This is a section of doi:10.7551/mitpress/14668.001.0001

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

2

JUST SUSTAINABILITIES AND THE DEBTS OF COMPUTING

Unless analyses of development begin not with the symptoms, environmental or economic instability, but with the cause, social injustice, then no development can be sustainable.

— Middleton and O'Keefe (2001, 16)

In "A Perfect Moral Storm," philosopher Stephen Gardiner emphasizes that climate change is a challenge of epic proportions because it is not only globally dispersed in space but also temporally dispersed across generations. The convergence of several characteristics causes an "ethical tragedy." The use of the word *tragedy* is poignant and intentional. The title references the popular book and film about a ship doomed to sink at the confluence of three major storms. In the case of climate change, the ship is us, and the storms are global, intergenerational, and theoretical. The global storm arises from the *spatial dispersion* of cause and effect in climate change. The intergenerational storm arises from the *temporal dispersion* of cause and effect. The theoretical storm, finally, arises from the inadequacy of current theories on "intergenerational ethics, international justice, scientific uncertainty, and the human relationship to animals and the rest of nature" (Gardiner 2014, 7).[1]

The global storm has somewhat familiar characteristics. Causes and effects are spatially dispersed, agency to cause change is fragmented, and vulnerability to the harmful effects of climate destruction is unevenly distributed. The intergenerational storm adds that climate change is resilient to change (at each point in time, current impact is irreversibly caused by the past), that it carries significant inertia (actions in the past have already committed us to significant warming), and that it is *substantially deferred* (the cumulative impact of climate change will be felt across a significant time span). But most seriously, the intergenerational storm introduces strong asymmetry. In principle, the under-resourced victims of climate change that are already alive can exert *some* influence over rich polluters, even if it is weak. But the same is not true at all for agents dispersed in time. Future generations can neither speak for themselves in the present nor coordinate with the past. Instead, each generation exerts influence over subsequent generations but has no way to hold the past accountable. This irreversible asymmetry in influence across generations creates "a new kind of collective action problem" that Gardiner (2014) terms "intergenerational buck-passing" (35). While it is collectively rational for each generation to act, it is individually rational for each generation to pass the buck. Because of *asymmetric vulnerability*, this is particularly challenging: The first generation is asked to make a sacrifice, and its incentive structure is set so that it reaps no reward. If it passes the buck, the subsequent generation faces the exact same situation, only worse. Each generation's inaction worsens the problem.

It is clear that under certain conditions, human societies have very successfully governed shared, common-pool resources for thousands of years. Favorable conditions include an alignment of incentives and timelines across the actors who share the resources, a democratic establishment of binding norms and infringement penalties, and transparency about rules and governance (Ostrom 2015; 2016).[2] But the dispersion of causes and effects across space and time fragments agency and renders existing institutions inadequate (Gardiner 2014). Against the backdrop of these storms, in a public sphere heavily distorted by the interference of fossil fuel lobbying (Oreskes 2010; Supran and Oreskes 2021), Gardiner (2014) warns of the danger of *moral corruption* and reminds us to be skeptical about our own abilities of judgment—we cannot rely on ourselves to simply "invoke and apply" the "correct theories" because "our assessment of theories and their consequences is not made in a neutral evaluative setting" (309). Hence, "an issue of intergenerational justice is at stake, and . . . we are likely—given the perfect storm—to be biased in our own favor" (Gardiner 2014, 430).

SUSTAINABILITY, MEET SOCIAL JUSTICE

Gardiner's emphasis on the moral and ethical nature of climate change drives home a central point: climate change is an issue of ethics and justice. As Agyeman, Bullard, and Evans (2003) write, "the issue of environmental quality is inextricably linked to that of human equality. Wherever in the world environmental despoilation and degradation is happening, it is almost always linked to questions of social justice, equity, rights and people's quality of life in its widest sense" (1).[3] They highlight three aspects. First, inequality within countries is strongly correlated with environmental destruction. Second, the damages inflicted by environmental destruction disproportionately affect those who are already disadvantaged. And third, they endorse Middleton and O'Keefe's epigraph and add:

> Sustainability, we argue, cannot be simply an "environmental" concern, important though "environmental" sustainability is. A truly sustainable society is one where wider questions of social needs and welfare, and economic opportunity, are integrally connected to environmental concerns. This emphasis upon greater equity as a desirable and just social goal, is intimately linked to a recognition that, unless society strives for a greater level of social and economic equity, both within and between nations, the long-term objective of a more sustainable world is unlikely to be secured. The basis for this view is that sustainability implies a more careful use of scarce resources and, in all probability, a change to the high-consumption lifestyles experienced by the affluent and aspired to by others. (Agyeman, Bullard, and Evans 2003, 2)

Sustainability for them is "at its very heart a political rather than a technical construct" (Agyeman, Bullard, and Evans 2003, 6). It is no coincidence that the emergence of the environmental justice movement in the US was spurred by evidence showing that hazardous waste disproportionately affected African American communities (US GAO 1983; Justice 1987), illustrating the pervasiveness of "environmental racism" (Bullard 1993). Environmental justice is also deeply implicated in gendered workplace struggles around health and pollution in the global tech industry (Pellow and Park 2002). As Dorceta Taylor (2000) shows, this context is significant in at least two ways. First, it was people of color who constructed the environmental justice paradigm, and second, this placed "concepts like autonomy, self-determination, access to resources, fairness and justice, and civil and human rights" into central positions of the discourse (534). As Agyeman et al. (2016) summarize, the paradigm "explicitly links the environment to race, class, gender, and social justice, effectively reframing environmental issues as injustice issues" (326). They attest partial success to the environmental justice movement, citing the prominent emphasis on equity and justice in the 2015 Paris Agreement as evidence and pointing out that the revised UN Sustainable Development Goals (SDGs) (United Nations 2015a; United Nations General Assembly 2015) "read like an institutional agenda for just sustainabilities" (Agyeman et al. 2016, 335). It is true that the seventeen SDGs range from gender equality and the reduction of inequalities to "justice for all" and include such metrics as 16.7, "responsive, inclusive, participatory and representative decision-making at all levels" (United Nations 2015b). Considering that as recently as 2013, *Nature* published an article on SDGs without the word "justice" that only mentions "equity" in terms of gender (Griggs et al. 2013), this is certainly progress. But critics of the framing of sustainable development remain highly skeptical of the willingness and ability of entrenched institutions to facilitate the kind of transformative change required to make these goals come true (e.g., N. Klein 2014; Escobar 2018; Kothari et al. 2019). Instead, they argue for the need to develop "alternatives *to* development, rather than development alternatives" (Escobar 2011, xiii), and we will meet their proposals later in this book. In parallel, "there remains an issue of the historic injustice imposed on the poor and the poor countries of the planet. There is a vast historic ecological debt" (Agyeman, Bullard, and Evans 2003, 30).

Agyeman, Bullard, and Evans's (2003) conception of *just sustainabilities* distills the convergence of social justice and environmental sustainability into an approach that aims to "ensure a better quality of life for all, now, and into the future, in a just and equitable manner, while living within the limits of supporting ecosystems" (2). The relationship between the environment and justice is complex and mutual: "A poor environment is not only a symptom of existing injustice; rather, a functioning environment provides the necessary conditions to achieve social justice" (Agyeman et al. 2016, 335). Others add that justice, equity, and inequality are in fact leverage points toward ecological sustainability:

> We are made to choose between human welfare or ecological stability—an impossible choice that nobody wants to face. But when we understand how inequality works, suddenly the choice becomes much easier: between living in a more equitable society, on the one hand, and risking ecological catastrophe on the other. Most people would have little difficulty choosing. Of course, achieving this will not be easy. It will require an enormous struggle against those who benefit so prodigiously from the status quo. And presumably this is why some are so eager that we avoid this course of action: they would prefer to sacrifice the planet in order to maintain the existing distribution of global income. (Hickel 2020, 144)

That sacrifice, after all, happens in a future that cannot reach back to hold them accountable. In other words, "justice is the antidote to the growth imperative—and key to solving the climate crisis" (Hickel 2020, 137). Historians, anthropologists, and economists alike have convincingly demonstrated that capitalist growth always comes at an enormous cost to others. Economic growth has to be fueled by resources that come *from somewhere* (Hickel 2020). In five hundred years of capitalism, growth has come variously from the enclosure of common public ground in Europe that forced commoners into labor contracts, from the extraction of raw materials at unfair conditions imposed by colonialism, from the forced unpaid labor of slaves, or from the extraction of fossil fuels. Throughout the centuries, capital has found the next frontier of growth whenever growth stalled, from deregulation to the extraction of data that describes the bodies and lives of people. The justification of continued growth and decoupling claims that our societies *need* economic growth and that everyone benefits from growth. But beyond a threshold long surpassed by the Global North, only a minuscule portion of our world's population actually benefits from economic growth.[4] The widely told story of growth leading to improved health and well-being has long been debunked. Instead, "progress in human welfare has been driven by progressive political movements and governments that have managed to harness economic resources to deliver robust public goods and fair wages. In fact, the historical record shows that in the absence of these forces, growth has quite often worked against social progress, not for it" (Hickel 2020, 128).

To think that the argument to abolish economic growth as an objective—degrowth—would be an argument for austerity and sacrifice would be a misinterpretation. It refers to a lean metabolism much rather than starvation:

> The term "degrowth" may be confusing if perceived as negative GDP growth. To be clear: degrowth is not negative growth. The goal of degrowth is not to make GDP growth negative. There is a name for that: "recession" or, when prolonged, "depression." In economics terms, degrowth refers to a trajectory where the "throughput" (energy, materials and waste flows) of an economy decreases while welfare, or well-being, improves. (Kallis 2018, 9)

The growth of economic activity, measured in GDP, is not a necessary condition for humanity to flourish; it is a necessary condition for capital to thrive and expand. Under the current economic arrangements, national economies and many companies are pressured to grow to survive their debt load, so a reduction in growth is felt through job cuts.[5] Degrowth restructures our societies' economic arrangements so that we do not *need economic* growth. It reverses the artificial scarcity on which capitalist economies depend. "Austerity calls for scarcity in order to generate more growth. Degrowth calls for abundance *in order to render growth unnecessary.*" (Hickel 2020, 168). The call is strongly linked to arguments to decolonize not just social and political relations but also the way we conceive of the relationship between us humans and the rest of nature, away from a view of nature as a lifeless resource, toward an understanding of interdependent relationships—a view that learns from Indigenous worldviews what Western science has pushed to the margins (Hickel 2020; Escobar 2018; Maturana and Varela 1992; Kimmerer 2013; Kothari et al. 2019). These perspectives offer new directions for reorienting and designing sustainable economic activity on our planet, including technology design.

SOCIAL JUSTICE IN COMPUTING AND SYSTEMS DESIGN

The term *social justice*—"justice in terms of the distribution of wealth, opportunities, and privileges within a society" (Oxford 2020e)—has only relatively recently been adopted as an explicit framing in computing discourse (Dombrowski, Harmon, and Fox 2016), but fields such as computer-supported collaborative work (CSCW) have long been invested in the issues it raises (Fox et al. 2017). According to Fox et al. (2016), "As a perspective, social justice explicitly takes into account how the historicity, situated context, and social issues (e.g., class, race, gender, ability, sexual orientation, health and wellness, food access, and so on) impact people's experiences . . . , including how technology is designed and developed, how policy impacts information and communication practices and experiences, and how marginalization and oppression impact people's experiences of and practices with technology" (3294). This perspective has been a central driver of recent work critiquing the booming industry that develops and applies machine learning algorithms at large scales to the classification of human beings in capitalist societies.

Patricia Hill Collins (1990) introduced the term *matrix of domination* in *Black Feminist Thought* to describe how dimensions such as gender, race and class that influence power differentials are connected and organized across four domains (structural, hegemonic, disciplinary, and interpersonal). In current views, the dimensions include gender, class, race, ability, sexual orientation, religion, geography, age, and other factors. The related concept of *intersectionality* (Crenshaw 1991) highlights how the position of individuals at specific locations—at intersections of these dimensions—affects power and privilege in ways that are not captured by an additive accounting of dimensions. Here, as elsewhere, as systems thinker and management scientist Russell Ackoff (1999b) put it, the system is not the sum of its parts but must be understood as the product of their interactions. Intersectional and Black feminist scholars who examine computing from this perspective emphasize that the uneven distribution of power and privilege in society should be conceived on many dimensions. It is notable and hardly a coincidence that not one of the books cited in the next section is authored by a White man.

When Google indexes the web, it shapes the web, and with its ranking of search results, it shapes our media landscape, market dynamics, and politics. The enormous influence of this and other algorithms has become an object of public discussion over the past years. In *Weapons of Math Destruction*, data scientist Cathy O'Neil (2016) explains and demonstrates how algorithms using big data reinforce inequity and existing power structures to the detriment of justice, fairness, and equity. The weapons of math destruction (WMD) metaphor arises from the combination of three characteristics often found in data science models: opacity, damage, and scale. Large-scale models create their own feedback loops, not merely reproducing the inequitable status quo but amplifying and accelerating it to inflict damage at scale. WMD afflict individual lives across all stages, from college admissions and credit scoring to job applications, credit loans, the criminal justice system, and parole decisions. In *Algorithms of Oppression*, Safiya Noble (2018) examines the role of search algorithms in particular and demonstrates how they reinforce racism and oppress women of color, in an excellent illustration of just how strong these reinforcing feedback loops can be. Latanya Sweeney (2013) had demonstrated racial discrimination in online ad delivery much earlier, but nothing had changed since then.

But algorithmic sorting, control, and discrimination are not limited to global platforms. Any computational system classifying its users raises comparable concerns that could in principle be addressed in its design but often aren't. In *Automating Inequality*, Virginia Eubanks (2018) shows, based on long-running in-depth investigations, how the design and use of software systems in social services profiles the poor across the US and reinforces their marginalized status. Inevitably, the material implications of algorithmic discrimination and oppression are unevenly distributed.

> I am often approached by engineers or data scientists who want to talk about the economic and social implications of their designs. I tell them to do a quick "gut check" by answering two questions: Does the tool increase the self-determination and agency of the poor? Would the tool be tolerated if it was targeted at non-poor people? Not one of the technologies I described in this book rises to this feeble standard. We must demand more. (Eubanks 2018, 173)

In the US context, the racist tendencies of machine learning algorithms reinforcing existing systemic inequalities and structural racism have become an object of public debate at last. In *Race after Technology*, Ruha Benjamin (2019) provides important historical context and traces the historical trajectory of technology from a Black feminist perspective. She shows how "technologies reflect and reproduce existing inequities but . . . are promoted and perceived as more objective or progressive than the discriminatory systems of a previous era" (15).

To grasp how this takes place and to change it, we need to combine a technical understanding of how they work with a critically appreciative understanding of their social construction, the political forces they are subject to, and their historical evolution. Computer science has been notoriously keen on avoiding this conversation. Notably, one of Eubanks' s (2018) pledges is to "design with history in mind," because to do otherwise is to be "complicit in the 'unintended' but terribly predictable consequences that arise when equity and good intentions are assumed as initial conditions" (174). The crucial point is that what the proponents of large-scale technology interventions often call "unintended" consequence or describe as "unforeseeable" is in fact not unpredictable at all from other standpoints.

These arguments have given significant visibility to the deeply problematic role of computing in maintaining societal structures of power and oppression, made the underlying historical trajectory visible, and explained it in accessible ways to a broad audience. They also call for different approaches to technology design, and recent work has moved from a focus on critique to speculating about new directions for design (Rosner 2018) and to developing generative principles for action. In *Design Justice*, Sasha Costanza-Chock (2020) draws on a wide array of practices and examples to illustrate how an intersectional awareness of the matrix of domination can be brought to bear directly on the practice of design (interpreted broadly with a focus on software systems). They highlight that "the matrix of domination is constantly hard-coded into designed objects and systems. This typically takes place not because designers are intentionally 'malicious' but through unintentional mechanisms, including assumptions about 'unmarked' end users, the use of systematically biased data sets to train algorithms using machine-learning techniques, and limited feedback loops" (Costanza-Chock 2020, 71).[6] The aim is explicitly to "retool design," drawing inspiration from activist and advocacy work from areas including disability justice.

Similar developments are under way in data science. In *Data Feminism*, Catherine D'Ignazio and Lauren F. Klein (2020) relocate and reconstitute data science on feminist ground: "Because the power of data is wielded unjustly, it must be challenged and changed" (14). Correspondingly, data feminism is built on seven principles: (1) examine power, (2) challenge power, (3) elevate emotion and embodiment, (4) rethink binaries and hierarchies, (5) embrace pluralism, (6) consider context, and (7) make labor visible. The authors emphasize that data feminism is not only about and for women, and not only about gender but about power (2020, 14). They discuss "What makes a project feminist? . . . a project may be feminist in *content*, in that it challenges power by choice of subject matter; in *form*, in that it challenges power by shifting the aesthetic and/or sensory registers; and/or in *process*, in that it challenges power by building participatory, inclusive processes of knowledge production" (2020, 18). This sense resonates strongly with the direction taken later in this book.

These arguments reflect a highly overdue confrontation in the computing discourse between its traditional modes of reasoning and a critically informed, intersectional feminist perspective.[7] A wave of practice in design and advocacy accompanies these writings too, with numerous groups forming to collectively address specific harms. Too numerous to list them all, these efforts range from research institutes such as the NYC-based AI Now Institute and researcher-driven efforts such as the Algorithmic Justice League founded by computer scientist and digital activist Joy Buolamwini to federated groups of researchers and practitioners such as the Design Justice Network and to organized labor, as in the Tech Worker Coalition. Such efforts of collective organizing and advocacy regularly perform successful interventions. In 2020, *Springer Nature* was forced to rescind an offer to publish a piece of work that can only be described as neophrenologist. It presented a machine learning algorithm that, its creators claimed, successfully "predicted" criminality based on photographs. The retracting came after a social media campaign gathered 2,500 signatories for a compellingly argued and heavily substantiated letter.

> Machine learning programs are not neutral; research agendas and the data sets they work with often inherit dominant cultural beliefs about the world. These research agendas reflect the incentives and perspectives of those in the privileged position of developing machine learning models, and the data on which they rely. The uncritical acceptance of default assumptions inevitably leads to discriminatory design in algorithmic systems, reproducing ideas which normalize social hierarchies and legitimize violence against marginalized groups . . . any effort to identify "criminal faces" is an application of machine learning to a problem domain it is not suited to investigate, a domain in which context and causality are essential and also fundamentally misinterpreted. (Coalition for Critical Technology 2020)

Because of such scholarly advocacy work, the racially biased and discriminatory role of technology in our societies is now well established. Still, despite these voices, the mainstream discussions of social justice in computing, especially in machine learning, continue to frame it narrowly and reductively. Instead of acknowledging the contextual, historical, sociotechnical nature of *algorithmic bias*, the issue is reframed as a computational problem. On that basis, "garbage in, garbage out"—the metaphor that bad data, especially biased data, leads to bad outcomes, for example biased classifiers—is used as an argument to deflect attention from broader issues and refocus it into a technical question. A satirical piece illustrates beautifully how impoverished and dangerous this is. *A Mulching Proposal* unflinchingly describes in technical terms a corporate effort to make completely fair, accountable, and transparent a machine learning algorithm that selects who out of a general human population should be picked up by a drone and turned into food. Initial concerns about the algorithm disproportionately targeting White people are quickly addressed by adding people of color to the training data. State-of-the-art metrics of computational fairness are successfully deployed to ensure fair outcomes: everyone is equally likely to be processed. The terrorizing project it describes is technically consistent with the revised ACM Code of Ethics. Since its latest revision, the ACM Code allows systems to be built that cause intentional harm, and the Code places responsibility for deciding whether to build these systems explicitly into the hands of the system designers, with no call for external legitimation (see chapter 12). That is exactly what *A Mulching Proposal* does, to striking effect (Keyes, Hutson, and Durbin 2019).

I want to highlight a few observations. First, many researchers and practitioners in computing reduce concepts with rich social history, such as *fairness* and *justice*, to algorithmic definitions. The well-meaning computer scientists at Dagstuhl whom I mentioned in the Introduction are not outliers; they are the norm, but their impoverished reframing of justice is utterly inadequate (Selbst et al. 2019; Jacobs and Wallach 2021). Only recently has this balance begun to shift at venues like the FAccT conference series on "fairness, accountability, and transparency in sociotechnical systems."[8] Many of these authors have emphasized the historical, political, and social nature of computing from a perspective located outside of computing, but in recent years, writers such as O'Neil, Timnit Gebru, Buolamwini, and others have written about these issues from a computing insider perspective as well. This has made it much harder for those in computing to ignore or silence their voices—though not for lack of trying (Hao 2021b; Vincent 2019). These courageous researchers have chosen the difficult route of writing *for computing* and working *in computing* for a better, more just, computing.

Second, justice has also been a concern in the sustainable development discourse, and as mentioned earlier, some work in ICT4S has mentioned the distributive justice aspect of technology development. But these concerns have remained at the margins. As late as 2022, the conference on ICT4S—the premier venue focused on such work—included no paper on environmental justice, social justice, or just sustainabilities. And even the socially conscious transformation mindset discussed in chapter 1 does not mention race, class, gender, or power. *Just sustainabilities* incorporate social justice into the sustainability discourse. Wider aspects of social justice remain outside of this framing.

Yet, based on the burgeoning literature on social justice issues and computing, and a range of communities working on issues of sustainability in computing, the nascent attention to the nexus of all three is apparent. For example, the *Smell Pittsburgh* project by Hsu and colleagues (Hsu and Nourbakhsh 2020; Hsu et al. 2017) adopted a community-centered approach to citizen science that empowered communities to hold corporate polluters accountable through collecting and visualizing air quality data. The project is situated explicitly within the landscape of sustainable human–computer interaction, and it makes direct reference to environmental justice. On a broader scale, the threat to the Environmental Protection Agency (EPA) that arose from the 2016 US election galvanized activists and academics into forming the Environmental Data & Governance Initiative (EDGI). The organization initially coordinated distributed efforts to save EPA data from impending erasure and continued to monitor the evolving situation to hold the US government accountable. The framework of *environmental data justice* fuses the value systems of environmental justice and data justice to pursue "public accessibility and continuity of environmental data and research, supported by networked open-source data infrastructure that can be modified, adapted, and supported by local communities" (Dillon et al. 2017). Vera et al. (2019) emphasize that environmental justice calls for approaches that "replace extractive data practices with community-based participatory research and community science projects," just what Hsu and colleagues aimed for, but as the authors argue further, the "desire for more data not only demands that harms be proven through a technically narrowed mode of evidence legible to dominant institutions, but also often requires communities to generate evidence of environmental violence from their own suffering" (1015). In their view, environmental data justice cannot resolve this tension but commits to ongoing critical reflection.

REALITY CHECK: THE DEBTS OF COMPUTING

If we could take stock of computing's balance with our worlds, what would we find? Like any other technology, software is never a neutral element of our societies. Software filters, sorts, and selects; it tracks, monitors, and intervenes; it initiates and shuts down and controls. In extreme cases, it passes judgment (Julia Angwin et al. 2016; Benjamin 2019; Eubanks 2018). It reinforces existing power dynamics, amplifies structural differences, and automates the reproduction of racism, inequality, and inequity. In doing so, it embeds the values and politics that underpin the choices of its designers and the context in which they make them.

But software is not just an element of oppression; it also is a tool of liberation. Airbnb was mentioned earlier as a software-based platform with far-reaching effects, deeply implicated in the erosion of housing availability. We know this because of data-driven investigations that equally rely on, and build, software technology in examining the role of Airbnb's platform in housing markets across the world from a critical angle. The work of urban geographer David Wachsmuth, for example, crucially involves the development of algorithmic de-obfuscation tools to make data analysis possible (Wachsmuth and Weisler 2018).

On a broader scale, civic movements across the world develop open-source software systems to securely and safely organize and coordinate peaceful protests; journalists rely on some of the same systems to report on human rights violations under dangerous conditions; and software systems are used to support environmental sustainability initiatives that range from supporting permaculture (Norton, Penzenstadler, and Tomlinson 2019) to community citizen science initiatives reporting pollution (Hsu and Nourbakhsh 2020) and safeguarding evidence of climate change from corrupt regimes (Vera et al. 2019). We might call these "algorithms of liberation" (Roberts et al. 2018). They demonstrate that, in principle, significant room for maneuver is available to those designing and developing software systems. The emergence of tech workers who take their disagreement with company policies into the open to protest unethical practices attests to an emerging social consciousness and a rising desire for such change in the IT industry (see chapter 12).

Despite these promising efforts and initiatives, however, it seems clear that the harmful direct and indirect effects of software technology development are not shrinking but rising. I call these the *debts of computing*. In fact, most of the initiatives mentioned here that strive for change are issues-centered efforts to address precisely some of those debts, to make them visible against the efforts of the tech industry, and to prevent the worst outcomes. The *Ledger of Harms* operated by the Center for Humane Technology lists the costs explicitly: “Under immense pressure to prioritize engagement and growth, technology platforms have created a race for human attention that’s unleashed invisible harms to society. Here are some of the costs that aren’t showing up on their balance sheets.” The harms listed in the evolving catalog range from algorithmic oppression and negative effects of tech on cognitive abilities to misinformation, social isolation, and threats to democratic governance (Center for Humane Technology 2020).

In software engineering, the term *debt* has received significant attention in the form of *technical debt*, which refers to “design or implementation constructs that are expedient in the short term, but . . . make future changes more costly or impossible” (Avgeriou et al. 2016, 112). In this view, “technical debt presents an actual or contingent liability whose impact is limited to internal system qualities” (112). Technical debt is a burgeoning research area with its own dedicated conference series. The term has been very effective because it explores the meaning of a domain that is hard to grasp (the temporally diffuse effects of design choices) by mapping key elements onto a relatively well-known and formalized domain (finance). That is the nature of metaphor (Lakoff and Johnson 1980; Lakoff 1993). Far from being merely figures of speech or poetic constructs we use to adorn our writing, metaphors are the conceptual mappings by which we view domains we cannot fully explain through the conceptual lens of domains we understand. This domain mapping is central to the structure of human thought and language, and to some degree we can trace them. Consider this statement: “Our new colleague gave a talk yesterday. It was full of interesting ideas, but I wasn't sure what to take away from it." A talk is an intangible, temporally bound event, hard to grasp as an object of thought. Yet we describe it as if it was a tangible object that could be handed from one person to another: Our colleague *gave us* the talk. More so, in the example the talk is treated as a container, with ideas being objects inside the container. Like marbles from a bowl, we assume these ideas can be taken out, and carried away, by the audience. Through layers of such conceptual mapping, metaphors ultimately ground our conceptual understanding of the world in our lived physical, bodily, and cultural experience. This applies not just to colloquial speech but equally to scientific theories and concepts (see also Morgan 2006; M. Jackson 2003), so it is worth reflecting more deeply about the implications of metaphors in design (Blackwell 2006; Becker 2018a).

In the case of technical debt, the metaphor enabled a community to talk in clearer terms about issues that had been bothering researchers and practitioners since the founding of the discipline (Naur and Randell 1969; Lehman 1979; Parnas 1994).[9] When it comes to technical debt, the conceptual mapping suggests that it may eventually be *paid back*, but it also carries additional implications. When formerly expedient design constructs get in the way of further work, the additional effort is understood as interest. This financial concept places the additional burden in a proportional relationship to the presumed benefit gained from the expedience of prior choices. Indeed, debt can be taken on strategically. The literature of technical debt has been explicit about this framing, developed an elaborate vocabulary around it, argued for the positive leverage to be gained from strategically incurring technical debt, and has even started to pull in mathematically sophisticated approaches from finance to address the problem of "managing technical debt" as if it were a financial portfolio. Technical debt is thus on its way to becoming a scientifically minded "theory" in software engineering.

One of the efforts made to leverage the debt metaphor further, *sustainability debt*, refers to "the hidden effects of past decisions about software-intensive systems that negatively affect [the] sustainability of the system under design" (Betz et al. 2015). This angle extends the established metaphor of technical debt outwards, drawing on the systemic effects of software systems visualized in figure 1.2. Where technical debt is about temporal dispersion of effects that remain firmly in the court of the design team, sustainability debt expands the timespan and adds social distance to the picture: The hidden effects of design decisions are not borne out by the designers in the future; they are inflicted upon the world at large. When an affluent business traveler pays a premium to book an inner-city Airbnb for a conference trip, they contribute to the displacement of affordable housing by highly profitable short-term rentals. This remains an individual choice, but it is enabled, shaped, and encouraged by the specific features in the design of Airbnb's platform. Reinforcing feedback loops between consumer behavior and investor capital make this option attractive for travelers and investors alike. The aggregate structural impact is felt by those urban populations who can no longer afford inner-city rents in busy global cities.

Perhaps the most egregious example of externalized debt is bitcoin. In 2022, bitcoin alone generated about 40,000 tons of electronic waste per year (Digiconomist 2022), comparable to the Netherlands (cf. de Vries and Stoll 2021). Already in 2018, "each $1 of Bitcoin value created was responsible for $0.49 in health and climate damages in the US and $0.37 in China" (Goodkind, Jones, and Berrens 2020), for an overall 1:1 matching of market valuation and material damage. Who carries the damage? Not those who invest spare resources in bitcoins or bitcoin mining, but those marginalized populations who are already disproportionately exposed to the health hazards of pollution and climate change.

COMPUTING HAS EXTERNALIZED ITS DEBTS

The metaphor of debt breaks down here at a crucial moment, for the debts of computing—the hidden, delayed, and remote effects of systems design decisions on the world—continue to be externalized. They are offloaded to and paid for by others. Computing is not paying back these debts to societies, nor does it carry any of the interest. Far from it. Instead, from the health costs paid by the copper miners to the higher credit rates paid by African Americans, those who bear the burden of interest for the generated profit, financial or otherwise, are those at disadvantaged locations of the matrix of domination and also often displaced in time or space. They are the least likely people to be involved in technology design themselves. This is evidently unjust. Critiques of this externalization are sometimes met with extraordinary cynicism from those in positions of power, going as far as inventing the term "underpollution" in the context of e-waste:

> In 1991, Larry Summers, then Chief Economist of the World Bank (and now President of Harvard University), spoke of the economic sense of exporting first world waste to developing countries (Summers, 1991). He argued that the countries with the lowest wages would lose the least productivity from "increased morbidity and mortality" since the cost to be recouped would be minimal; the least developed countries, specifically those in Africa, were seriously underpolluted and thus could stand to benefit from pollution trading schemes as they have air and water to spare; and that environmental protection for "health and aesthetic reasons" is essentially a luxury of the rich, as mortality is such a great problem in these developing countries that the relatively minimal effects of increased pollution would pale in comparison to the problems these areas already face. (Widmer et al. 2005, 437–438)

This is not a book about the material extraction and supply chains of IT (Frankel, Chavez, and Ribas 2016; Reid-Henry 2012; Crawford and Joler 2018; Crawford 2021) or the continued colonial offshoring of environmental debt (Fuchs, Brown, and Rounsevell 2020) and e-waste (Wang et al. 2020). But it is important to bring this context to attention before we examine how computing's way of framing its domains of interest pushes this destructive impact into the background of systems design practice.

CONCLUSIONS: JUST SUSTAINABILITY DESIGN CHALLENGES

The questions of sustainability cannot be treated as detached from social justice concerns. The framing of *just sustainabilities* serves as a reminder that the two are inextricably entangled, and that their conceptualization must be pluralist, making space for diverging meanings in local contexts (Agyeman 2013). These are important shifts in understanding. Not all social justice concerns are about environmental sustainability, however. Systems design in computing today always must be aware of its never-neutral role in shifting power and causing harm to those at the margins of the process. Because harms are often dispersed temporally and spatially, and because the vulnerability to these harms is always asymmetric, it is just too easy for those who design to disregard the debts that their choices incur. Economically, this is because the debts are externalized: foisted on others. Cognitively, the actors who carry them are located at a psychological distance. Ethically, those who design are typically located at privileged positions in the matrix of domination, so they are especially prone to the dangers of moral corruption. Just sustainability design must attend to the challenges that arise from the multifaceted nature of these concerns.

Modeling challenges arise out of the difficult interrelationships that bridge social, technical, economic, human, and nonhuman domains. Systemic approaches to grappling with these interrelationships have made major strides over the past decades and allow us to articulate in increasingly nuanced ways the dynamic causal and correlational factors of just sustainability in given design contexts.

Cognitive challenges are complicated by our limited understanding of individual and social cognition. We do not have a robust theory to explain how systems design decisions are made in situations where their outcomes are uncertain, ambiguous, and located far away in time and space at a psychological distance to those involved in design. How do teams handle these cognitive challenges, where delayed and accumulating impact comes only clearly into vision as distance decreases? We cannot hope to develop effective methods that prescribe how to design for sustainability and justice without such a theory.

Ethical challenges, finally, arise out of the spatial and especially temporal dispersal of effects, combined with asymmetric vulnerability and the inadequacy of theories. Together, they raise the specter of moral corruption and remind us that we're literally "biased in our own favor" (Gardiner 2014, 430). Most importantly, they remind us that science and technology do not provide the conceptual tools for their own ethical justification. For all its power, the monological form of deductive reasoning manifested by the scientific method has no access to moral judgment. "In this way, science loses sight of its potential role in the search for nonoppressive forms of culture and society. It cannot even enter into dialogue with other forms of knowledge given its de facto claim to have the monopoly on knowledge, compassion, and ethics" (Escobar 2018, 89). Instead of insisting on the formalization of ethics, we have to defer for these judgments to other perspectives and forms of reasoning. But which?

Given the contributions of computational thinking and modeling in conjunction with system dynamics and related approaches, systems design in computing is relatively well equipped to deal with the modeling challenges to anticipate dispersed direct and indirect outcomes at multiple scales. Many difficulties and challenges remain (Kienzle et al. 2020), but these challenges on their own are of a type of complexity that computing is well-equipped to handle. A myriad of collaborations is under way to tackle these types of concerns. I will not attempt to contribute much to advancing these issues but instead focus on the cognitive and ethical challenges. This is because they are so dominant in their influence that progress in the first area will hardly make a dent on sustainability and justice unless we gain more effective insights on them. The computing discourse has little to offer on these questions, so we need to look elsewhere.

In the meantime, computing has managed to externalize its debts, and so far, it has gotten away with it. Part of what is happening here is the corruption of discourse itself. The reduction of "wet" concepts like environmental protection and justice to a "dry" computational problem conveniently allows many in computing to avoid facing an uncomfortable realization: that the present state of technology is a living historical reality that continues to reproduce historically grown power structures, with all the privileges, inequities, and injustices they incorporate. Aside from the incentive structures that make it possible and desirable for those in positions of privilege to pursue this framing, the next chapter will show that the pervasive resistance to acknowledging the wider issues is supported by the metaphors and narratives that structure what computing considers its domain of reasoning.

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