1

Why Design Is Everything

YOU ARE IN TOTAL CONTROL. At least, that's the story you've been told. You get to choose what you share online, whom you share it with, and how it is shared. Tech companies are bending over backward to give you more perceived control over your personal information. It is no wonder the common narrative is that victims of privacy breaches only have themselves to blame. But the truth is this: all the focus on control distracts you from what really affects your privacy in the modern age. The most important decisions regarding your privacy were made long before you picked up your phone or walked out of your house. It is all in the design.

Design affects how something is perceived, functions, and is used. Technologies are great examples of the power of design. Consider this marketing photo depicting the user interface for the media service Snapchat:

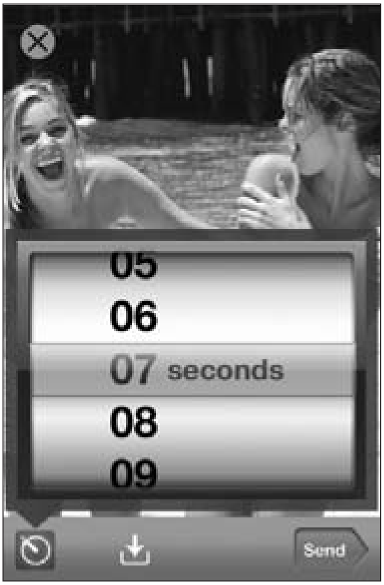

The selectable "seconds" connected to a clock icon, the Send button, and the photo carefully cropped to show only nude skin all suggest that this software is designed to allow users to temporarily share intimate or suggestive photos.[1] All of this is conveyed without an explicit promise or disclaimer—just design. New users would be justified in thinking that naughty photos, or "Snaps," would completely disappear.

Except that's not how it works. The Snaps simply become invisible to recipients. Copies of the photo still exist. Most modern phones are designed to let users take screenshots—a tactic that is regularly used to "capture" Snaps. Data forensics experts are able to recover copies of photos still lingering in storage. There is even third-party software that allows users to save Snaps before they disappear. That's how nineteen-year-old student Zeeshan Aqsar saved a nude photo sent by a fifteen-year-old schoolgirl via Snapchat, which he then used to blackmail her for more photos and money.[2] Like many Snapchat users, the young girl thought the photo she sent would disappear. As previously mentioned, Snapchat initially failed to ensure that only its own software client could access its application programming interface. Design could have been leveraged to better shape users' expectations and make saving photos more difficult. For a while, it was not.

Even when design is not shaping our perceptions of a technology, it is in the background, shaping what happens to us. Users usually cannot tell what kinds of personal information the websites and apps they visit are collecting. Every website and mobile application collects some kind of arguably personal data, such as Internet protocol (IP) addresses, browser type, and even browsing activities. Did you know that when you click on a link, your browser tells the website that you just loaded where you came from? For example, if you are reading a blog about sexual fetishes and click on a link to a book on Amazon, the online merchant would know which salacious website you were reading before you clicked that link. This information is contained in what is known as the hypertext transfer protocol (HTTP) referrer header. We are awash in invisible data streams. Modern information technologies are specifically designed to collect more, more, more.

In this chapter I'll show how design affects your privacy. I'll make three simple points. First, privacy-relevant design is everywhere. It's part of every action we take using our phones, laptops, and tablets. It's also a force upon us as we interact in the physical world. The best way to spot privacy-relevant design is to look for the ways in which technologies *signal* information about their function or operation or how technologies make tasks easier or harder though *transaction costs*. Signals and transaction costs shape our mental models of how technologies work and form the constraints that guide our behavior in particular ways.[3]

Second, I'll show that design is power. Every design decision makes a certain reality more or less likely. Designers therefore wield a certain amount of control over others. Because people react to signals and constraints in predictable ways, the design of consumer technologies can manipulate its users into making certain decisions. Design affects our perceptions of relationships and risk. It also affects our behavior: when design makes things easy, we are more likely to act; when design makes things hard, we are more likely to give up and look elsewhere. The power of design makes it dangerous to ignore.

I will conclude this chapter by setting up a conflict at the heart of this book: the misconception that design is neutral. A popular argument in technology law circles is that the creation of technologies that are used in harmful ways is less deserving of legal scrutiny than the act of collecting, using, or sharing another's personal information. Those who see great promise in "big data" often argue in favor of rules regulating use of data rather than limitations on the ability to collect personal information.[4] Critics of legal action against companies for having poor data security liken it to the government punishing victims of burglary.[5] Advocates of immunity for services that host harmful content point to the fact that they are just the messenger.[6] By calling attention to the most immediate or proximate source of privacy harm, companies downplay how design affects our privacy. Those seeking to avoid scrutiny for the things they build often critique calls to regulate design by arguing that design is neutral.[7] They're wrong. Design is never neutral. It is political. And it should be a key part of our information policy.

Design Is Everywhere

This book is concerned with the design of two particular kinds of consumer technologies: those used by people and those used directly upon them. Most of the technologies we use *mediate* our experiences. Browsers, mobile apps, social media, messaging software, and the like all act as media through which we find and consume information and communicate with others. Meanwhile, surveillance technologies like license plate readers, drones, and facial recognition software are used upon us. We don't interact with these kinds of technologies, yet they have a profound effect on our privacy.

You can see the impact of design on our privacy almost everywhere you look. In the physical world, doors, walls, closed-circuit television, modesty panels for lecture hall tables, and any countless number of other design features shape our notions of privacy. Structural protections, such as our homes, create "reasonable expectations of privacy," which is the law's touchstone metric. Structure can also erode our privacy when it facilitates surveillance and information misuse. There is a reason bathroom stalls in the workplace are not made of transparent glass, yet many conference room walls are. All the better to see you with, my dear. And any book on privacy and design must make at least a passing, obligatory reference to the most famous of privacy-related structures: the panopticon.

Philosopher and social theorist Jeremy Bentham designed the panopticon as a prison comprising a circular structure with a guard tower at its center, thus allowing a single guard (or a small group of guards) to observe all the prisoners at any given time.[8] The design obscures the view of the guard tower from the cells, and prisoners cannot know when they are being watched or whether a guard is even present; as a result, they are compelled to act as though they are always being watched. This design has become a metaphor, pioneered by philosopher Michel Foucault, for the modern surveillance state and the cause of surveillance paranoia.[9]

Design is just as critical to our privacy in online, mediated environments. For example, app developer GoldenShores Technologies designed the Brightest Flashlight app to collect a mobile user's geolocation data.[10] There is no particular reason a flashlight app needs geolocation data to function. It's a flashlight app. It just needs access to a flash. Yet because "data is the new oil" the company couldn't resist the opportunity made available by the architecture and capabilities presented by the little surveillance devices we all keep in our pockets. GoldenShores probably designed the app to collect our location data because it could and because it was financially advantageous for it to do so. The Federal Trade Commission (FTC) alleged that GoldenShores' design did not comply with its privacy promises and that the company was engaged in an unfair and deceptive trade practice. In Chapter 4, we will revisit this complaint as an example of how lawmakers might better consider privacy-related design.

Compare GoldenShores' design decision with that of tech media outlet Gizmodo, which configured its services to not store the IP addresses of those who visit its website. In explaining the company's decision, Annalee Newitz writes, "if we received [a] subpoena tomorrow, demanding that we hand over personally-identifying information about you because you went into comments and spilled some corporate secrets about Apple—well, we couldn't. Our technical systems team couldn't. We don't log that data, period. And we chose to set up our commenting system that way because we want to protect your right to speak freely and anonymously."[11] Gizmodo's decision demonstrates how design can advance privacy-related values like obscurity and anonymity. These values are furthered by making certain things easier or more difficult to accomplish—or, in the case of Gizmodo, impossible. Of course, engineering plausible deniability cuts both ways, as it can protect wrongdoing as well as provide a space for free expression and human flourishing. In any event, the design here is consequential and value laden.

Apple has designed its encryption system for mobile devices to similarly protect the information on its phones. Because of its design choices, Apple cannot disclose the information of those using its encrypted devices.[12] The decision reflects a commitment to secure, private communications and technology use. That commitment can also be seen in Apple's decision in early 2016 to resist the FBI's request for a customized operating system that would cripple a security failsafe keeping the government from accessing the phone of a terrorism suspect in the shootings in San Bernardino, California.[13] This dispute also demonstrates the very real cost of protective design, which can frustrate law enforcement and intelligence agencies in their duties, reduce cybersecurity in some instances, and facilitate harassment and abuse. For example, law enforcement might be unable to obtain information important in solving a crime. Intelligence agencies might miss important information. Yet these designs protect people's most personal communications and personal information from abuse of process by government and direct attack from hackers. They provide the freedom and autonomy necessary for human development, commerce, infrastructure management, and our own national security.

Because we use digital technologies every day, we tend not to think about the role of design. It's just always there. To understand why design is so important, let's first explore the function of design.

What Does Design Do?

Broadly speaking, particular designs of technologies *communicate* information and *enable* (or *hinder*) activities. Often, design does both, as when labels on online clickboxes and drop-down menu options channel user choice. Password requirements are great examples. Consider a company's practice of securing sensitive electronic documents like paystubs and tax forms with a password. When employers protect PDFs with a password, they simultaneously *hinder* access by ensuring only those who were given the password can view the file, and they also communicate to recipients that the contents of the file are not for general consumption. While the exact nature of what is being communicated is open to interpretation, the design choice of an authentication requirement sends a signal that this information is not for just anyone. Design can act as a medium, communicating *on behalf of* both designers and users. It can also act *upon* users, constraining or enabling them in particular ways. By focusing on signals and barriers, we can see how design affects people and begin to understand the values at stake.

Design Provides Signals

Imagine that you just rented a laptop from a consumer electronics store while yours is in the shop. You boot up your rental and see this pop-up window[14]:

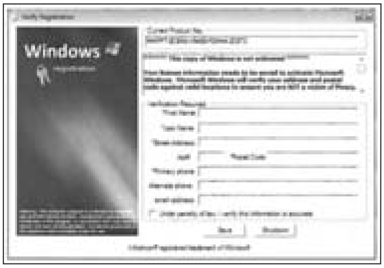

You recognize these kinds of pop-up windows because they are common with new pieces of software. (This is not your first computer, after all.) The window has the famous Microsoft logo and some keys that represent security and access. You also see what would appear to be the product key, which is a signal that this software is legitimate. The prompt asks routine questions about personal information. There's even an official-looking requirement for you to check that "Under penalty of law, I verify this information is accurate." Serious stuff. Every aspect of the design of this pop-up window communicates to you that Microsoft has created this interface to facilitate activating the software on your laptop.

The problem is that this window is a lie. It was not created by Microsoft but by a software company called DesignerWare. It is part of a pernicious piece of spyware called PC Rental Agent to be used by companies that rent laptops.[15] The software gives companies the ability to track laptops and disable them if renters are late in making their payments. This particular window is part of the software's "detective mode," which can log user keystrokes and hijack the laptop's webcam to record anyone in view.

The pop-up window is a false front. The registration window registers nothing. Instead, it captures all of the personal information entered into the prompt boxes and transmits it back to DesignerWare and to the store renting the laptop. This is no hypothetical. In one of the few instances of privacy law taking design seriously, the FTC alleged this software to be a deceptive practice based on the representations made via the fake registration screen.[16]

I give this example to demonstrate that one of the principal functions of design is to *communicate*. It provides signals to people. The communication can be a message from the designer, information from or about other users, information about how the technology functions, what a user's options are, or simply an aesthetic choice. In the case of the fake registration page, the communication is evident. The designers used text and graphic design to signal authority, security, and protection to manipulate users into sharing their personal information.

Sometimes design communicates in subtler or even subconscious ways. For example, Apple made the decision that the background color for free text messages sent between two iPhones would be a soothing blue, whereas text messages between an iPhone and another type of phone that counted against your text limit would be a harsher green.[17] This difference in color communicated several very subtle messages, including when a user would be charged for texts (communicating with a competitor's phone). The distinction might also mean that texting with people who use a competitor's phone is slightly more irritating than people who have calmer, blue buttons. The green color also communicates that the recipient's phone is not an iPhone.

Through signals, design helps define our relationships and our risk calculus when dealing with others. Design affects our expectations about how things work and the context within which we are acting. Consider the padlock icons that are ubiquitous online. We seem them in our browsers, on our phones, and on social media. They're everywhere.

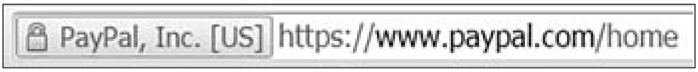
PayPal, Inc. [US] https://www.paypal.com/home

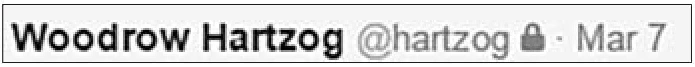
Woodrow Hartzog @hartzog · Mar 7

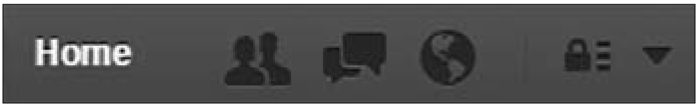

The padlock icon is one of the most well-known privacy-related design signals, and it is used whenever a company or designer presumably wants to convey some sense of security or protection—like its real-life counterpart, which limits access to rooms and boxes to those with the key. In the case of my Internet browser, the padlock icon lets me know when there's a certificate in place to provide encrypted HTTPS communications versus standard HTTP (the *S* stands for "secure").[18] HTTPS protects against eavesdropping and "man-in-the-middle attacks," whereby attackers secretly insert themselves between actors who think they are communicating directly to each other. The padlock icon on Twitter represents a "protected" account, which requires users to approve followers before accessing the profile. The padlock on Facebook is the button that directs the user to Facebook's privacy settings.

Each padlock icon invites reliance upon a technology or service. It signals safety. In the case of Twitter and Facebook, the locks can affect peoples' expectations regarding their relationship with a company. Technology users rely upon signals such as this to shape our expectations regarding what companies are going to do with our personal information. As I'll discuss later in this chapter, when signals like padlock icons and privacy settings cause us to perceive a low risk of harm for disclosing information, we are more likely to disclose more because we are comfortable with the perceived odds that we will be negatively affected by sharing.

Design Affects Transaction Costs

When design hinders or facilitates action, it affects the transaction cost of a certain activity. In economic theory, transaction costs refer to a range of expenses required for participating in market exchanges.[19] But the concept can be expanded to cover the expense required to do anything. For example, time and effort are valuable resources. We often evaluate how desirable possibilities are by calculating how much time they will take or how much effort we'll need to spend to accomplish them.

Consider the practice of putting bars on the window of a house. Bars do not make it impossible to burglarize a home. A burglar with sufficient motivation could get a steel file or even more high-powered machinery to remove the bars. But this usually doesn't happen, because the cost of doing so is too high. Filing the bars down would take far too long and machines that can cut steel are too loud, visible, and expensive. Bars don't provide perfect security. They are just good enough for most houses and stores. Online examples are usually not this dramatic, of course. Instead, transaction costs almost imperceptibly guide choice and expectations online.

Transaction costs play a key role in the design of digital technologies. User interface design and user experience design focus on ease of use. For example, in Facebook's early days, there was no central news feed that aggregated all of your friends' activities in one location. Facebook users had to visit the profile of each friend if they wanted to see their posts and activity. The news feed made it easier for people to see what their friends were up to because they didn't have to go through the exercise of first thinking of a friend they wanted to look up, entering that friend's name into the search bar, and navigating to their profile to find them, then repeating the process for each friend. Design made this task *easier.*

In fact, the entire data economy is founded on design that makes tasks easier. Digital data itself is the result of design that makes the recall of information easier because it is preserved in a persistent, searchable state at marginal cost. Because of databases and communication technologies, the sum of humankind's knowledge is available within seconds to anyone with an Internet connection and the proper log-in credentials. This has largely replaced having to travel to libraries and other research areas around the world—that is, assuming you even knew what you were looking for. Talk about making things easier!

Design can also make tasks *more difficult.* For example, Amnesty International helped design a digital "mutant" font for Internet users who wanted to make sure their writing was read by humans only rather than computer bots. Many people using the Internet would prefer to be invisible to online trackers, which seek out and process text through the shape of characters or the source of a font. To accommodate those people, Mutant Font's design includes "small graphic interventions that will block machines from viewing its shapes. It comprises seven different fonts that generate thousands of codes to confuse tracking. Its algorithm is also shuffled every 24 hours to impede automatic scanning."[20]

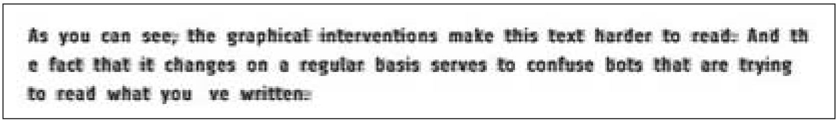

As you can see, the graphical interventions make this text harder to read. And th
e fact that it changes on a regular basis serves to confuse bots that are trying
to read what you ve written.

The practical effect of the "mutant font" design is to make the task of widespread processing of data more difficult. Those who seek to track others must dedicate human resources and time, which is more difficult than deploying software bots that can track others without fatigue and at low cost. Sometimes design can make certain tasks practically impossible. For example, the strength of encryption is measured by how long it would take a theoretical attacker to guess all the possible key combinations or passwords until the right one is found (known as a "brute force" attack). At one point experts estimated that it would take a supercomputer around one *billion* billion years (that's a billion billions!) to crack the 128-bit Advanced Encryption Standard key.[21]

Even very slight costs can discourage certain behavior. Facebook again serves as a good example here. The social media service has a reply feature for its private messages, but no forwarding feature.[22] While you could cut and paste your private conversation with one of your friends into a separate message box for other users, that takes more time and effort than simply pressing a Forward button and typing in a name. It is a transaction cost. While the cost might be slight, it adds up over time and works as a nudge against sharing. It also shows how design practically protects the obscurity of these conversations.

Design is capable of other things, of course. It creates aesthetic value, for example. Do you remember how ugly everyone's home pages were in the early days of the World Wide Web? Bad design. But design's most important feature with respect to information and privacy is that it provides signals to people and affects transaction costs.

How to Identify and Think About Design

Few people have clearly articulated the importance of design in our day-to-day lives better than psychologist Don Norman. In his lauded book *The Design of Everyday Things* Norman explains why design is so important for the things we use. He also shows how to design objects in a user-friendly way. "Well-designed objects are easy to interpret and understand," he notes. "They contain visible clues to their operation. Poorly designed objects can be difficult and frustrating to use. They provide no clues—or sometimes false cues. They trap the user and thwart the normal process of interpretation and understanding."[23]

Norman bases his theory of design on the way humans process visual and tactile clues. He theorizes that the fundamental principles of designing for people are to "(1) provide a good conceptual model [allowing users to mentally simulate an object's operation] and (2) make things visible."[24] He argues that the keys to these two principles were in the visible structure of objects, specifically their *affordances, constraints,* and *mappings.*

According to Norman, *affordances* are "the perceived and actual property of the thing, primarily those fundamental properties that determine just how the thing could possibly be used." Affordances "provide strong clues to the operations of things. Plates [on doors] are for pushing. Knobs are for turning. Slots are for inserting things into. Balls are for throwing or bouncing. When affordances are taken advantage of, the user knows what to do just by looking: no picture, label, or instruction is required."[25] The concept of affordances, which was pioneered by psychologist James Gibson, is quite useful in understanding technological design as well.[26] On web pages, boxes are to be checked, down arrows are to be clicked to see more, and the cursor shifting to a pointy finger indicates a mouse click or hyperlink.

Mapping refers to "the relationship between two things, in this case between the controls and their movements and the results in the world."[27] Norman's example is that when steering wheels are turned clockwise, cars turn right. Mapping helps drivers understand the relationship between the control (the steering wheel) and the result (turning). It is easy to learn how to steer a car because the wheel is visible, a clockwise motion is closely related to the desired outcome of a right turn, and the motion provides immediate feedback. Mapping is also crucial for information technologies.

Cursors on the screen mimic the movements of your mouse or touchpad. Knobs, switches, and icons all rely upon mapping to tell us how to use an interface.

Constraints are the final key to good user design. Norman writes, "The physical properties of objects constrain possible operations: the order in which parts can go together and the ways in which an object can be moved, picked up, or otherwise manipulated. Each object has physical features—projections, depressions, screw threads, appendages—that limit its relationships to other objects, operations that can be performed with it, what can be attached to it, and so on." Social norms—what Norman calls "cultural conventions"—act as a constraint as well; he observes, "Because of these natural and artificial constraints, the number of alternatives for any particular situation is reduced, as are the amount and specificity of knowledge required within human memory."[28] Encryption and password prompts, for example, are technological design constraints that prevent third parties from accessing information. Affordances and constraints can work together in design to clearly guide users to the proper course of action, even when they are encountering an object for the first time.

Norman articulates seven principles for using design to make tasks simpler and more intuitive:

1. Use both knowledge in the world and in your head (in other words, be thoughtful and do research).
2. Simplify the structure of tasks (for example, require users to take fewer steps).
3. Make things visible (i.e., make an object's use obvious from its visual elements, and make important design elements obvious).
4. Get the mappings right (i.e., make tasks performed on or with the object correspond intuitively with their results in the world).
5. Exploit the power of constraints, both natural and artificial.
6. Design for (human) error.
7. When all else fails, standardize (i.e., use universally recognizable signals).[29]

This insight is not just useful for designers. It can help everyone understand when and why design has failed us or is exploiting us.

We're all familiar with at least one design feature that makes human error more likely: the Reply All button, which can cause us to send an email to a number of people rather than just the one we intend it for. The Reply All and other accidental email disasters have claimed more than their fair share of victims. The most memorable such incident in my mind is the unfortunate young law student who meant to send an email update to his best friend about his summer internship at a prestigious law firm: "I'm busy doing jack shit. Went to a nice 2hr sushi lunch today at Sushi Zen. Nice place. Spent the rest of the day typing e-mails and bullshitting with people. . . . I should really peruse these materials and not be a fuckup. . . . So yeah, Corporate Love hasn't worn off yet. . . . But just give me time."[30] Unfortunately, the summer associate had accidentally emailed the firm's entire underwriting group. Bad times. There are hundreds of thousands of similar Reply All and listserv disaster stories. Yet they still keep happening. While we can always be more careful online, errors of this frequency point to a design flaw. Companies are failing to design software to help avoid obvious human error. Sometimes companies learn from their mistakes. Thankfully, relatively new features like pop-up warnings and better layout have been introduced to mitigate the scourge of the Reply All button.

We can analyze privacy design failures the same way. In 2012, computer scientists from Columbia University conducted a usability study of the privacy settings on Facebook. The results were disturbing: they found that "overwhelmingly, privacy settings do not match [users'] sharing intentions." Social media users are not sharing or hiding information the way they think there are. What's worse, "a majority of participants indicated that they could not or would not fix the problems. The prevalence of such errors—every participant had at least one incorrect setting—suggests the current approach to privacy controls is deeply and fundamentally flawed and cannot be fixed. A completely different approach is needed."[31] At worst, this version of Facebook's privacy settings may have been intentionally confusing to encourage more sharing. At best, Facebook got the mapping wrong—the settings users were asked to choose from did not correspond logically with their intentions. This study echoes the findings of many researchers that design is a major obstacle to managing our online privacy.[32]

All of this research leads us to two conclusions. First, design is incredibly important. The hardwiring in our brains makes us susceptible to design, which affects us at every turn. Technological design can determine how we communicate and share information, how often we use a particular technology, and how we relate to the businesses and other entities we deal with online. It can make us feel in control or it can frustrate us to the point of giving up—and it can trick us. Second, design is difficult to get right. It requires methodical consideration of how we interact with the things around us and what we expect from those interactions. Design must anticipate human error and respond to feedback. In short, technology needs to be designed for *people.* People that design for a living get this. Researchers rigorously study design, and companies invest heavily in it. Companies have every incentive to use the power of design to their advantage. Yet, as we will see, privacy law has failed to take design seriously.

Design Is Power

If the first truth of design is that it is everywhere, the second truth of design is that it is also a form of power. Power has been defined as "the capacity or ability to direct or influence the behaviour of others or the course of events."[33] Given how design can shape our perceptions, behavior, and values, *power* and *design* often feel like synonyms. Design is power because people react to design in predictable ways. This means that with the right knowledge, design can impose some form of order on chaos.

I want to emphasize that I'm not arguing that design completely dictates our privacy and that nothing else is relevant. Such an extreme argument is a misguided form of technological determinism, which is the idea that technology makes cultures what they are and is the exclusive key to change.[34] But technology does affect us in powerful and tangible ways.

Governments and industry have long known that design is critical to accomplishing basically any significant endeavor. They have always leveraged design to achieve particular ends. This is evident from wide-scale city planning efforts that facilitate the movement of cars and people via traffic circles, grid systems, and road dimensions. It can also be seen in something as small as a park bench, with armrests spaced evenly across the bench to prevent people from lying flat and sleeping on it.

Entire bodies of literature have been dedicated to understanding and harnessing the power of design in policy and industry. The entire field of engineering involves leveraging design to accomplish something. Architects match design with purpose in order to make "structure express ideas," in the words of Frank Lloyd Wright.[35] Urban planners use design to accomplish the goals of a city, improve public welfare, and protect the environment. But the study of design goes far beyond the construction of cities, buildings, and machines. Researchers also explore design's social effects.[36] Rooms built to expose occupants to natural lighting can increase workplace performance.[37] Elevator banks have reflective surfaces to give you something to look at while you wait for an elevator, which can reduce anxiety and anger. The walk from your airplane to the baggage claim area is intentionally long because you are less conscious of the time it takes to claim your baggage if you are walking rather than standing and waiting.[38] Placebo buttons—buttons that do nothing and exist entirely to make you feel better—are everywhere. The "door close" button on the elevators? Probably fake. Crosswalk button? Probably fake. Your office thermostat? Probably fake. I was surprised too.[39]

Economics professor Richard Thaler and law professor Cass Sunstein have pioneered the concept of "nudging," which leverages design to improve people's lives through what Thaler and Sunstein call "choice architecture." Choice architects are people who have "the responsibility for organizing the context in which people make decisions." A nudge is "any aspect of the choice architecture that alters people's behavior in a predictable way without forbidding any options or significantly changing their economic incentives."[40] Designers and engineers are choice architects. Thaler and Sunstein give numerous examples of design nudges, including alarm clocks that run away from you, forcing you to chase them to turn them off (and thus ensure you are awake), increasingly grouped white lines on dangerous curves to make drivers feel as though they are going faster (and thus reflexively slowing themselves down), and smiling and frowning emoticons on power bills to encourage people to use less energy. Nudging through choice architecture does not give designers total control over people. However, it can give designers control at the margins.

Design is powerful because we as people are more easily manipulated than we'd like to think. We often fail to act in our own self-interest. We like to think of ourselves as rational and autonomous actors, but the fact is that we are not. Decades ago, psychologists Daniel Kahneman and Amos Tversky began to unravel the leading "rational actor" model of human decision making by demonstrating the mental shortcuts (rules of thumb known as heuristics) and biases that dominate our mental process and judgment. They hypothesized that while sometimes we call upon a reflective system of logic and reasoning to form judgments, more often we use rules of thumb to automatically form judgments without the cognitive burden of measured, contemplative thought. In his foundational book *Thinking Fast and Slow* Kahneman called the fast, instinctive, and emotional method System 1, and the more deliberative and logical method System 2. System 1 is the one we use most but, unfortunately, it's the troublemaker. It frequently guides us to make decisions that are less than optimal.

While we might not be "rational" in the sense that we often fail to act in our own self-interest, we are pretty consistent. Behavioral economists and psychologists have demonstrated that people have a consistent bias in making routine judgments.[41] For example, we have a tendency to be influenced by irrelevant numbers, often the first ones that come to mind (a process known as the anchoring effect). When trying to gauge the probability of some event happening, we rely far too much on the information that is the easiest to recall in our minds, rather than the most relevant information (a process known as the availability heuristic). We have a tendency to be far more optimistic than we should be, even in the face of contrary evidence (a process known as optimism bias). We regularly keep investing in things even though it's a bad idea, or "throw good money after bad," just because we instinctually want to avoid regret over how much we've already spent (a process known as the sunk cost fallacy). As we will see later in this chapter, we change our attitudes about things based solely on the way facts are presented, even if the facts stay the same (a process known as framing). We consistently choose a "smaller-sooner" reward rather than a "larger-later" one because it will occur earlier in time (a process known as hyperbolic discounting).[42] We also tend to search for, interpret, favor, and recall information in a way that confirms our preexisting beliefs and give less consideration to alternatives that do not (a process known as confirmation bias).[43] Because these effects are consistent and predictable, design can be adjusted to leverage all of these biases at a designer's will.

Design shapes our privacy perceptions, which in turn shape how we use and respond to technologies. Because privacy is so difficult to pin down and the harms are so diverse and often remote, we crave privacy guidance. Design gives it to us. In summarizing the robust literature around privacy and decision making, Alessandro Acquisti, Laura Brandimarte, and George Loewenstein note that three themes have emerged from the literature that explain our vulnerability to design and other external forces that shape how we disclose personal information and make choices regarding our privacy.[44]

First, people are *uncertain* about the nature of privacy trade-offs and about what kinds of trade-offs they prefer. Information asymmetries keep people from properly assessing risk, and even when privacy consequences are clear, people are uncertain about their preferences. Second, our privacy preferences are almost entirely *context dependent*. As Acquisti and colleagues note, "The same person can, in some situations, be oblivious to, but in other situations be acutely concerned about, issues of privacy."[45] Finally, our privacy preferences are incredibly *malleable*—that is, they are subject to influence from others who have better insight into what will make us act.

So while people may be unaware of the forces that modulate their privacy concerns, companies that rely upon personal information are not. Acquisti and colleagues found evidence of the manipulation of subtle factors that activate or suppress privacy concern "in myriad realms, such as the choice of sharing defaults on social networks, or the provision of greater control on social media, which creates an illusion of safety and encourages greater sharing."[46]

Our uncertainty, malleability, and dependence upon context all work together. We might be so dependent upon context because of how uncertain we are about the outcomes of sharing personal information. We need lots of clues to give us a hint as to what to do. And our privacy preferences and behaviors are malleable and subject to influence because we are so dependent upon context. Because our privacy intuitions are so malleable, we are vulnerable to those who would manipulate context to their own advantage.

Enter design. Recall Norman's theory of good design being a function of mental mapping, technical and normative constraints, and affordances. The concept of affordances—the fundamental properties that determine how a thing can be used—is useful because it gives us a framework for understanding how people interpret and then interact with an object or environment. James Gibson has theorized that although we all interact with the same objects and environment, our abilities and limitations cause us to perceive these things differently. People perceive cliffs as dangerous; birds perceive cliffs as irrelevant, or as great places to build nests.

Affordances can be negative or positive, depending upon how they are perceived, and they can also be seen as subject to change. Stairs can be carved into a steep hill to provide a walkable surface. Yet stairs themselves are a negative affordance for those who use wheelchairs. Whether an affordance is true or false, perceptible or hidden, also affects how people interact with objects and environments. Hidden affordances do people no good. Checkboxes allowing people to opt out of information collection are worthless when they're tucked away at the bottom of a screen where they can't be easily seen. False affordances can trick people into relying on protections and escape routes that do not exist. Because people often rely on affordances, an opt-out checkbox that actually does nothing is worse than no affordance to opt out at all.

Language can also be incorporated into design to shape our perceptions, and the way we receive a message may be just as important as the message itself. Even if the substance of a communication remains constant, the presentation of that substance can significantly affect how we perceive it. In communications, sociology, psychology, and related disciplines, this concept is known as framing.[47] Framing theory holds that even small changes in the presentation of an issue or event can produce significant changes of opinion.[48] For example, older studies have shown that some people are more willing to tolerate rallies by controversial hate groups when such rallies are framed as an exercising of free speech rather than a disruption of the public order.[49]

Consider two questions: Do companies with data use policies protect your privacy? And, do companies with privacy policies protect your privacy? While these questions are constructed differently, they are essentially asking the same thing. But the choice of which question to ask could predetermine the answer. In 2005, Joseph Turow led a research effort uncovering that 59 percent of people falsely believe that websites with a privacy policy cannot sell personal information without consent.[50] Meanwhile the term *data use policy,* as used by Facebook and others, carries no implicit marker or promise of respect for privacy.

Judges, lawmakers, and the public all use and are influenced by frames.[51] According to Robert Entman, "To frame is to *select some aspects of a perceived reality and make them more salient in a communicating text, in such a way as to promote a particular problem definition, causal interpreta-*

tion, moral evaluation, and/or treatment recommendation for the item described," and framing offers a way to articulate the "power of a communicating text." By increasing the salience of certain bits of information, frames enhance the probability that receivers will interpret the information in a certain way, discern a particular meaning, and process it accordingly.[52]

While frames do not guarantee an influence on audience thinking, frames that comport with the existing schemata in a receiver's belief system can be particularly effective.[53] Kahneman and Tversky have offered what is now likely the most well-known example of how framing works by highlighting some features while omitting others. In an experiment, the researchers asked test subjects to consider the following hypothetical:

> Imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimates of the programs are as follows:
> If Program A is adopted, 200 people will be saved. . . .
> If Program B is adopted, there is a one-third probability that 600 people will be saved and a two-thirds probability that no people will be saved. . . .
> Which of the two programs would you favor?

Here, 72 percent chose Program A. Kahneman and Tversky followed this experiment with another that offered mathematically *identical options* to treating the same situation, but the programs were framed in terms of *likely deaths* rather than *lives saved:*

> If Program C is adopted, 400 people will die. . . .
> If Program D is adopted, there is a one-third probability that nobody will die and a two-thirds probability that 600 people will die.

With this alternative framing, 22 percent chose Program C, even though 72 percent of the previous experimental group selected Program A, Program C's mathematical twin.[54] In short, the alternative framing resulted in what was essentially a reversal of the percentages.

The frame in this experiment determined whether most people noticed a problem and how they understood and remembered it. The frame also determined how people evaluated and chose to act on the problem.[55] Perhaps one of the most important functions of frames is that by calling attention to particular aspects of a described reality they, by construction, direct attention away from other facets. This logical sleight of hand means the power of most frames lies in that what they omit as well as include. Omissions of things like definitions of problems, explanations, evaluations, and recommendations might be just as important as what a frame includes in guiding the audience.[56]

Design and language work together to make frames that affect our perceptions of how a technology works and, ultimately, decisions that affect our privacy. For example, contrary to the common assumption in both the literature and popular culture that people have stable, coherent preferences with respect to privacy, Leslie K. John, Alessandro Acquisti, and George Loewenstein have found that "concern about privacy, measured by divulgence of private information, is highly sensitive to contextual factors," including design.[57] In a series of experiments, John and colleagues manipulated the design of interfaces to make privacy concerns more or less salient to certain participants. In each experiment, participants provided identifying information (email addresses) and then indicated whether they had engaged in a series of sensitive, and in some cases illegal, behaviors.

In one experiment the researchers found that participants were more likely to divulge sensitive personal information in a survey with a more frivolous, less professional design than in one with a more formal interface. John and colleagues write, "The study was inspired by news stories about postings on the Facebook group '20 reasons why a girl should call it a night' in which young women voluntarily post compromising pictures of themselves—pictures that, in most other contexts, they would be mortified to share."[58] The researchers were interested in whether the frivolous nature of the site encouraged self-revelation and suppressed concern for privacy. They tested people's responses to two different designs of online surveys that asked people for very sensitive disclosures such as, "Have you ever smoked marijuana (i.e., pot, weed)?" The first design was intended to downplay privacy concerns. It titled the survey "How BAD are U??" and looked lighthearted.

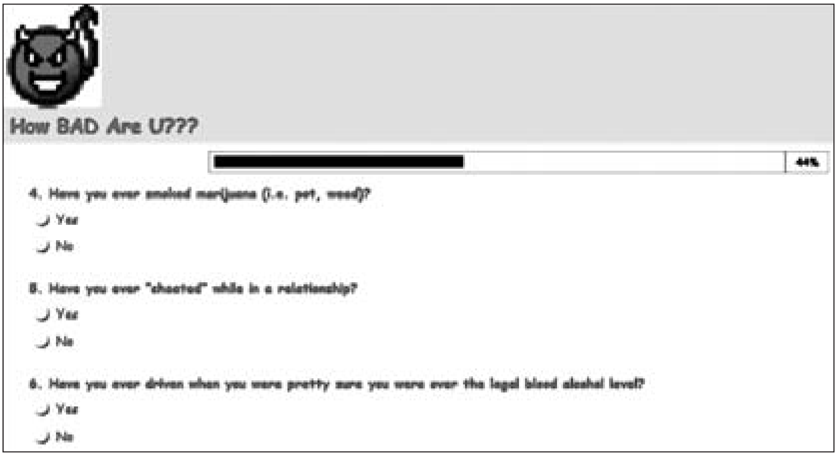

How BAD Are U???

44%

4. Have you ever smoked marijuana (i.e. pot, weed)?

○ Yes
○ No

5. Have you ever "cheated" while in a relationship?

○ Yes
○ No

6. Have you ever driven when you were pretty sure you were over the legal blood alcohol level?

○ Yes
○ No

The second control design set the survey within a professional context and was titled "Carnegie Mellon University Survey of Ethical Standards." This was less frivolous in appearance and was thus hypothesized not to minimize privacy concerns in the same way as the more frivolous interface might.

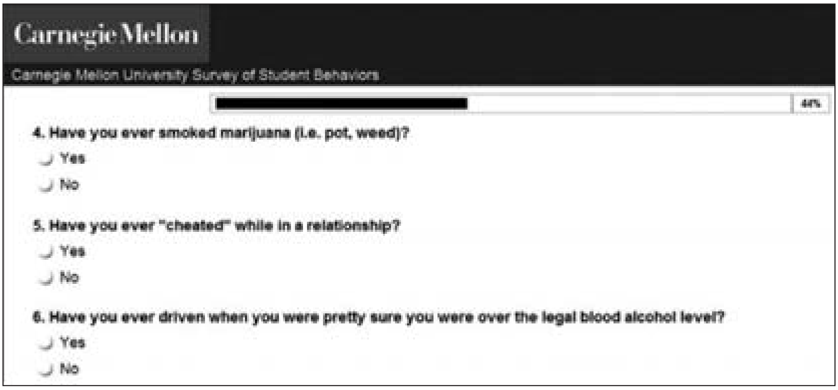

Carnegie Mellon

Carnegie Mellon University Survey of Student Behaviors

44%

4. Have you ever smoked marijuana (i.e. pot, weed)?

○ Yes
○ No

5. Have you ever "cheated" while in a relationship?

○ Yes
○ No

6. Have you ever driven when you were pretty sure you were over the legal blood alcohol level?

○ Yes
○ No

Relative to the nonfrivolous interface, participants in the frivolous-looking survey that asked identical questions were on average 1.7 times more likely to admit to having engaged in risky behaviors. The researchers found that "a participant in the [frivolous-looking survey] was on average 2.03 times more likely to admit to having 'ever taken nude pictures of [him]self or a partner' People, it seems, feel more comfortable providing personal information on unprofessional sites that are arguably particularly likely to misuse it." The authors conclude, "By illustrating that disclosure of private information is influenced by contextual factors that have little, if any, normative justification, the current research casts serious doubt on whether individuals may be able to navigate this complexity according to their best interests."[59]

Once design affects our perceptions, it begins to shape our behavior. Once it shapes our behavior, it can be used to control us because it shapes what we perceive as normal. And once norms are established, they are difficult to change. Studies have explored how framing and design affect our privacy-related behavior. In multiple experiments in 2014, researchers changed the design, framing, and presentation of online privacy choices to users. For example, one group was exposed to "privacy settings" while another was exposed to "survey settings." They found that "[a]cross three experiments, the choice of the privacy protective options for objectively identical choices ranged from 17% to 55% depending on choice framing."[60]

The researchers concluded that given the value of personal information, it would be surprising if companies did not strategically leverage framing to get people to disclose more. And market forces are unlikely to help consumers much here, because people are unlikely to notice the subtle manipulations like the ones studied by the researchers. Furthermore, people are unlikely to notice the impact of those manipulations on their behavior. We just are not that self-aware. Even when design purports to give people more "control" over their information, we often have a false sense of security leading to risky personal disclosures.[61]

These studies are just the tip of the iceberg for understanding how design affects our privacy perceptions. Defaults, malicious interfaces, intentionally not ringing "privacy alarm bells," and a host of other design choices can all be used to manufacture personal disclosures. Acquisti, Brandimarte, and Loewenstein argue that policy approaches that rely exclusively on informing or empowering people will not work. Even "control" and "transparency," two hallmarks for privacy protection, are demonstrably vulnerable and in many cases radically insufficient when used apart from other principles of data protection. They argue that "People need assistance and even protection to aid in balancing what is otherwise a very uneven playing field."[62] Design can create power disparities, but it can also better distribute power within relationships. The authors conclude, "To be effective, privacy policy should protect the naïve, the uncertain, and the vulnerable. It should be sufficiently flexible to evolve with the emerging, unpredictable complexities of the information age."[63]

Design Affects Our Values

We are what we create, so we must be very careful with what we create.[64] Design reflects the values of its creators and can either support or undermine human values more generally.[65] In Langdon Winner's work on the politics of artifacts, he tells the story of Robert Moses, expert builder of roads, parks, bridges, and other public works from the 1920s to the 1970s in New York. Moses designed overpasses "to specifications that would discourage the presence of buses on his parkways."[66] Stripped of context, this decision seems odd. But when you consider the fact that minorities and the poor in this community and time disproportionately relied upon buses for transportation, this design is revealed as malicious: "According to evidence provided by Moses' biographer, Robert A. Caro, the reasons [for the low overpasses] reflect Moses' social class bias and racial prejudice. Automobile-owning whites of 'upper' and 'comfortable middle' classes, as he called them, would be free to use the parkways for recreation and commuting."[67] No buses, no poor people. (Or at least fewer poor people, given the transportation hardship).

If true, Moses's values were reflected in those bridges, and their design undermined larger human values of openness and equality, to say nothing of efficient transportation. In some contexts, we should pay even more attention to the values furthered by design because of lock-in effects. Winner writes, "For generations after Moses' death and the alliances he forged have fallen apart, his public works, especially the highways and bridges he built to favor the use of the automobile over the development of mass transit, will continue to shape that city. . . . As New York planner Lee Koppleman told Caro about the low bridges, . . . 'The old son-of-a-gun had made sure that buses would *never* be able to use his goddamned parkways.' "[68]

Like that of bridges, technology design can be imbued with social values. Consider the infamous wearable technology Google Glass; envisioned as the first major wearable ubiquitous computing device, it was released in 2012 to widespread criticism (fans of the device came to be known by some as Glassholes).[69] The technology essentially looked like eyeglasses with a small computer and lens slightly covering one eye. Ideally, Google Glass was a way for people to look at a screen without using their hands.

But one design decision may have been too much for us all to handle. Google included a camera on Google Glass. It would have been useful without a camera, but the public has become accustomed to being able to take a picture at any moment, and image sensors are quite useful. It turns out the camera may have been a bridge too far. People need time to adjust to new technologies. Without a camera, we might have been able to wrap our minds around always having little computers in front of our eyes. But a somewhat surreptitious camera—one that could snap a photo at any moment without the user having to hold up a conspicuous device—represented surveillance and a threat to our privacy and autonomy.

Anyone in range of a Google Glass device was put on notice of being watched. Conversations that were casual and free were subject to preservation and mass publication within seconds. The camera embodied everyone's worst fears about surveillance in the modern world.[70] The visible indicator light meant to alert bystanders that the device was recording and the company's lack of support for facial recognition apps for Google Glass could not save its reputation as a privacy-invasive technology.[71] Compare this with the cameras on smartphones. Adding a camera to a phone happened more gradually, which has made the adjustment to living in a world where people are constantly taking and sharing pictures a little easier. Smartphones also live in pockets, purses, and bags and require at least some effort to take out and use. Google Glass shows how both design and culture work together to shape how technologies are used and whether they are accepted. It also shows how design implicates values.

Design Is Political

For the past few years I have spoken with people in industry, academia, government, and civil society about the importance of privacy-related design. In these talks I've heard one argument come up repeatedly: the notion that it is users, not technologies, who are to blame for privacy violations. When I tell people the thesis of this book, sometimes they respond with some form of the argument, "There are no bad technologies. Only bad technology users."

According to this line of thought, it is not technology itself that matters but people—and the social and economic system they exist in. We are accustomed to thinking of tools that can be used well or poorly, for good or for evil, or somewhere in the middle of all of that.[72] This instrumentalist conception of technology is regularly used in ethics and policy discussions about the appropriate use of a technology to focus attention on the actor using the technology rather than the technology itself. For example, intellectual property law gives a pass to technologies with "substantial non-infringing uses."[73] The idea is to not blame the technology when it's really the user that's causing the harm. In this light, technologies are largely innocent bystanders.

This concept of neutral technologies might seem intuitive. Why blame technologies for harms carried out by people? After all, technologies are just tools. They are not self-executing; they require someone (or many people) to bring their intended purpose to fruition. In order to operate your car, you must ignite the engine, press the gas pedal, and turn the steering wheel. Your computer requires you to push the power button, move and click the mouse, and depress letters on your keyboard to make words. Autonomous technologies are still reliant upon designers and operators for initial execution. Even simple, docile objects like your refrigerator must be positioned and plugged in to keep your food cold. Given this subservient role of technology, it is not surprising that a dominant strain of "technological neutrality" exists in policy, industry, and our social lives.

The technological neutrality argument is evocative of a slogan sometimes used by gun rights activists: "Guns don't kill people. People kill people."[74] Technological neutrality arguments have political appeal and have been used effectively to shift lawmakers' focus away from technologies and toward bad actors using those technologies. But this sort of technological neutrality argument gives short shrift to the political and practical power of technology. Every single technology instantiates values by virtue of its design that makes certain realities more or less likely.

At this point you might be thinking, whatever happened to personal responsibility? Indeed, in the absence of a default setting, iPhones do not save someone else's explicit photos on their own. Facial recognition software that is used to find and harass people cannot reflect on its own wrongdoing. These technologies merely serve the people who make harmful decisions.

One way to downplay the importance of design is to call attention to the many different theoretical ways a technology could be used. People can use drones to film birds rather than peep through windows. People can let intimate Snaps they receive fade into the ether instead of saving and posting them to pornography and revenge websites. Framed this way, people's harmful uses of technology seem much more blameworthy than the design of the technology itself.[75]

But it's not that simple. Consider the Pinhole Spy Toothbrush Hidden Camera. The toothbrush is designed with no wires, a long battery life, no external memory card, and a very small lens to make it look as much like a regular electric toothbrush as possible. The website selling this spycam markets it as ideal for illegal behavior and, in an astonishing display of hubris, disclaims that "we are not responsible if this camera is used for illegal activities, this is a home security camera and should be treated as such, with responsibility."[76]

Of course, this technology could be used many different ways and for many legitimate ends. People could use it to brush their teeth. They could use it to take video of their family vacation on the beach. They could use it as a doorstop or play catch with it. Or perhaps they could just mount it on their mantel for aesthetic purposes.

But of course, all those uses are ridiculous. To suggest otherwise is to ignore how this technology embodies behavior-shaping values. To paraphrase Evan Selinger, the design of the toothbrush indicates the preferred ends to which it "should" be used—namely, surveilling others in secret. People aren't likely to spend $230 on a simple toothbrush, toy, or sculpture. And using a toothbrush camera to record your family vacation would just be awkward; better and cheaper cameras that won't make you look like a weirdo are available. The design of this technology facilitates bathroom voyeurism, pure and simple. There is nothing neutral about it.

To be sure, the notion of technological neutrality can be useful. It can help us avoid uncritically blaming objects and technologies for problems without taking into account the social and economic system in which such technologies are embedded. Winner argues that ideas like this, which have been referred to as the "social determination of technology," serve as a "needed corrective to those who . . . fail to look behind technical things to notice the social circumstances of their development, deployment and use. This view provides an antidote to naive technological determinism—the idea that technology develops as the sole result of an internal dynamic, and then, unmediated by any other influence, molds society to fit its patterns."[77]

But we must not overinvest in the belief of technological neutrality. Winner argues, "taken literally, [the social determination of technology] suggests that technical *things* do not matter at all." He instead suggests that artifacts are laden with *politics*—that is, "arrangements of power and authority in human associations as well as the activities that take place within those arrangements." Sometimes, the design or arrangement of a specific technical device or system "becomes a way of settling an issue in a particular community."[78] Winner notices that "we usually do not stop to inquire whether a given device might have been designed and built in such a way that it produces a set of consequences logically and temporally *prior to any of its professed uses*" and contends, "If our moral and political language for evaluating technology includes only categories having to do with tools and uses, if it does not include attention to the meaning of the designs and arrangements of our artifacts, then we will be blinded to much that is intellectually and practically crucial."[79]

He is right. When we fail to pay proper attention to the design of technologies, we end up focusing too much on human conduct. We ask too much of ourselves and others to overcompensate for design that pulled or pushed us in a direction adverse to our interests. We blame ourselves and others for being clumsy or mechanically incompetent. Norman has noticed people's tendency to think they are at fault when they fail in using an everyday object.[80]

On social media, we even have a term for our own "privacy fails"; we call it "oversharing."[81] The Internet is littered with articles like "The 30 Absolute Worst Facebook Overshares" and "Facebook Overshare: 7 Things You Might Not Realize You're Telling the World."[82] These posts describe tales of lost virginity, menstrual awkwardness, kidney stones (including pictures!), sexually transmitted diseases, and almost any other explicit details that one might simply rather not know about someone else. Even random clicks—for example, "liking" someone's photo that was posted three years ago (revealing you might be spending a little too much time "Facebook stalking" someone)—can be considered oversharing.

These obvious examples make it seem like only idiots and people with no dignity overshare. But that's not true. If we consider oversharing to simply be "sharing more than we should" or "sharing more than we would in offline, nonmediated settings," we all overshare regularly. As I will discuss in Chapter 6, social media are disclosure-manufacturing machines. Their entire purpose is to facilitate sharing early and often. For some social media, the more personal the information, the better. It is common for those of us with social media accounts to share our birthdays, illnesses, travel plans (i.e., when we will be away from home), intimate thoughts, and pictures of us in our skimpy swim trunks and bikinis. And sometimes we forget who can see our posts or we simply make a mistake. It's okay to admit it. No one has perfect judgment.[83]

And when people overshare, we almost always say they are to blame. For example, in an article titled "You're as Much to Blame for Facebook's Privacy Woes as Mark Zuckerberg," columnist Farhad Manjoo calls out social media users as the cause for their own privacy concerns, noting, "You should approach Facebook as cautiously as you would approach your open bedroom window" and, regardless of your privacy settings, "you should imagine that everything that you post on Facebook will be available for public consumption forever."[84] This is the only failsafe for sharing on social media. Manjoo also had a column titled "How To Stay Private on Facebook in One Easy Step" that consisted of only two words: "Quit Facebook."[85] To Manjoo, oversharing is ultimately your fault. To prevent it, just shut up. You are not forced to post anything, after all.

This user-fault mentality is also prominent in the context of nonconsensual pornography, sometimes called revenge pornography. As Danielle Citron writes in her book *Hate Crimes in Cyberspace,* "Revenge porn victims have told me that the overwhelming response to their struggles is that they are responsible for their nude photos appearing online." She quotes the operator of the revenge porn site Texxxan.com, who told the press, "When you take a nude photograph of yourself and you send it to this other person, or when you allow someone to take a nude photograph of you, by your very actions you have reduced your expectation of privacy."[86]

Putting aside the self-eroding and problematic nature of the notion of a "reasonable expectation of privacy," this myopic focus on only the most proximate and direct causes of harm ignores the entire structure and context that facilitated it. Of course, sometimes people are careless in their interactions with others. Other times they are betrayed by trusted confidantes, or their actions were just enough to tip delicately balanced scales.

We've all seen a scenario in movies or television where an out-of-control car comes to a stop, dangling perilously on the edge of a cliff, when one of the passengers leans forward slightly, causing it to go over. As any first-year law student will tell you, the person leaning over may have technically "caused" the fall, but she is not the only one at fault.

The notion that human choice, and not design, is what matters most in protecting our privacy is firmly entrenched in U.S. law. In this way, privacy law reflects a kind of neutral or even protectionist approach to technology. Courts generally refuse to hold technology companies liable in tort for unsecure computer code.[87] The four privacy torts all restrict people's conduct—people cannot publicly disclose another's private facts, intrude upon another's seclusion, commercially misappropriate another's name or likeness, or depict another in a false light. Surveillance law in the United States prohibits unauthorized interception of communications but has little to say about the design of surveillance technologies (with one notable exception for spyware, which we will return to in Chapters 5 and 7). With exceptions for data security, statutory regimes like the Family Educational Rights and Privacy Act, the Gramm-Leach-Bliley Act, and the Health Insurance Portability and Accountability Act largely focus on people keeping a confidence but ignore the tools of disclosure.

Even in our day-to-day lives, we often fall prey to the myth of technological neutrality and miss how technological design affects us. Consider the password. Many data breaches are the result of an attacker successfully figuring out someone's password. This has led to the standard advice that we've all heard: "Select a long and complex password. Use upper and lower case, letters and numerals, and special characters. Commit it to memory. Change it frequently. And do this for every account you have."

Most people don't do this. How could they? The most popular passwords are still words like "password" or other simple things that even a bad hacker can crack in seconds. People reuse their passwords, write them down on sticky notes next to their computers, or carry them in scraps of paper in their wallets. And people don't change their passwords very often.

And when accounts with poor passwords are breached, our first reaction is often to blame the user: "This idiot used the password 12345." When my own personal email account was compromised by a spammer a few years ago, I blamed myself. I had used a relatively short, recycled password, and I knew better. I was embarrassed because I write about privacy and data security for a living. How could I be so dense?

What I didn't think about at the time was how design helped precipitate the breach. The system has set us all up to fail. Our memories can handle at most a few long and complex passwords, but not a lot of them and not if they must be changed frequently. According to one study, consumers have an average of twenty-four online accounts. For those who use the Internet more robustly, that number is much higher. It's just too much to handle, so it's no surprise that people fail to follow good password protocol.[88]

If we say "Systems don't facilitate data breaches, people do," we obscure the orchestration of design choices that give people little practical say in the matter. Several different design decisions could have been made to anticipate human limitations and predictable error. For example, designers could have required two-factor authentication, the essence of which is simple: in order to log in, you must have something you know (usually a password), as well as one additional factor, usually something you have (usually your cell phone). Additional factors could be mixed and matched to avoid relying on the sole password. Under a "two-channel" authentication scheme, companies would not authenticate users until they actually hear back from them on the second channel (such as the cell phone) dedicated to authentication. Other factors can include having a friend vouch for you or sign your name. None of these are foolproof (and some are certainly better than others), but generally two factors are much better than one for important accounts.

I am not saying that people should be able to do anything they want online and avoid the consequences of their behavior. That is a ridiculous notion. We should all act reasonably online. We should not be reckless. When we are, we must accept the consequences of our actions. But we must move past the notion that the *only* thing that matters in the privacy debate is collection, use, and disclosure of personal information. The current trajectory of privacy law and discourse cannot continue to marginalize the role of design. Design picks data winners and privacy losers.

This is true even when we cannot characterize outcomes as "intended" or "unintended," which are often oversimplified categories. Winner notes that sometimes "the very process of technical development is so thoroughly biased in a particular direction that it regularly produces results counted as wonderful breakthroughs by some social interests and crushing setbacks by others."[89] Consider the data collection infrastructure that feeds the phenomenon we know as *big data,* an amorphous term usually best thought of as shorthand for powerful data analytics. Scientists rejoice over big data. It holds the potential to answer previously unanswerable questions regarding our health and well-being. Insurers and advertises also love big data. It can allow for better, more efficient personal profiles and more profitable decisions regarding who gets coverage and what ads we see. But big data is also dangerous to those being profiled. It can be used to give us all scores—similar to credit scores—for nearly any aspect of our lives.[90] While the scored society promises many advantages, it also threatens to be a ruthlessly efficient discrimination machine that can curtail our autonomy, disproportionately burden minority communities, and demand ever more personal information to feed the beast. In the world of big data, more is always better.

Winner writes that for most technologies, "it is neither correct nor insightful to say, 'Someone intended to do somebody else harm.' Rather one must say that the technological deck has been stacked long in advance to favor certain social interests, and that some people were bound to receive a better hand than others."[91] To properly understand the role design plays in our privacy, we must understand how, why, and for whom the deck has been stacked.

Design Should Be Policy

Because design can allocate power to people and industries, it is inherently political. Advocates for technological innovation often say "the law should stay out of technology." This is sometimes perceived as a call for neutrality, leaving market forces and norms to shape technology. There are often very good reasons why the law should not directly address the design of technology, but such a decision is not neutral. To not address design is to sanction the power of creators to determine the context in which people make decisions. Most of the time this is the right path. The law usually should not dictate the minutiae of design. It would be a nightmare if every little design choice for every digital technology was micromanaged and had to be preapproved by some regulator.

But there is a difference between zealously *dictating* design and *addressing* it. The law can be responsive to the power of design by articulating boundaries, guidance, and goals. The law can simply better recognize the role that design plays in shaping our privacy expectations and reality. It is an ethical lapse to ignore design altogether under the auspices of neutrality. Because design and architecture affect every choice we make, even silence matters.

Others argue that controlling design itself is ethically problematic, because, stripped bare, design regulations are attempts to manipulate people into making particular choices. In addressing the ethics of nudging, Cass Sunstein writes, “When people make decisions, they do so against a background consisting of choice architecture. A cafeteria has a design, and the design will affect what people choose to eat. The same is true of websites. Department stores have architectures, and they can be designed so as to promote or discourage certain choices by shoppers (such as leaving without making a purchase).”[92] This architecture will affect people’s choices even if the effect is not intended by designers. For example, people have a tendency to buy things they encounter first, so even if you are not trying to sell a lot of a sweater vests, putting them at the entrance to your store will probably move more vests than if you tuck them away in some corner.

We cannot avoid choice architecture. As Sunstein notes, “Human beings (or dogs or cats or horses) cannot wish [choice architecture] away. Any store has a design; some products are seen first, and others are not. Any menu places options at various locations. Television stations come with different numbers, and strikingly, numbers matter, even when the costs of switching are vanishingly low; people tend to choose the station at the lower number, so that channel 3 will obtain more viewers than channel 53.”[93] A website has a design, and that design will affect how users will navigate the site and how long they will stay. When lawmakers and judges ignore design they implicitly endorse the status quo and the unchecked use of power.

Consider default rules and settings, which are the preselected choices that will take effect anytime a person does nothing or fails to select an option. Sometimes it is impossible to avoid creating a default. Consider the on / off toggle button that is so common on our smartphones:

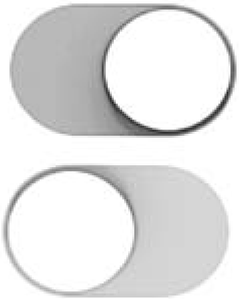

There are only two possible choices here: *on* or *off.* Designers must choose to preselect the default position for the button. This choice cannot be avoided, because even some halfway choice in a binary decision would basically function as *off.* Because we know defaults are sticky and it would take the user's scarce resources of time and attention to change the setting, the default decision reflects a value. If the default for location tracking and sharing were set to *on,* we should expect the exact location of many more phone users to be tracked and shared. Many of us might not mind, but political dissidents and those seeking refuge from domestic abuse might be jeopardized by this default. The danger of defaults is acute when people are unaware of them or when the panel of selections looks like this cascading array of options:

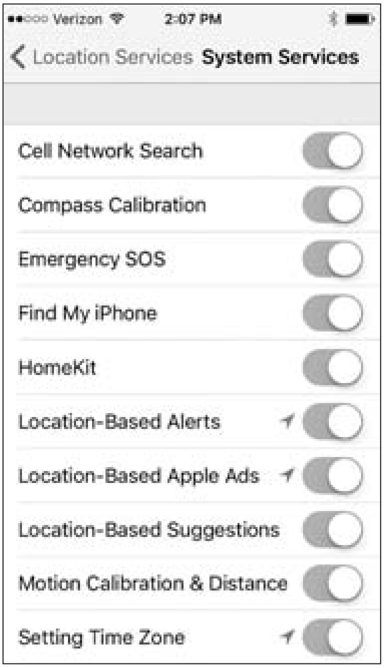

Cell Network Search

Compass Calibration

Emergency SOS

Find My iPhone

HomeKit

Location-Based Alerts

Location-Based Apple Ads

Location-Based Suggestions

Motion Calibration & Distance

Setting Time Zone

The default for many social media services is set to maximize sharing at the expense of discretion. For example, at one point the mobile payment service Venmo publicly shared its users' transactions by default.[94] The process for users to limit access to all of their transactions is not very intuitive. The Default Audience setting only affects the visibility of the charges and payments initiated by the user. The transactions initiated by friends who have not changed the default privacy setting will still be shared publicly. To obscure all transactions, a user must hunt down and change the inconspicuous Transactions Involving You setting.[95] These kinds of defaults matter. According to a 2015 study from Kaspersky Lab, 28 percent of social media users neglect privacy settings and leave all of their posts, photos, and videos accessible to the public.[96] There are hundreds of buttons under the hood of our technologies. Each default setting affects our choices.

And we users are severely outmatched. Technology users operate as isolated individuals making decisions about things we don't know much about, like the risk of disclosing information and agreeing to confusing legalese. We're up against the brightest graduates from Ivy League schools who spend their days figuring out how to get us to disclose personal data.

It's like we are disorganized little tribes fighting against an organized army. Large tech companies claim to be running thousands of A / B tests at a time, and these tests compare two different versions of an interface to see which one works better. How many of those tests are about maximizing how much personal information can be extracted from people? This is why progressive new laws like Article 25 of the European Union's new General Data Protection Regulation titled "data protection by design and by default," which requires that core data protection principles be integrated into the design and development of data technologies, are so important.[97]

A regulator's choice to address design or to ignore it is also a kind of default, which comes with its own set of consequences. This is actually what the law does—it provides the default rules for how we deal with each other. Even if a government claims to be firmly committed to free markets, private property, and keeping regulation out of technology, it cannot simply refrain from acting—or, in Thaler and Sunstein's word, "nudging." Even a free market, competition-based system requires a legal framework. As Sunstein writes, "the rules of contract (as well as property and tort) provide a form of choice architecture for social ordering." Design can, of course, maintain freedom of choice where it is important, and it can also consciously pursue neutrality in important ways. But, Sunstein notes, "choice architecture itself is inevitable, which means that it is pointless to object to it on ethical grounds."[98]

In fact, ethics compel lawmakers to take design seriously. Designers facilitate morality because the decisions they make in creating an object have morally relevant outcomes. Philosopher Peter-Paul Verbeek notes that technologies facilitate morally "good" and "bad" outcomes; speed bumps, for example, "help us make the moral decision not to drive too fast near a school." Technologies help shape our actions and experiences, and in so doing facilitate our ethical decisions. For example, social media design can nudge us to increase our civic participation or to heckle or harass others. According to Verbeek, because technologies are "bursting with morality," we must develop an ethical framework to conceptualize the moral relevance of technology.[99]

Lawmakers must acknowledge that the law inevitably influences design. Once lawmakers, policy creators, and judges realize that choice architecture is unavoidable and that the design of systems has moral implications, they have a choice to make: they can keep the status quo and basically ignore the design of technologies when addressing privacy law, or they can confront the important ways design shapes morally relevant outcomes. Lawmakers are justified if, upon assessing the effect of design on desired outcomes, they intentionally decline to address design in their laws and regulations and instead preserve the status quo. But failure to act out of a desire for neutrality or pure randomness is not justified, because such outcomes are impossible. Because of the great capacity for design to facilitate human flourishing or inflict suffering, no principal justifies refusing to form a deliberate legal framework, be it permissive or intervening, over the design of technologies. We cannot ethically ignore design.

In this chapter we explored how design affects our privacy because it is everywhere, it is powerful, and it is political. In other words, it is a force that should not be ignored by policy makers. In Chapter 2, I will make the case that privacy-related design deserves much more attention in law and policy.