

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 environ-

ment. 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

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.