Insolvent

How to Reorient Computing for Just Sustainability

By: Christoph Becker

Citation:

Insolvent: How to Reorient Computing for Just Sustainability
By: Christoph Becker
DOI: 10.7551/mitpress/14668.001.0001
ISBN (electronic): 9780262374668
Publisher: The MIT Press
Published: 2023

The MIT Press

12

A SILICONE RING

SOCIAL RESPONSIBILITY AND COLLECTIVE ACTION

From the outset, therefore, the engineer was at the service of capital, and, not surprisingly, its laws were to him as natural as the laws of science. If some political economists drew a distinction between technology and capitalism, that distinction collapsed in the person of the engineer and in his work, engineering.
—D. Noble (1977, 34)

If democracy is going to mean anything, it is the ability to all agree to arrange things in a different way.
—Graeber (2011, 390)

At 5:31 p.m. on August 29, 1907, in the span of fifteen seconds, the world's largest cantilever bridge collapsed. Thousands of tons of steel, twisted and torn like melted plastic, sank seventy meters into the Saint Lawrence stream near Quebec City. The bridge was not completed yet, and eighty-six workers fell with it. They were about to leave the site for the day when the iron under their feet gave way. Seventy-six workers died, half of them buried at the bottom of the river, never to be recovered. In the days before, several had raised concerns about the obvious bending of the iron chords holding up the structure, but their fears had been dismissed. Subsequent inquiries into the causes of the tragedy absolved all involved parties from negligence. But a reexamination of evidence suggested that the leading engineers could have averted the disaster had they paid attention. "There was no trace of humility in any of the senior engineers involved with the design of the Quebec Bridge. Their arrogance and absolute confidence in their work prevented them from realizing what the workers building the bridge already knew: that the bridge was failing under its own weight" (Levert 2020).

The collapse left a mark on the psyche of Canadian engineering. Not long after the tragic disaster, a group of engineers and engineering professors introduced a ritual into the profession that persists to this day. Upon graduation, young engineers complete a private ceremony, away from public sight and celebration, in which they take an Obligation and a ring. Both remain unchanged since 1925. The Obligation reminds them of their "assured failures and derelictions" (see Raymond Francis 2006). It is not quite a Hippocratic Oath, but close. The Iron Ring bestowed on the engineer is worn on the little finger of the working hand so that it rubs on working materials as a constant reminder of humility and dedication.

We will never know if the Obligation oath and the wearing of the ring have materially improved the safety of Canadian infrastructure. There is no doubt, however, that the Obligation and its material symbol impress the minds of engineers and serve as a powerful reminder of the responsibilities they carry (Levert 2020) as they use specialized technical expertise to shape the world around us.

HOW (NOT) TO AVOID HARM

In 2018, I still had hope. Hope as in confidence in the ability of institutional action to establish effective guides for ethical conduct. I felt confident that the organization I thought best represented me as a professional and researcher—the Association for Computing Machinery (ACM)—would take meaningful action to address the social responsibility of computing in the twenty-first century. ACM was finally revamping its outdated code of ethics (ACM Committee on Professional Ethics 1992), which was in dire need of overhaul and improvement. It spoke of "avoiding harm," of course—that is considered the absolute minimum for a code of ethics, after all. But its examples for harm were tellingly narrow: "'Harm' means injury or negative consequences, such as undesirable loss of information, loss of property, property damage, or unwanted environmental impacts. . . .

Harmful actions include intentional destruction or modification of files and programs leading to serious loss of resources or unnecessary expenditure of human resources such as the time and effort required to purge systems of 'computer viruses'" (ACM Committee on Professional Ethics 1992). In its focus on property and material damage, the ethics code was, as always, a product of the times. As ACM president Cherri Pancake wrote in 2018, "in 1992, many of us saw computing work as purely technical" (ACM Committee on Professional Ethics 2018, 1).

The process for updating the code looked promising. Significant consultation seemed built in. Under the leadership of an interdisciplinary taskforce, multiple rounds of public feedback were to shape the new edition of the code, informed by critical insights into the conflicted role of computing in our societies. Multiple drafts were released for comments. Many excellent changes were introduced to update and strengthen the code and consider feedback of the community. When I read the final draft, however, alarm bells went off. The code contained crucial changes to a central principle that weakened it substantially and fundamentally. The last public draft of *Principle 1.2 Avoid Harm* looks harmless enough at first glance, but note the highlighted passages.

> In this document, "harm" means negative consequences to any stakeholder, especially when those consequences are significant and unjust. Examples of harm include *unjustified* physical or mental injury, *unjustified* destruction or disclosure of information, and *unjustified* damage to property, reputation, and the environment.

Why is each example of harm an *unjustified* damage? Is it only harm when it isn't justified? By whom?

> Well-intended actions, including those that accomplish assigned duties, may lead to harm. *When that harm is unintended*, those responsible are obligated to undo or mitigate the harm as much as possible. Avoiding harm begins with careful consideration of potential impacts on all those affected by decisions.

How actionable is the distinction based on intentions in today's computing landscape? How are those responsible supposed to mitigate harm that they have not intended in the first place? Why are they not instructed to avoid harm?

> *When harm is an intentional part of the system*, those responsible are obligated to ensure that the harm is ethically justified and to minimize *unintended* harm. To

minimize the possibility of indirectly harming others, computing professionals should *follow generally accepted best practices*. Additionally, the consequences of emergent systems and data aggregation should be carefully analyzed. (Gotterbarn et al. 2017, 124)

The closer we look, the more dubious it sounds. When harm is intentional, it does not have to be minimized? How are computing professionals supposed to ethically justify the harm caused by their work, and to whom? And since when are "generally accepted best practices" sufficient guidance when they lead us to a state of insolvency?

In ethics, the doctrine of *double effect* allows negative harm for a positive outcome under certain conditions. But "applications of double effect always presuppose that some kind of proportionality condition has been satisfied. Traditional formulations of the proportionality condition require that the value of promoting the good end outweigh the disvalue of the harmful side effect" (McIntyre 2014). For example, a surgeon may cause significant harm to a patient's skin to save their heart. In academic research, whenever I collect personally identifiable information, I need to justify the value of doing so in relationship to the risk posed to participants. Crucially, I am not allowed to make that judgment myself: an institutional review board makes it for me.

The new code introduces the double effect without addressing proportionality at all. The text above instructs those responsible to minimize unintentional harm without requiring them to minimize intentional harm. The word "unjustified" hints at a proportionality condition but wraps it up within the examples, without opening a conversation about it. Its use to articulate all examples suggests that justified negative consequences would not count as harm at all. At no point does the text address the crucial question of accountability: the act of justifying is left to those who are involved. The result is a principle that spells "avoid harm" while expressly condoning intentional harm, suggesting that justified harm is no harm at all, *and* leaving it up to the designers to decide which is which.

Consider the effect on two examples, one benign, one murderous. In the process of coordinating a collective response, my colleagues and I discussed the action of alerting all members of a dormant community email list to the danger posed by the proposed principle. This involved minor harm (the annoyance of yet another unsolicited email) for a greater public good (the wider discussion of an ethics code). It may have also involved unintentional harms of some minor kind, but we are unable to anticipate those with any certainty. Maybe these could be some people's possible irritation about a standpoint they do not share or the carbon footprint of emails. The draft Principle 1.2 required us to minimize these unintentional harms—which we could not feasibly anticipate—but it did not require us to minimize the intentional harm caused by yet another unsolicited email. It seems obvious and common sense, though, that we should aim to minimize that, for example by taking care to explain the rationale behind the message.

The principles fail more strikingly in an example more typical of ethics literature. Consider a team that designs a bomb to end a war. They face a choice between developing a bomb that would neutralize electronic equipment and developing a hydrogen bomb. According to Principle 1.2, the responsible agents would have to minimize unintentional harm, but not intentional harm. This would make the hydrogen bomb the more appealing solution *according to the ethics code*. Yes, it is absurd.

To understand how this text came to be, I examined the publicly available documentation. The distinction between intentional and unintentional harm had been introduced in the final draft to expand the code's applicability to arms manufacturers (pardon, "the defense industry") (Gotterbarn et al. 2017). To voice my concerns, I published a Twitter thread and reached out to colleagues, who helped me draft an analysis (Becker 2018b). The ethics committee swiftly responded on social media, publicly asserting they would meet with me to discuss the concerns.[1] They did not. Instead, I was informed in writing that the text had already been approved by ACM Council. The committee's response did not address my concerns, but eventually, the final text *was* changed to address one of the issues my colleagues and I had objected to: The final principle states "In either case, ensure that all harm is minimized" (ACM Committee on Professional Ethics 2018), just as I had proposed. It was a bittersweet moment to see that bare minimum change implemented. No other issue was addressed.

ETHICS WON'T SAVE US

I tell this story as a reminder that vested interests capture ethical concepts under the cover of the objectivist illusion. Engineers tend to look to codes of ethics as a seal of approval to absolve them from responsibility for misconduct. From the orthodox standpoint of insolvent computing, the "ethical problem" posed by just sustainability is solved at one clearly identified point: the code of ethics. The codes of ethics of professional associations like ACM and IEEE provide the principles of actions that guide ethical conduct. The development of ethics codes is the central point at which governance and value judgments are admitted into a clearly circumscribed location of this picture. Within the expressly political process of developing a new code, ethical and moral philosophy are translated into prescriptions by formalizing supposedly democratic decisions into ethical principles. Outside the process, the resulting codes formalize the criteria on which the ethicality of design can be adjudicated. Methods further translate these principles into actionable guidelines. If those methods are followed properly, so they say, then the outcomes will be ethical. By implication, failure can be blamed on the individual (Metcalf, Moss, and boyd 2019). This is why it is so important that the process carry the democratic appearance of legitimacy. When I raised concerns at a late stage, the institution ACM was more concerned with that appearance than with the substantive questions I raised. The episode made it clear to me personally that these codes have very dubious legitimacy.

Artificial intelligence (AI) ethics is a booming business today, and still, some of the public discussion evolves around the content of ethics codes. With much less attention paid to the question of *how* ethical principles and directives come to be, and how they work in practice, it is common for industry to shop around for the set of ethics codes that best fits their existing business practices, to use these ethics codes to make exaggerated claims about the business's social responsibilities, to lobby for self-imposed ethics regulation over externally imposed legislation, and to remove unethical components of computational systems, such as the invisible labor behind so-called AI, from highly regulated countries, only to re-import the resulting technologies (Floridi 2019).

This is why ethics is often critically seen as an instrument of control and complacency. Western moral philosophy has long tried to pretend that its reasoning takes place in a vacuum. While there is still value to moral philosophy in AI ethics (Bietti 2020), the critical friends highlight that counter to the "ideal" tradition that dominates it, ethical reasoning always already happens on a ground of inequity and injustice

(A. Hoffmann 2020). On this ground, industry sometimes hires critical researchers with public visibility to conduct critical work with the intent of using them as "legitimation workers" (Whittaker 2020). The central purpose of their employment is not to conduct substantive critical research, but to perform the appearance of critical research. Many noticed what happened when Timnit Gebru and Margaret Mitchell, two such legitimation workers, dared to interpret their job at Google differently: they got fired, "revealing just how much companies will prioritize profit over self-policing" (Hao 2021b, 51).

Some pointed out that I had been naïve to expect anything but moral bankruptcy from an institution like the ACM, and they may be right. ACM has had a long history of entanglement with the military industrial complex, after all. It is hard to draw the line between the principles discussed earlier and a satirical comment such as this one: "*impact assessment* (ph)—A review that you do yourself of your company or AI system to show your willingness to consider its downsides without changing anything" (Hao 2021a).

The theory of economic regulation distinguishes between two forms of *regulatory capture*. Financial capture refers to material incentives, often of dubious legality, which shape the behavior of a regulating entity. Cultural or *cognitive capture* refers to the shaping of regulatory thought by industry. Both types are at play in the capture of AI regulation by Big Tech. Tech companies do not just operate the revolving door connecting their offices to the halls of government, they are also "startlingly well positioned to shape what we do—and do not—know about AI and the business behind it, at the same time that their AI products are working to shape our lives and institutions . . . [they] control the tooling, development environments, languages, and software that define the AI research process—they make the water in which AI research swims" (Whittaker 2021). In this context, institutional declarations of goodwill mean little if the actions of the institution do not correspond. And we know that to identify an organization's purpose, we must look at its behavior, not at what it declares its purpose to be (D. H. Meadows 2008, 14).

Amid regulatory capture, ethics washing, and the inadequacy of ethical codes to guide responsible design, it is clear that even good ethics codes alone won't save us. So what might?

BIG TECH WON'T SAVE US

Not to worry about sustainability and the fate of the global poor, say the founders of Big Tech. We don't need ethics codes to do good. In fact, they just hold us back. Step aside and let us save the world. We know what we are doing. After all, we have already proven to you how capable we are, and we have nothing else to prove. We'll save you. Bill Gates has no geoengineering credentials but has published an entire book on it, in which he proudly declares: "I think more like an engineer than a political scientist" (Milman and Rushe 2021). Elon Musk markets autonomous electric cars to rescue us from fossil fuel dependency while his company exploits lithium reserves in the Global South. Anyway, he plans to retire on Mars. And Jeff Bezos claims that Amazon will be net zero by 2040—which would be too little, too late—but for now, Amazon continues to market its cloud services to the fossil fuel industry (Palmer 2020). It does not take a lot of scrutiny to become suspicious of these claims. But a common perception is that they must have some point. After all, they are highly successful entrepreneurs, and that means they have judiciously taken high, calculated risks and demonstrated visionary leadership in developing innovative technological breakthroughs with real value to the world. Or have they?

We often hear that origin myth of Big Tech. It tells us that innovation begins at places such as Silicon Valley, where technological entrepreneurs devise ingenious and entirely novel ideas—ground-breaking innovations that lead to breakthrough technologies. Shouldn't computing then deserve the profits it earns and the soaring stock market evaluation of its shares? While no one disputes that innovative entrepreneurs should make money, economic studies show that it is the state, not private entrepreneurs, who takes the biggest risk, funds the real breakthroughs, and takes the long-term view necessary to make these breakthroughs happen. Economist Mariana Mazzucato (2013) has compiled overwhelming evidence showing that private investors, focused on short-term profits, are more risk-averse than the state and rarely fund the kind of technological breakthrough that enables the businesses they invest in. Instead, they come in afterwards to turn these technological innovations into profits. This is true especially for the iconic successes of Silicon Valley, such as the iPhone. From the touchscreen to the internet, from GPS to voice recognition and from battery technology to

HTML, each component that makes it work was made possible by risky, strategic, long-term state funding: "far from getting out of the way of private innovation the State paves the way" (Barendregt et al. 2021). There is a long history in which government funding enables true breakthrough innovation in IT and private business harvests the profits.[2]

Instead of placing our bets on Big Tech getting it right, "we should be striving for a much more radical agenda that envisages the wholesale transformation of the computing profession, putting an end to the technological solutionism of Silicon Valley, turning it into a humble enterprise that places human dignity first," as my colleague Steve Easterbrook puts it (2021). As I put it with a group of coauthors in a new platform we launched after a series of workshops on collective action:

> Now is the time to radically redirect the future of tech, by reclaiming the purposes of technology development, and redistributing the associated responsibilities and benefits, in the service of our collective and sustainable well being . . . to redirect Big Tech's excessive revenue flow, we must transform the conditions and funding structures that enable it. The aim is to free up resources to support a wide range of socially beneficial ends, not least community-based and community-oriented initiatives to develop digital infrastructures that better serve the public interest. While we are not calling for the demise of Big Tech, we are calling for radical reform. (TechOtherwise Collective 2021)

In fact, defunding Big Tech is not as radical a position as it may appear. It simply means that societies ought to collectively reap the rewards of their collective investments. To achieve that, every participant in our societies must exercise our agency to reduce the outsized influence of corporate tech giants over how our lives on this planet are organized and for whose benefit. As Graeber's epigraph puts it so well: democracy should mean that we can all agree to change course. The proposal offers steps for tech companies, policymakers, tech professionals in industry and academia, civil society organizations, and individuals. One of these groups gives me the most hope.

TECH WORKERS FOR SOCIAL RESPONSIBILITY

While the executives at Amazon & Co are busy brokering deals with the oil and gas industry, employees at the same companies are organizing for sustainability and justice. In 2019, more than 8,000 Amazon employees signed an open letter and supported a shareholder proposal to accelerate the company's commitment to carbon neutrality. Throughout the pandemic, however, while Amazon's carbon intensity sank, its overall environmental footprint grew. Google workers may have made the biggest headlines. In spring 2018, thousands joined a walkout in protest over a Pentagon weapons contract (Tarnoff 2020). Later that year, walkouts from offices around the globe protested the company's mishandling of sexual misconduct. With success: the company stopped the Pentagon project that year and updated its processes for handling misconduct.

Protest organizers faced retaliations, and prominent organizers like Meredith Whittaker left the company, explaining: "The reasons I'm leaving aren't a mystery. I'm committed . . . to organizing for an accountable tech industry—and it's clear Google isn't a place where I can continue this work." (Vincent 2019). Ben Tarnoff, who has studied the "tech worker movement" for years, highlights that tech specialists such as software engineers occupy a middle ground. They are located between executives and the outsourced service workers who keep the posh corporate campuses running and who also should be understood as tech workers. In fact, the latter were the first to organize, including on issues of environmental justice (Pellow and Park 2002). The experience of their organizing was profoundly influential on how IT specialists understood their own role.

> By virtue of their position between labor and capital, they inhabit what Erik Olin Wright famously called "contradictory class locations" . . . they are pulled in two directions. . . . The tech worker movement offers a fascinating illustration of the latter phenomenon. It involves many members of the middle layers coming to see themselves as workers. This new identity has in turn enabled those individuals to act collectively as workers: to use their leverage over the spaces where profit is made in order to demand more control of their work and their workplaces. They are wielding the weapons of working-class struggle while speaking its language. Most significantly, they are building relationships of solidarity with their working-class colleagues in the tech industry and coordinating their efforts in order to advance their campaigns together. (Tarnoff 2020)

During the pandemic, tech workers around the world have organized, lobbied, protested, and unionized. The wave of collective organizing in tech companies that we are witnessing today demonstrates that many tech workers are keenly aware of the tension between corporate values and their own sense of professional social responsibility, and that they are ready to act collectively on their commitments, even if this puts their financial security at risk. This tension is not all new. As briefly seen in chapter 10 and recalled in the chapter epigraph, the room for action available to specialists in technology fields has always been shaped and constrained by their employers' interests—directly, through incentives, responsibilities, and rules, and indirectly, through the production and reproduction of ideas through education and the cultural imagination. These tensions emerged together with the engineering disciplines and professions themselves.[3]

As the chapter's epigraph illustrates, Noble provides a sharp analysis of this history. Strikingly, the leaders of previous centuries did not find it necessary to hide their intentions. One told engineering students that "the dollar is the final term in every engineering equation." Another spoke to engineers about the role of the engineer, saying "he must always be subservient to those who represent the money invested in the enterprise" (D. Noble 1977, 34–35).

For over a century, some tech workers have disagreed very strongly. Already in the 1910s, some progressive engineers in the United States rejected the idea that the dollar should always have the last word. Most were committed to corporate liberal reform, however, and others remained peripheral (D. Noble 1977, 50–63; Wisnioski 2012, 5). By the heady 1960s, the growing power of computing and fields such as operations research fueled the idea that engineering could be applied to society at large. The scaling up of engineering education had also shifted the labor dynamic: by the 1960s, the number of engineers in the US had risen from a very small cadre of highly respected engineers to a workforce exceeding 800,000. With that came a shift into the dynamics of a nameless workforce and the values of team play and compliance (Wisnioski 2012, 31). The growth in number also entailed ever-increasing degrees of specializations. And for each subdivision of specialized expert labor, politics are supposedly just outside. As Feenberg highlights, "the differentiation of specializations gives specialists the illusion of pure, rational autonomy. This illusion masks a more complex reality. In reality, they represent the interests which presided over the underdetermined technical choices that lie in the past of their profession" (Feenberg 1999, 139). But at the same time, activism grew in many professions, including engineering fields (Wisnioski 2012; Hoffman 1989).

> As traditionally defined, professional work is the work of "experts"—those who apply a given body of abstract knowledge to specific problems . . . the traditional model asserts that the dependence of the client, the knowledge of the expert, and the importance of the task, make it necessary to maintain an impersonal orientation to the task at hand, a hierarchically structured relation between client and professional, and professional control of most if not all aspects of professional work. For critics, on the other hand, professionalism is not a "best solution" to an inherent set of dilemmas, but a form of practice that maximizes professional control to maximize professional self-interest. (Hoffman 1989, 6)

Hoffmann's study of political activism focuses on medicine and urban planning, but her characterization of professionals' struggles and strategies apply to other fields too. Some professionals changed how services were delivered to address deficiencies in bureaucratic institutions. Some aimed to create movements for social change by workplace organizing and intellectual leadership. And some worked on empowering people:

> Professions were seen as bureaucratic actors in organizations which served dominant interests. Professionalism, defined as a system of dominance and dependence, was the culprit, and deprofessionalization, the solution. The problem for minorities and the poor was not lack of experts and services per se, but awe of professionals and service bureaucracies. The activist roles were to mobilize client communities to help themselves and to democratize service delivery by transferring expertise. (Hoffman 1989, 8)

At the same time, "a small and vocal minority attempted to redefine engineering by rethinking the nature of technology through collective action" (Wisnioski 2012, 6). In the 1980s, computing professionals and academics followed suit and organized as *Computer Professionals for Social Responsibility* (CPSR 2007), galvanized by concerns about the role of computing in nuclear war (Ornstein, Smith, and Suchman 1984). Despite early connections to immigrant labor organizations concerned with working conditions, health hazards, and environmental justice, the CPSR focused on technical arguments about the limits to reliability inherent in computing and initially argued for limitations on computing's reach. This technical focus lent credibility to the moral argument its members made (Finn and DuPont 2020). As the threat of nuclear war faded in the 1990s and the interests of its leadership shifted, the organization struggled to reorient itself. It tried to integrate issues such as privacy, elections, and community networks into its portfolio. For a number of reasons, including financial insecurity, lack of focus, and diminished enthusiasm on the part of the leadership, its membership dwindled, and the organization eventually folded in 2013.

An important shift takes place among all these dissenters: from professional responsibility, defined and shaped by professional obligations and corporate institutions, to a social responsibility informed by social critique and understood as standing in tension to profit motives. In a small study of requirements professionals, many were similarly aware of the tension between what they considered their broader social responsibilities for sustainable IT and the corporate incentive structures they work in today (Chitchyan et al. 2016). Alternative approaches to design have often emerged from outside this corporate context. For example, the practices that crystallized into design justice principles developed in community contexts (Costanza-Chock 2020). Their approach aligns, in key aspects, with Hoffmann's characterization of professional activists. One effect is that it is proving very difficult to transfer them into professionalized corporate contexts (Spitzberg et al. 2020). But we need to introduce different approaches inside the mainstream of systems design practice because that is still where the majority of design happens.

That this shift is happening despite all structural obstacles should remind us: Hobbes was wrong. Mutual destruction is not a foregone conclusion. Since the 2010s, concerns over computing's many disasters are shared by a wider audience than ever, spanning all groups: tech workers, users, executives, researchers, NGOs, lawmakers, and everyone else. These concerns are taken up by new groups springing up all over the globe. There are too many to list, and there is no doubt that these voices are stronger, better coordinated, and sharper in their critique than ever before. Take the successful international campaign against the use of facial recognition technology, which combined the technical expertise of computer scientists with legal scholarship and critical social theory, organizing effectively to intervene in publishing, in policy, and in product development. In a workshop held to debate whether computing needs a new CPSR (Becker, Light, et al. 2020), the view prevailed that computing already has many voices speaking up in critical conscience on substantive matters.

What will we see in the 2020s? In North America, the tech worker movement is growing strong. Summer schools now support tech workers in social responsibilities and organizing (Logic School 2021). Ifeoma Ozoma, who blew the whistle on Pinterest's racism, created a handbook of resources to guide tech workers in speaking up about public issues and helped introduce legislation to protect whistleblowers in California (2021). And in April 2022 as I was completing this manuscript, workers at Amazon achieved a major milestone as they formed their first union. "We want to thank Jeff Bezos for going to space because while he was up there, we were signing people up," lead organizer Chris Smalls said with a triumphant smile. Meanwhile, the #TechWontBuildIt hashtag unites thousands of tech workers rejecting job offers on the ground of moral objections to the potential employers' practices. Will we soon see toolkits allowing socially conscious system designers to vet employers for ethical practices?

One way to think of this shift in how tech workers understand their profile is by contrasting it to conventional ideals of the perfect tech worker touted by business. The most frequently repeated of these ideals is the T-shaped professional, who has deep subject expertise in one narrowly defined technical area and broad shallow understanding of a range of fields (Hansen and Oetinger 2001; Neeley and Steffensen 2018); and the Π-shaped professionals, which combine deep technical expertise with deep domain expertise for some application domain of interest. I think of responsible tech workers instead as W-shaped designers (figure 12.1), who combine technical expertise with social theory, bridged by a critically systemic view, which enables them to work together with community partners and simultaneously form critical friendships to make sure they learn those lessons that may be hard to accept with humility.

A SILICONE RING

Rumor had it that the original iron rings for engineers in Canada were manufactured from the wreckage of the collapsed Quebec Bridge. There is no evidence that this is true, but it's a powerful image. The persistent rumor shows how much we care about the idea that we can be humble and learn from our mistakes.

Computing has already collected too much wreckage. Books on safety, reliability, and security have no shortage of examples on offer. The Cambridge Analytica scandal and others stand for privacy. Weapons of Math Destruction can be found across many areas of algorithmic

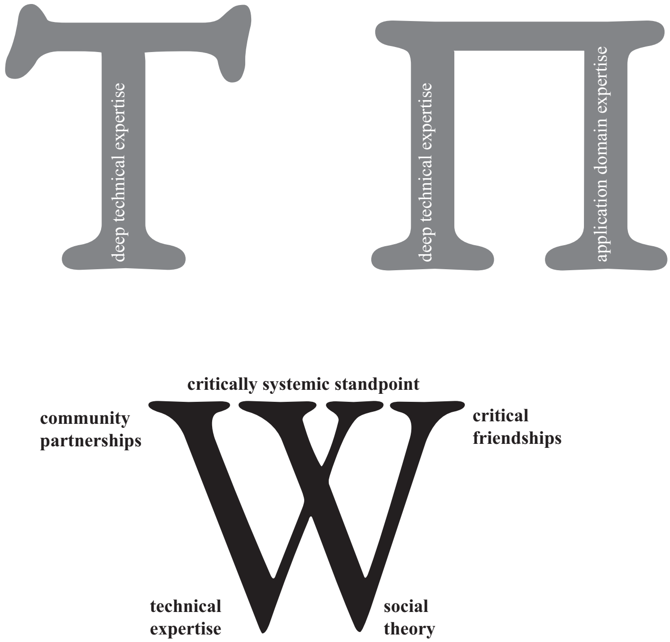

12.1 The W-shaped systems designer.

decision-making, from college admissions and recidivism assessment to credit scores. Facial recognition is so obviously toxic that it's called the "plutonium of AI" (Stark 2019). But because the nature of IT often removes the visible disaster from the time or place where and when a system is designed, constructed, and operated, it is harder to point to a particular smoking gun that spells out "just sustainability." Our collective computing debts are mounting, and we cannot wait, counting it, until all that impact has materialized.

I propose that we draw inspiration from the Iron Ring of Engineers, that we consider what kind of ritual would be appropriate for the twenty-first century. That we begin to develop such a ritual as a symbol of responsibility and a reminder of humility. After all, "it is the prevailing culture of ideas that shapes the direction of a society, that determines what is thinkable and unthinkable, what is possible and impossible. Yes, factors like economic structures, political systems and technology all play vital roles, but never underestimate the power of ideas" (Krznaric 2020, 19).

In 2025, one hundred years after the first Iron Ring graced an engineer's hand, I imagine IT workers here and there sporting a Silicone Ring on their pinky fingers, responsibly made from recycled wafers and CPUs formerly used for bitcoin mining. It means something slightly different for each of them, but carries a shared meaning, too. The Silicone Ring represents not professional, but social responsibility. The allegiance it reminds them of is to our planet. It reminds them that its flourishing is morally loftier than the client and the corporate authority. Those who decide to wear the ring are reminded that information technology is never neutral. Whenever they design, it reminds them that their work is not neutral either. They carry an obligation in the twenty-first century, after all, in which we are facing the biggest crises humanity has ever had to overcome. The ring reminds them to be aware of asymmetric vulnerability and the danger of moral corruption; to reflect on the limits of their technical knowledge; to consider how its exercise shifts power; to design with, not for, those whose knowledge is different but who must live with the consequences of the design; and to critically reflect on the legitimacy of their expertise. Some see it as a reminder of the simple fact that there is no Planet B and that all tech comes from somewhere and goes somewhere. Others simply see it as a symbol for collective action. As Roman Krznaric writes, "the key question is not 'how can I make a difference?' but 'how can *we* make a difference?' A mere shift of pronouns has the power to change the world. The urgency of our current crises demands strategies of change based on collective action directed at those in power more than isolated personal actions" (2020, 205).

CONCLUSIONS

Computing's dominant responses to the climate crises, injustice, and ethical challenges have been to double down on rationalist reasoning and the myths of computing. But codes of ethics will not materially shift or even reorient the practices and impacts of computing. Instead, tech workers are called on to find their collective and individual leverage points and organize collective actions.

We are beginning to witness the emergence of a greater sense of collective responsibility and collective action in tech workers of all kinds, driven by their growing awareness of the urgency and importance of changing social course and by a renewed sense of solidarity and "the collective, the communal, the commons, the civil, and the civic" (N. Klein 2014, 460).

You too can reorient your work: your design, your coding, your modeling, your development, your engineering, your research, your teaching, your organizing, and your volunteering. Everywhere you can, find the levers to apply force to amplify the small changes you have control over. And, in all of this, find allies and make critical friends and work together.

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Becker, Christoph (Director of the Digital Curation Institute), author.
Title: Insolvent : how to reorient computing for just sustainability / Christoph Becker.
Description: Cambridge, Massachusetts : The MIT Press, [2023] | Includes bibliographical references and index.
Identifiers: LCCN 2022038283 (print) | LCCN 2022038284 (ebook) | ISBN 9780262545600 | ISBN 9780262374651 (epub) | ISBN 9780262374668 (pdf)
Subjects: LCSH: Electronic data processing—Social aspects. | Computer systems—Environmental aspects. | Information technology—Social aspects. | Sustainable development.
Classification: LCC QA76.9.C66 B435 2023 (print) | LCC QA76.9.C66 (ebook) | DDC 303.48/34—dc23/eng/20221121
LC record available at https://lccn.loc.gov/2022038283
LC ebook record available at https://lccn.loc.gov/2022038284