

Introduction

In September 2010, eighteen-year-old Tyler Clementi, a freshman at Rutgers University, was filmed secretly through his bedroom webcam while making out with another man. The video, taken by his roommate, was then distributed through Twitter and among students at his school. When Tyler had a second date, a group of students organized a viewing party to watch footage from a hidden camera.¹ Tyler was ridiculed. His sexuality was exploited for fun. After the incidents, he jumped to his death from a bridge into the Hudson River.

It's tempting to misconstrue the harm caused to Tyler as the result of an individual actor – his roommate. But it's inaccurate.² The harm was collective and it was enabled by a slew of digital services and devices we use daily, provided for profit. While blameworthy, Tyler's roommate and classmates were part of a digital ecosystem that, dismissing value in people's privacy, creates enormous tangible and intangible harm.³ From 2010 to today, this system of unfettered collection and sharing only got larger, more sophisticated, more profitable, and more harmful.

All your interactions, movements, and decisions are collected in real time and attached to profiles used by advertisers to compete for your attention. Not because they think you're special or because they're interested in learning about you for the sake of getting to know you, but because, regardless of your age, gender, or country of origin, you're monetizable. When combined, these little pieces of ourselves fuel a trillion-dollar industry that threatens livelihoods, lives, and democratic institutions.

The worst part is not that we get little in exchange. It's that, much like companies that pollute the atmosphere or that offshore production to places where they can violate workers' human rights, every step of the data industry creates losses and harms that are opaque but real. Companies that collect, process, and sell our personal information create harms that are out of sight but have dire consequences for those affected.

Clara Sorrenti experienced firsthand the consequences of data harms. She received death threats, had her home address found and shared, saw intimate documents about her family revealed, and was "swatted" – the practice of falsely reporting a police emergency to send armed units to an innocent person's home,

an experience that for Clara ended up with an assault rifle pointed at her head.⁴ Clara was a victim of Kiwi Farms, a platform that coordinates the gathering of information available online to target trans people.⁵ For the law to consider that you harassed someone, you need to contact them several times. So, if a platform pools information and coordinates people who each contact a victim once, as was the case for Clara, it produces harassment while avoiding its legal definition. Describing her experience, Clara explained: “When you get your own thread on Kiwi Farms it means there are enough people who are interested in engaging in a long-term harassment campaign against you.”⁶ She left her home after the swatting incident, but Kiwi Farms found her by comparing hotel bedsheet patterns from a picture she took with information available online.⁷ Clara fled the country to escape abuse and Kiwi Farms found her again.⁸

Viewing Clara’s harassment as the work of a few bad individuals ignores a broader systemic problem. In our digital ecosystem, it’s easy to obtain and use our data in ways that inflict harm on us – like a roommate exploiting a teenager’s sexuality for entertainment or a website exploiting trans women’s physical safety for dollars. Because Tyler and Clara’s harms weren’t just a result of individual trolls, but rather emblematic of an ecosystem that enables and magnifies data harms, Tyler and Clara’s situations are not exceptional.⁹ Part of what’s shocking about their stories is that they faced enormous harm from something as common as webcams and blogs.

Data harms differ. Some are visible and affect people such as Clara on an individualized basis, with immediate consequences on their livelihood. Daily victims include women facing online harassment and abuse, racialized individuals experiencing magnified systemic discrimination, and anyone going through identity theft because their financial information was taken without their knowledge.¹⁰ Most harms in the information economy, however, are opaque and widely dispersed. Examples are online manipulation to make personal and financial choices against our best interests (called “dark patterns”) and the normalization of surveillance to constantly extract personal data (called “data mining”).¹¹

Tyler and Clara were pulled into the information economy – the trillion-dollar industry fueled by the collection, processing, and sharing of personal information to produce digital products and services.¹² When we look at data interactions, we sometimes forget that it’s there. For example, in nonconsensual distribution of intimate images, there’s a tendency to concentrate all blame on the first perpetrator. But when intimate photos go viral, that means hundreds of people reposted them and websites derived ad or subscription profit from them.¹³ Victims are harmed because there’s a data ecosystem that facilitates and encourages it. The information economy enables corporations and individuals to instrumentalize others for their own gain – and it amplifies them.¹⁴ So cases of one perpetrator and one victim barely exist. We lack accountability over what happens with our data and what harms happen to us because of our data. By having a better picture of that data ecosystem, laws can better reduce and repair data harms.

This book builds on academic and policy critiques to privacy law, which is the body of law that governs the collection, processing, and sharing of personal information. It explores these critiques' consequences to explain where privacy laws fall short when it comes to the information economy and how their shortcomings relate. It then proposes what we could do about it.

The central problem is the following: privacy law across the world – including in the United States (US), the European Union (EU), and countries that modeled their privacy laws on either of them – is based on regulations designed for contract-like relationships. The foundations of privacy law therefore rest on two critical assumptions: that people can freely and rationally make data decisions that increase their wellbeing and that legislators can design rules that anticipate and prevent data harms. Neither of these assumptions is true. Further, privacy law fails us because it relies on false assumptions about how people behave and what people believe regarding their privacy.

Facebook serves as an example. In 2016, shortly before the Cambridge Analytica scandal, journalists uncovered that Facebook had been facilitating housing discrimination. Facebook's software made it possible for advertisers to filter who saw their ads by race, gender, nationality, and other protected characteristics. Marketers used this feature to avoid showing housing ads to racialized users and have whiter tenants.¹⁵ The US National Fair Housing Alliance sued, supported by law enforcement.¹⁶ After it did, Facebook agreed to remove its "ethnic affinity" filter.¹⁷ But the information that Facebook has about its users is so detailed and nuanced that this hardly made a difference. For example, advertisers can't filter by who's Latinx, but they can filter by who likes Telemundo. Advertisers can't filter by who's gay, but they can filter by who likes gay tourism websites. And advertisers can filter by "multicultural affinity."¹⁸ Facebook continues to classify its users by over 5,000 categories, some of which enable indirect discrimination.¹⁹ The company didn't eliminate discrimination; it just hid it. Because of the host of information it collects and infers, the company has enough power to discriminate while complying with antidiscrimination and privacy law. The issue isn't unique to housing or to Facebook, but common to platforms that can weaponize information about us to selectively expose us to opportunities, turning our information against us.²⁰

The international tendency to base privacy law on consent models from contract law is most extreme in the US. Omri Ben-Shahar and Lior Strahilevitz once described the tendency as "a quiet legal transformation whereby the entire area of data privacy law has been subsumed by consumer contract law."²¹ In the EU and countries with EU-inspired data protection laws, similarly, laws hinge on individual consent and individual control – as if the relationships were in a market.²² In both cases, laws' framework is founded on the notion of bilateral commercial relationships. The underlying dynamic for the contractual view is that there's a trade in which people agree to give up their personal information in exchange for a service. There's not.

In contrast to privacy law's assumptions, we mechanically click "I agree" on documents that would be unhelpful to us even if we read and understood them.²³ We do so for a wide range of corporations, such as websites, apps, and internet service providers (ISPs) that profit from our data. To supplement these agreements, governments make long checklists for corporations to tick to achieve legal compliance. But, as Facebook did when enabling discrimination, corporations cause enormous individual and social harm, often while remaining compliant.²⁴ Corporations obtain meaningless "I agree" clicks, performatively comply with checklists, and continue business as usual.²⁵

This book's core premise is that, rather than grounding privacy law on concepts from contract law, which sets the rules for voluntary agreements, we need to ground it on concepts from tort law, which sets the rules for harms caused to others.²⁶ This premise may sound technical, but the reasons justifying it are intuitive because they respond to the social reality we all live in. Contract law works well for standard exchanges and agreements, like when we buy groceries or hire the services of a dry-cleaner. But the mutual understanding and agreement on the specifics of an interaction central to contract law (what legal scholars call a "meeting of the minds") doesn't exist in privacy. Privacy agreements' subject matter is opaque to the people they involve. We don't know what we give up in data interactions – like Clara couldn't predict being found from a nondescriptive picture.²⁷ And many companies that hold our data never interacted with us in the first place – like Tyler, who had no Twitter account.²⁸ By exerting power over people both within and beyond empty agreements, corporations do mass harms, including but not limited to their users.

The result of the mismatch between laws' assumptions and social reality is that corporations are free to exploit people whose information they collect, process, and share. They can do so by misusing their information for financial gain (data misconduct) and profiting from people's data without keeping it safe (lack of data security). Technologies that make it easier to analyze large amounts of data facilitate this type of exploitation. The increasing reliance on artificial intelligence (AI) for processing data and our increasing dependence on data-mediated social and economic interactions make exploitation a serious concern for what our society may soon become. The corollary is that solutions must involve substantive reform. We need to rethink the building blocks of privacy protections in the private sector.

Given the failure of the current model, the book proposes a program for building meaningful accountability into the information economy through liability for individual and group harms. Law's framework should move to one of compensating harms that occur outside mutually beneficial agreements. This program departs from existing laws and liability proposals, which focus liability on breaches of procedural rules or individual agreements. These breaches are too narrow to capture the different and unpredictable ways in which people can be harmed and exploited. To overcome these problems, the book proposes a theory of harm and exploitation that addresses common concerns with liability, such as standing, causation, class certification, and compensation.

The needed changes extend to regulatory reform. Regulations should complement harm-based liability regimes by focusing on systemic risks. Regulations must match the right type of underlying relationships and power dynamics. The changes this book proposes would depart from the current focus on individual rights for controlling personal information. Regulations that reinforce individual control are unhelpful because they implicitly rely on a contractual “meeting of the minds,” and laws can’t reinforce something when it doesn’t exist. Individual choices and individual control rights that provide people with options can’t meet this imperative because they elide an array of social harms.²⁹ Patching the current system with additions that uphold its underlying contractual logic, like the right to data portability or the right to be forgotten, is a band-aid solution, rather than a cure. Regulators are well positioned to reduce systemic risk and the magnitude of widespread harms, which requires looking into and moderating the power dynamics embedded in data practices’ business models.³⁰

The type of liability developed is crucial for achieving accountability. Liability proposals so far suggest compensation when corporations break a promise made in their terms of agreement or undertake an activity prohibited by a procedural rule. These forms of liability are contract-like, similar to liability arising from breach of contract or breach of a legislated mandatory contractual clause. They fail to reduce harm because they rely on similarly flawed assumptions over underlying dynamics. Data harms are different; they resemble mass harms addressed by modern tort law, such as environmental harms.³¹ To address them, privacy law needs to hold corporations accountable through tort-type liability. Protection requires that corporations are held accountable for the consequences of their data practices – not for the checklists they complete or the notices they send.

The pervasive data harms that exist in the information economy show that this type of accountability needs to be at the center of the protection system, rather than an add-on to the system’s enforcement. Recently, privacy scholars developed other calls for consequence-focused meaningful accountability, such as information fiduciaries, privacy by design, and relationships of trust.³² This proposal builds on their motivations and is compatible with their implementation. They all respond to a social phenomenon that took off just under twenty years ago.

If you’re old enough to remember one of the first-ever cases of viral information sharing, you may remember that it happened because hundreds of entities exploited and humiliated a woman for private gain. Twenty-four-year-old Monica Lewinsky found herself at the center of the news over her affair with President Bill Clinton. Lewinsky discovered the cost of artificially inflated shame for profit.³³ The more shame and scandal created, the more clicks they received. And the more clicks, the more ad revenue. Clinton’s infidelity may have been newsworthy, but hundreds of memes, posts, photos, and commentaries made it about her. Lewinsky wasn’t a public figure, but rather an intern in a relationship characterized by an exorbitant power differential. As she explains, people “plastered photos of me all over to sell newspapers, banners online, and to keep people plastered to

the tv ... the attention and judgment I personally received was unprecedented.”³⁴ Lewinsky, like Clara, was not harmed by a specific individual, but rather by an aggregation that the information economy’s incentives structure enables and fosters. Back then, there was barely a name to designate what she went through. We now call what hundreds of individuals did to her “online harassment.” We still lack a name for the systemic effect.

Traditionally, when laws and courts address privacy issues, they focus on tangible consequences. This is true whether the problem is a data breach, like when a company is hacked, or the violation of a data right, such as the right to know what information a company has about you. This important but insufficient conception contemplates financial harm such as loss of money, loss of reputation that damages one’s employment relationships, and, in some cases, physical consequences, such as harm to one’s health or safety. Privacy laws attempt to foresee and prevent these harms.

Modern data practices changed things. They introduced complicated power dynamics where corporations use people’s information, often with the help of AI, to make decisions about their opportunities and experiences. Modern data practices also allow harms to arise between parties who never interacted with one another, such as harms from data brokers, who buy your data to aggregate it and sell a profile about you to others.³⁵ Through these power dynamics, modern data practices introduced and fuel informational exploitation, a different type of data harm that involves profiting from people’s information with disregard for the harm that it causes them. Informational exploitation differs from other data harms in that it’s systemic, it’s opaque, and it facilitates, while simultaneously hiding, other harms. Informational exploitation is the systemic effect that Lewinsky was put through.

Surveillance that facilitates exploitation is easier, cheaper, more pervasive, and less evident than ever before. Practically every time you interact with a screen, your clicks are monitored, what you look at is recorded, your activity is surreptitiously linked to your identity, your information is traded, and all of it is aggregated with information from others. Most significantly, statistical inferences are constantly made about you and the groups you belong to. This dynamic gives hundreds of corporate entities power over you.

To address the systemic effects that new relationships of power produce, we must identify privacy-violating data practices by connecting them with the reasons for which we value privacy. Privacy is a social value, so it’s about more than preventing negative tangible consequences.³⁶ Protecting privacy is important for building trust, preserving autonomy, and maintaining relationships. It protects us from emotional harm, such as distress and anxiety. Numerous theories of privacy explain what privacy is and why we protect it, underscoring its relationship with intimacy, autonomy, personhood, and trust. These theories show that privacy has intrinsic and instrumental value: it has independent social value and it protects people from other harms, such as financial fraud and physical violence.

This book doesn't advance a new concept of privacy. Rather, it builds on these concepts of privacy, together with lessons from behavioral science, economics, psychology, and sociology, to better design the system that protects people. Its proposals apply to different conceptions of privacy, which can be advantageous for advancing policy arguments in light of differing views.³⁷ The problems of the traditional protection system, as well as the need for accountability for data practices' consequences, apply across privacy theories because all those theories recognize that there's something in privacy worth protecting.³⁸ An accountability program based on privacy harm just requires recognizing that there's something valuable in privacy that can be subjected to systemic loss and harm.

The privacy fallacy causes us to miss that value. This fallacy refers to the disjuncture between a notion that privacy has value in and of itself and the conviction that, at the moment of protection, only tangible consequences – like physical or economic harm – are real, concluding that privacy can be sufficiently protected by preventing those outcomes. It contains a contradiction, because the idea that privacy has intrinsic value implies that there's a value in privacy that can be harmed, even absent physical or economic consequences. Opinion leaders succumb to the privacy fallacy when they solely address privacy's instrumental consequences in a particular issue and subsequently claim to have successfully protected privacy, dismissing the loss of privacy's social value. Thinking that preventing Tyler Clementi's suicide would have solved his invasion of privacy, for example, would be falling for the privacy fallacy. Regulators and industry members do so when they understand and endorse the value of privacy in theory, but forget about it in practice. Authorities fall into the fallacy when their protection regimes only recognize the tangible consequences of privacy losses, while politicians repeatedly remind people of the value of privacy. People fall into the fallacy when they say that, even though privacy is important in general, you shouldn't worry about it in a specific situation if you have "nothing to hide."

In its most popular and most dangerous form, the privacy fallacy is used to argue that each individual should protect themselves from those tangible consequences. It overlooks the loss of privacy's social value and how, in any information interaction, we affect each other. Traditional laws buy into the privacy fallacy by committing to the idea that it must treat people as hypothetically rational and perfectly informed entities and that, absent physical harm or financial fraud, their own choices will protect them from harm in the information economy. Public policy efforts buy into this fallacy by building on the mistaken belief that, by adding procedural requirements, people will at some point take control over their data. This approach pays lip service to privacy. It creates the illusion that we're moving forward and legislators are placing strict requirements on corporations that will, one day, achieve individual control. Though, even if regulators did provide individual control for people to prevent tangible consequences, the privacy values they claim to protect would remain unprotected.

* * *

The book develops this argument in seven chapters.

Chapter 1 ties together problems in central elements of privacy law: the individual choice-based system, the fair information principles that originated it, the view that privacy is about secrecy, and dichotomies such as public versus private. We don't have actual choices about our data beyond mechanically agreeing to privacy policies because we lack outside options and information, such as what each choice means and what risk we're taking on by agreeing. The choice-based approach creates a false binary of secret versus public information when, in reality, privacy is a spectrum. The idea that someone, at any given time, has either total privacy or no privacy at all is unfounded. Additionally, data are bundled: you can't reveal just one thing without letting companies learn other things. Reckoning with this reality defeats the popular "I have nothing to hide" argument, which traces back to Joseph Goebbels.

Chapter 2 shows the falseness of two ideas that underlie the central elements of privacy law: that people make fully rational privacy choices and that they don't care about their privacy. These notions create a dissonance between law and reality, which prevents laws from providing meaningful protection. Contrary to rationality, context has an outsized impact on our privacy decisions and we can't understand what risks are involved in our privacy "choices," particularly with AI inferences. The notion that we're apathetic is prevalent in popular discourse about how much people share online and the academic literature about "the privacy paradox." Dismantling the myth of apathy shows that there's no privacy paradox. People simply face uncertainty and unknowable risks. People make privacy choices in a context of anti-privacy design, such as dark patterns. In this process, we're manipulated by corporations, who are more aware of our biases than regulators are.

Chapter 3 shows why the contracts model doesn't work: consent in the information economy is an illusion. Inferences, relational data, and de-identified data aren't captured by consent provisions. Consent is unattainable in the information economy more broadly because the dynamic between corporations and users is plagued with uneven knowledge, inequality, and a lack of choices. Privacy harm can't be seen as a risk that people accept in exchange for a service. Data harms are collective and unknowable, making individual choices to reduce them impossible. Worse, privacy has a moral hazard problem: corporations have incentives to behave against our best interests, creating profitable harms after obtaining agreements. Privacy's moral hazard leads to informational exploitation. A manifestation of valid privacy consent is consent refusals among individuals. We can consider them by thinking of people's data as part of them, as their bodies are.

Chapter 4 delves into two modern efforts to reinforce individual consent: opt-in and informed choice. It illustrates why, in the information economy, they also fail. Power asymmetries enable systemic manipulation in the design of digital products and services. Manipulation by design thwarts improved consent provisions, interfering with people's decision-making. People's choices regarding their privacy are determined by the designs of the systems with which they interact. European and

American attempts to regulate manipulation by changing tracking from opt-out to opt-in and reinforcing information crash against the illusion of consent. Contract law doctrines that aim to reduce manipulation are unsuitable because they assume mutually beneficial agreements, and privacy policies are neither mutually beneficial or agreements. Best efforts to strengthen meaningful consent and choice, even where policies are specifically intended to protect users, are ultimately insufficient because of the environment in which privacy “decisions” take place.

Chapter 5 examines traditional data protection law’s regulatory structure in light of these considerations. It shows why data protection rights and rules, while desirable, don’t address the core problems of the contracts model and can’t work well without the liability model. Data protection rights unintendedly impose administrative burdens on those they protect. Mandatory rules address power asymmetries and manipulation better than defaults. But our procedural rules overregulate while they underprotect: they benefit large players by adversely affecting new players and they allow companies to comply merely by following box-ticking exercises. Against this backdrop, laws legitimize exploitation that can be executed while remaining compliant. Risk-reduction approaches based on standards can reduce informational exploitation more effectively.

Chapter 6 explores a different path: building privacy law on liability. Liability for tangible and intangible privacy harm would improve our protection systems. To achieve meaningful liability, though, laws must compensate privacy harm, not just the tangible consequences that stem from it. Compensation for financial and physical harms produced by the collection, processing, or sharing of data is important but insufficient. The proposed liability framework would address informational exploitation by making companies internalize risk. It would deter and remedy socially detrimental data practices, rather than chasing elusive individual control aims. Applying it, courts and regulators can distinguish harmful losses from benign ones by examining them on the basis of contextual and normative social values. By focusing on harm, privacy liability would overcome its current problems of causation quagmires and frivolous lawsuits.

Chapter 7 proposes how the liability framework should be implemented. Harm liability can flow from a statutory standard or local tort law. This focus allows liability to complement, rather than replicate, public enforcement. The quantum of liability should depend on the harm incurred by the victim, rather than on the wrongfulness of the perpetrator’s conduct or the consequences that the perpetrator foresaw. Because harms are often dispersed, privacy liability is most effective as part of a mechanism of collective redress, such as class actions. Considering privacy problems at scale, we need a framework recognizing mass privacy effects for regulators and courts. A robust notion of loss and intrinsic harm can address problems of insufficient compensation and uncertainties in class certification.

* * *

The information economy is formed by entities that profit from people’s data and the people whose data those entities profit from. The corporate entities in the information economy are varied. They’re websites, apps, advertising companies, product designers, social networks, data brokers, search engines, and manufacturers of Internet of Things devices, among others. Many, but not all, are tech giants. Often, I’ll refer to a specific type of entity, but when discussing all of them I’ll refer to “corporations” or “companies.” Similarly, I’ll refer to “users” as a shorthand when referring to those who use an app or platform. But the more accurate term would be “affected parties” because sometimes we’re part of the information economy without using any service at all. Sometimes we engage in the information economy as “consumers,” but we’re not only affected while consuming something. For example, someone else can share data about us. EU law uses the term “data subjects” to stress that individuals are the key actors. But “data subjects” are seen in discrete ways, making it seem as if we’re detached individuals with different interests, as opposed to an interrelated group whose actions affect each other because they’re connected by data.³⁹ Because everyone’s data is part of the information economy, most times I’ll refer to “people.”

The world changed significantly since 1973, when privacy law was conceptualized. Our new environment necessitates a different approach to privacy than what was conceived back then. Privacy law’s challenge is no longer regulating individual choices, but rather regulating relationships of power.⁴⁰ And addressing power doesn’t require presenting choices to the powerless. Addressing power requires holding the powerful accountable for the consequences of what they do.⁴¹

There’s an old Silicon Valley mantra to “move fast and break things,” encouraging disruption regardless of risk.⁴² Privacy harms today are more pervasive and significant than those that took place back then. Data’s central role in our economy and the increasing role of AI inferences in our daily lives will continue to accelerate this drift. The law needs a solution that allows for technological, economic, and social progress while protecting people from being turned into collateral damage. To move privacy law forward, we must abandon the old contractual paradigm and try something more difficult: holding corporations responsible for the things (and people) they break. Taking harm and exploitation in the information economy seriously is overdue.