilla.org/Security/Contextual_Identity_Project/Containers.

Munn, Nancy. 1986. *The Fame of Gawa: A Symbolic Study of Value Transformation in a Massim Society (Papua New Guinea)*. Durham, NC: Duke Univ. Press

Nakassis, Constantine V. 2020. “Deixis and the Linguistic Anthropology of Cinema.” *Semiotic Review* 9. https://www.semioticreview.com/ojs/index.php/sr/article/view/65.

———. 2018. “Indexicality’s Ambivalent Ground.” *Signs and Society* 6(1): 281–304. https://doi.org/10.1086/694753.

———. 2013. “Brands and Their Surfeits.” *Cultural Anthropology* 28 (1): 111–26.

Narayanan, Arvind. 2013. “What Happened to the Crypto Dream?, Part 1.” IEEE Security and Privacy 11 (2): 75–76. https://doi.org/10.1109/MSP.2013.45.

National Research Council. 1996. *Cryptography’s Role in Securing the Information Society*. Washington, DC: The National Academies Press. https://doi.org/10.17226/5131.

Nelson, Deborah. 2002. *Pursuing Privacy in Cold War America*. New York: Columbia University Press.

Neocleous, Mark. 2007. “Security, Liberty and the Myth of Balance: Towards a Critique of Security Politics.” Contemporary Political Theory 6 (2): 131–49. https://doi.org/10.1057/palgrave.cpt.9300301.

Nguyen, Nick. 2018. “Latest Firefox Quantum Release Available with Faster, Always-on Privacy with Opt-in Tracking Protection and New Features.” The Mozilla Blog (blog). January 23, 2018. https://blog.mozilla.org/en/products/firefox/latest-firefox-quantum-release-now-available-with-new-features/.

Nielson, Jakob. 2009. “Powers of 10: Time Scales in User Experience.” Accessed March 22, 2022. https://www.nngroup.com/articles/powers-of-10-time-scales-in-ux/.

———. 1993. *Usability Engineering*. Cambridge, Ma.: Academic Press, Inc.

Nippert-Eng, Christena. 2010. *Islands of Privacy*. Chicago: The University of Chicago Press.

Nissenbaum, Helen. 2010. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford, Calif: Stanford Law Books.

———. “Where Computer Security Meets National Security.” Ethics and Information Technology 7: 61–73. https://doi.org/10.1007/s10676-005-4582-3

Novak, Chelsea. 2018. “Browse without baggage in Firefox: Set Tracking Protection to always on” Accessed March 22, 2022. https://blog.mozilla.org/products/firefox/tracking-protection-always-on/.

O'Reilly, Tim. 1996. “Publishing models for Internet commerce.” *Communications of the ACM* 39: 79-86.

Oreskes, Naomi, and Erik M. Conway. 2010. *Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming*. New York: Bloomsbury Press.

Osterwalder, Alexander & Yves Pigneur Yves. 2003. “Modeling value propositions in e-Business.” *ACM International Conference Proceeding Series* 50: 429-436.

Palen, Leysia, and Paul Dourish. 2003. “Unpacking ‘Privacy’ for a Networked World.” In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 129–36. CHI ’03. New York, NY: Association for Computing Machinery. https://doi.org/10.1145/642611.642635.

Paley, Julia. 2008. “Introduction.” *Democracy: Anthropological Approaches*. Julia Paley, ed. Santa Fe, NM: School for Advanced Research Press.

Pedersen, Morten Axel and Martin Holbraad. 2013. “Introduction: Times of Security.”*Times of Security: Ethnographies of Fear, Protest and the Future*. Martin Holbraad and Morten Axel Pedersen, eds. New York: Routledge.

Perlroth, Nicole, Jeff Larson, and Scott Shane. 2013. “N.S.A. Able to Foil Basic Safeguards of Privacy on Web.” The New York Times, September 5, 2013, sec. U.S. https://www.nytimes.com/2013/09/06/us/nsa-foils-much-internet-encryption.html.

Pfaffenberger, Bryan. 1992. “Technological Dramas.” *Science, Technology, & Human Values* 17(3): 282–312. https://doi.org/10.1177/016224399201700302.

Pold, Søren. 2008. “Button.” *Software Studies: A Lexicon*. Matthew Fuller, ed. Leonardo. Cambridge, Mass: MIT Press.

Porter, Michael. 2001. “Strategy and the Internet.” *Harvard Business Review* 79: 62-78.

Porter, Theodore M. 1995. *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life. Princeton*, N.J: Princeton University Press. https://catalog.lib.uchicago.edu/vufind/Record/1742250.

Povinelli, Elizabeth A. 2011. *Economies of Abandonment: Social Belonging and Endurance in Late Liberalism*. Durham: Duke University Press.

———. 2006. The Empire of Love: Toward a Theory of Intimacy, Genealogy, and Carnality. Durham: Duke University Press.

———. 2001. “Radical Worlds: The Anthropology of Incommensurability and Inconceivability.” *Annual Review of Anthropology* 30: 319–34.

Power, Michael. 1994. *The audit explosion*. London: Demos.

Prosser, William L. 1960. “Privacy.” *California Law Review* 48 (3): 383–423. https://doi.org/10.2307/3478805.

Rabinow, Paul, George E. Marcus, James D. Faubion, and Tobias Rees. 2008. *Designs for an Anthropology of the Contemporary*. Durham: Duke University Press.

Rachels, James. 1975. “Why Privacy Is Important.” *Philosophy and Public Affairs* 4(4): 323–33.

Rainie, Lee, Sara Kiesler, Ruogu Kang, and Mary Madden. 2013. “Anonymity, Privacy, and Security Online.” Internet and American Life Project. Pew Research Center. https://www.pewresearch.org/internet/2013/09/05/anonymity-privacy-and-security-online/.

Reidenberg, Joel R. 1997. “Lex Informatica: The Formulation of Information Policy Rules through Technology,” *Texas Law Review* 76(553).

Riles, Annelise. 2005. “A New Agenda for the Cultural Study of Law: Taking on the Technicalities.” *Buffalo Law Review* 53(3).

Rosaldo, Michelle Z. 1982. “The Things We Do with Words: Ilongot Speech Acts and Speech Act Theory in Philosophy.” *Language in Society* 11(2): 203–37.

Rose, Nikolas. 1996. “Governing Enterprising Individuals.” In *Inventing Our Selves: Psychology, Power, and Personhood*, 150–68. Cambridge Studies in the History of Psychology. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9780511752179.008.

Sahlins, Marshall. 2013. *What kinship is—and is not*. Chicago: University of Chicago Press

Schneier, Bruce. 2013a. “Surveillance as a Business Model,” Schneier on Security. Accessed March 22, 2022. https://www.schneier.com/blog/archives/2013/11/surveillance_as_1.html.

———. 2013b. “The US Government Has Betrayed the Internet. We Need to Take It Back.” *The Guardian*, September 5, 2013, sec. Opinion. https://www.theguardian.com/commentisfree/2013/sep/05/government-betrayed-internet-nsa-spying.

———. 2013c. “Internet Hardening.” Presented at the IETF Technical Plenary, Vancouver, BC, Canada, November 6. https://www.ietf.org/proceedings/88/minutes/minutes-88-iab-techplenary.

Schüll, Natasha Dow. 2014. *Addiction by Design: Machine Gambling in Las Vegas*. Princeton, N.J.: Princeton University Press.

Seibel, Michael. 2016. “How to Pitch Your Company: Fundraising.” *YC Startup Library* (blog). https://www.ycombinator.com/library/4b-how-to-pitch-your-company.

———. 2014. “Sales and Marketing; How to Talk to Investors.” CS183B: How to Start a Startup. Class lecture at Stanford University, Palo Alto, C.A., Fall Quarter. https://www.youtube.com/watch?v=SHAh6WKBgiE&t=1s.

Shankland, Stephen. 2017. “Firefox Fights Back.” Accessed April 27, 2022. https://www.cnet.com/special-reports/mozilla-firefox-fights-back-against-google-chrome/.

Shapin, Steven, and Simon Schaffer. 2011. *Leviathan and the Air-Pump: Hobbes, Boyle, and the Experimental Life*. Princeton, N.J: Princeton University Press. https://catalog.lib.uchicago.edu/vufind/Record/8547333.

Shryock, Andrew, ed. 2004. Off Stage/On Display: Intimacy and Ethnography in the Age of Public Culture. 1 edition. Stanford, Calif: Stanford University Press.

Silverstein, Michael. 2004. “‘Cultural’ Concepts and the Language-Culture Nexus.” *Current Anthropology* 45 (5): 621–52. https://doi.org/10.1086/423971.

———. 2003. “Indexical Order and the Dialectics of Social Life.” *Language & Communication*. 23: 193-229. 10.1016/S0271-5309(03)00013-2.

———. 1992. “The Indeterminacy of Contextualization: When Is Enough Enough.” In *The Contextualization of Language*. Peter Auer and Aldo Di Luzio, eds. Amsterdam; Philadelphia: J. Benjamins.

———. 1976. “Shifters, Linguistic Categories, and Cultural Description.” In *Meaning in Anthropology*. Keith H. Basso and Henry A. Selby, eds. School of American Research Advanced Seminar Series. Albuquerque: University of New Mexico Press.

Singel, Ryan. 2007. “Spying in the Death Star: The AT&T Whistle-Blower Tells His Story.” *Wired* May 10, 2007. https://www.wired.com/2007/05/spying-in-the-death-star-the-att-whistle-blower-tells-his-story/.

Slater, Don. 2002a “Markets, materiality and the ‘new economy.’” In: Metcalfe, Stanley and Warde, Alan, (eds.) *Market Relations and the Competitive Process: New dynamics of innovation & competition*. Manchester: Manchester University Press, 95-113.

———. 2002b. “From calculation to alienation: disentangling economic abstractions.” *Economy and Society* 31(2): 234-249.

———. 2000. “Consumption without scarcity: exchange and normativity in an internet setting.” In *Commercial Cultures: Economics, Practices, Spaces*. Oxford: Berg.

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

Souders, Steve. 2013. “Moving beyond Window.Onload() | High Performance Web Sites.” Steversourders.Com (blog). May 13, 2013. https://www.stevesouders.com/blog/2013/05/13/moving-beyond-window-onload/.

Strathern, Marilyn. 1988. *The Gender of the Gift: Problems with Women and Problems with Society in Melanesia*. Studies in Melanesian Anthropology 6. Berkeley: University of California Press.

Strohmeier, Dominik and Harald Kirschner. 2017. "Designing for performance: A data-informed approach for Quantum development." *Mozilla Hacks* (blog). Accessed on April 27, 2022. https://hacks.mozilla.org/2017/06/designing-for-performance-a-data-informed-approach-for-quantum-development/

Stowsky, Jay. 2004. "Secrets to shield or share? new dilemmas for military R&D policy in the digital age." *Research Policy*, 33(2): 257-69.

Stuntz, William J. 1995. "Privacy's Problem and the Law of Criminal Procedure," Michigan Law Review 93: 1016.

Suchman, Lucy. 2006. *Human-Machine Reconfigurations: Plans and Situated Actions*. 2 edition. Cambridge; New York: Cambridge University Press.

Suchman, Lucy, Carolina Follis, and Jutta Weber. 2017. "Tracking and Targeting: Sociotechnologies of (In)security." *Science, Technology, & Human Values* 42(6): 983–1002. https://doi.org/10.1177/0162243917731524

Sullivan, J. 2013. "Personalization with Respect." *The Mozilla Blog* May 10. Accessed March 20, 2022. https://blog.mozilla.org/products/firefox/personalization-with-respect/.

Tang, Stanley. 2014. "How to Get Started." CS183B: How to Start a Startup. Class lecture at Stanford University, Palo Alto, C.A., Fall Quarter. https://www.youtube.com/watch?v=oQOC-qy-GDY&t=80s.

Taussig, Michael T. 1999. *Defacement: Public Secrecy and the Labor of the Negative*. Stanford, Calif: Stanford University Press.

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

Teece, David J. 2009. *Dynamic Capabilities and Strategic Management*. Oxford; New York: Oxford University Press, 2009.

Tene, Omer and Polonetsky, Jules. 2013. "A Theory of Creepy: Technology, Privacy and Shifting Social Norms." *Yale Journal of Law & Technology* 16(1): 59-102.

The Chertoff Group. 2016. *The Ground Truth About Encryption And The Consequences of Extraordinary Access*.

Terranova, Tiziana. 2000. “Free Labor: Producing Culture for the Digital Economy.” *Social Text* 18(2): 33–58.

Thiel, Peter. 2014. “Competition is for Losers.” CS183B: How to Start a Startup. Class lecture at Stanford University, Palo Alto, C.A., Fall Quarter.

Thrift, Nigel. 2006. “Re-Inventing Invention: New Tendencies in Capitalist Commodification.” *Economy and Society* 35(2): 279–306. https://doi.org/10.1080/03085140600635755.

———. 2001. “‘It’s the Romance, Not the Finance, That Makes the Business Worth Pursuing’: Disclosing a New Market Culture.” *Economy and Society* 30 (4): 412–32.

Tribe, Laurence H. 1985. “Constitutional Calculus: Equal Justice or Economic Efficiency?” Harvard Law Review 98 (3): 592–621. https://doi.org/10.2307/1340870.

———. 1973. “Technology Assessment and the Fourth Discontinuity: The Limits of Instrumental Rationality.” *Southern California Law Review* 46(3): 617-60.

———. 1972. “Policy Science: Analysis or Ideology?” *Philosophy & Public Affairs* 2 (1): 66–110.

Tschofenig, Hannes. 2014. “Improving Security on the Internet.” Position paper, W3C/IAB Workshop on Strengthening the Internet Against Pervasive Monitoring, London.

Thomson, Judith Jarvis. 1984. “The Right to Privacy.” Philosophical Dimensions of the Right to Privacy: An Anthology.” Ferdinand David Schoeman, ed. Cambridge, UK: Cambridge University Press.

Tufekci, Zeynep. 2018. “Opinion | Facebook’s Surveillance Machine.” *The New York Times*, March 19, 2018, sec. Opinion. https://www.nytimes.com/2018/03/19/opinion/facebook-cambridge-analytica.html.

Turkle, Sherry. 1995. *Life on the Screen: Identity in the Age of the Internet*. New York: Simon & Schuster.

Turner, Fred. 2013. *The Democratic Surround: Multimedia and American Liberalism from World War II to the Psychedelic Sixties*. Chicago; London: University Of Chicago Press.

Turner, Fred. 2006. *From Counterculture to Cyberculture: Stewart Brand, the Whole Earth Network, and the Rise of Digital Utopianism*. Chicago: University of Chicago Press.

Ullman, Ellen. 2017. *Life in Code: A Personal History of Technology*. New York: Farrar, Straus and Giroux.

Vyas, Tanvi. 2016. "Contextual Identities on the Web." *The Mozilla Blog* (blog). June 16, 2016. https://blog.mozilla.org/tanvi/2016/06/16/contextual-identities-on-the-web/.

Vyas, Tanvi, Andrea Marchesini, and Christoph Kerschbaumer C. 2017. "Extending the Same Origin Policy with Origin Attributes." In *Proceedings of the 3rd International Conference on Information Systems Security and Privacy - Volume 1:* 464-73. DOI: 10.5220/0006210404640473.

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

Warner, Michael. 2002. Publics and Counterpublics. New York: Zone Books.

———. 1990. The Letters of the Republic: Publication and the Public Sphere in Eighteenth-Century America. Cambridge, Mass: Harvard University Press.

Weiser, Mark. 1993. "Ubiquitous Computing." *Computer* 26(10): 71–2. https://doi.org/10.1109/2.237456.

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

Whitman, James Q. 2004. "The Two Western Cultures of Privacy: Dignity versus Liberty." *Yale Law School Faculty Scholarship Series* Paper 649.

Williams, James. 2016. "The Clickbait Candidate." *Quillette*. October 3, 2016. https://quillette.com/2016/10/03/the-clickbait-candidate/.

Williams, Jamie and Seth Schoen. 2015. "Crypto Is For Everyone—and American History Proves It." Electronic Frontier Foundation (blog). October 30, 2015. https://www.eff.org/deeplinks/2015/10/crypto-everyone-and-american-history-proves-it.

Winton, David, dir. 2000. *Code Rush*. Winton duPont Films and KTEH.

Wirtz, Bernd, Oliver Schilke & Sebastian Ullrich. 2010. "Strategic Development of Business Models: Implications of the Web 2.0 for Creating Value on the Internet." *Long Range Planning* 43: 272-290.

Wolfe, Gary. 1994. "The (Second Phase of the) Revolution Has Begun." *Wired*. October 1, 1994. https://www.wired.com/1994/10/mosaic/.

Wu, Timothy S. 1997. "Cyberspace Sovereignty? – The Internet and the International System." Harvard Journal of Law & Technology 10(3): 647-666.

Zelizer, Viviana. 2011. "Circuits within Capitalism." *Economic Lives: How Culture Shapes the Economy*. Princeton: Princeton University Press. https://catalog.lib.uchicago.edu/vufind/Record/10364321.

———. 2005. *The Purchase of Intimacy*. Princeton, N.J.: Princeton University Press.

Ziewitz, Malte. 2016. "Governing Algorithms: Myth, Mess, and Methods." *Science, Technology, & Human Values* 41(1): 3–16.

Zittrain, Jonathan L., Matthew G. Olsen, David O'Brien, and Bruce Schneier. 2016. "Don't Panic: Making Progress on the 'Going Dark' Debate." Berkman Center Research Publication 2016-1.

Zuckerman, Ethan. 2014. "The Internet's Original Sin." *The Atlantic*, August 14. Accessed March 20, 2022. https://www.theatlantic.com/technology/archive/2014/08/advertising-is-the-internets-original-sin/376041/.