Introduction

There's a scene in the HBO series *Silicon Valley* where a character tries to show off by purchasing a $14,000 smart refrigerator that can identify when it's running out of beer and when someone is about to put expired yogurt on one of its shelves. Hilarity ensues when another character hacks into it and sets the start-up screen to an inappropriate and looping video. The absurdity is a welcome relief for skeptics like us who bristle at the breathless media coverage of "smart" gadget rollouts as paving the path towards interconnected utopia. Unfortunately, techno-social engineering, the main subject of this book, is no laughing matter. How we engineer ourselves and are engineered by others is one of the most important questions of the twenty-first century.

The companies, organizations, and institutions that use and design smart technology are our leading techno-social engineers. They seduce us by promising smart tools will make our lives easier and better. But like all narratives about pure progress, this isn't the whole story. As we collectively race down the path toward smart techno-social systems that efficiently govern more and more of our lives, we run the risk of losing ourselves along the way. We risk becoming increasingly predictable, and, worse, programmable, like mere cogs in a machine.

Critics often claim that new technologies dehumanize, especially in recent decades with the widespread adoption of computers, the Internet, and, more recently, smartphones. But the public generally takes such claims to be alarmist, and so the claims remain untested and ultimately drowned out by rampant enthusiasm for new technology. Yet techno-social engineering of humans exists on an unprecedented scale and scope, and it is only growing more pervasive as we embed networked sensors in our public and private spaces, our devices, our clothing, and ourselves.

To get a clear sense of where on the path we are, let's play a game, the type that philosophers call a thought experiment. Imagine that an "evil, tech-phobic monarch" forced everyone to stop using products

and services from the major technology companies: Amazon, Apple, Facebook, Microsoft, and Alphabet (the parent company of Google), a.k.a. the "Frightful Five."[1] No more Instagram. No more email. No more searching the Internet. If you had to stop cold turkey, you wouldn't like it, would you? It might feel like the end of the world – the technological apocalypse.

But how, exactly, would our lives "deteriorate" by pulling out these cords?[2] It's hard to imagine the specifics because we depend so deeply on each one of these companies. Consider Amazon, the supreme retailer whose stock "has been rising at nearly 42% a year."[3] Amazon began by selling books online and offering customers a simple way to share book reviews and get automated recommendations for books we might like to read. Then it expanded and expanded some more. And the company just keeps on creeping along, pursuing "dominance, comprehensiveness, and the pursuit of monopoly,"[4] edging ever-closer towards a "United States of Amazon."[5]

Without Amazon, we'd lose one-click, fast delivery ordering of everything from diapers to breakfast cereal – purchases that are so easy to make that there's effectively "no thinking required."[6] If Amazon couldn't deliver vast libraries of streaming television shows, movies, and music to our desktops, laptops, phones, and tablets, the loss of entertainment would sting. And if, suddenly, we couldn't walk into brick-and-mortar stores and compare their prices to Amazon's, we'd feel like bargain shopping died.

It's become clichéd to say that the future is already here but not evenly distributed, but Amazon proved that this is so back in 2014. That's when the company drew headlines for acquiring a patent for "anticipatory" shipping. This is exactly what it sounds like – a patent for a system that can predict what customers want to buy before they even know they want to make the purchases. Amazon's goal is to "box and ship products it expects customers to buy pre-emptively, based on previous searches and purchases, wish lists, and how long the user's cursor hovers over an item online."[7] Amazon's "significant" stake in cloud computing – essentially running and renting space for other online businesses – means that the company is prepared to "power the public infrastructure that keeps the world running," once self-driving cars go mainstream and run on smart grids that are "underpinned by cloud computing networks."[8] In light of all that Amazon offers, does, and will do, *New York Times* technology writer Farhad Manjoo characterizes the company as his "keeper of lists, a provider of food and culture, an entertainer and educator and handmaiden to my children."[9]

When Manjoo describes his personal experience of getting sucked into Amazon world, he notes that the vortex intensified significantly when the company rolled out the Echo. Echo is a hands-free device. It interacts with users through a digital, voice-activated assistant named Alexa that "is designed to get smarter every day" by "adapting to its users' speech patterns, vocabulary and personal preferences."[10] All Manjoo needs to do is ask, and Alexa will perform a range of tasks for him. She'll look up the weather for him, turn on his favorite music playlists, and place his Amazon orders. And that's just the beginning. Manjoo notes that Echo has become such an integral part of the "most mundane moments" of his day that the device is "well on its way to becoming" his "household's brain, a kind of butler in the sky that runs the place for" him.[11]

Notice what Manjoo is saying about how Amazon instills a can't-live-without-you mindset. Alexa directly mediates Manjoo's everyday experiences and habituates him to think and act in collaboration with the device, and Alexa persuades him by design to fundamentally change how he performs household tasks and makes consumer choices.[12] Manjoo will get some benefits from this "relationship," but he probably won't recognize all the subtle and profound ways that Alexa is programmed to program him. The folks at Amazon knew exactly what they were doing when they gave Echo a human name and a human-sounding voice. These are two anthropomorphic features, giving the illusion of humanity. And as the research shows, both incline people to bond and empathize with inanimate technology.[13]

While the "brain" and "butler" comparisons suggest that Manjoo is using a networked device that is, at once, both master and servant, the reality is that Alexa doesn't present evenly balanced powers. What Manjoo identifies is the beginning of a path where powerful companies use smart technologies to gain control over us by framing our choices and nudging us towards programmed lives of convenience and cheap bliss. Cheap bliss is addictive. If it weren't, you could stop eating after you had exactly one potato chip. Or one bite of any of the other foods that are engineered to get us to come back for more, and more, and more . . . And so, Manjoo appears to say that if technology companies can deliver cheap bliss by optimizing his life, he's all for it. He's even willing to pay for their services with agency and self-determination.

Manjoo's desires are not unique. We are all like Manjoo. Consumer demand for various kinds of digital assistants is growing, and during the much-touted Amazon Prime Day, we considered purchasing the deeply discounted Echos.[14] Hal Varian, chief economist for Google, goes so far as

to declare: "Centuries ago, rich people had servants, and in the future, we will all have cyberservants."[15] Apparently, in the future everywhere we go, technological valets will track and assess our behavior, steer us towards our anticipated goals, and take care of our predicted needs.

You might well wonder, what's the harm in technology companies making shopping easier for us? Or making it easier for us to communicate with our friends? Or making it easier to get valuable information like directions for how to get to a meeting across town during rush hour traffic? These all seem like good things that enhance our lives. That's why it would feel catastrophic to lose the technological services that we've grown accustomed to. At the same time, however, we're being sold a misleading vision of cyberservants and digital assistants. These tools don't just do our bidding. They're also smart enough to get us to do theirs.[16]

Our discussion of Amazon reveals a piece of a larger puzzle, a blueprint for building a world that's filled with ubiquitous smart programming. Such a world will be dramatically different from our own. And that's why it's important to take a step back and critically consider the human-level implications of being programmed by the environments that are being designed for us to live, work, and play in.

Such programming was on full display during the 2016 US presidential race, in what's come to be known as the fake news election.[17] While it remains debatable just how much fake news helped Donald Trump get elected, one thing is certain: propaganda campaigns let loose highly automated networks of social media bots. The software posed as real people – regular folks offering earnest, special-interest-free, political opinions – and masked their real agenda of being tools designed to sway votes and circulate calculated talking points. Even though disinformation campaigns have been going on for a long time and attack ads have become a political staple, the bot situation is especially troubling. In a polarized world, when bots are designed to look and sound like us, our neighbors, and our friends, it can be hard to know who – or better yet, what – is engineered to follow a deviously programmed script. This problem, the growing hold Amazon and other technology giants have on us, and many other related issues in the personal, social, and political spheres all concern *twenty-first-century techno-social engineering*.

Techno-social engineering refers to processes where technologies and social forces align and impact how we think, perceive, and act. That's the "techno" and "social" components of the term. "Engineer" is quite close in meaning to "construct," "influence," "shape," "manipulate," and "make," and we might have selected any of these terms.[18] After due consideration, however,

"engineer" won out for two reasons. First, the practice of engineering is directed at designing and using tools to achieve desired ends. Second, the term "engineer" lends itself to analysis that draws parallels between designing environments and designing the people who live in them.[19]

Techno-social engineering has many components. An especially potent one is surveillance. We live in a surveillance society now, and while some people, groups, and even nations resist, most of us are being conditioned to accept surveillance expanding in scale and scope. Business leaders, policy-makers, and consumers are clamoring for a world with smart technology embedded in everything. And that world can't function without always-on people interacting with always-observing, always-analyzing, and always-acting technological systems.

Consider a few examples of techno-social engineering from your everyday life. Have you ever been relentlessly pursued by targeted advertising across the Internet – perhaps a pair of shoes or a jacket that you once considered buying pop up wherever you browse and won't go away? That's done to wear you down. The more you need to exercise will-power when considering whether to buy something, the more your will-power depletes.[20] Or, have you ever clicked "I agree" and accepted the terms of service for online contracts that you didn't bother reading? We all have. Those contracts are designed so that there's no point in reading the fine print. See it, click it, stimulus-response. Or, have you ever been in social situations where you shouldn't check your phone but you do because you just can't help yourself? That's addiction by design, and it cuts both ways.[21] It also accounts for why other people annoy us when they can't leave their digital tethers behind.[22]

Then there's social media. Ever intend to bare your soul or engage in a reasoned debate but end up sticking to the widely used expressions that the interfaces promote – clicking "like," "retweet," or "heart" instead of formulating more thoughtful responses? We have. And that's because social media platforms are optimized to get users to communicate this way. The platforms profit from this style of communication.

Let's not forget the games. Billions of dollars are spent each year on mobile games that are free to download. Free to download, however, doesn't necessarily mean free to play. Gamers pay with their time, attention, and data. They make in-app purchases and get sucked into playing during the time programmers select when they heighten their control over players by limiting when special rewards and challenges are offered.[23]

These experiences and many others reveal that powerful techno-social engineering is occurring everywhere and that a common theme runs throughout them: We are being conditioned to obey. More precisely, we're being conditioned to want to obey.

One extreme scenario that's worth considering is that the smart programming of the future will require us to automatically accept the shots that algorithms call. Perhaps the only way we'll be able to do all the things that smart systems require will be for humans to accept a new lot in life and behave like simple machines. That's the dark side to twenty-first-century techno-social engineering.

Should such a future arise, it will be a long way off. But before the programming deepens, it's crucial to get a clear sense of how decisions that are made today can impact the world of tomorrow. Conventional wisdom says we've made tremendous technological progress in the past century and that it's been driven by the rational behavior of producers and users who develop, deploy, adopt, and use innovative technologies to satisfy consumer preferences and pursue happiness. The conventional wisdom obscures the truth and engineers complacency.[24] Our preferences are increasingly manufactured rather than freely adopted, thanks to techno-social engineering calling the shots. The worst, perhaps, is yet to come.[25]

Welcome to the Experience Machine n.0

Farhad Manjoo's thought experiment about how contemporary technology companies are shaping our values reminds us of a different thought experiment – one that the philosopher Robert Nozick first articulated over forty years ago, long before the invention of the commercial Internet.[26] Nozick didn't seem to have much interest in being a futurist. As an exercise in theorizing about well-being, he wondered whether he or anyone else would choose to plug into a hypothetical "experience machine" that could convincingly simulate any desired experience. In the blink of an eye, the experience machine would let you take on the role of a renowned novelist, a caring father, an ascetic saint, or any of the myriad of other possibilities, like rock star, brilliant scientist, or world-class athlete.[27]

Nozick seemed to imagine the experience machine as a huge mainframe computer. By now, it seems safe to say that he envisioned the wrong type of machine. If a contemporary experience machine were to be built, it wouldn't be anything like a 1970s-era mainframe computer that one plugs into with a cord.

Nozick wasn't far off in other respects. He imagined neuropsychologists would supply us with the sensations we desire and ostensibly crave. Today, technologists, entrepreneurs, and policy-makers are importing scientific insights about how minds work and can be manipulated into their engineering projects and business plans. Knowledge from cognitive science, psychology, and behavioral economics guides how technologists design contemporary computer programs, architect technical systems, and create human-computer interfaces.

Extrapolating from the present to the near future, trends point toward the possibility of creating distributed experience machines, comprised of interconnected sensor networks and big-data-driven automation of socio-technical systems around, about, on, and in human beings. In the final iteration, the distributed experience machine would be ubiquitous and all-encompassing. In this imagined future, our entire environment would be a host of interconnected experience machines, what we'll call Experience Machine n.0 for short. Deployed and integrated incrementally over decades, people will be gradually prepared for and conditioned to accept how it reshapes our entire world and ultimately us.

If the Experience Machine n.0 strikes you as unrealistic, remember we're using it as a metaphor. It represents the combined effects of several real technological developments – all of which are gaining momentum today. We're not claiming that an actual variation of Nozick's thought experiment will be built. The dynamic relationships between social and technological tools and the complex systems within which they are nested and deployed are not easily reduced to a linear series of cause and effect relationships.[28] Nevertheless, reports ranging from the White House's "Internet of Things: Examining Opportunities and Challenges" to the Pew Center report "The Internet of Things Will Thrive by 2025" suggest that the Experience Machine n.0 metaphor dovetails closely with projected projects and scenarios.

Nozick invented the thought experiment to challenge hedonism. This theory stipulates that what matters most in evaluating the quality of our lives is our subjective experience of happiness. Many who have engaged his hypothetical have assumed people would only enter the experience machine if they freely choose to – that is, if they willingly embraced hedonism. The presumption of choice, however, deserves more scrutiny in the context of the Experience Machine n.0. It's hardly a "choice" to plug in anymore. It's almost a practical necessity. Fighting for the freedom to be off will be one of the most important battles of the twenty-first century.

How could the Experience Machine n.0 get built? In an essay titled "Utopia?" we identify several pathways.[29]

- One possibility is a *slippery slope*. Slippery slope refers to the process by which incremental steps down a sloped path can lead to tipping point – a slip and fall, so to speak.
- Another possibility is *engineered complacency*. Engineered complacency refers to one of the mechanisms for accelerating slippage down the slope. If we're engineered to avoid critically questioning innovation, it's hard for us to pay attention to whether change accords with values we deem important or to deliberate about strategies for avoiding change that threatens our values.
- Another possibility is the *aggregation of trillions of perfectly rational choices*. The aggregation of trillions of perfectly rational choices refers to the idea that the incremental steps we take down the slippery-sloped path often will be perfectly rational when evaluated one-by-one on their own seemingly independent terms. This frame evokes the tragedy of the commons, which we'll revisit momentarily.
- Yet another possibility is the *ubiquitous deployment of "smart" techno-social resource management systems for the purposes of maximizing human happiness at minimal social cost*. This possibility links means with ends specifying what type of infrastructure could support Experience Machine n.0.[30]

Each possibility captures part of the techno-social engineering story. Collectively, they highlight the key features of the path we seem to be on.

Humanity's Techno-Social Dilemma

Let's consider in more detail how the path towards Experience Machine n.0 could be fueled by the aggregation of trillions of perfectly rational choices. A helpful comparison is the tragedy of the commons, a famous environmental allegory. In ecologist and philosopher Garrett Hardin's original formulation, the tragedy of the commons involves a dilemma faced by a community of sheep herders who share a common pasture. The herders create a disaster by thinking and acting selfishly. Each one wants to use limited land to feed her own sheep. And so, each individual proceeds under the assumption that it's rational to increase the size of her own herd to capture the benefits of a pasture that everyone shares while only bearing a fraction of the costs that accrue as the common resource gets

exhausted. These externalized costs add up, however, and over time the mad rush for resources leads to massive depletion.

Many believe that things could work out differently if the herders adopt a different outlook. To avoid disaster, they need to better understand their relationships to each other and their shared resources and develop governance strategies for cooperatively bringing about sustainable well-being.

The tragedy of the commons is shorthand for describing many problems that involve a shared resource, a lack of governance, rational, selfish behavior, external costs, and incremental individual actions that aggregate over time to disastrous, often irreversible, social consequences. One of the most pressing examples is climate change – a super-sized, global tragedy of the commons. Remarkably, it has taken decades for the public to appreciate that a large-scale climate change problem exists that humans bear responsibility for creating. Despite widespread scientific consensus for years, the mainstream media only recently have come around to gloomy portrayals of our greenhouse gas crisis. How to understand the relationships between key factors and how to respond to the problem remain highly contentious works-in-progress.

In the context of techno-social engineering of humans, we're calling the tragedy-of-the-commons-like problem *humanity's techno-social dilemma*. Like climate change, there are an incredible variety of small-scale decisions we each make about technology that seem, on their own terms, rational and unproblematic. Yet the increments aggregate, and, like individual herders who need to decide whether to add another sheep to their flock, we suffer if we fail to account for the systemic effects of our decisions, including the production of negative externalities and the impacts on ourselves and future generations.

Just because techno-social engineering is old news doesn't mean we've got a handle on it. Think about our dependence on carbon-based fuels. Relying on them has induced status quo bias (the tendency to accept how things currently are) and made it hard for many people to acknowledge that climate change poses an existential threat. It is hard to accept that lifestyles, industries, and politics need to change. Similar things can be said of technology and techno-social engineering.

One of the ways that humanity's techno-social dilemma differs from the tragedy of the commons is that we're frequently unsure if the problems associated with techno-social engineering are being imposed on us, whether we're electing to behave in short-sighted and insufficiently reflective ways, or whether both factors are in play.[31] Companies, institutions, and designers regularly treat us as *programmable objects* through hyper-

personalized technologies that are attuned to our personal histories, present behaviors and feelings, and predicted futures. Although some finger wagging at powerful corporations is justified, let's not fool ourselves into believing we're innocent victims. The overly simplistic "us vs. them" dichotomy is an ideological trap. There's not always a bright line dividing either, and even when there is, we can't blame "them" fully. We're at fault, too. We choose to participate or choose not to choose and simply follow laid out plans as our default orientation. We adopt technology and mindlessly bind ourselves to the terms and conditions offered. We carry, wear, and attach devices to ourselves and our children, maintaining a connection and increasing our dependence. In doing so, we leash ourselves. As we feed on incremental satisfactions, curiosities, updates, and attention, we treat ourselves as grazing sheep and make ourselves more susceptible to conditioning. We outsource memory, decision-making, and even our interpersonal relations, among many other things. In constructing many different aspects of ourselves, ranging from intelligence to fitness, attentiveness to sociality, we rely on the techno-social engineers' tools to train ourselves, and, in doing so, let ourselves be trained. We both herd ourselves and get herded by others.[32]

Take social robots – think of an embodied and upgraded form of Alexa. When they go mainstream, our new "companions" will engage in highly intimate forms of techno-social engineering by inviting us to change our habits and altering how we relate with others. Will those changes be good or bad for us? It's hard to know without possessing a framework for identifying the central components of techno-social engineering and evaluating some of its normative consequences. We create that framework as we analyze the fundamental ideas associated with techno-social engineering, develop a theory about what makes contemporary techno-social engineering more troubling than previous versions, and propose tests to measure the impact of techno-social engineering upon our capabilities and dispositions. Finally, we offer suggestions for how to minimize undesirable techno-social engineering in the age of smart systems.

The Structure of the Book

This book is divided into four parts. In the first part, we use contemporary observations, thought experiments, and theoretical analysis of creep phenomena and slippery slope arguments to reflect on why it's so hard to understand techno-social engineering and come to grips with humanity's techno-social dilemma. Some reasons concern the difficulty of identifying