USER FRIENDLY

How the hidden rules of design are changing the way we live, work, and play

Cliff Kuang with Robert Fabricant

MCD
Farrar, Straus and Giroux
120 Broadway, New York 10271

Copyright © 2019 by Cliff Kuang and Robert Fabricant
All rights reserved
Printed in the United States of America
First edition, 2019

Library of Congress Cataloging-in-Publication Data
Names: Kuang, Cliff, 1978– author. | Fabricant, Robert, 1966–
Title: User friendly : how the hidden rules of design are changing the way we live, work, and play / Cliff Kuang with Robert Fabricant.
Description: First edition. | New York : MCD/Farrar, Straus and Giroux, 2019. | Includes bibliographical references and index.
Identifiers: LCCN 2019022184 | ISBN 9780374279752 (hardcover)
Subjects: LCSH: Product design—Social aspects.
Classification: LCC TS171.4 .K83 2019 | DDC 658.5—dc23
LC record available at https://lccn.loc.gov/2019022184

Designed by Richard Oriolo
Illustrations by Chris Allen

Our books may be purchased in bulk for promotional, educational, or business use. Please contact your local bookseller or the Macmillan Corporate and Premium Sales Department at 1-800-221-7945, extension 5442, or by e-mail at MacmillanSpecialMarkets@macmillan.com.

www.mcdbooks.com · www.fsgbooks.com
Follow us on Twitter, Facebook, and Instagram at @mcdbooks

10 9 8 7 6 5 4 3 2 1

For my wife and daughter

—CLIFF

To my family and friends
Hopefully this will explain once and for all
what I do every day, and why it matters

—ROBERT

Contents

Introduction:
The Empire of User-Friendliness 3

Part I: Easy to Use

1. Confusion 15
2. Industry 49
3. Error 75
4. Trust 99
5. Metaphor 129

Part II: Easy to Want

6. Empathy 161
7. Humanity 187
8. Personalization 215
9. Peril 247
10. Promise 279

Afterword:
Seeing the World Through User-Friendly Eyes 301

Appendix:
A Brief History of "User Friendly" 331

Notes 349

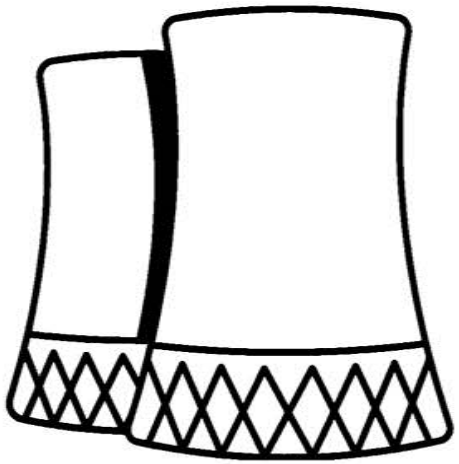
Bibliography 373

Acknowledgments 387

Index 393

Part I

Easy to Use



Three Mile Island cooling towers (1978)

1

Confusion

Wednesday, March 28, 1979: The worst nuclear accident in American history begins with a clog in the basement at a lonely hour. One of the specialists, Fred Scheimann, huffs from the control room down eight flights of stairs into the sprawling guts of the Three Mile Island facility. Scheimann knows every pump and pipe and gauge that he passes while crossing the basement's central walkway, which stretches nearly the length of a football field. He finds his way to Tank 7, where his night shift workers have been gathered since before midnight. Then he clambers on a giant pipe running along its flank so that he can peer in through the tank's sight glass. It's jungle hot, noisy with clanking pumps and sighing valves. Far down on the

other side of the complex, the plant's five-hundred-ton turbine—a block long and spinning thirty times a second—issues a piercing squeal.[1] Scheimann takes his glasses off to get a better look, then mops his brow. *This goddamn clog.* "Hey, Miller . . ." he calls out. But before Miller can answer, the entire assembly begins to shudder. The gathered men all feel a surge of water, "like a freight train,"[2] barreling through the giant pipe Scheimann is standing on. Scheimann hurls himself off just before the pipe bursts from its moorings, then cracks, showering the spot where he'd just been standing with a geyser that would have stripped off his skin.

Still, this is a tiny leak. A plant like this has been designed to protect itself, and the men can hear the plant's thousands of subsystems shudder into action. Hundreds of yards away in the heart of the complex, the reactor core begins an automatic shutdown. High above them all, the cooling towers, thirty stories tall, release a million pounds of steam into the predawn skies above the slow-moving Susquehanna River. A farmer across the river in Goldsboro later recalls stopping short under the light of his barn, listening to what sounded like the whoosh of a jet engine.[3]

Scheimann picks himself up off the ground and then runs back to the control room at a dead sprint. The place is laid out like the bridge of a ship. In fact, almost everyone in the room is ex-Navy, having served stints aboard nuclear submarines or aircraft carriers. In the center of the room is a massive console; behind that, there's another wall of control panels arrayed in an arc ninety feet long and reaching up to the ceiling.[4] All told, there are eleven hundred dials, gauges, and switch indicators, and more than six hundred warning lights. At this moment, it seems like every one of them is wailing. The room is buried in noise. Here, at a critical moment, the machine is generating not just noise but chaos in the minds of its operators. The taint of that chaos will linger for hours.[5]

What on earth does any of it mean? How do you find the one thing wrong when the system is telling you there are hundreds? Scheimann starts riffling through the emergency manuals, making sure that every procedure is followed to a T. Reactor trips are a pain, but they aren't uncommon. With hundreds of fail-safes layered atop one another, the chances of a meltdown seem vanishingly small. Barring human intervention, the plant will shut itself down at any hint of danger. But you shut down one light and another goes off; the way it's all designed—or isn't designed—makes it impossible to imagine how it's all linked, or how one missed signal can cascade.

Every single one of the reactor's systems is designed to serve two abiding purposes: to either create heat or contain it. The core itself is made of thousands of uranium pellets the size of a finger. The heat flows from the uranium as its atoms split apart, spewing heat and neutrons. The neutrons cause more uranium to split, powering a chain reaction that grows upon itself exponentially—the power of that chain reaction is such that each pellet can produce as much heat as a ton of coal.[6] While all that happens, the heat must be controlled, and that requires massive amounts of cool water, flowing over and past the reactor, carrying heat from the core, driven by pumps two stories tall that are powerful enough to reverse the flow of the Colorado River. The water flows over the core, transferring the heat and carrying it away. That water, now hot, is used to make steam; the steam spins a giant turbine, which generates enough electricity for a small city.

The men up in the control room first make sure that the core has all the water it needs by switching on pumps and monitoring boilers and turbines. Then something strange happens—a critical slippage between what the men understand and what the machine is telling them. The water level in the pulmonary loop that cools the reactor is actually going *down*, even though the emergency pumps are going

full blast. Scheimann is still at the manuals, shouting out every protocol one step at a time, then nodding as someone shouts back that it's been done. Then the water level stops sliding. It starts to hold. The emergency pumps seem to finally be filling the system back up with water. Relief washes across the room. Minutes later, that relief evaporates. There has to be pressure in the system. That's what tells you there's enough water. But there's a Goldilocks range, and the reactor loop seems to be shooting past it. The pressure begins to rise slowly at first, then quicker. *What the hell is happening?* To 160 inches, 180 inches, 190, 200. And then a spike: 350, higher than they'd ever seen. Worry starts creeping through the room's cool professionalism.[7]

"Okay, we're going solid!"

This is everyone's greatest fear. "Going solid" means that the reactor loop is filling completely with water, so the pressure will just build until the pipes burst, draining the reactor. Hurriedly, the men switch off the emergency pumps, to keep them from adding more water to the core.[8] It will turn out to be the day's single worst decision.

While all this is happening, the temperature in the core is continuing to rise. That shouldn't be possible. If there's so much water in the system, why isn't the reactor cooling down? *Maybe there's a valve open somewhere, and all the water is simply leaking out?* There's a gauge in the control room that should have the answer. It's hard to find, squirreled away out of sight on the back of a control panel on the other side of the room. The man sent to check the gauge finds it, and sees that it looks fine. But he's looked at the wrong one. So he walks away and tells everyone else that the valves are closed: the system's not leaking water, and it's time for everyone to start over again and look for another answer. No one knows the core is beginning to destroy itself.[9]

Just a few hours later, at 6:00 a.m., the catastrophe has taken on a sickening momentum, though no one seems close to figuring out

what's really happening. Pete Velez walks into the control room of Reactor 2 to begin his usual shift. There's typically only a couple of men there, gliding about among the institutional-green control panels and thousands of placidly glowing lights, enjoying the special silence of supreme order. But today Velez can see that something terrible is in motion. People are everywhere amid the telltale signs of swallowed panic: coffee cups strewn about, stacks of safety manuals piled high, men tearing at them while sweat stains spread under their arms. Velez has encountered some of the bigwigs now milling around only as names atop a memo—regional managers and their managers, in from headquarters in Ohio, all of them trying to decipher what the hell has gone wrong. He plucks a fresh green notebook from his pocket and jots down his first entry of the day: *Aw shit.*[10]



Velez and everyone here knew the dangers that came with working in a nuclear plant. It's what they trained for until the dangers were familiar enough to feel routine. But Velez was also uniquely intimate with a gruesome kind of calculation: As the radiation-protection foreman, it was his job to know exactly how much nuclear exposure a worker could handle. Typically, he shouldn't be exposed to more than 3 rem of radiation in three months. (This being 1979, at a nuclear plant, the workers are almost all men.) Emergencies were different. Say you had to send a man to save some critical piece of equipment. In a month, he'd probably be fine if he'd soaked up 25 rem. No lasting effects. You could justify that tiny personal risk against the chance of a great catastrophe. Any more radiation than that, things got complicated. The rule of thumb was that saving another man's life was worth the risk at no more than 100 rem. Any higher—say 120 rem—and it became a decision no one else could make but you. Could you let someone just suffer there alone?

None of this is abstract anymore when the next day, March 29, Velez finds himself cracking a door to steal a peek inside a room that could kill him. He huddles with Ed Hauser, the chemistry foreman in charge of monitoring the water that constantly cools the reactor, saving it from its own ferocious heat. They're both covered head to toe in coveralls, wet suits, gloves, boots, and masks, with every seam taped up. In that split second, he can see that when the hazard klaxons had started screaming, the men on duty had literally dropped what they were doing and ran: hats and coats still on the rack, telephones off the hook, a pot of coffee scalding the table.[11] At the back of the room there are about twenty-five faucets above a hooded sink. That's what Velez and Hauser are here for: to understand just what's happening in the reactor, to understand how bad the situation has become, because instruments in the control room aren't making sense.

No one knows how much radiation is leaking from the reactor core. These twenty-five valves connect to twenty-five pipes running thousands of feet through the building's unseen creases. There is one pipe in particular, no bigger around than a man's finger, connecting to the next building. It is a live wire to the nuclear core, which might already be melting.

Working like this, the men share the exposure and share the risk. Each man does his part, and no one gets more radiation than the other. Hauser is glad for it. Just yesterday, running a test elsewhere in the plant as everything turned to shit, he'd soaked up 600 rem within minutes. Yet he's still here, on duty again, planning with Velez. It becomes clear that Velez doesn't know how all the valves are ordered. Hauser does. So despite his prior exposure, he's the first man in, 600 rem and counting.

Velez looks to Hauser, checks his watch, and marks the time. *Go!* Hauser races into the room, straight for the valves at the sink, tearing open fifteen of them in exactly the right sequence. He's barely got his hand on the last one before he's turning heel, bursting back into the hallway. Now they wait. It takes an agonizing forty minutes for the water from Reactor 2 to travel the thousands of feet to this place and then sputter into the sink. There's still more valves to be turned. Velez has every excuse not to rush in—he doesn't know which ones. But instead, he presses Hauser for a description of the room so that he can take a dose of the exposure, so that Hauser doesn't get any more than necessary. Velez rushes in to turn the last valve. Now it's Hauser's turn to finish up the job. He rushes into the room, up to the sink, and catches a water sample in a vial. It seethes like a witches' brew, yellowed with chemicals engineered to soak up radioactive isotopes. Hauser raises his dosimeter to the sample. It spikes to 1,250 rem: so high that if you touched that vial with a bare hand, your fingertips would tingle.[12]

Miraculously, both Hauser and Velez survived the ordeal. It would be tempting to say that neither would have had to risk his life, but for one misread light on the back of some poorly laid out control panel—or, preceding that, a chain of misunderstanding which simply didn't go as it should have. But when we step back to see what went wrong—when we step back to imagine those eleven hundred dials and six hundred alarms, blaring at once—it's not merely that a machine broke down, or that a human failed to do what he was supposed to. The machine might have been made differently, with some greater awareness of how too much information and too little meaning can overwhelm the humans who are supposed to be in control. Instead, the machine and the human couldn't speak to each other in a language that each could understand. They were opposed in ways that no one in the moment could appreciate. The story of that opposition carries on today.

The only reason I'd ever thought to delve into the history of Three Mile Island was a hunch: that when you look hard enough at monumental machine disasters, you can usually find a design problem. It's almost always the case when planes crash. In fact, a misread signal at the worst possible time was responsible for the burning of Notre-Dame in 2019: A state-of-the-art fire system with inscrutable controls led to a bungled inspection while the blaze grew unchecked for thirty minutes.[13] Disasters always mirror the way things *should* work. So what would Three Mile Island reveal about the ways in which humans and machines should interact?

I expected to tell that story and why it mattered through analogies and metaphor—rhetorical sleight of hand, if I'm honest. But then, buried in a report about the disaster, I found a passing reference to another investigation into what had gone wrong, commissioned by Congress and coauthored by one Donald A. Norman. Could it be *that* Don Norman, the guy who'd actually invented the term "user experience" in the 1990s?[14] Seeing Norman's name mentioned in that report suddenly made it seem that the tenuous thread connecting Three Mile Island to the problems of the present day was instead a steel cable, buried but already in place.

In the era before user experience came to define digital life in the twenty-first century, Norman was the Moses of product design: In 1988, he published perhaps the only mainstream bestseller about the field, *The Design of Everyday Things*, which documented all the ways in which the fodder of everyday life failed us, from door handles to thermostats. His books became desk references for a generation of interaction designers. He'd tried to retire in the early 1990s, before Apple lured him there. He started by creating a panel of usability gurus, whom he dubbed "user-experience professionals" and who were meant to track every product as it was developed; in doing so, he became an early champion of the recently hired Jony Ive, who would go on to design the iPod, the iMac, and the iPhone.[15] Yet when I thumbed through Norman's books and all their footnotes, I found only passing references to nuclear reactor design—and none of them seemed to mention Three Mile Island at all. How had that catastrophe shaped the godfather of modern design?

Norman is slight, around five foot three, with stooped shoulders and a thin waist and a daily uniform of black turtleneck, jeans, and a gray newsboy cap. He stays fit by walking from his home to his office, up and over the steep hills that run across the UC San Diego campus, which straddles a series of scenic gorges. I visited him at the Design Lab, which he'd founded, on a typically warm, sunny December afternoon when the air was bright with the scent of eucalyptus and wild rosemary.

We were sequestered in a tiny meeting room with self-consciously kooky green shag carpeting. I was seated in a low-slung patio chair as Norman stood above me, pacing as he warmed up to the rhythm of a lecture. He began to describe the project that was to be his last great work, begun just six months prior. It was his second time coming out of retirement, and he would soon turn seventy-nine, on Christmas Day. "The university, you see, is filled with people who analyze in great depth," he said in his high, gnomish voice. "Designers don't analyze, they put together. This lab is an opportunity to put together all the knowledge at this university to solve problems in the environment, aging, health care. These are the kinds of problems we want to solve."[16]

I looked out into the lab beyond: just a few desks so far, filled by a handful of graduate students tapping at lines of code. The Design Lab held a prime location on campus, at the corner of a shiny new

postmodern building. Like a lot of newer college campuses, UCSD is an open-air timeline of modern architectural fads, beginning with the library, done in the brutalist style popular in the 1960s, when the university was established, and continuing through a few ironic studies in classical forms that marked 1980s postmodernism. This building, composed of jagged bands of steel and glass, was the newest wave.

Norman's sprawling vision for design's power doesn't belong to him alone—it has, in fact, spread among the designers themselves. "I was in Shanghai visiting Frog, where they were very proud to tell me that their firm owned product design," Norman says. "Then I visited IDEO and told them what their competitor had said. And they said, 'We don't care. Singapore came and asked us to design their whole city.'" That wasn't a boast: Today, "design thinking"—the processes that inform modern design—has spread far beyond the design firm IDEO, which was a pioneer in marketing the movement. Design thinking is now marshaled to solve myriad problems at every scale. What was once a niche profession more commonly associated with chairs is now talked of as a solution to the world's ills, simply because of a shift in perspective.

Norman has a tendency to insert long pauses between thoughts or before answering questions—the kind of thing you grow comfortable doing only after many decades of people hanging on your every word. When I finally had the opening, I asked him, What do you remember about Three Mile Island? He described it as a break in his career, between the wonky academic research he'd been doing and the wider world. Early in his career, Norman spent years classifying the many ways people err when a task is set before them. What he discovered at Three Mile Island was notable in that it revealed just how little *other* people seemed to know about what he'd been working on. "The problem was that they spent so much time

designing the technical parts, and none on understanding what it was like to work there, what was going on for people," Norman recalled. "The control room was done last, almost an afterthought when there wasn't time or money left."

He told a terrifying story of just how ingrained that myopia was: Reactors were almost always built in pairs; at some point, someone had realized that rather than customize two separate control rooms, it was cheaper to build one and then build its mirror image. Thus, the staff would have to work one day in one control room, and the next in a literal bizarro world, where everything was reversed. Those examples and others made Norman "realize that there wasn't any understanding of technology combined with psychology. We were building technology for people, but the technologists didn't understand people."

That myopia was reflected in the culture at large—a fracture between academics like him, who'd been studying how humans used the machines around them, and the people who created those machines. "The really good work in human error started in the Second World War, but it didn't become a major item for everyday people. Researchers like me didn't know who else was doing the same work," recalled Norman. And meanwhile, the designers "used to come from art schools or advertising, so it was all about style without any substance." Norman didn't know any designers himself at the time; he was an outsider to a profession that had leaped ahead, stumbling into a new world without a compass or a guide. Norman's books thus adopt an abiding tone of bemusement: *Dear God, will these people please listen to me?*

Norman's grand emphasis on such complex problems—the environment, period—belied the fact that he was famous, in large part, for the modesty of his insights as a design guru, thinking about doorknobs and teakettles. In his books, Norman is like Job, always

being tested by some uncaring design god. He pushes on doors when they should be pulled, he has a hard time turning on the lights in his house, he's constantly being scalded in the shower. But in his thoroughgoing confusion, he's us.

The most consequential assumption behind all his work is that even if human error is to blame, it is hard to imagine any human not making these errors. Humans might fail—but they are not wrong. And if you try to mirror their thinking a little, even the stupidest and strangest things that people do have their own indelible logic. You have to know why people behave as they do—and design around their foibles and limitations, rather than some ideal. His great insight was that no matter how complex the technology, or how familiar, our expectations for it remain the same. Norman's discipline, cognitive psychology, wasn't so much about the nuances of buttons and control panels—though there's plenty of that, if you want to look—but rather the ways in which humans *assume* their environment should work, how they learn about it, how they make sense of it. This is what you have to understand if you are to design an app that people can use the first time they try it, or a plane that humans won't crash, or a nuclear reactor that humans can't cause to melt through the continental shelf.

All these lessons might have remained the obscure province of professors and anonymous designers and engineers if not for a coming wave of technological change: the profusion of computers and electrical gadgets in our everyday lives, driven by the rise of cheap transistors and silicon. Starting in the 1980s, the complex problems you might have found at Three Mile Island became consumer problems, having to do with making buttons work on gadgets such as VCRs and computers; the nuances of designing such devices went on to be expressed in the smartphone. It's no surprise

then that the reasons a bad app drives you crazy have a direct relationship to the reasons that Three Mile Island almost melted into the earth. The problems that caused Three Mile Island are similar to the ones that frustrate you when you're trying to turn off the notifications on your smartphone; the inscrutability of a poorly designed light switch shares the same cause as your inscrutable cable box: a button that seems misplaced, a pop-up message that vanishes before you can figure out what it means, the sense that you did something but you don't know what. The presiding notion that you don't know how something works.

It was perhaps only natural that as the smartphone came to take over our everyday lives, the principles that had created it would come to seem like the answers not just to problems of the moment (How do I get people to understand this app?) but to problems of the era (How do I get people to understand their health care?). It made perfect sense if you believed that all these problems came down to the way that the machines failed the people who used them—and knowing that those failures revealed a truth about how people made sense of the world around them, and how they expected the artifacts of everyday life to behave.



Back to March 28, 1979. The graveyard shift was nearly over and the sun just starting to rise when the men finally solved what they thought was the biggest problem—too much water flooding the system, going solid. To stop that, they had made a fateful choice: to shut off the emergency water pumps. Once those pumps were off, they watched as the water levels fell in the pressurizer that fed the reactor's circulatory loop. Relief once again filled the room. Once again, it didn't last. Even though there seemed to be so much water

in the system that it threatened to burst—so much water that the reactor should have been cooling down—the temperature just kept creeping up with a sickening steadiness.

As you watch this disaster movie unfold in your mind's eye, let the camera open upon the reactor control room, tracking slowly to the control panel at the center. Then let the camera pan across all those blazing warning lights. And then let it stop upon just one of the lights, a great big red one with a tag taped below it, saying exactly how to read it. It's one of the few lights along the panel that isn't lit, and that's a good thing. The men must have glanced over at it hundreds of times just to make sure. But that light was lying.

Its importance came from what it was supposed to be connected to—a manual release valve at the top of the reactor. The valve works like the whistling spout of a teakettle, venting steam whenever the pressure inside the reactor gets too high. If it was open, then that would mean the reactor had a massive leak at the top. Yet, as investigators later learned, the so-called PORV (pilot-operated release valve) light was designed around a deep conceptual error: it turned off when someone flipped the switch controlling the valve—not when the valve *actually closed*. Put another way, the light was merely marking intent, not action. If the light was off, that might mean the operator had done the right thing, closing the valve—or it might mean that the operator had done the right thing, but the switch wasn't working.[17] The misdesigned light could only ever say things were just fine.

In fact, the reactor's circulatory system had a huge hole in it, and no one could know, simply because the switch's feedback was meaningless. As the men called for help from more and more people across the country, and as the country began learning that something was going wrong at Three Mile Island, the temperature in the core kept rising. The machines had reached a limit to what

they could reveal. The computer readouts, relaying the core temperature, stopped at 700 degrees. Now all they were saying was "???"[18] The systems were mute about what was actually happening. In fact, the core had reached an astounding 4,300 degrees. At just 700 degrees more, the 150-ton uranium core would have melted, searing though the eight-inch-thick steel containment vessel, then the twenty-foot-thick concrete foundation, not stopping until it hit bedrock beneath the Susquehanna River, blowing radioactive geysers straight into the sky.

It was nearly three hours after the clog in the basement caused the reactor's systems to lurch into a shutdown that the hole was finally plugged, by a man on the next shift who came on with fresh eyes and a hunch about what might have been missed. He'd shut the backup to the PORV, just to make sure. Then, hours after that, one of the system's original engineers finally commanded that the emergency cooling system be turned back on, ending the disaster. It was only later that anyone discovered that a complete meltdown of the reactor core had been as little as thirty minutes away.[19]

The Three Mile Island disaster happened less than two weeks after the release of *The China Syndrome*, a Jane Fonda blockbuster about a cover-up at a nuclear power plant. The movie's title came from the urban legend that an American reactor meltdown could bore a path through the earth's core, to China. Pop-culture fantasy and a real-world disaster together killed the growth of America's nuclear energy industry.[20] Plans for around eighty plants were scrapped; not a single new reactor was approved until 2012.[21] Today, what some experts argue is the safest, cheapest, and most reliable source of renewable energy remains clouded by fear. So, measured by what might have been, it's reasonable to call Three Mile Island the biggest design failure in American history. It's also the most instructive. The failures at TMI are a mirror image of the user-friendly age we live in, a tally

of all the underlying principles that allow smartphones and touchscreens and apps to blend into our lives.

What Norman and the investigative team discovered at Three Mile Island was terrifying because it all seemed so obvious. The catastrophe unfolded over two days. In that time, hundreds of eyes pored over the system. If someone had closed the right valve sooner, or if, at any time, someone had thought to turn the emergency pumps back on, the place would have been saved. These men were not stupid. And yet even today, the few reports on TMI still blame "equipment failures and operator errors."[22] That's not it at all. At Three Mile Island, there were no grand equipment failures. The staff were some of the industry's best, and, incredibly, they never panicked.

In fact, the plant mostly behaved like the beautifully engineered machine it was meant to be. It would have saved itself if the men had left it alone.[23] What happened instead was that the men, thanks to catastrophically bad control-room design, were unable to understand what was going wrong. Swaddled in a fog of misdirection, they made catastrophic choices. The plant and the men were talking past each other: The plant hadn't been designed to anticipate the imaginations of men; the men couldn't imagine the workings of a machine.

Begin with the lights of the control panel. Despite every appearance of industrial precision, they had no abiding logic that a user could understand. Yes, a lit-up red bulb meant that a valve was presumably open. But not every valve was meant to be open or closed. Thus, normal operation was a hodgepodge of conflicting indicators, instead of them all being just one color when things were fine. The investigators who descended upon TMI in the wake of the accident reported that there were *fourteen* different meanings for red, and *eleven* for green. The consistency we now expect in countless

rounded buttons and red warning lights was totally absent at Three Mile Island.

Sometimes the lights were above the control they corresponded to, sometimes off to the side. They weren't even grouped in a way that made sense: On the very same panel that would warn of water leaking from the reactor were alarm lights indicating elevator trouble. It was as if someone had taken a map of the reactor, cut it up into pieces, thrown it into the air, then taped it all back together. With such a map, you'd never be able to navigate. One reason we find apps easy to understand even if we've never used them before is that *navigability* and *consistency* are so ingrained into the patterns of app design today. Menus all largely behave the same way; so do swipes and taps.

There was also no indicator that would tell the men that the circulatory system was empty, like the gas gauge on your car. They were so fixated on the first panic—of filling the whole thing with too much water—that their imaginations simply failed. This, too, was a design failing that we now consider a given. When something works well enough for you to predict what it'll do next, you eventually form a *mental model* of it. That mental model can be deep or shallow—it might vary from just a sense that this button does that, to a picture in your head about how your hybrid car charges its battery. But those mental models are knowingly crafted by the designers who put interfaces in front of you.

Meaningless alarms, information clustered nonsensically, no consistency anywhere—these things translated to no mapping, no navigability, no mental models. These are tenets that everyone who owns a smartphone today takes for granted. These are the principles that make the user-friendly world work. You need all these principles in place to master a machine, whether that machine is a nuclear

reactor or a child's toy. For you to master how a machine works, it needs to adhere to a pattern language.

But there was one essential thing whose failure loomed largest at TMI, one essential thing that we demand of any gadget in our lives: feedback. When the light was lying, when the temperature readouts were printing "???," and when there was no indicator telling anyone the system's total water levels, the machine just wasn't telling the men what they needed to know. With every little thing they tried, they grabbed on to the wrong feedback, focusing on the wrong things.

Feedback that works surrounds us every day, so we rarely think about it. It's feedback that defines how a product behaves in response to what you want. It's feedback that allows designers to communicate to their users in a language without words. Feedback is the keystone of the user-friendly world. In fact, the importance of feedback for both mankind and machines was a founding insight of both neuroscience and artificial intelligence. It was pioneered in 1940 by Norbert Wiener, a mathematical genius teaching at MIT. At the height of World War II, the German Luftwaffe had unveiled new warplanes faster than anything that had come before; they bombed British cities with impunity, banking too fast for any gunner to react—retaliatory artillery shells exploded in empty skies. Wiener thought to invent an algorithm that might automatically take radar data about a warplane's position, add in the flight time of an artillery shell, and spit out an anticipated vector where a gun could be pointed. The idea was to identify a brief window in time and space where an attacking plane would probably be, given the incoming radar signals. As new radar signals came in, that window would shift, creating a feedback loop.

Wiener and his collaborator, Julian Bigelow, realized that they'd stumbled onto something bigger. Imagine picking up a pencil: You

form the idea in your head. You start to move your arm. And as you do, your brain must make an infinity of tiny corrections, using your eyes, muscles, and fingertips. Wiener knew, from a neuroscientist friend, that this was precisely what went wrong in certain hand tremors: The brain, having overshot its mark, would get stuck in a ricochet of overcorrections—just as Wiener's equations predicted. Wiener and Bigelow realized that feedback was "required by *any* voluntary action."[24] Feedback is what links the ineffable stuff in our minds—the things we want—with the machinery of our bodies and the information from our environment. As the anthropologist Gregory Bateson later marveled, "The central problem of Greek philosophy—the problem of purpose, unsolved for twenty-five hundred years—came within range of rigorous analysis." Feedback is what allows information to become action—and not just at the level of data, neurons, and nerves.

When you swing an ax to chop a piece of wood, the wood either splits or it doesn't. If it doesn't, you set the wood upright and swing again. When you put your bread in the toaster, you push the lever, and it clicks when you've pressed it far enough to turn the toaster on. Then you hear the filaments start to hum with electric current, a sign that the toaster has in fact turned on. You're getting feedback all along the way that the toaster has done what you'd wanted it to do. There was the click of the button, which had to be designed and engineered. And there was the sound of the wires heating up, which is simply a useful by-product of the toaster's physics. Without all those signals along the way, you'd just be endlessly fiddling, trying to understand whether the toaster was working.

The natural world is filled with feedback; in the man-made world, that feedback has to be designed. When you push a button, does the button actually affect the thing it's supposed to? The world of everyday life is so densely layered with information that it can be

hard to realize how much information—how much feedback—we have to re-create in the world of design. And yet feedback is what turns any man-made creation into an object that you relate to, one that might evoke feelings of ease or ire, satisfaction or frustration. These are the bones of our relationship with the world around us.

Is there any problem in which our behavior doesn't match how we'd like to live that isn't a feedback problem? When we eat too much or eat the wrong things, it's a problem of not realizing in the moment how that tiny choice might affect our future. In the United States, doctors rarely track what happens *after* they've prescribed a drug or procedure, and so they just keep prescribing both to new patients, aiming to try everything since they can't see if any one thing actually works. And so we spend more and more each year on medical costs. Even climate change can be seen as a feedback problem. We cannot see our everyday contributions to carbon emissions, and the timeline is too long for us to see their effects. Imagine if carbon emissions had no other effects than they do now, but that carbon accumulation turned the sky from blue to green. In a world like that, it's hard to believe that we'd still be arguing about whether mankind was having an effect on the climate. We might instead be arguing about what to do about it. These are all problems of not feeling the stakes. Until the worst has happened, there is no feedback about what the effects of our actions are, and by then it's too late.[25] There may be no greater design challenge for the twenty-first century than creating better, tighter feedback loops in places where they don't exist, be they in the environment, health care, or government.

Feedback already defines the world we live in today. For example, we tend to assume that the internet's great revolution was connecting people. That's partly true. But consider the birth of buyer/seller feedback. eBay was an unknown startup until it rolled

out a feature in which buyers and sellers could rate one another. Today, buyer/seller feedback is what has made us comfortable with the online economy—from buying products that we've never seen before on Amazon to staying in the homes of people we've never met, through Airbnb. In a previous era, we used brands to create trust—when you saw a toothpaste stamped with Colgate, you knew it was the product of a big, stable company whose long-term success depended on good products. Today, we have feedback from people who've tried out something we might like; even if you don't know them, you put your faith in there being a lot of them. As the economist Tim Harford has mused, without feedback, internet commerce might not be like it is now, with strangers trusting one another. It might be more like hitchhiking, something done only by people willing to take a risk.[26] Even the biggest startup of the last fifteen years, Facebook, was a company formed because of feedback. The Like button offered nothing less than a new way to send and receive affirmation, and in so doing, it rewired the social fabric of one-third of the world. (We'll see more of what that world has wrought in chapter 9.)

New technology improves the kind of feedback we can get, and how fast, allowing us to be more efficient and to act on new types of information. When you think about futuristic new technologies, you're often thinking about feedback that doesn't exist yet. From slickly designed custom nutrition regimens tailored to your metabolism, to public buses that are rerouted in real time, according to demand—these are all products predicated on bringing new feedback to the market. One of the most significant technologies of the twenty-first century, artificial intelligence, rests on feedback: Put simply, AI and machine learning are a collection of methods that allow algorithms to gauge how well they've performed, and then tweak their own parameters until they perform better. AI's chief

breakthrough was in allowing algorithms to process feedback. (The very first "neural networks" were proposed by Warren McCulloch and Walter Pitts; McCulloch was inspired during one of Norbert Wiener's first lectures on feedback.)[27]

While the goal of most feedback is just to reassure us that something has gone as we expected, there are higher values and needs that feedback can address, whether they be soothing us or making us anxious or spurring our competitive instincts. For example, the Facebook Like button allowed us to attach a number to the loose uncertainty of our social bonds; it created a lighter, more fleeting definition of what counted as a relationship. And differing approaches to feedback lie behind two of the most successful startups in recent history: Instagram and Snapchat.

Instagram came first, in 2010. The app initially just let you share photos and see the photos of friends to whom you were linked. Soon after it launched came the commonsense policy, popularized by Facebook, of letting your friends like your posts (and, of course, seeing how many likes your friends' posts had gotten). That was the simple act of affirmation that the app revolved around. How many likes had you gotten? How many had your friends? Snapchat was built differently. It, too, was meant as a basic photo-sharing app. But it differed in two crucial ways: For one, photos you posted would disappear after twenty-four hours. And two, the photos did not earn likes. If you saw the photos your friend had posted, the only option available was to message your friend in response. Put another way, the only feedback you might get from a photo you posted was a direct conversation with a friend. The idea was that you could share without consequences—that sending a message of yourself looking sad and writing "BAD DAY" was enough, that you wouldn't be judged for it. That it would just be for the people you cared about. That it would just be *you*.

By 2016, Instagram had noticed that its users were talking differently about the product. It had been intended as a service for sharing photos in the moment—but users didn't talk that way about it anymore. They talked about the worry of wanting to curate what they showed other people. They talked about wanting to present the best possible picture. A feedback loop had set in, and reinforced itself: Likes made users more self-conscious about what they were posting, at the same time as improving cameras and Instagram stars were steadily raising the bar of what counted as a great Instagram post. "It started feeling like Instagram was for highlights rather than what was going on now," said Robby Stein, a product manager at Instagram.[28] This might sound innocuous, but it was potentially disastrous. The fear was that self-conscious users, because they posted less and less, would drift ever so slightly away from the app. If you were Instagram, that was a mechanism that could end your business altogether. And so Instagram ruthlessly adopted Snapchat's feedback philosophy, inserting a strip of "stories" at the top of the app, which let users share photos without the possibility of getting likes, and with the only way of replying being to send a direct message. It worked. Instagram Stories was wildly successful, used by 400 million people a day within a year of launch, helping Instagram nudge its average user time ever upward.

You could talk about Snapchat and Instagram as the story of two billion-dollar apps that gave people new ways to entertain themselves, and one another. Instagram had started to become the equivalent of an art gallery. Snapchat had become a spontaneous goof for your friends. The difference between those two businesses was really the story of the differing experiences that feedback could produce. Forty years after Three Mile Island, feedback is more than just what makes machines intelligible. When feedback is tied not merely to the way machines work but instead to the things we

value most—our social circles, our self-image—it can become the map by which we chart our lives. It can determine how the experiences around us *feel*. In an era when how a product feels to use is the measure of how much we'll use it, this is everything.

The scariest things are often the easiest to forget. Sometimes, the forgetting is a reflex meant to keep us safe in our routines. Other times, it's a more willful erasure that carries an agenda. In the case of Three Mile Island, both seemed to be true. Three decades after the accident that nearly ended America's investment in nuclear energy, the industry has been utterly quiet, apparently afraid to draw any attention to itself whatsoever.

In the wake of the accident in 1979, Reactor 2 at Three Mile Island was shuttered and sealed. Meanwhile, Reactor 1 quietly continued to operate for another forty years. After Norman's commission suggested a number of changes to the design of all American reactors, it was retrofitted to be easier to operate. You might think that the very fact of Reactor 1's continued operation would be a success to celebrate, and yet that isn't the case for an industry covered fearfully in the press and misunderstood by the public. I wanted to learn exactly how it had been redesigned—what had been done to prevent the fog of confusion seen at Reactor 2. It took me months to convince the power plant's PR minders that I wanted to visit not to see what remained dangerous but to see what had been *fixed*. Finally, I went—not long before it would finally close, in September 2019.

I approached the Three Mile Island power plant on a wooded two-lane road, where the trees would briefly open up to offer flickers of a riverbank. Then, finally, I saw two gargantuan cooling towers on a tiny island of their own. They were three hundred feet tall, out of scale with everything around. Steam rose from just Reactor 1, in

cottony plumes at the shambling tempo of a work song. Next to it, the tower of Reactor 2 stood quiet and streaked with rust, a dead sentry. It produced an eerie, strangely beautiful effect. Impossibly expensive to tear down, Reactor 2 loomed like a monumental postmodern sculpture. It was the commentary shadowing the reality.

Across the road was the ho-hum, squat brick building where workers trained so that the remaining reactor at Three Mile Island would never again come close to failing. Inside, it was appointed in institutional drab: orange-brown-gray carpet, public-school-style furniture made from chrome and chipboard.

The only reason this plant came to Pennsylvania at all was because of organized crime. In the 1970s, Metropolitan Edison had at first tried to build it in New Jersey. The Mob, which ran thick in the local unions, threatened to sabotage the work site unless it got the customary 1 percent kickback on the total building cost—which amounted to $7 million against $700 million. The power company pressed ahead anyway and began laying the foundations. Then, while a crane was lowering the 700-ton reactor core into the ground, some unknown construction worker dropped a literal wrench into the crane's workings. The message was clear: Pay up, or else the plant would be sabotaged in ways no one would ever know until it was too late. The power company promptly abandoned the site and decamped for a little spit of land in Pennsylvania. As a result, Reactor 2 was reconfigured at Three Mile Island within a mere ninety days, for a site it was never meant to occupy. For those who worked there, Reactor 1 always performed beautifully. Reactor 2 remained a tetchy, temperamental beast.[29] Indirectly, the Mob truly had sabotaged the thing.

The simulated control room inside the training facility—with control panels painted an industrial green and rows of lights

shrouded in protective cowls—resembled a movie set for Apollo Mission Control. My tour guide that day was the man who ran the room, a friendly, slightly built engineer with wire-rim glasses and decades on the job. He wore sturdy brown walking shoes and clothes that were scrupulously beige, making him look like an extra from a period movie about the Space Race. He was in charge of simulating the beginning of the end of the world, to see how the workers would respond. The room was an exact replica, down to the switch, of everything within the control room at Reactor 1.

In the wake of the accident at Reactor 2, a slew of subtle but powerful changes have worked their way across the industry. Standing in the facsimile control room, I could see them all. For one, from the back of the control room, everything was visible—there were no hidden indicator lights to be forgotten behind a panel. The room was easily navigable. The lights here were consistent: When everything checked out, they would all glow blue.

Normal wasn't why we were there, though, and it wasn't why this room was designed. The engineer stole away into an observation room at the back, and then came out to announce that the reactor had shut itself down. Just like it would have when the clog happened at Reactor 2.[30] A bank of lights went off, but the effect was muted. Contained. I pushed the button to quiet everything so that we could isolate what was happening.

For the people who were assembled in the control room at Three Mile Island in 1979, all these problems—erroneous feedback, controls that were inconsistent and impossible to navigate—added up to a greater problem. The men on duty literally couldn't imagine what was going wrong, because the machines wouldn't let them. They had no mental model showing how all these disparate and strange events might be connected, which would have helped them deduce what was going on.[31]

Mental models are nothing more and nothing less than the intuitions we have about how something works—how its pieces and functions fit together. They're based on the things we've used before; you might describe the entire task of user experience as the challenge of fitting a new product to our mental models of how things *should* work. To take one simple example, we have expectations of how a book "works": It has pages of information, laid out one after the other in a sequence; to get more information, you turn the page. One key to the enduring success of the touchscreen Amazon Kindle lies in how well it has remapped that mental model; just as you turn a page in a book, you "turn" a page by swiping at an e-book.[32]

When we can't assume how a gadget works, we use feedback—in the form of trial and error—to form a hazy mental model of its logic. But the most literal way to develop a mental model is to draw a picture. Looking around the Reactor 1 control room, I could see how it had been remade to create a mental model of the entire reactor. Even for a neophyte like myself, the major pieces of the system were easy to imagine. The room simply mirrored the reactor's design. Each control panel represented a discrete system—for example, the secondary circulation system or the reactor core—so that when I surveyed the room, I could see how those systems linked up, each flowing from one to the next. The reactor had been mapped to the room—just the same way you find the burners on your stove mapped to their corresponding dials, or the controls on your car's driver's seat mapped to buttons that resemble the seat's parts.[33] All of it was meant to create a durable picture in the minds of its operators, one whose steadiness could keep confusion at bay.

But the most curious thing I noticed was the way that workers were trained to interact inside this little bubble of precision. When workers would go to confirm some crucial reading, they went in pairs:

One would do an action, the next one would confirm the action; the first one would confirm the action had been done right, then the second one would as well. This process was meant to eliminate the kind of error that happened in 1979, when one worker went to the back of a control panel and misidentified the gauge that might have revealed an open valve. Instead, the workers would now follow the same steps built into any working button: After it's been pushed, the button issues feedback to confirm what's been done. In the reactor control room, that feedback came verbally, from the second worker. It's the same idea.

It is a strange kind of world we live in, where to make sure that men make no mayhem with a machine, they're made to behave like buttons. But then, it's maybe not surprising on deeper reflection: As Norbert Wiener discovered in his pioneering work designing feedback algorithms for shooting down German bombers, feedback is what turns information into action. Buttons, in turn, have become the connection point between our will and the user-friendly world. Embedded in them is a fundamental truth about how our minds make sense of the world. As banal as buttons may seem, properly viewed they can also seem like everything. The point arrives from surprising places, all the time. For me, the strangest was when my wife told me that her psychologist had said that the secret to having a productive argument with your spouse is to listen to what she has to say, *repeat* what you just heard, then finally have your spouse *confirm* that's what she meant. Push the button, provide feedback, confirm the action. Like a button. The creation of a shared understanding precedes any influence we might wield upon the world. Design is nothing more—and nothing less—than creating artifacts imbued with such shared understanding, legible to their users.

What happened at Three Mile Island was a transition between two eras: from one in which machines doomed their operators to make mistakes, to one that is the user-friendly world we live in today. Consider that Three Mile Island was designed in complete ignorance of the ways in which man and machine interact, which had been studied with excruciating care in a thirty-year period beginning after World War II. Whatever lessons should have been learned were instead squirreled away. And yet just five years after the Three Mile Island accident, the first Mac ads began to appear, touting a machine made to be utterly intuitive. It would be too much to say that without Three Mile Island, you'd never have the iPhone, but the two are bound together in a great chain of influence. The lights that didn't work, the gauges placed in baffling places—their opposite lives on in gadgets that get these ideas right, so that we never have to bother about figuring it out. You tap a button on the screen, it depresses just a little to show it's been pushed, then a new screen pops up in response—success. You encounter some new screen filled with new buttons on a menu in your phone, yet you can still figure out what to do, because the options are all laid out right before you. In all these cases, the opposite example—important because it illustrates how things go wrong without some shared understanding—is Reactor 2, with its lights and dials hidden from view and its fateful buttons that never confirmed what they were meant to.

Not long after I spent the afternoon with Harlan Crowder, the retired engineer living near Apple's new campus, I took a walk around downtown Palo Alto, trying out Google Lens. It was hard not to feel as if I were holding the future in my hand. Simply by using the camera on my smartphone, I could glean information about what was around me. I could train the camera on the marquee outside a theater and, at the bottom of the screen, buttons

would appear for finding out showtimes, booking tickets, or reading reviews about the movie. I could hold my phone up to a tree and be told its species.

It was as if the internet had been seamlessly superimposed on the real world. Or, more precisely: Google Lens is what you get when you transform the company's familiar search box into a spyglass. What powers it is some of the most advanced artificial intelligence on the planet, capable of recognizing words and objects, then guessing what you are wondering about based solely on what your phone is seeing. It represents billions of dollars in research, reaching all the way from neural networks—simulated clusters of neurons, dwelling in an algorithm—that can distinguish cats from dogs, and Korean from Japanese, to the very servers that house those calculations, whose circuitry has been custom-engineered for the computations required. Yet the end product of all that technology is not much more complicated to use than a magnifying glass. Here is one of the most advanced technologies in the world and it needs no instruction manual.

To be clear, there's a difference between a nuclear reactor and a smartphone. One is meant to be used by specialists, the other meant to be used by everyone. But there is a broad trend in our culture driving the most specialized things to become easier to use.[34] You can spy it from many angles at once. We expect everything from our washing machines to our deodorant to be "professional grade"—powerful enough for the most extreme application, even if we're not running a laundromat or an NFL team. Meanwhile, our economy is increasingly built upon systems in which the human professional is interchangeable—transient, contract workers whose value doesn't lie in their training or specialization but in their "hustle." Uber drivers don't need to know the city they drive in; factory

workers, if they still have jobs at all, have a precarious hold on the tasks that machines cannot do yet.

We now expect almost every aspect of our lives to work as simply as something on our smartphones—and it has to, for a rapidly changing workforce that often can't catch up to the pace of technological change. Indeed, because of the smartphone's ubiquity and immediacy, heavy work is now done on the move. Via smartphone, we can diagnose fleets of jet engines, exotic forms of cancer, and tetchy wind turbines. The user-friendly world is encroaching into successively tinier niches—and the nature of human expertise is shifting. Today, we take it for granted that the most advanced technologies ever created should never need any explaining whatsoever. How is that possible? How are we able to understand what something should do without ever being told?

Before I left Donald Norman at his fancy office, I remarked to him that it seemed like a kind of magic that some insight about the way we use a nuclear reactor control panel, or reach for a doorknob, might tell us so much about the shape of tools we haven't yet dreamed up. I meant to open a crack so that Norman might reflect a little about all the connections he'd seen. Wasn't it like a kind of magic? "But it's not!" Norman exclaimed. "It's science! That's what science does. You try to find general principles that hold true." Norman said that it seemed to him obvious that people weren't to blame for the errors they made. He believed that minds were knowable by science—that the frailties he was discovering weren't a feature of things going wrong, but were an essential facet of us, and what we expect of the world, whether we're swinging an ax at a log or pushing a button. We just expect things to work. Almost all of design stems from making sure that a user can figure out what to do, and can tell what's going on. The beauty and difficulty lie in what

happens when the object at hand is new, but needs to feel familiar so that its newness isn't baffling.

I asked if he had any designs that he lived with that he admired. Norman looked around, and then his gaze settled on his watch: a black Braun, whose clean lines and haute-utilitarian face were descended from the famous models designed by Dieter Rams and Dietrich Lubs in the 1960s. Except the watch face was marred by one detail: In addition to the analog hands, there was a honking LCD readout. It told the time twice in two different ways, for no apparent reason at all. For a long time, Norman said, he hated that watch; when he ordered it, he hadn't known it would come with a double readout. It was a silliness that was remarkable insofar as it was unique: You don't see telephones with two sets of keypads, or cars with two dashboards. But it grew on him, the way the watch seemed at war with itself. The way it seemed designed by two competing minds. "I love the conflict of my watch between beauty and function," Norman said. "I love that tension. This beautiful elegant dial that's ruined by this ugly LCD." There are many ways to make something work—deciding which one is right is another thing entirely. There's a psychology to design, but there's also art in it, and a culture, too. Design presumes that we can make objects humane, but doing so requires a different way of seeing the world. Three Mile Island showed us how machines should behave; by itself, this explains only part of the user-friendly world. Left unsaid are the motivations we bring to the things we make.