1

The Traditionalist Approach to Privacy

Imagine you notice someone is following you on your way to work one morning. You find it concerning, but brush it off. Then another stranger follows you to the gym in the afternoon. You get worried, but carry on with your day. You eventually find out that the two share notes with each other, and become shocked. You finally become scared when you find out that everyone you spoke with that day also took notes about you during casual conversations and reported back to the strangers. This, essentially, already happens. Only it's devices like your phone, laptop, and smart home devices that do the tracking. And it's every written communication that gets recorded – unless you have devices with mics, and then it's the spoken ones too.

Privacy law emerged without the Internet or AI and evolved without revisiting its core assumptions. As a result, it's stuck in time. Core concepts in privacy law no longer correspond with daily social interactions in the information economy.

Privacy law across the world is grounded on ideas from nineteenth-century neoclassical economics of contracts – what I call "the traditionalist approach to privacy." Neoclassical economics makes assumptions about how people behave in market exchanges: it assumes people behave rationally, optimizing choices for their own wellbeing based on available information. These assumptions permeate how the law addresses commercial interactions. In many contexts, such as in mergers and acquisitions, the stock market, and most commercial contracts, these assumptions are helpful. In other contexts, such as in parent–child caregiving, less so. When the law uses the wrong assumptions, placing weight on them can impede it from protecting the vulnerable parties that it's meant to protect.

These assumptions don't reflect the reality of contemporary data interactions.[1] Yet the law places enormous weight on them. They dictate the law's worldview about how people make privacy choices (rationally, in an informed way), how people use their privacy (to keep secrets), what activities underly (bilateral commercial transactions), and how people's privacy ought to be protected (by providing more choices).

This book explores the myths that the neoclassical contracts conception creates and how privacy law can and should overcome their obstacles. It argues that the traditionalist approach led privacy law to ineffectively build on concepts from contract

https://doi.org/10.1017/9781108995825.002 Published online by Cambridge University Press

law and shows how it can and should build on concepts from tort law instead. It attempts to chart how to change the foundations of privacy law to move toward a paradigm that protects real-life people in the twenty-first-century economy.

A FORCING PEOPLE TO CHOOSE

In their song "Freewill," rock band Rush says that "If you choose not to decide, you still have made a choice."[2] The idea comes from a famous quote from philosopher Jean-Paul Sartre, who emphasized that not choosing is, in itself, an important choice.[3] The information economy deprives us of this choice. Every day, we must make decisions about our personal information that we're not prepared to make.

Notices, Choices, and Self-management

As you diligently read the Amazon Web Services Terms & Conditions before agreeing to them, you probably noticed a curious clause in its gaming section. The clause indicates that a limitation won't apply "in the event of the occurrence (certified by the United States Centers for Disease Control or successor body) of a widespread viral infection transmitted via bites or contact with bodily fluids that causes human corpses to reanimate and seek to consume living human flesh, blood, brain or nerve tissue and is likely to result in the fall of organised civilization."[4] Technically, you consented to a way of certifying a zombie apocalypse.

The idea of valid consent (often called meaningful or informed consent) is pivotal in privacy law.[5] With limited exceptions, over the past fifty years individuals' consent has been the main basis to collect, process, or share their personal information, forming the bedrock of corporate privacy practices.[6]

Legislatures around the world are guided by the primacy of individual consent when establishing the default legal basis for collecting, processing, and sharing people's personal information. When discussing how to update data protection law, EU Justice Commissioner Julia Fioretti asserted that "[c]itizens should have more possibilities, more chances to be the masters of their personal data, to be informed on what somebody does with their personal data."[7] In 2012, the US Federal Trade Commission (FTC) proposed changes in US privacy law to give people "the ability to make decisions about their data at a relevant time and context."[8] A White House effort that year aimed for people to have "clear and simple choices, presented at times and in ways that enable consumers to make meaningful decisions about personal data collection, use, and disclosure."[9] As early as the 1990s, the Canadian government held that notice and consent "are the core values in any personal information code."[10]

Recent modernization efforts are also guided by individual consent as their gold standard. Press releases of the European Parliament state that people "should have full control over their data and be empowered to take decisions about it."[11] The interpretative authority for the EU Data Protection Directive repeatedly stated that

https://doi.org/10.1017/9781108995825.002 Published online by Cambridge University Press

control over personal information is central to data protection, where control is achieved through consent.[12] In 2021, the Canadian government proposed to overhaul its private-sector privacy regime to "enhance consumer control by requiring organizations to get meaningful consent from Canadians."[13] The Australian Information Commissioner's website states that "Consumer consent for the collection, use, and disclosure of their data is the [law's] foundation … ensuring they can direct where their data goes to obtain the most value."[14]

The main way for companies to obtain our consent and for us to manage our privacy is through "privacy notices."[15] These notices are privacy policies or terms of service that corporations share with their users to explain how they collect, process, and disclose personal information, asking their users to agree to them. Privacy notices capture individual consent as the key to unlocking the data practices described in them.

Privacy notices, and the promises companies make in them, are central to privacy law globally. The global popularity of this practice may be linked to how it embraces a common regulatory approach: give people control (in this case, over their information) so they take care of themselves.[16] After all, that's how the law deals with most of our possessions, from apples to non-fungible tokens (NFTs). The history of how this practice took over our information dates back to 1973.

The Fair Information Principles

Every time you download a new app on your phone, you're asked to agree to its terms of service. The reason dates back to the early days when the world worried about the digitization of personal data and developed the 1973 Fair Information Principles (or Fair Information Practice Principles, usually referred to as the FIPs) to address it.

The FIPs have slowly become synonymous with privacy law. As their name indicates, the FIPs aim to make practices relating to peoples' personal information fairer.[17] They were initially principles developed in an American advisory committee report as guidelines for the private sector.[18] Rapidly growing out of that report, they became FTC guidelines, Organisation for Economic Co-operation and Development (OECD) international guidelines, and eventually law.[19] Today, they're the backbone of privacy and data protection legislation around the world.[20] They're the basis of privacy and data protection laws, for example, in Argentina, Australia, Canada, the EU, Japan, Malaysia, New Zealand, the Philippines, Singapore, South Korea, the United Kingdom (UK), and the US, among many others.[21]

The FIPs have many permutations and one commonality. For example, the FTC lists notice, choice, access, security, and enforcement as the principles the private sector should abide by.[22] Europe's General Data Protection Regulation (GDPR) lists eight.[23] The OECD proposes eight others.[24] The lists of principles go on.[25] Despite their differences, though, all FIPs permutations have one thing in common. They aim to increase people's control over their personal information.[26]

Apps ask you to agree to their terms of service because individual control is mainly afforded through opportunities to consent (or not) to the collection, processing, and distribution of your personal information.[27] And, according to the FIPs, consent requires two things: giving you notice and giving you a choice. For people to have notice, they must know how an organization will collect, use, and share their personal information. For people to have a choice, they must be able to decide whether to agree with the collection, use, or sharing by considering whether the benefits they may get from it outweigh the risks.[28]

The primacy of notice and choice is most marked in the US, where a regulatory peculiarity elevates privacy policies to a mainstay of privacy governance. The federal US agency tasked with regulating and enforcing privacy is the FTC. Established to protect consumers, the FTC's mandate is to investigate and pursue "unfair and deceptive" practices.[29] The agency ensures that corporate promises (if any) are fulfilled and sanctions companies when they fail to notify consumers of a practice – or, occasionally, for improper conduct, such as maintaining inadequate cybersecurity measures.[30] What corporations promise they'll do isn't the main object of scrutiny; whether they did what they promised is.[31]

The US emphasis on promises was part of a broader regulatory strategy. From the 1970s to the 1990s, Congress and US regulatory agencies prioritized disclosure schemes such as notices to achieve regulatory goals, rather than designing substantive regulation.[32] In privacy law, this strategy stuck.[33] Although in theory the FTC's privacy enforcement is guided by all FIPs, in practice it prioritizes the principle of notice.[34] The FTC refers to notice as the "most fundamental principle."[35] Ensuring proper description and adherence to data practices lies at the core of the FTC's role as a privacy regulator.[36] Functionally, the FTC mostly enforces private agreements.

The tech industry lobbies for notice-and-choice.[37] Mandating notices is a much lighter regulatory intervention than mandating or forbidding data practices. Mandating notices, rather than developing substantive regulation, reduces regulatory costs.[38] Notices take a market-style approach that intervenes without actually intervening: they're in line with the approach of regulating by giving "choices" to people and, instead of mandating or forbidding practices, letting people decide which ones they'll accept.[39]

Notices are also easy for agencies to enforce.[40] They're easier to develop than substantive regulation because they place the onus on each individual to decide what's OK and what's not.[41] In the face of different business practices, technologies, and processes that affect people's privacy, the easiest thing regulators can do is to verify that each corporation adequately describes its data practices and adheres to them. Privacy regulators bind corporations to their privacy policies by punishing them for breaking the promises made in them outside the US too.[42]

The FIPs' goal of individual control over information fails because we're not given the means to make those choices.[43] As Woodrow Hartzog puts it, "privacy law is in a bit of a pickle thanks to our love of the Fair Information Practice Principles."[44] The pickle is that the FIPs have become synonymous with privacy protection. Initially

designed as guiding principles and not specific provisions, they turned into provisions that regulate personal data around the world. Designed at a time before people had computers at home, let alone the Internet, the FIPs aimed to protect people in a vastly different environment than the current AI-driven information economy. Since then, privacy scholars have heavily criticized them.[45]

The FIPs' failure to protect people's privacy isn't their own doing. As this book explores, it's rather the failure of the paradigm on which they're developed and implemented. The appropriate solution isn't to just change the FIPs. It's to change the building blocks that support them.

Notices that Don't Inform

Privacy policies are in a predicament.

In the early 2000s, Aleecia McDonald and Lorrie Cranor had an unorthodox idea: to check how long it would take the average person to read the privacy policies of every website and app that she uses for a year. The answer was astonishing: 244 hours per year, or six full-time working weeks.[46] In the decade and a half since the study, this number can only have increased. Another study found that less than 5 percent of people read them, a result that may be optimistic.[47] Not even the sitting Chief Justice of the US Supreme Court reads them.[48]

The no-reading problem is only the tip of the iceberg of privacy policies' issues. Even when we read a privacy policy, we can't understand it.[49] Their meaning is lost to people navigating passages that are too detailed or ambiguous to be helpful.[50] The result is that readers trying to penetrate the obscure content of the one document that's supposed to explain how they're being surveilled are left with either a sense of confusion or a false sense of understanding.

Many call for more user notices to increase transparency so that people can make informed choices – doubling down on the traditionalist view.[51] But others suggest that privacy notices are ineffective at increasing user awareness.[52] Empirical evidence shows that simplifying their language doesn't make people understand them better, improve people's awareness of data practices, or lead people to make different choices.[53] These findings suggest that privacy notices haven't only been consistently ineffective, but they're also likely to continue being ineffective for the foreseeable future.[54]

Researchers at the University of Michigan developed an algorithm, called Polisis, that uses AI to visualize privacy policies.[55] If you go through the representations generated by the algorithm, though, you'll notice they're somewhat unhelpful to understand what's going on with your data. Their limitations illustrate that the real issue is not that you don't read your privacy policies. It's that they're uninformative – even after recruiting the help of AI.

Because no one reads them, people don't choose one product or service over another based on its privacy policy. So corporations have incentives to have privacy policies that are the most beneficial for them and the least beneficial for their users.

And what's most beneficial to a company changes from one to another. This reality leads to a peculiarity. Although privacy policies are unified in their unhelpfulness, they're dissimilar in their content.[56] Reading one or two of them won't provide insight into the content of the others to reduce the no-reading problem.

Privacy policies have a no-reading problem, a comprehension problem, and an indistinction problem. These problems make them uninformative: we learn close to nothing from them. Privacy policies' uninformativeness leads many experts to believe that they don't matter,[57] or that making them the target of regulatory efforts is a red herring.[58]

Choices with No Options

The information economy eviscerates the idea of people having choices over what happens with their information. Beyond the insufficient yet unfulfilled aim to inform people as the paramount means of protection, people are rarely afforded genuine choice to do anything other than agree with them.

Our notice and choice model, inspired by neoclassical contract theory, was conceived fifty years ago, when today's Internet was unimaginable.[59] The model emerged in a context where personal information transfers took place between few and easily identifiable parties, for discrete purposes as part of a business exchange, and in relatively predictable and transparent ways. Data transfers happened, for example, when stores requested customers' phone numbers to inform them of a product's arrival or when banks needed their clients' social security numbers to log their financial information. Back then, it was far easier to know with whom you were interacting, what information they collected about you, how it would be used, and whether it would be shared with anyone.

The information economy, defined by multiparty data exchanges, is fundamentally different.[60] Today, corporate use of personal information includes data sharing, data mining, data trading on the back end, and profiling based on inferred data. Even a simple interaction, like buying shoes at your favorite store, includes the possibility of the other party selling your information to data brokers, who aggregate it with other information about you and sell it.[61] Other parties, such as your bank, are obligated to report your information to credit reporting agencies, whose job is to aggregate information about you to probabilistically infer your trustworthiness as a borrower through your credit score.[62]

The information economy's paradigm shift makes it impossible for people to understand who has what information about them and what purpose they may use it for (let alone how it got there). Third parties collect and use an unprecedented amount of personal information beyond people's knowledge and understanding.[63] In this context, we don't know what we're saying yes to when we tick the "I agree" box.[64] The shift in the collection, use, and sharing of people's data from fifty years ago to today's information economy that makes notices difficult also makes choices impossible.

The rationale behind the choice model is that, in theory, it could allow people to manage their privacy risks.[65] Privacy self-management was thought of as a way to avoid paternalism by making each individual decide what data risks they find acceptable to incur, when, and with whom.[66] In theory, it accounts for the fact that people's privacy preferences may differ and their preferences may vary from one context to another.[67]

Choice assumes knowledge and understanding of risks. It's impossible to make a real choice if you don't know what the choice is and what its consequences can be – what risk you're taking on by agreeing. So the failed informativeness of privacy policies is key to the failure of choice.

Two Forms of Individual Agreement

Despite privacy policies' long history, legal scholars and courts disagree about what kind of legal document they are: notices or contracts.[68] I find this disagreement puzzling. Privacy policies are corporations' main vehicle for informing their users, so they have incentives to clarify what kind of document they are. And, particularly but not exclusively in the US, regulators use privacy policies to oversee corporate data practices, making privacy policies an important mechanism for protecting privacy. So regulators have incentives to clarify it too.

Legally speaking, a notice is a tool to convey to someone else what they can do based on your property rights. For example, "no shirt no service" notifies patrons that a business will exercise its right to refuse service to anyone who doesn't wear a shirt. "Entry beyond this point is trespass" aims to notify that an area is private and anyone who enters it is liable.

A notice can shift liability only when informing someone is relevant for determining liability. For example, a warning label on a product can free a manufacturer from liability if an injury results from an improper use that the label said to avoid. But notices can't expand the preexisting rights of the notice-giver.[69] For example, you can put a sign on your fence informing others that walking beyond the fence is trespass, but you can't decide the punishment for trespassing. Signs indicating that trespassers will be shot don't actually establish homeowners' right to shoot trespassers. A notice can allocate risks, but it can't give or take away rights. To allocate rights, one needs a contract.

Treating a description of data practices as a notice implies that the notifying corporation has the right to do whatever such notice contains – and is simply informing us of what it will do. Treating the document as a privacy contract implies that the corporation lacks the right to do what's in the document unless it obtains the consent of each user.

For many legal scholars and courts, privacy policies are more akin to "privacy contracts" than to "privacy notices."[70] For example, Facebook's Terms of Use include a forum selection clause indicating that any dispute will be resolved by California

courts, which is something only a contract can do.[71] Treating privacy policies as notices made sense when there were no restrictions about what corporations could do with our data. The contracts lens, rather than the notices lens, better reflects that corporations aren't free to do as they please with our data with a mere obligation to let us know. Instead, corporations must obtain agreement.

Privacy policies have further similarities with contracts. Some contracts, called standard form contracts, share a set of problems with privacy policies. They're written in complicated language that people must often agree to without someone to clarify the terms – and often without reading them. Consumers rarely know everything they're agreeing to, and there's no room for negotiation because it's a take-it-or-leave-it offer.[72] However, these are only a subset of the problems that privacy policies have.

The contracts model ultimately also fails to reflect twenty-first-century personal data interactions. In the information economy, privacy policies differ from contracts in that there's no "meeting of the minds": the mutual understanding and agreement on the specifics of an interaction that's essential in contract law.

Even in the most egregious contracts, consumer standard form contracts, there's a meeting of the minds. We may not read standard form contracts, but at least we know what their object is: we know what we're giving up and receiving in exchange. If you purchase a cellphone plan with AT&T, you know you're giving money in exchange for a cellphone service. But in privacy interactions, we don't know what we're giving up.[73] Standard form contracts can be valid, even if some of their terms are invalid, as long as there's a core agreement between parties, such as trading a good or service for a price.[74] This core agreement doesn't extend to data practices. Often, there's not even a trade involved. Standard form contracts must have sufficient notice and a chance to read and understand the terms before agreeing.[75] But privacy policies can even be changed unilaterally.[76] In standard form contracts, we can choose whether to complete a transaction, but many companies that hold our data are entities we never heard of. Treating privacy policies as contracts mistakenly situates them in relationships of mutually chosen trade – a more consequential misconception than believing people read them.

Ultimately, persuading courts to treat privacy policies as standard form contracts doesn't solve the problems posed by the notice-and-choice regime. In Canada, for example, where courts routinely treat privacy policies as consumer contracts, scholars critique its privacy law for characterizing privacy in market terms, thus placing disproportionate importance on business interests.[77] Notices- and contracts-based models equally reinforce the idea of privacy self-management, which mistakenly sees the relationship between corporations and their users as series of bilateral market transactions.

What are people managing when they self-manage their privacy, according to the traditionalist view? The next section addresses this question. The short answer is only the secrets that they want to hide from the entire world.

B THE BINARY BLINDERS

The law was deeply unfair to Pamela Anderson. She and Tommy Lee filed an invasion of privacy lawsuit against *Penthouse* magazine in 1997 for publishing intimate photos of the couple.[78] The photos were stills from a tape that had been stolen from their home and posted online without their knowledge.[79] *Penthouse*, Anderson explains, offered to pay the couple for the photos, but they refused and asked the magazine to destroy them, explaining they didn't want people to see them.[80] The magazine published them anyway, exploiting the couple's intimacy for profit.[81] The judge overseeing the case dismissed it, arguing that because intimate material of the couple had been previously published, they had forfeited their privacy.[82]

Thinking that Anderson lost all privacy over the pictures once someone shared them, and that she lost nothing by the subsequent publications, is a result of the binary blinders. It results from thinking that once someone's personal information is disclosed for the first time, all privacy interests over that information are gone. The binary approach misconstrues privacy's value because it disregards the context in which disclosures occur and that further disclosures generate new harm. Privacy isn't binary as this notion assumes. It sits on a spectrum. People's privacy can decrease by different magnitudes, depending on the informativeness and sensitivity of what other people learn or infer about them. Recognizing this spectrum is more important than ever.

Bracketed into a Binary

The traditionalist approach to the information economy is built on a worldview of bilateral commercial exchanges that leaves out people in situations like Anderson's. This binary worldview results in the notice-and-choice system that privacy laws across the globe incorporate. Under this view, you either have privacy or you don't – just like you either fulfill a contract or you don't, with no in-between.

In the age of algorithms, recidivist privacy invaders permeate daily social interactions in an unprecedented way. A binary conception of privacy may have been adequate (it probably wasn't) in a world of one-time bilateral intrusions. In that simplified world, a person could open only one of your letters (a single intrusion) and publicize its contents (a single disclosure), but it would be unlikely to go beyond that. That same person was unlikely to open and disclose many more of your letters because it would be difficult for them to have the resources to do so. By contrast, in the information economy we're involved in repeated and ongoing interactions with actors that reduce our privacy, from social networks we're too familiar with to data aggregators we never heard of. Getting stuck in the idea that one either "has privacy" over something or one doesn't prevents one from capturing this context.

The story of Holly Jacobs, who founded the Cyber Civil Rights Initiative, illustrates the pitfalls of reducing privacy to that dichotomy.[83] Dr. Jacobs had exchanged intimate

photos with her boyfriend while they were in a long-distance relationship. Eventually, he posted the pictures online, where multiple websites reposted them – most of them deriving ad or subscription profit.[84] Jacobs spent months sending takedown notices and, after monumental efforts, got them scrubbed. But they reappeared on about 300 more websites. The police told her there was nothing they could do.[85] Telling Jacobs that, once her ex-boyfriend posted the pictures, it didn't matter how many websites reposted them, like the court told Anderson, would have been detached from reality.

Courts and policymakers often engage in this type of poor privacy reasoning. In a case against the city of Petersburg, for example, employees were required to answer a questionnaire asking about the criminal histories of their family members, their complete marital history, their children, and their financial status. The court dismissed the claim that their privacy was violated, reasoning that there was no privacy interest in the information because it was already available in other records.[86] The binary view brackets courts like this one to only two possible readings of the world: a person either "lost" their privacy or they didn't. It leads to an unreasonably high bar for harm and makes privacy claims unfathomably difficult to prove in today's context of multiparty data exchanges.

A continuous concept of privacy loss is paramount for understanding the information economy. Recognizing that privacy exists on a spectrum captures intuitions about privacy better than binary views. When a company like Alphabet (Google) gains more knowledge about one of its users, it's false to say that the user no longer has *any* privacy – just as it's false to say that they had perfect privacy before. It's also incorrect to say that nothing happened to their privacy. Privacy loss is about the user's level of privacy dropping from one level to another.[87] Viewing people's privacy as a spectrum better captures the reality that they face regarding their privacy losses.

Determining any rights violation is a binary exercise in one broad sense: in a trial, courts have to rule whether there was a violation or there wasn't. Recognizing degrees of losses, however, is essential to identifying those privacy violations correctly. Likewise, when estimating "reasonable care," courts consider degrees of care and apply a cut-off. The estimation mistake is overlooking that privacy losses, like levels of care, exist in degrees.

Accounting for nuance in privacy through a spectrum of losses and gains is key because privacy violations that get to court involve grey areas: they involve different gradations of privacy losses.[88] Rejecting binary perspectives in favor of an understanding that privacy losses exist in a continuum is necessary for developing sensible laws for the information economy.

Counting Only Secrets

In its worst form, privacy viewed through the binary blinders is reduced to secrets. Under the secrecy view, once you reveal information to someone in any way that makes it possible for others to see or know it, you abandoned all privacy over it.

The shortcomings of the secrecy view are clear in the painfully frequent scenario of nonconsensual distribution of intimate material, such as Dr. Jacobs' story. It's eerily common to hear that, because a victim shared the material with someone, it wasn't a secret anymore and the victim assumed the risk of its distribution.[89] This misconception sometimes extends to courts.[90] For example, in 2015, a woman called Dana sent intimate pictures to an ex-boyfriend, and someone else who saw them on his computer plastered them over a public Facebook page. The Vermont Supreme Court said that Dana chose to abandon her privacy over the pictures when she sent them to someone with whom she wasn't in a relationship.[91] This type of dismissal occurs even when courts rule in the victim's favor, but still frame the victim's harm as a cautionary tale about their excessive or irresponsible risk-taking.[92]

Positing that people abandon privacy expectations over information whenever they share it with one person, as the secrecy view does, is mistaken. Dana retained some privacy over the images when she shared them with one person and lost significant privacy when they were shared with the world. So did Anderson. This view is worse than victim-blaming.[93] It also implies that the victim didn't have her rights breached at all – that she wasn't even a victim.[94]

The dynamic at play in these cases follows us into our daily lives. Their dynamic is replicated when corporations acquire massive amounts of information that are deemed public, such as taking pictures of us on the street or gathering our online profile photos to train facial recognition software that can identify us.[95] As a result, online interactions are plagued with surveillance, harassment, and risks of violence.[96]

Secrecy is a uniquely problematic aspect of the traditionalist approach because it further narrows privacy protections from privacy self-management's "let people make choices about their privacy" into one specific choice: hiding information about oneself from others. The flawed secrecy conception permeates the notice-and-choice principle, indicating that people chose to abandon privacy over information when they chose to disclose it. This fundamental error illustrates why notice and choice fails as a privacy framework and, worse, leads to people bearing the risks of corporate data practices.

From a secrecy perspective, privacy also disappears when a person moves from private to public spaces.[97] Those notions of public information and public spaces that nullify privacy are often defined too broadly or not defined at all.[98] Secrecy leads to the belief that, as Scott Skinner-Thomson puts it, "the right to privacy while in public is nearly nonexistent, that privacy is more or less 'dead' once you walk out of your front door."[99]

Maintaining a privacy claim under the secrecy paradigm means having to keep information to oneself.[100] However, keeping any digital record in absolute secrecy in the information economy is beyond impractical; it's impossible.[101] By requiring people to do so, this view of privacy inordinately disadvantages the disadvantaged: those without property, those without a home who need to use public spaces, and those who belong to communities that are disproportionately surveilled.[102]

Secrecy-based views of privacy require a binary conception because they zealously abandon privacy expectations and protections when information is revealed. But not all binary views of privacy are secrecy-based. One could (misguidedly) believe that a person either has complete control or absolute lack of control over their information – failing to capture that one usually controls some aspects of it but not others. Binary views, whether they're about secrecy or control, mistakenly pose that when we share something in one context we lose our privacy over it in all contexts.[103]

By inferring preferences solely from behavior (someone revealed information to a platform so they must not care about privacy), the traditionalist narrative weaponizes the binary blinders into deregulatory efforts. The most common consequence is the argument that, if you have nothing to hide, you have nothing to lose.[104]

"You Have Nothing to Hide"

Former Google chief executive officer (CEO) Eric Schmidt once famously said that "if you have something that you don't want anyone to know, maybe you shouldn't be doing it in the first place."[105] A strong statement for a data profiteer.

Schmidt's infamous statement is an example of the most widespread consequence of the secrecy conception: the "you have nothing to hide" argument, which myopically equates privacy with hiding terrible secrets.[106] The statement illustrates how viewing privacy as secrecy leads to a mistaken understanding of choices, even when making choices is possible. Often, the argument is used to present policymakers with a false all-or-nothing choice between privacy and another social value, such as national security or public health.[107]

The idea that only people with "something to hide" care about privacy is the most pervasive argument against privacy that one can find.[108] Anyone who has conversations about privacy has heard someone else indicate that if they have nothing (bad) to hide, they have nothing to lose. The argument gets repeated by industry members, regulators, and community members.[109] With the nothing to hide argument, the secrecy view reduces privacy to something merely instrumental. In this view, privacy exists solely for trickery.

Variations of this argument appear regularly in statements by politicians and government entities. A defense of the British public surveillance system by its Conservative Party was "if you've got nothing to hide, you've got nothing to fear" – a quote originally from Joseph Goebbels.[110] In the US, the argument harkens back to McCarthyism, when it was used to pressure witnesses to confess to endorsing communism.[111] Still today, the argument appears during legislative debates amid claims that only "criminals" should be concerned with their privacy.[112]

People accused of crimes aren't the only ones who find themselves on the wrong end of the nothing to hide argument. Everyone in the information economy does.[113] People are constantly surveilled in their digital lives to show them more accurate

ads.[114] The data harvested through this private surveillance system are the source of tremendous profits.[115] Given these economic incentives, it's unsurprising that the nothing to hide argument found its way into the information economy – and Google's CEO.

An influential version of the nothing to hide argument comes from Judge Richard Posner. Based on neoclassical economics, Judge Posner argues that people desire privacy because they "want more power to conceal information about themselves that others might use to their disadvantage."[116] In this view, privacy is the "right to conceal discreditable facts" about oneself.[117]

Judge Posner's argument begins with the premise that there are people with bad traits and people with good traits, and that people with bad traits want to hide them while people with good traits want to show them. Privacy, his argument proceeds, allows the bad types to hide their bad traits by reducing the information available in the market, making themselves indistinguishable from the good types.[118] From this market-based perspective, privacy creates an information asymmetry. This asymmetry, according to the argument, advantages bad types because others are more likely to engage with them if they can't see their bad traits. And it disadvantages the people examining information to choose whom to engage with (the "information receivers"). The notice-and-choice system, allowing people to waive privacy whenever they're asked, is its logical consequence because it doesn't see anything worth protecting beyond deceit.

The nothing to hide argument mischaracterizes what privacy is about. Resting on the secrecy conception, it views privacy as chicanery: it assumes the only valid reason to hold personal information private is to strategically deceive others. It's not.

Judge Posner presented the argument in the context of employment, which was an appropriate choice. Employers are part of the information economy when they profit, directly or indirectly, from their employees' personal data and not only from their work.[119] Uber, for example, gathers valuable information from its drivers when they don't have passengers.[120] Employment analogies, though, are more broadly illustrative of dynamics in the information economy. In both employment interactions and the information economy, corporations hold significant power to make decisions that impact people's lives based on their information.[121] It's clear that providing employees with a take-it-or-leave-it option under a power imbalance deprives them of real choices. Employees, like people in the information economy, have something to lose.[122]

When You Have Something to Lose

In 2000, economists Claudia Goldin and Cecilia Rouse developed a test for gender-biased hiring in symphony orchestra auditions.[123] They noticed a gap between the proportion of women in elite music schools and elite orchestras, so they held auditions behind a curtain, preventing those hosting auditions from knowing the

auditionee's gender. Female orchestra hires immediately increased by one-third.[124] The curtains increased the probability of each woman passing the initial round by 50 percent.[125] Imagine the value of this privacy protection for those women, who had absolutely nothing to hide.

Privacy protections often help prevent discrimination. The logic behind these protections is that decision-makers can't discriminate if they lack the information needed to do it.[126] With or without formal protections, people rely on privacy to avoid discrimination in interactions with their employers, landlords, healthcare providers, schools, and banks.[127] This protection is crucial in the information economy, where AI algorithms make decisions about us based on our data.[128]

This privacy protection wasn't afforded to Carter Brown, who was fired from his job in Texas in 2018 after a coworker outed him to management as trans.[129] Carter Brown didn't have anything nefarious to hide but, when his employer learned information that he would have liked to keep private, he lost his livelihood.[130] This protection wasn't afforded either to April Cox, who was refused by human resources giant Randstad when a drug test revealed her medicinal use of methadone to recover from a former addiction.[131] Brown and Cox aren't alone in their efforts to keep legitimate aspects of their life private.

The value of privacy protection extends to interactions without discriminatory intent.[132] To continue with the employment analogy, consider drug testing in the workplace, which illustrates dynamics of power and bundled information that extend to the information economy. Many workplaces employ random drug testing during employment relationships or as a condition for hiring.[133] Presumably, employers do it because they believe that it provides them with valuable information about their employees. Some employers might test if they don't care about recreational drug use itself but treat it as a proxy for something else. For example, they may mistakenly believe that people who use recreational drugs may experience more frequent health issues leading to higher absenteeism, that they're more likely to disobey other rules, or that they tend to be less conscientious or hard-working.[134]

General drug tests, which check for various substances, reveal other information besides what employers try to gauge based on drug use. They can reveal information irrelevant to employers about someone's health status by detecting prescribed drugs. They can reveal information about what someone does during weekends that's irrelevant to their productivity on weekdays. They can reveal methadone use, uncovering that someone is recovered or recovering from a former addiction, as was the case for April Cox. Or, if additional tests are run on the sample, they can reveal sex assigned at birth, the information that harmed Carter Brown.[135] What an employee seeks to keep private may be unrelated to what their employer wants or deserves to know.

One should consider privacy concerns over information attached to drug tests as evidence of harm that flows from revealing bundled information. This privacy harm can lead someone who doesn't use recreational drugs to avoid testing. Many who are

harmed by bundled information from mandatory drug testing (who have something to lose) aren't recreational drug users. They're harmed by the other information revealed by the test.

Because drug use data are bundled with other data, one can't infer drug use from test refusal. Contrary to the neoclassical economics of privacy and the nothing to hide argument, people like Brown and Cox may refuse the test not because of how nefarious what they're not revealing is, but despite it. Many who would have passed a test will either decline it, foregoing a job that they would be qualified for, or accept it and bear associated harm. Because employers can't know how much each employee is harmed by the bundled information, they can't distinguish between those who refuse because they use illegal drugs at work (those with something to hide) and those who refuse because it would reveal bundled information (those with nothing to hide, but something to lose).[136]

Bundled information justifies protecting people with nothing to hide from mandatory testing if there are other ways – beyond testing – for them to convey productivity information. These methods could include providing a list of contacts for recommendations or establishing a trial period. Allowing people to choose among means to convey information improves social welfare because it reduces privacy harm.[137] In the US, for example, three states allow for voluntary compliance with workplace drug testing.[138] Other jurisdictions should consider doing the same. In the meantime, employers in jurisdictions that don't would receive better employees by breaking away from the secrecy conception and understanding that resistance to testing doesn't mean that employees have anything to hide – they may just have something to lose.

Drug testing illustrates how the nothing to hide argument relies on the secrecy conception. From the perspective of the receiver, the information is binary: either an employee passes the test and gets the positive signal that their employer attaches to it, or they don't. Because employers believe that the test is informative about productivity (otherwise they wouldn't require it), they're likely to use it for promotion and retention decisions. They're wrong to believe that, because learning about productivity is legitimate, employees can't suffer privacy harm from how they learn about it.

Requiring people to disclose bundled personal information to obtain a benefit imposes social costs beyond employment.[139] Because information is bundled, people in the information economy often have to reveal irrelevant data to convey relevant ones, like when you need to provide your phone number to make an online purchase, hotels scan your passport to verify your identity at check-in, or you're recorded at a store for security purposes. The risks from those forced disclosures represent a social loss together with an individual one, as the consequences of surveillance constitute a social harm beyond individual harm. For example, people behave differently when they know they're under observation, regardless of whether they're trying to hide a wrongful activity.[140] Allowing people to keep bundles of information private, rather than forcing them to reveal them to access products and services, is beneficial to them and to society.[141]

https://doi.org/10.1017/9781108995825.002 Published online by Cambridge University Press

Considering that people value their privacy for its own sake, and not just to deceive others, shows that the secrecy-based view of privacy is inadequate. Personal information is bundled by nature because disclosing anything to someone often implies disclosing other things too. Sometimes, we want to keep information private not because of the information itself, but because of everything that's bundled with it. For orchestra auditioners, their physical appearance was bundled with their gender. For Brown, his past appearance was bundled with his membership in the Queer community. For Cox, her nonuse of recreational drugs was bundled with former use prior to recovery. For everyone in the information economy, information about their online activity is bundled with information about their characteristics, behaviors, and preferences.

The neoclassical economics conception of privacy is one-dimensional. In this one-dimensional conception, the only reason someone would value their privacy is to hide something that should be relevant to someone else so they can deceive. This view has erroneous microfoundations, meant for bilateral commercial transactions in which, if I know less about you or can't speak about you, you're competitively advantaged toward me.

The nothing to hide argument makes it seem that privacy is about hiding nefarious secrets from others. The value of privacy, however, is social.[142] Recognizing privacy's intrinsic value means moving to a nuanced worldview where we recognize someone might want to keep information from others that is or should be irrelevant to them.

* * *

In 1987, President Ronald Regan nominated Robert Bork, a fierce opposer of privacy rights on traditionalist grounds, to the US Supreme Court. Unbeknownst to Bork, during the debate over his nomination, a reporter walked into his video rental store, asked for, obtained, and published Bork's entire videotape rental history – something that, in 1987, was quite informative about what one watched.[143] His viewing history turned out to be unremarkable and his nomination unsuccessful. But the process changed American law for decades because it led members of Congress to write and pass the Video Privacy Protection Act, forbidding video stores from disclosing rental histories, at record speed.[144] Possibly worried about their own viewing histories, members of Congress were similarly situated to millions of people in the information economy. Their choices about what video store to rent from alone didn't protect them, they would lose more privacy from having their whole rental history revealed than from the parts they revealed to others, and they most likely had nothing to hide in their perfectly legal rentals, but had something to lose.

The notice-and-choice principle and the privacy self-management system that it underpins share mistaken assumptions about privacy interactions. First, they assume privacy interactions exist in a context of trade, where parties to a contract have the opportunity to notify each other and make choices. Second, they assume privacy is

binary, as shown through the popular nothing to hide argument and the widespread notion of privacy as secrecy. Through the binary blinders, they pigeonhole people as hermitic or impassive. This privacy paradigm was designed for a simple commercial context that no longer aligns with reality. While the paradigm's concepts about privacy interactions may have been relevant at the time, today they're stretched into contexts where they no longer make sense.

The privacy fallacy, operating in this paradigm, reduces privacy to an instrumental dimension. It creates a blind spot. Just because people occasionally seek something instrumentally, such as privacy, that doesn't exclude that people also value it for its own sake.[145] This reduction makes secrecy-based arguments such as nothing to hide mistaken in their own terms. Privacy laws fail when they silo social effects into instrumental individual choices. The instrumentalist view overlooks distributional aspects that inform privacy's social value: people with fewer resources who lack power to say no and data from one person that conveys information about others.[146] Decision-makers who fall into the privacy fallacy fail to capture real privacy interactions. They leave out negative effects on oneself and others.[147] Most privacy protections, in theory and in practice, don't protect bad people's dark secrets.

There's a reason why serious cases of privacy invasions abound. The traditionalist approach is built on the idea that people make rational and informed choices about their privacy when they're given notice, similar to how they decide to buy apples, a shirt, or an apartment. In the next chapter, I call this the "myth of rationality."[148] The approach is also built on the idea that if someone "chooses" to give information to someone, that means they don't care about keeping that information private. In the next chapter, I call this the "myth of apathy."[149]

The first idea, that people make rational and informed choices about their privacy, is the bedrock of privacy self-management. It's the same foundation on which contract law rests. But there are good reasons to qualify and depart from it for the information economy. The second idea, that if someone "chooses" to give information away they don't care about keeping that information private, is specific to privacy. But anyone interested in privacy enough to open this book faces the idea routinely. You face these supposed choices every time you open a website in a rush and click "I agree to cookies" without looking for the hidden "read without agreeing to cookies," or open a website in Incognito mode not knowing that you're still giving your browsing information to the website, the browser, and your Internet service provider.