CAMBRIDGE
Cambridge University Press
978-1-108-49242-3 — Industry Unbound
Ari Ezra Waldman
Excerpt
More Information

INTRODUCTION

We are told by companies, by lawmakers, and by civil society that privacy law is getting stronger. The European Union's General Data Protection Regulation (GDPR) has been called "comprehensive" and "one of the strictest privacy laws in the world." Between 2018 and 2020, nine proposals for comprehensive privacy legislation were introduced in the US Congress; two ballot initiatives and forty-two proposals were introduced in twenty-eight states in that same time. Several of those became law, including the California Consumer Privacy Act (CCPA). The Federal Trade Commission (FTC), the de facto privacy enforcer in the United States, is putting limits on the collection, use, and manipulation of personal information with unprecedented billion-dollar fines. The US Supreme Court has started to reclaim the Fourth Amendment's historical commitment to curtailing pervasive police surveillance. And the European Union's Court of Justice has challenged the cross-border transfer of European citizens' data, protecting the privacy of European citizens in the process.[1]

Even the information industry – the ecosystem of companies that collect, process, and use our data for profit – seems to be getting on the privacy bandwagon. Apple markets its iPhones as privacy protective. Facebook, today's dominant online social network, promises to build a future of "private, encrypted messages." Google "build[s] privacy that works" for us. These tech giants, and many far smaller ones, have spent millions of dollars and hired thousands of privacy professionals, privacy lawyers, data protection officers, and privacy vendors. Even some software engineers, historically not known for their concern for privacy, are coming to work recognizing that they have to do better to protect our privacy.[2]

These workers are at the front line of the social practice of privacy law, by which I mean the practices, behaviors, and performances that implement or navigate privacy law and privacy design on the ground. They write policies, interpret and apply the law for their employers, translate legal requirements into technical specifications, set up internal systems, write litigation briefs, build products, consult with coders and programmers, coordinate across departments, conduct internal assessments, manage audits, answer questions about privacy, and then some. When you talk to them, read their resumes, or scan their profiles on LinkedIn, they will tell you they are "committed" to privacy. Many are eager to "work with companies who care about" making "data protection a central part" of their business. Privacy professionals are focused on helping "businesses manage privacy risks" and dedicated to "compliance done right." Privacy engineers are committed to "bringing privacy into design."

And yet, every day, our privacy is slipping away. Face surveillance, DNA-testing kits, "smart" devices, gratuitous location tracking, and manipulative "dark patterns" are increasingly commonplace. Platforms like Google, Facebook, Snapchat, Pokémon Go, and Zoom, not to mention in-home assistants like Amazon's Alexa, are designed without our privacy in mind. Other products – like Uber's mobile app – were maliciously designed to invade our privacy, while social platforms like Facebook are designed to manipulate us into disclosure. Despite repeatedly promising to do better after privacy scandals, muckraking investigations revealing invasive designs, and a growing backlash to surveillance practices, very little has changed on the ground. There are some bright spots: Mozilla's Firefox, the DuckDuckGo search engine, the Signal messaging app, and Apple's decision to notify iPhone users when geolocation tracking is on. But these are exceptions, not the rule. Our websites still track us. Our apps still follow us. Our choices are limited. Our privacy is disappearing.[3]

Microsoft may have an experienced privacy team, but Windows 10 gathers information about us when we use its apps and when we browse the web, tracking our locations, our various machines, our home, our place of work, and our travel routes. Google makes much of its integrated approach to privacy (as have scholars), but the company's suite of products are data collection juggernauts. Google has been sued for its surveillance overreach in tracking university students but only because the company's practices violated a contract, not because they violated privacy law. And Facebook, a repeat offender in the privacy space, tracks not just our interactions on Facebook, Instagram, and WhatsApp, but almost everywhere else on the internet. The information industry is one big surveillance machine with powerful financial interests in commodifying our behavior. And it designs products to serve those interests.[4]

We have opposing interests. Some of us want to remain obscure from prying eyes and share intimate secrets with our families, our friends, and even strangers, all interests that are weakened by widespread surveillance. We want to make free and autonomous decisions and not be coerced by deceit or hidden manipulation. We have interests in equality and social justice that are undermined when our data is used to discriminate. And yes, although we also have interests in frictionless access to platforms that let us socialize or buy the products we need and we don't all think about privacy the same way, we all have individualized, collective, and dignitary interests in privacy that sometimes run counter to the data-hungry interests of the information industry.[5]

The Questions

How can privacy law and corporate commitments to privacy be on the rise without it having a significant effect on the designs of new technologies? We are supposed to be protected by privacy laws. Are they ineffectual? Are they not being enforced? Are they not even being implemented? We are supposed to have privacy advocates on the inside: armies of privacy professionals, privacy lawyers, and privacy engineers that are supposed to be changing the corporate culture. Are they all just marketing gimmicks? Are they undermining our privacy while misleading us about their privacy work? Are their employers' promises about doing better hollow? And what about laws like the GDPR that explicitly rely on in-house privacy professionals to do the ongoing work of interpretation, monitoring, and compliance? Is this approach to privacy law at all effective?

Stories about good versus evil are frequent in fiction, yet rarer in reality. It would be easy to dismiss corporate surveillance as

Cambridge University Press
978-1-108-49242-3 — Industry Unbound
Ari Ezra Waldman
Excerpt
More Information

the product of greed. And don't get me wrong, there *is* greed. A lot of it. But far more complicated structural and subtle forces are at play. Greed, for lack of a better word, is a bad look. Data-extractive capitalists who brag about how little they care about privacy are pilloried in the press. Even more importantly, greed cannot be the only motivator when there are armies of privacy professionals working inside technology companies with the earnest goal of protecting privacy. What's more, law is supposed to rein in the excesses of capitalism. But privacy laws on the books are not being translated into privacy protection on the ground. I went inside the information industry to explain why.

This book is about how actors in the information industry *do* privacy, or more specifically, how they can earnestly say they care about our privacy while simultaneously undermining it in practice. It focuses on what companies and their employees do behind the veil, behind the marketing, and behind the puffery when they actually translate the requirements of privacy law into their legal, organizational, and discursive behavior.

My findings are based on nearly four years of fieldwork, including interviews with current and former employees of large and medium-sized technology companies, interviews with start-up entrepreneurs and their venture capital backers, surveys of privacy professionals and software engineers, inductively coded analyses of industry literatures, interviews with in-house lawyers and their colleagues in private practice, reviews of internal protocols and procedures, analyses of legal arguments and public statements to legislative bodies and the press, observations of design processes, and listening and learning from those doing the work of privacy on the ground. The book consciously attempts to understand social phenomena from the ground up. And to do so, it relies on a diverse, interrelated set of conceptual models from law and the social sciences, including Foucauldian discourse theory, actor-network theory, science and technology studies (STS), performance theory, and critical sociolegal studies. I went into this project with an open mind, conscious that the truth always lies somewhere between cartoonish villainy and false heroism.[6]

The information industry often presents itself as contrite and dedicated to doing better at taking its privacy responsibilities seriously. Privacy scandals and new statutes trigger earnest rethinking among lawyers, privacy professionals, and technologists. Policy makers have passed new privacy laws that do more than require privacy policies that no one reads. But this work rarely has more than marginal effect on technology design. Why?

The Argument

The organizational, technological, and discursive system is stacked against privacy. Our privacy is at risk because of two related social forces operating within the information industry: coercive bureaucracy and normalizing performances. Tech companies use the tools of coercive bureaucracies to routinize antiprivacy norms and practices in privacy discourse, compliance, and design. Those bureaucracies constrain workers directly by focusing their work on corporate-friendly approaches to privacy. As information industry workers perform these antiprivacy routines and practices, those practices become habituated, inuring employees to data extraction, even as they earnestly profess to be privacy advocates. The result is a system in which the rank and file have been conscripted into serving the information industry's surveillant interests, and in which the meaning of privacy has been subtly changed, often without them even realizing what's happened.

Coercive Bureaucracy

Norms and routines inside corporations are the products of several internal and external influences. Situated within a sociolegal context, corporations are influenced by the web of laws, court decisions, rules, and real and threatened litigation and investigations that constitute the regulatory environment in which they, and their competitors, operate. They are also influenced by public opinion, market forces, and the behavior of their competitors. As a collection of individuals, corporations are also influenced by endogenous factors, including corporate structure and the embodied experiences of the real people doing the real work in the company's name. Of course, many of these influences overlap, but each works together to develop routines and embed norms throughout the corporation. Given this, some corporations inside the information industry work hard manipulate law, scholarship, structure, and the workplace experience to embed antiprivacy norms and routines wherever they can.[7]

When privacy professionals and privacy lawyers approach their work, they do so with background assumptions and understandings about what privacy means and how to protect it. Those ideas are heavily influenced by the values of informational capitalism. Industry leaders seek to influence how we think about privacy not just to erode our interest in and capacity to enact robust privacy laws, but to entrench corporate-friendly ideas as common sense and mainstream among their workers. This is the power of *discourse*.

The dominant privacy discourse today, from Silicon Valley to Washington, DC, centers around notions of choice, consent, and control; in other words, that privacy is about making our own choices about what to disclose, when, and to whom. It is a vision of privacy so narrow that it allows companies, their employees, and their allies to honestly say they care about privacy and still do little to improve privacy protections for their customers. But when you talk to people on the ground – the privacy professionals, the privacy lawyers, and the public policy shops – their independent self-reported views on privacy, though rarely hostile to corporate interests, are far more diverse than the party lines. Yet, few of those pro-privacy voices have any impact where it matters. In legal briefs, public reports, new products, press comments, and in testimony to Congress, discourses that protect corporate interests remain dominant.[8]

That is because tech companies use subtle and, at times, invisible strategies to silence pro-privacy voices and channel their employees' work to suit their ends. Industry executives set agendas for their privacy teams, require prepublication approval of academic research, control academic discussions through external funding of research, threaten researchers with restrictions on future access to data, and perpetuate false narratives about the efficacy of data-hungry AI tools. Lead in-house counsel and partners at private firms help determine legal strategies based on myopic definitions of privacy and enlist their subordinates in the effort while systematically denying them opportunities to voice alternative viewpoints. At the same time, many employees have internalized corporate cultures that encourage, value, and reward deference to leadership. This pushes them to incorporate the views of their bosses into their own,

Cambridge University Press
978-1-108-49242-3 — Industry Unbound
Ari Ezra Waldman
Excerpt
More Information

marginalizing their own views in the process. Together, these discursive tactics inculcate tech company workers with notions of privacy that perpetuate corporate power.

With these ideas in place, privacy law is undermined by narrow definitions of privacy that put few obligations on companies to actually protect privacy. And the privacy professionals on whom we depend to implement the law on the ground are so steeped in corporate-friendly privacy discourses and constrained by organizational structures that they end up weakening the laws we have even further. This is the power of *compliance*.

The newest privacy laws and proposals, from the GDPR to the CCPA and the roughly fifty new proposals for comprehensive privacy law in the United States, rely on internal organizational structures for implementation and ongoing monitoring. It reflects an incomplete form of collaborative or "new" governance. Reviews of these compliance structures, interviews with privacy professionals and lawyers, and observations of compliance operations in practice suggest that the largest and most entrenched companies in the information industry have built a house of cards of compliance structures. Tech companies make it look like they are following the law, but in truth, they are reframing it to achieve their own data-extractive ends. This happens in part organically. If you filter legal requirements through a corporate or managerial lens, you're going to get corporate and managerialized law. Management amplifies these effects by subtly manipulating and undercutting privacy professionals, reallocating budgets away from privacy, placing privacy advocates in stifling reporting hierarchies, siloing their departments, and leveraging workplace pressures and threats to silence pro-privacy voices.[9]

It is in this environment that the information industry designs its data-extractive products. When we share images with our friends on Instagram, buy products on Amazon, or conduct teleconferences on Zoom, we do so on their terms, not ours. Companies use this power of design to serve their data collection goals while making it difficult for us to realize our own privacy preferences. They build products that collect information without any benefit to us. They gather data from our clicks, our browsing, even where we move our cursor. They trick us into sharing information, hide opt-out buttons, make privacy navigation inordinately difficult, and trigger cognitive biases that constrain our choices.

It's fashionable to blame engineers for all of this; technologists are not trained to think about or design for privacy. That may be true, but programmers also work within corporate hierarchies that constrain and channel their work. Their work is just as much the product of manipulative social structures within organizations as it is the product of code. This is the power of *design*.

There are undoubtedly software engineers and compliance professionals and lawyers who are not only indifferent to privacy but see it as an impediment. The real story, however, is more nuanced. Corporate power over law and design processes means that companies can leverage internal organization, hierarchies, and policies to systematically devalue privacy and maximize data extraction in how they interpret the law and how they design new products, making antiprivacy designs more likely. Within this structure, even those software engineers and privacy professionals aware of their power and cognizant of privacy issues would nevertheless have a hard time pushing back against privacy-invasive corporate behavior. The result is a process that makes it difficult for privacy to establish a beachhead in design.

At the heart of this story is something more dangerous than mere bad faith: By influencing privacy discourse, undermining privacy law compliance, and constraining design processes, tech companies have not only unshackled themselves, they have created a corporate culture and environment in which all work, regardless of worker motivations and intentions, serve corporate antiprivacy goals. And the law not only allows this; it explicitly welcomes it.

Normalizing Antiprivacy Practices among Privacy Professionals

Throughout my research, I was struck by a disconnect between the stated motivations of information industry executives and those of their employees. Leaders are primarily motivated by capitalistic interests, which are almost always data extractive and, therefore, independent of or in conflict with privacy. "Too much privacy," one head of product development told me, "means our products won't be able to give people what they want: access, fast access, convenience, connection, and fun." A start-up executive in New York admitted that "privacy isn't at front of mind when you're trying to make it, like make it in this business. Data means better

Cambridge University Press
978-1-108-49242-3 — Industry Unbound
Ari Ezra Waldman
Excerpt
More Information

targeting, which means more money, and more investments. I need that data if I'm going to succeed in this environment." "Restrictive privacy law is bad for us," another executive admitted. These comments aren't isolated. Some information industry executives are on record asserting that privacy is in conflict with their businesses' success. Even those that recognize the importance of privacy do so in the interests of profit.[10]

But those doing the work of design and privacy law on the ground profess to find motivation elsewhere. Software engineers want to solve exciting engineering problems (even privacy ones!), to think of ways to make products more efficient, and to develop exciting new technologies; privacy professionals, for the most part, want to pragmatically facilitate privacy compliance. These motivations aren't capitalistic per se. They are technical, vocational, and careerist, just like many of the reasons we all go into our careers of choice. And yet, despite these differing motivations – some of which seem at odds with one another – the information industry realizes its surveillance goals in the end: Technologies are far more likely to ignore or violate our privacy than accommodate our interest in it.

The information industry perpetuates its power through a process of normalization. Tech companies focus their privacy work on narrow privacy discourses. They create compliance mechanisms that reduce privacy law to check boxes and outsourced technologies that allow industry to escape accountability. Companies then design new surveillant technologies while silencing pro-privacy voices. At each stage, corporate data extraction and the actions that facilitate it are normalized as ordinary and routine. With every privacy assignment focused on notice or security, employees become conditioned to thinking about privacy in narrow, underinclusive ways. With compliance focused on paper trails, checkboxes, and prefilled reports from outsourced vendors, privacy professionals come to confuse mere symbols of compliance with actual adherence to privacy law. And with every small engineering team focused on a narrow engineering problem, software engineers become accustomed to thinking that privacy is someone else's responsibility. These antiprivacy practices then become common sense for information industry employees, making it difficult for them to see privacy law as anything but what they have been doing.[11]

Cambridge University Press
978-1-108-49242-3 — Industry Unbound
Ari Ezra Waldman
Excerpt
More Information

These practices are performative. By “performative,” I don’t mean that they’re fake, ersatz, or cynical false fronts. They can be. But performances in this context are actions and behaviors that communicate something to the self and others. Performances are performa*tive* when their repetition and iteration socially construct and define our identities, our realities, and, I argue, privacy law. The social practices of privacy law feed and perpetuate themselves, constructing a reality where discourses of control and symbolic compliance *are* privacy law. And that’s why they don’t work. A regulatory regime that relies on regulated entities to flesh out the details of the law in practice performatively constructs a privacy law that is not only weak, but counterproductive. Symbolic compliance legitimizes what is really a con game.[12]

Privacy professionals, privacy lawyers, and software engineers perform privacy within constraining organizational bureaucracies. As they do, as they repeat practices described in Figure 1 that end up marginalizing privacy voices, their performances become routine, their routines become habits, and their habits become part of how they conceptualize privacy law. In other words, as they repeat corporate-friendly privacy practices, they normalize them and that normalization feeds back into a bureaucratic system built to drive corporate-friendly privacy discourses, compliance, and design. Therefore, our surveillant technological landscape is less the result of corporate shills than it is the product of organizational structures where antiprivacy work is habituated and normalized by ongoing performance.

Discourse	Compliance	Design
Discursive practices normalize corporate-friendly discourses of privacy – particularly, privacy-as-control – as common sense. These discourses form the backdrop for legal and technical work.	Constrained by corporate bureaucracies, privacy offices routinize compliance practices that normalize symbolic compliance.	With weak privacy discourses and symbolic compliance, bureaucratic practices not only take advantage of designer disinterest in privacy, but also remove opportunities for privacy entrepreneurship in the design process.

Figure 1 Privacy performances.