



The Tools of Their Tools

Evan Selinger and Jathan Sadowski

When people write about Nicholas Carr, they usually characterize him as a technology journalist. To make the point that he is especially good at his craft, they might say that he is a gifted storyteller and a Pulitzer Prize finalist. But Carr's latest book, *The Glass Cage*, confirms that these labels are not quite right and that the standard alternative, "technology critic," won't do either. Nicholas Carr deserves to be seen as a philosopher of technology.

The Glass Cage, like Carr's previous books on technology—*The Big Switch: Rewiring the World from Edison to Google* (2008) and *The Shallows: What the Internet Is Doing to Our Brains* (2010)—takes a critical look at recent technological developments, drawing from political economy and empirical research on human psychology. For the purposes of this review, however, we wish to focus on the phenomenological aspect of Carr's analysis—that is, the parts of his argument concerned with the perception, use, and experience of technology. Carr digs past the surface level of most contemporary

discussions of technology, exposing the subtle yet far-reaching ways that technology shapes who we are and what we do.

Carr's topic in the new book is automation. Although the word can ultimately be traced back to the Greek *automatos*, typically rendered "self-moving" or "self-acting," Carr notes that our English word "automation" is of surprisingly recent vintage: engineers at the Ford Motor Company reportedly coined the term in 1946 after struggling to refer to the new machinery churning out cars on the assembly lines. A little over a decade later, the word had already become freighted with the hopes and anxieties of the age. Carr tells of a Harvard business professor who wrote in 1958, "It has been used as a technological rallying cry, a manufacturing goal, an engineering challenge, an advertising slogan, a labor campaign banner, and as the symbol of ominous technological progress." Carr aims to investigate automation in all these variegated senses, and more.

The conventional wisdom about technology—or at least one popular,

*The Glass Cage:
Automation and Us*
By Nicholas Carr
Norton ~ 2014 ~ 276 pp.
\$26.95 (cloth)

mainstream view—holds that new technologies almost always better our lives. Carr, however, thinks that the changes we take to be improvements in our lives can obscure more nuanced and ambiguous changes, and that the dominant narrative of inevitable technological progress misconstrues our real relationship with technology. Philosophers of technology, as Albert Borgmann said in a 2003 interview, tend not to celebrate beneficial technological developments, “because they get celebrated all the time. Philosophers point out the liabilities—what happens when technology moves beyond lifting genuine burdens and starts freeing us from burdens that we should not want to be rid of.” A true philosopher of technology, Carr argues that the liabilities associated with automation threaten to impair the conditions required for meaningful work and action, and ultimately for leading meaningful lives.

In the spirit of finding the future where it already exists in the present, Carr directs our attention to the world of aviation, where automation has been pervasive for quite some time. As he points out, nearly all parties involved—airplane manufacturers, airlines, civil aviation agencies, and the military—have proven particularly keen at harnessing technologies to automate tasks previously done by human beings. He recounts the tales of two recent plane crashes, highlighting the tragedy that can

ensue when automation degrades skills and cognition.

In February 2009, a Bombardier Q400 turboprop embarked on a routine trip from Newark, New Jersey to Buffalo, New York. The quick hop should have been no problem at all. And it wasn’t, until the plane began its approach into the Buffalo airport. “The plane’s ‘stick shaker’ had activated, a signal that the turboprop was losing lift and risked going into an aerodynamic stall,” Carr recounts. This caused the autopilot to disconnect—as it was programmed to do—and relinquish all control to the captain and first officer. The captain reacted and grabbed onto the yoke, “but he did precisely the wrong thing.” Instead of pushing the yoke forward, he yanked back on it—even as the plane’s stall-avoidance system attempted to push forward. “Rather than prevent a stall,” Carr writes, the captain “caused one.” The Q400 plummeted to the ground, killing all 49 people on board and one person in the house it crashed into.

A few months later, “an eerily similar disaster, with far more casualties” occurred. An Airbus A330 with 228 people on board was making the long overnight journey from Rio de Janeiro to Paris and, again, the autopilot disengaged, this time due to the air-speed sensors becoming caked with ice. The first officer tried to regain control by pulling back on the stick. Even as stall warnings blared, he continued on. This caused the jet

to climb sharply before losing the velocity needed to stay airborne. The plane dropped into the ocean, leaving no survivors.

The investigations of both tragedies reached similar conclusions. In the case of Q400, the National Transportation Safety Board said that pilot error caused the accident; the evidence pointed to “a significant breakdown in [the captain and first officer’s] monitoring responsibilities.” In the case of A300, French investigators said the flight crew suffered from “a total loss of cognitive control of the situation.” In both crashes, it seemed that human beings did not hold up their end of the bargain as parts of a complex system.

Carr argues that we cannot just blame “human error” and wholly absolve automated systems of responsibility without digging further into the causes underlying both tragedies. He notes that pilots have become increasingly dependent on automation, and that over time autopilot technology has gone from being an aid that offered pilots relief from taxing workloads to being a plane’s primary controller. Evidence collected over decades has shown that trip after trip of sitting in the cockpit effectively relegated to monitoring duties can atrophy the psychomotor and cognitive skills needed to fly. It is thus unsurprising that disasters and near misses sometimes occur when automation systems disengage or malfunction—as inevitably happens—and

manual control is forced back into the hands of pilots whom these very systems have deskilled.

Pilots are not ignorant of the negative effects of automation. “They’ve always been wary about ceding responsibility to machinery,” Carr writes. To be sure, automated systems have tended to make aviation safer. But to see that fact as a refutation of Carr’s larger argument, as some reviewers have done, is to miss his point about the altered relation between plane and pilot, tool and user. Pilots themselves are concerned about the change: “Even as they praise the enormous gains in flight technology, and acknowledge the safety and efficiency benefits,” Carr notes, “they worry about the erosion of their talents.” By placing an intermediary between us and the activities we perform, automation can dull the skills and awareness we need to make our way through the world.

Carr’s foray into aviation is not an isolated case study. He sees it as a window into a future in which automation becomes increasingly pervasive. “As we begin to live our lives inside glass cockpits,” he warns, “we seem fated to discover what pilots already know: a glass cockpit can also be a glass cage.” The consequences will not always be so catastrophic and the systems will not always be so totalizing. But that does not make the range of technologies any less worthy of the crucial questions that Carr asks: “Am I the master of the

machine, or its servant? Am I an actor in the world, or an observer? Am I an agent, or an object?"

In order to explain the general problems of automation, Carr lays out some of the factors that make it difficult for us to recognize those problems in the first place. (This process of clearing away the biases that prevent people from seeing what needs to be confronted is sometimes called "phenomenological reduction," although Carr does not use the term.)

First, the most important things people stand to lose by letting automation go too far are hard to measure: an active sense of agency, a robust experience of autonomy, and the capacity to execute skills that add meaning to our lives. While these are all real, they feel completely subjective and possibly ineffable. Because we cannot quantify losses in these domains, it is easy to underestimate the significance of what is slipping away and how much diminution is occurring at any moment. Moreover, because the dwindling takes place through the gradual integration of new technologies rather than happening at a single, decisive, and overwhelming moment where a particular engagement with technology turns us into mush, we underrate the cumulative effect. The things we stand to lose, Carr laments, "are the kinds of shadowy, intangible things that we rarely appreciate until after they are gone."

Second, it is difficult to figure out how to use automation wisely and protect ourselves against future losses when technology changes more quickly than our ability to understand its effects on us. "Whereas computers sprint forward at the pace of Moore's law," Carr writes, "our own innate abilities creep ahead with the tortoise-like tread of Darwin's law." Rather than confronting this temporal disparity head on, we are too often inclined to believe optimistically that we will simply adapt to, or just muddle through, whatever comes our way. After all, the human species has made it this far, and despite recurring panics about alienation, we haven't become soulless automata yet, right? But the belief that we can always just adapt to technological change can blind us to the need occasionally to set boundaries, to draw limits, to protect aspects of the human condition that we should deem inviolable.

Third, we are biased by a "substitution myth"—a fallacious assumption that inclines us to believe that when a labor-saving device is used, it offers a simple "substitute for some isolated component of a job." In reality, something more holistic and far-reaching can occur. Automating an activity sometimes transforms "the character of an entire task, including the roles, attitudes, and skills of the people who take part in it." That is, replacing one part of a system can cause a ripple effect that changes other parts.

Eventually, a “degeneration effect” takes hold where the resulting technological dependency leaves us less able to adapt to new situations and make our way in the world without the crutch of automation: “we naturally come to rely more on the software and less on our own smarts.”

Fourth, it is difficult to acknowledge downsides to automation when we are smitten with longstanding ideologies that construe technology as a great liberator. Many of us are enculturated to believe that the more we automate the grunt work pervading our private lives, the freer we will become. Unfortunately, we do not always appreciate how much fulfillment we find from seemingly menial tasks that demand concentration and skill. “One of the most remarkable things about us,” Carr declares, “is also one of the easiest to overlook: each time we collide with the real”—doing manual labor or cooking a meal from scratch or writing a quick message to a friend—“we deepen our understanding of the world and become more fully a part of it.”

Fifth, we hold a prejudice that technology ultimately amounts to nothing more than a collection of tools that we, their creators, can master. This attitude dates back to the ancient Greeks. As Carr reminds us, in the *Politics* Aristotle posits an equivalence between slaves and tools, “the former acting as ‘animate instruments’ and the latter as ‘inan-

imate instruments.” At the same time, many people also hold to a kind of technological determinism that amounts to a belief that we are the slaves of our creations. These extreme, opposing views obscure two important facts: technologies are not neutral, but often come in forms that incline us to behave in very specific ways; and when technologies do limit our scope of action, it is usually because of how they are employed by people or institutions for their own ends—often to enhance their own power—rather than because of the technologies’ inherent properties.

Once we have removed, or at least become aware of, these blinders, it becomes possible to appreciate how automation can induce what Carr calls two “cognitive ailments”: *automation complacency* and *automation bias*. We slip into automation complacency when we sacrifice our own attentiveness by treating automation technology as an unfaltering supervisor. “We become so confident that the machine will work flawlessly, handling any challenge that may arise, that we allow our attention to drift.” This makes it easier to miss warning signs when the technology malfunctions. The related ailment of automation bias comes into play when “people give undue weight to the information coming through their monitors.” Our devices sometimes give us wrong or misleading information but we often continue to place our trust in it, even when

it conflicts with “other sources of information, including [our] own senses.”

Both of these cognitive ailments ultimately originate from the same place: “limitations in our ability to pay attention.” And they both “tend to become more severe as the quality and reliability of an automated system improve” because we get lulled into disengaged laziness. Such sleepwalking can drastically increase the danger of a situation that would otherwise have required only minor corrections had an alert person been attuned to the warning signs. Carr quotes a human-factors expert who refers to “a growing body of research” showing that automated systems can paradoxically “increase workload and create unsafe working conditions.” When people overly reliant on automation systems get thrown into a situation where those systems are absent or broken, they become overwhelmed.

Automation complacency and bias have entered the popular consciousness, in part because of examples familiar in our everyday lives. Some of these have become fodder for comedy, as in a well-known slapstick scene from the U.S. version of the television show *The Office*. It features two characters, Michael Scott (Steve Carell) and Dwight Schrute (Rainn Wilson), driving down unfamiliar country roads and relying on GPS directions. Although the GPS gives a dubious-sounding instruction,

Michael still defers to it. Resolutely trusting the electronic voice over his own senses and over Dwight’s advice, Michael plunges the car into a shallow lake, all the while yelling, “The machine knows where it’s going—the machine knows!”

Alas, succumbing to automation complacency and bias is not always so obviously foolish or relatively harmless. Take, for instance, Carr’s account of a Norwegian-owned ocean liner named the *Royal Majesty*. During the last part of a weeklong journey in the Atlantic in 1995, the GPS antenna for the automated navigation system became damaged, causing the computer to give inaccurate readings. Unbeknownst to the captain and crew, for over thirty hours “the ship slowly drifted off its appointed route...despite clear signs that the system had failed.” They continued to trust that the navigation system was guiding them along the correct path. Until, that is, the ship ran aground on a sandbar. “No one was hurt, fortunately, though the cruise company suffered millions in damages.”

Throughout *The Glass Cage*, Carr identifies many cases of automation complacency and automation bias, including in the fields of medicine and finance. The examples suggest the presence of what we might dub “automation creep”—the expansion and increasing sophistication of automation technologies. Automation is poised to continue

expanding within industry and professional settings and also to permeate ever more aspects of everyday life.

One crucial area where we need to be vigilant against automation creep is the domain of digital consumer goods that mediate our relation to information. While Carr does not focus much attention on these products, his selective remarks offer clues for what the future might look like.

In Silicon Valley there is a strong commitment to creating “frictionless” environments—a commitment that arguably amounts to an ideology. Friction, in this sense, is synonymous with inefficiencies that waste our time, slow the pace of innovation, and prevent optimal performance—both from machines and us. For incumbent firms and startups alike, friction has become taboo. In these circles, it is widely assumed that the easier it is for consumers to create, locate, and share information, the better off they will be.

But contrary to the prevailing platitudes, minimizing friction doesn’t simply remove speed bumps from our paths. While getting rid of obstacles may seem like a process of subtraction, frictionless design cannot be advanced without first introducing new devices and systems, and then promoting them as superior alternatives to and replacements for older ones. After they are rolled out, these technologies subject us to new kinds of “choice architecture” (a

term from the literature on nudging) that can modify our behavior. Over time, as more and more people incorporate frictionless technology into their daily lives, cultural norms shift. Automation may promote technical values at the expense of humanistic ones; it may make it harder to choose deliberative practices that require attentiveness and conscientiousness.

To get a clearer sense of how automation creep can modify our actions—both the means and ends—consider three phases of automation in a common technologically mediated process: automated writing. In the first phase, we started using automated spell checkers. Carr notes their primary function was to act as “tutors” that “highlight possible errors, calling your attention to them and, in the process, giving you a little spelling lesson. You learned as you used them.” Automation served a spotlighting function, letting you know when something might be amiss but ultimately relying on your deliberation and judgment for corrections.

Over time, however, the spell checker was made bolder, given new abilities that made it a new type of cyber-servant: autocorrect. At this point, the technology “instantly and surreptitiously” fixed mistakes without providing any feedback. You would “see nothing,” Carr observes, but also “learn nothing.” In the abstract, autocorrect seems great: Who needs to worry about learning how to spell

words correctly when the technology is ubiquitous and enables you to write well-proofread letters and papers? Especially in educational or professional settings, when the stakes for rapid and error-free writing can be high, you might be considered foolish for failing to take as much assistance as technology can offer.

As a further advantage, when misspellings occur during smartphone correspondence, you have a good shot of being immunized from blame. It turns out that people are relatively forgiving of spelling mistakes in messages sent from smartphones; the little “Sent from my iPhone” tagline works like a “built-in forgiveness clause,” as a 2013 article in *The Wire* noted. (One of the authors of the present essay has customized his smartphone disclaimer to read “Mistakes courtesy of iPhone.”)

But all of these benefits may come at a price. Speaking anecdotally, we have noticed that over time, having been habituated to using spell check and autocorrect programs, we have also become increasingly worse spellers when put into situations that require handwritten prose. We have even caught ourselves needing to look up words online with our smartphones, trying to avoid the embarrassment of making basic misspellings. Of course, whether such a tradeoff really matters is a point of disagreement.

In the third phase of its evolution, the spell checker has gone from serv-

ing a relatively trivial cognitive function to a more intimate and interpersonal one. The purview of autocorrect has expanded and morphed into the presumptuous function of predictive texting. The iOS 8 software update, for example, promises to deliver the iPhone’s “smartest keyboard ever,” one that “predicts what you’ll likely say next. No matter whom you’re saying it to.” The “QuickType” software—the name alone emphasizes the value of speedy, optimized, frictionless communication—is supposed to switch style and tone depending on whom you are talking to (your boss or your best friend) and what app you are using (text message or e-mail). The more you type, the more it learns about you.

Of course, QuickType doesn’t force you to accept its recommendations. It’s still up to users to decide if they want to endorse the suggestions. But it may prove difficult to resist relying on it, even when its recommendations are imperfect. And when you have predictive type spitting you back to you, “you” become a facsimile of yourself, a set of personalized clichés. Whereas clichés are typically overused expressions that are generic to a community—“How are you?” “Fine.” “Awesome.”—predictive-type clichés will be idiosyncratic to individuals. This is much worse than degraded spelling skills: we will give others a stand-in version of ourselves as we get lost among hollowed-out messages. Interpersonal communication

will be suffused with personalized banality.

We can only speculate about how far automation creep will go. A possible fourth phase in the automation of communication is hinted at in a patent application that Google has filed. The software it describes would learn how you respond to social media posts and recommend updates and replies you can make, effectively encouraging you to outsource future “interactions” with friends and followers. The automation in this software does more than incentivize you to become a cliché. Generating whole-cloth responses—rather than partially recycled pieces of text—is ventriloquism in which we are the dummies. At the extreme, bot proxies could end up “talking” to other bot proxies, the direct human presence fading away.

When looked at by themselves, these changes in how we communicate may seem insignificant, too removed from all the other challenges of life to adversely impact who we are and how we interact with others. But these examples help us to see bigger trends, in which automation systems are becoming ever more sophisticated, subtle, and ubiquitous. “At some point,” Carr writes, “automation reaches a critical mass. It begins to shape society’s norms, assumptions, and ethics. People see themselves and their relations to others in a different light, and they adjust their sense of personal agen-

cy and responsibility to account for technology’s expanding role. They behave differently too.”

Max Weber, the sociologist and woebegone appraiser of modernity, famously described the modern economic order as an “iron cage.” An irresistible drive towards rationalization and bureaucratization, Weber thought, displaced old forms of kinship and undermined the freedom of people to pursue anything other than material goods. There is more than a token resemblance between Carr’s title and Weber’s lament. While the likeness appears unintentional, as Carr never mentions Weber, it is not a mere coincidence. Both thinkers are, among other things, studying a tendency to render the world more efficient and manageable, whether through technology or social, political, or economic organization. Moreover, both Carr and Weber try to understand the effect of these changes on the subjective side of human behavior.

In this respect, a comparison of the cage metaphors is telling. Whereas Weber’s iron cage suggests a harsh confinement that constantly impinges on our awareness, Carr’s glass cage seems less confining but also harder to perceive. Carr thinks that we risk being beguiled by the wondrous vistas that technology provides—too beguiled, that is, to notice that the very glass transmitting the vistas divorces us from their source.

As *The Glass Cage* makes clear, automation systems can not only change how we interact with the built world and its artifacts but also can transform how we interact with other people—encouraging us to treat them as objects rather than subjects who warrant consideration, care, and effort. We may do fewer of the things required to demonstrate caring, the things that show people they are connected to us by the bonds of family, friendship, and love. As we choose the ease and comfort of automation over our own autonomy and agency, we may tend more often to treat other people as instrumental material to manipulate and objectify instead of as subjects worthy of respect.

Moreover, thanks to automation complacency and bias, we are likely to be blind and susceptible to the systems' influence on us over time. Even if we can intellectually recog-

nize the Faustian bargain of convenience, it will be hard to refuse that bargain, and once we develop new habits and internalize new norms, it will be practically very difficult to change them. In the end, whether or not you agree with the values embedded in these automation systems is incidental to the fact that they ought to be held to scrutiny and not simply designed, implemented, and allowed to change us and our society with impunity.

Evan Selinger is an associate professor of philosophy at Rochester Institute of Technology. He advised the author of the book here reviewed on matters related to ethically assessing technology, commented on the manuscript before publication, and is mentioned in the book's Acknowledgements. Jathan Sadowski is a Ph.D. candidate in the human and social dimensions of science and technology at Arizona State University.