Chapter 7
Toward a Praxis of Information Justice

Abstract This chapter summarizes the arguments of this book, situating them amidst the booming literature on information ethics that has emerge over the (too) long process of writing. Unfortunately, nothing like a full theory of information justice has emerged from this, but we can now see important considerations for how we might think about information within what we already know about justice. That presents several possibilities for theoretically-informed action and action-oriented theory. I also suggest a range of possible principles, policies, practices, and technologies that are worthy of a deeper look that can engage data scientists, citizens, and governments. Ultimately, however, information justice (like political justice generally) is not likely to be something that can be established solely by easily executable principles. It will necessarily involve an information justice movement.

The central argument of this book has been the need to view information ethics questions as matters of justice. After establishing a critical-constructive understanding of technology in Chaps. 1 and 2 I turned to the politics of information by studying two cases in which information ethics questions are prominent: the effort to open data to public access and the use of predictive analytics in higher education, showing that both present questions traditionally understood to be questions of political and social justice. In Chaps. 3 and 4 I examined information technologies as political practices, showing that data systems exist as socio-technical translation regimes transform observations into data states constructively rather transcribing them objectively. Those translations include not just the atomizing, normalizing, and unifying translations of characteristics but the translation of the subjects tracked in data systems into mere bundles of information, termed "inforgs" by Floridi. Chapter 4 showed that data systems practice politics, first by encoding data through the translation regimes but then by decoding and institutionalizing metrics using the encoded data. In both cases, social factors are at least as important to the ultimate form of the information stored in and extracted from data systems as technical ones. The institutionalized metrics then play very traditional political roles, distributing moral and material goods and structuring the power of the state and political actors. Chapters 5 and 6 explored how philosophical conceptions of justice help us understand this political view of information technologies. Chapter 5 demonstrated that both instrumental and distributive views of justice are helpful to understanding information

J.A. Johnson, *Toward Information Justice*, Public Administration and Information Technology 33, https://doi.org/10.1007/978-3-319-70894-2_7

privacy, but encounter challenges at the well-known limits of distributive justice generally. Chapter 6 showed that understanding information justice necessarily has a structural component that considers how information both shapes and is shaped by the social structures that support self-determination and self-development.

At the outset of this book, I noted that I did not set out to fully develop a theory of information justice, so I have no shame in admitting that I didn't arrive at one.[1] But this book has clearly shown that one is needed. While it is tempting to say that we need only to think about how information is political and how what we already know about social and political justice applies to that, I think the remarkable success of the environmental justice field shows that there is much to be gained by building a coherent theory of information justice as well. In this conclusion I look toward that, considering it from two directions. The first looks at the burgeoning literature on information and data ethics and justice that emerged late in the process of developing the arguments in this book. Powerful social critiques of contemporary information technologies have led to a scholarly consensus against a neutral, realist understanding of data. Those critiques have spurred the development of several other efforts to develop a theory of information justice. The other direction that we need to consider is praxis: what can we do to bring about information justice? I consider principles and practices that could be developed further: making information politics explicit, protecting normative and contextual validity, encouraging information participation and foundational open data, understanding information pluralistically, and developing processes for information federation. These and other such approaches are increasingly institutionalized in codes of professional ethics for data scientists. Taken together, these two directions posit four principles for the praxis of information justice. But ultimately I conclude that the nature of information technologies is such that the challenge of information justice is a political challenge in which principles and policies are supportive tools for an ongoing information justice movement.

7.1 A Theory of Information Justice

This project has suggested the broad outlines of the issues a theory of information justice needs to address, but it has by no means arrived at such a theory, and even if it did it would not be the only approach to information justice. Since this project began, what we know about information technologies as social and political practices—even just what has been published in peer-reviewed form, let alone the blogs, conferences, and tweets where the most innovative discussion is taking place—has been growing faster than it can be consumed. Incorporating that evolving literature into the argument of the entire book would delay the book even further than it has

[1] That is, of course, artistic license on an author's part. I wrote the preface the day before I wrote the conclusion, so it was easy to make 4 years of writing appear to cohere nicely. Both, however, genuinely do reflect my intentions in writing this book.

been; at some point, an author must draw a line and just write both for productivity's sake and to preserve the sanity of one's editor. But a decent respect for one's colleagues demands at the least that I situate my work in relation to this emerging work, especially when it is of such uniformly high quality. Two themes have emerged in that literature that influence where the arguments I've made should move forward: the consensus on the subjective and constructed nature of data and the elaboration of what it means to consider information from the perspective of justice. I consider all of these quite promising, indeed exciting, developments in moving toward information justice.

It should now be accepted as consensus that data is subjective rather than an objective representation of reality. A growing body of work on data, metrics, and algorithms has approached information from views consistent with the critical-constructive perspective that I developed in the early chapters of this book, rejecting the neutrality thesis and suggesting that data is inherently constructed. Many big data applications show that algorithms' content reflects a political economy dominated by large corporate interests (Pasquale 2015) and the spaces of contestation in which they operate (Crawford 2015). The problems, knowledge, and actors algorithmic actions include are mutually constitutive rather than independent as the realist view of data would suggest (Introna 2015). The "city of visualised facts" that comes from such a realist and instrumental view of data obscures the assemblages that constitute metrics, benchmarks, and dashboards (Kitchin et al. 2015). Information is seen as a political tool, unique to bureaucratic forms of government and rooted such regimes' need to make their subjects legible to the apparatus of authority by transforming an underdetermined reality into standardized, aggregatable, static facts that are capable of consistent documentation (Scott 1998, pp. 80–81). This is, of course, a central requirement of the processes of rationalization that McMillan Cottom and Tuchman (2015) address, especially of accountability regimes, that are emerging in contemporary higher education. Ultimately, predictive analytics has been described as "opinions embedded in math" that undermine democracy and equality (O'Neil 2016) and as "money laundering for bias" (Cegłowski 2016). In all of these views, the process by which data comes to exist is driven by social factors and riddled with at best unexamined assumptions and values and at worst dangerously prejudicial biases that remain hidden behind claims of objectivity and neutrality. My work, and more so that of these colleagues (much of which is far more rigorous that I have done here), should put to rest the idea that one can't argue with the data. Indeed, one must argue with the data in order to use data meaningfully and effectively.

The most exciting recent development, however, comes in the form of other theories of information justice. I am enthralled to say that I am no longer the world's leading expert on information justice by default. Several excellent frameworks have emerged that are defining this emergent field.[2] Prinsloo and Slade (2017) argue that

[2] In addition to the published approaches I discuss here, there are approaches that are currently in working form, such as Taylor (2017), that I look forward to seeing in final form in the near future that, out of fairness to a work in progress, I will refrain from critiquing here.

an ethics of justice for big data in higher education must be complemented by an ethics of care that cannot be subsumed into the former. The note in particular that concepts of justice, especially those related to fairness and desert, are already central parts of many big data applications sorting people one normative as much as empirical categories (e.g., self-motivated or unmotivated). As such, they extend the view of information justice that I have previously articulated (2014b) by unpacking the actors involved in claims to information justice:

> In considering information justice as a useful heuristic for engaging with the complexities of the collection, analysis and use of student data, it is crucial to also raise (if not address) the issue of "whose justice" is served by our definition of information justice—students, faculty, the institution or society. In the context of the asymmetrical power relationship between students and the providing higher education institution, it is a real possibility that an institution's perception of information justice is determined by the reporting and compliance regimes of various regulatory and legal frameworks. (Prinsloo and Slade 2017, pp. 113–114)

This compels consideration of students' rights to receive care, and thus the need to adopt a complementary ethics of care in the pursuit of information justice lest an ethics of justice alone lead to gross injustice through overreliance on rules and an emphasis on achieving justice through sameness. This ethics of care implies both a relational understanding of learning data and an obligation to care for students based on their unique qualities rather than on universal rational claims. It leads to four propositions designed to identify the limits of an ethics of justice and apply an ethics of care: that justice and care are contextual; that they are multidimensional, dynamic, and permeable; that care in education must be scalable to not be unjust; and that care must be distinguished from pity.

This approach parallels many of the arguments that I have made since 2014, particularly the argument for a structural component of information justice. Prinsloo's and Slade's opposition of justice and care relies on a rather narrow vision of justice as reductionist, rationalist, universal, and rule-driven. While this certainly captures the dominant, especially distributive, views of justice, it adequately captures neither structural views such as Young's nor alternatives such as restorative justice or capabilities approaches. To a significant extent, the two can be reconciled through these alternative views of justice; Young's view of justice is at least sympathetic to an ethics of care, and more recent approaches to ethics of care have built explicitly on her work (Clifford 2013). Certainly, a structural justice approach leads just as strongly to the four propositions that Prinsloo and Slade propose, which I support enthusiastically as considerations in a just practice of information.

This concern with power relationships is also evident in Dencik's et al.'s (2016) work on data justice and the surveillance state. Their approach is specific to the practices of data-driven surveillance in the contemporary capitalist political economy, but it remains focused on structural concepts of justice. They define data justice as "the implications that data-driven processes at the core of surveillance capitalism have for the pursuit of substantive social and economic justice claims" (2016, p. 9). Data justice in this sense explores the application of claims of social justice to surveillance capitalism specifically, both challenging the interests, power

relations, and political agenda behind these data practices and developing an ideal relationship between social organization and digital infrastructures. "Advancing this agenda," they argue, "would transform surveillance from a special-interest 'issue' into a core dimension of social, political, cultural, ecological and economic justice, and thus respond to the central position of data-driven processes in contemporary capitalism," concluding that "concerns with the collection, use and analysis of data need to be integrated into activists' agendas, not just to protect themselves, but also to achieve the social change they want to make," a more promising approach than simple "techno-legal solutionism" (2016, p. 9).

Heeks and Renken (2016) approach data justice from the perspective of development, arguing that there can be no development justice without data justice because data has become a "primary, public good" central to decision-making. The general direction of Heeks and Redken's argument suggests that in saying data is a "primary, public good" they do not mean to argue that data is a "primary social good" in the sense that the phrase is often used by theorists of justice.[3] Rather, I take "primary, public good" to mean a good that is both primary and public in the sense of being central to the operation of public authority and resolution of commons problems in contemporary societies. This conception of data as a public good on the level of security or infrastructure alone is important for understanding why information justice is such a critical and challenging issue, putting conflicts among public problems, private (and state) data ownership, and state interests in the forefront of the information justice debate. Information is not a luxury or a business advantage; it is necessary to addressing the kinds of problems that John Dewey (1954) held gives rise to polities to begin with. In such a society, information justice is an essential dimension of social justice. Of course, that raises challenging questions about the nature of publics and goods, such as those taken up by Taylor (2016).

Given this conception of the problem, Heeks and Renken are unsurprisingly dissatisfied with the instrumental and distributive approaches that I addressed in Chap. 5, as well as procedural conceptions of justice that focus on fair data management and use practices. To be use, they do see some value in distributive approaches, which they understand less as distributing rights and more as using rights to distribute data. That is inconsistent with the typical conceptions of rights as a distributive form of justice but soundly within Young's critique of the use of distributive concepts to allocate moral goods (1990, pp. 24–30). They also quite effectively broaden the question of distributive justice from questions about the distribution of information (as I have characterized privacy) or of rights to questions about the distribution of the benefits of information. This is an area I haven't systematically addressed here, and one that must be considered much more seriously in the future. Nonetheless,

[3] Rarely do I parse punctuation, as authorial intent isn't not usually an issue that is critical for the kinds of philosophical questions I ask. But in this case the comma really does appear to matter. Rawls (2005), for example, argues that primary social goods such as rights and liberties are to be distributed according to the greatest equal liberty principle, while material goods are to be distributed according to the difference principle. It would be hard to understand how data is a primary social good. Though one might argue that privacy is, that place emphasis on restricting data flow and thus would be inconsistent with Heeks and Redken's arguments about the ubiquity of data.

Heeks and Renken move in the same direction that I have and that Prinsloo and Slade have: It is insufficient to rely on distributive information justice alone.

Heeks and Renken build instead on concepts of "small data" and on Amartya Sen's capabilities approach to justice to complement the distributive approach. Small data orients information justice toward supporting individuals and communities in using "the data that people need in order to live the life they value" rather than the large datasets held by enterprise users (Heeks and Renken 2016, p. 6). Sen's approach brings the focus to justice-in-practice, considering the abilities one had to access justice as an essential feature of it. This makes one's capability to use data to further one's values as important to justice as the distribution of data, rights, or benefits. These criticisms make the case for understanding information justice structurally far more compelling, as small-data capabilities are strongly influenced by social structure. Hence, Heeks and Renken argue,

> the foundation of data justice must be structural data justice, which we can define as "the degree to which society contains and supports the data-related institutions, relations and knowledge systems necessary for realisation of the values comprised in a good life". (Heeks and Renken 2016, p. 7)

deliberately paralleling Young's definition of structural justice. But they go beyond Young's foundations, offering three alternative forms of structural data justice: a cosmopolitan form in which the focus is on the nature and structure of social networks, an approach rooted in the emerging field of critical data studies in which the focus is on information and information systems as sites of politico-economic contestation (particularly a "critical modernist" form that maintains a space for agency, and Sen's capabilities approach to justice.

Heeks and Renken identify and help resolve a significant gap in the approach to information justice that I've presented here, one that is apparent especially in Chap. 6. Young's understanding of justice as structural is vitally important and provides a very strong basis for critique in the "Drown the Bunnies" case. But Young's approach is ultimately a theory of *in*justice: Social systems are justifiably blameworthy to the extent that they interfere with self-development and self-determination. I will not dismiss the power of starting from a claim of injustice when dealing with actually existing social problems. But that chapter struggles to go beyond the claims of injustice. What would a just system of student analytics look like? It is not clear that it would simply be a system that promotes self-determination and self-development; there are many possibilities there, and many potential conflicts as well, enough so that one might well be tempted to suggest that justice and injustice are not opposite ends of a spectrum for Young's approach. A claim of structural injustice, as I make in Chap. 6, does not necessarily help us build justice. Perhaps Sen's approach—or an ethics of care or a restorative justice approach, which together draw attention to the values, intents, and actions of the students affected by predictive analytics and to their relationships with their universities—might well prove as useful a guide in building as Young does in critiquing.

7.2 Information Justice in Practice

The possibilities for implementing information justice are myriad, and nothing in this book has, so far, made a conclusive case for any single practice being either necessary or sufficient to achieve it. Perhaps a fully developed theory will do that, but I doubt it; practice emerges too quickly with too much variation. But we can look to some practices that might prove generally useful at least. I can think of nothing more important to the pursuit of information justice than making information politics explicit. The political background and consequences of data must be consciously considered in the practice of information. Data scientists routinely speak of the "data provenance," the origin, source, and process that accounts for the data (hopefully through a series of records). Data provenance needs to be analyzed not just for its technical aspects (e.g., how reliable and valid the data is) but for its social aspects as well (e.g., the justification for coding the data the way that it was). Any claim that data is objective, realist, value-free, or apolitical must be seen as a political claim itself. And normative assumptions must be considered as important as empirical ones in understanding the soundness of information systems.

If data is indeed a moral and political practice, attention must be paid to the normative arguments that support the inferences and interventions based on it. Models are often unable to substantiate their own value assertions where they are external to the model, taken either from elsewhere in the problem-model-intervention nexus or part of the model's set of substantive and methodological axia. This is a familiar problem to empirical researchers in higher education: the problem of validity.

> The term "validation" and to a lesser extent the term 'validity' tend to have two distinct but closely related usages in discussions of measurement. In the first usage, 'validation' involves the development of evidence to support the proposed interpretations and uses . . . . In the second usage, "validation" is associated with an evaluation of the extent to which the proposed interpretations and uses are plausible and appropriate. (Kane 2006, p. 17)

One can think of scientism and the uncritical assumption of values as an attitude that compromises (or, at the least, assumes rather than demonstrates) the normative validity of the problem-model-intervention nexus.

If this way of understanding scientism is correct, it suggests that researchers can address these problems much as researchers would address empirical validity. Kane (2006) presents an approach to validating measures based on a series of inferences from observation to construct. While the specifics of Kane's approach vary widely according to the particular type of measurement, the basic principle of ensuring a sound path of inferences throughout the research process—including the point of taking action based on the research—can serve as a model for data mining applications. In developing or applying a data mining process, institutional researchers should ask themselves if the chain of inference from problem to model to implementation is sound, both scientifically and normatively. Where it is, ethical problems originating in scientism are likely to be averted. Where it is clearly flawed, researchers should change their approaches. But most importantly, where there are gaps in the reasoning researchers should identify the assumptions that allowed those

gaps to be bridged uncritically and then subject those assumptions to critical analysis. Practiced iteratively, this approach can minimize especially the effects of scientism in data science, and likely improves the achievement of information justice generally.

Another central problem of information justice is exclusivity: individuals, their experiences, their values, and their interests are left out of information systems by the data collection process, the dissemination process, or the operation of the system as a whole. It seems likely, then, that a practice of information justice will be built around forms of pluralism. Information pluralism would embrace, rather than problematize, the "messiness" of data. Rather than seeing conflicting data as inherently erroneous it would encourage information systems to be designed to incorporate and highlight differences in data, identifying them as moments of conflict among assumptions and values to be resolved through social rather than algorithmic solutions. It could take advantage of big data's increasing abilities to process narrative and unstructured data and to solve for solutions built on the diversity of individual cases rather than the central tendency of the dataset. And it could incorporate the myriad values that compete for the attention of technologists: openness, efficiency, privacy, security, benefit. This would be joined to a kind of participative pluralism, where information systems are designed with the participation of all actors who are part of the system, including those who will serve as the data points and as the objects of decisions based on the information. Such a system would reflect concepts of "deliberative development" or "collaborative transparency," where concerns with transparency are mediated by the countervailing power of public participation (Donovan 2012).

Especially important to information pluralism is encouraging participation in the development of data: what one might call "foundational open data." This approach recognizes the virtues of open data, and in particular the need for open data as a condition of examining the politics of an information system or practice. As long as the data is closed and the algorithms black-boxed, it is very difficult to examine the processes, assumptions, and biases of the system. But opening data is, I suggested in Sect. 2.1, often a path toward exacerbating the injustices built into the data. A more promising process would be to make the development of the data itself an open process in which the subjects of the data are included in its development. Making data open at its foundation rather than after its development would at the least allow those challenged by information initiatives to expose the politics in the process to examination, and may well provide inputs that lead to more just data systems.

For these approaches to become widespread, however, they must become central to the practice of data science generally, themselves acting as social institutions. Data scientists have begun to recognize the ethical challenges involved in big data and predictive analytics and in response have begun developing codes of professional ethics that go beyond information privacy. Higher education learning analytics especially has been the subject of robust analysis, and several ethical frameworks for the field have been developed to varying extents. Much of this effort draws on the excellent work of Sharon Slade and Paul Prinsloo (2013). Slade and Prinsloo

take issue with the assumption that information is inherently helpful to learners, independently offering arguments similar to those I've made here. In particular, they find that the problems of data provenance and interpretation, privacy and consent, and data management are inherently connected to the power structures within which they operate; Slade and Prinsloo pay particular attention to Foucauldian power structures that have received scant attention here but very much should be considered as part of any political analysis of information technologies. They argue strongly for

> viewing learning analytics as moral practice, recognizing students as participatory agents with developmental and temporal identities and learning trajectories and the need for reciprocal transparency. Learning analytics as moral practice functions as a counternarrative to using student data in service of neoliberal consumer-driven market ideologies. (Slade and Prinsloo 2013, pp. 1511–1512)

This view of students rests, insightfully, on a conception of students' identities as both transient and pluralistic, changing over time and including multiple, often conflicting, dimensions simultaneously (e.g., learner of critical thinking and adherent to authoritative religious practices).

Slade and Prinsloo argue for an approach in which "an institution's use of learning analytics is going to be based on its understanding of the scope, role, and boundaries of learning analytics and a set of moral beliefs founded on the respective regulatory and legal, cultural, geopolitical, and socioeconomic contexts" (Slade and Prinsloo 2013, p. 1518). They identify six principles:

1. Learning analytics must be understood as a moral practice.
2. Students must be understood as agents.
3. Identity and performance must be understood as dynamic constructs rather than essential characteristics.
4. Success must be understood as complex and multidimensional.
5. Universities must be transparent about their purposes.
6. Universities must use learning analytics to improve outcomes.

These are quite valuable, certainly consistent with the understanding of information justice presented here and, as Slade and Prinsloo operationalize them, valuable for guiding practice in learning analytics.

Considerations such as Slade and Prinsloo offer are the basis for a growing number of guidelines and frameworks seeking to establish a professional ethics or codes of conduct in learning analytics, often built on regulatory regimes in other areas of information practice. Based on the Association for Institutional Research Statement of Aspirational Practice and the findings of a working group of researchers and vendors, Rachel Boon (2016) identified seven steps for sharing data with students that promote both transparency and shared understandings of data collection. Jisc, which provides technology services to the UK higher education sector, identified eight principles for post-secondary education institutions use of learning analytics, including responsibility, transparency and consent, privacy, validity, access, enabling positive interventions, minimizing adverse impacts, and stewardship of data (Sclater

and Bailey 2015). The New America foundation, which was a major force in developing the completion agenda, recently proposed five guiding principles for learning analytics (Ekowo and Palmer 2017) that are more operationally oriented than, for example, Slade and Prinsloo.

Perhaps the best known set of standards are those established by The Open University in the UK (2014). The policy sets out in detail both the business case for using analytics and the context and concerns its use presents; identifies specific data that is and is not expected to be used in learning analytics; incorporates existing university policy, oversight processes; and identifies eight principles for using student data ethically to provide student support:

> Principle 1: Learning analytics is an ethical practice that should align with core organisational principles, such as open entry to undergraduate level study.
> Principle 2: The OU has a responsibility to all stakeholders to use and extract meaning from student data for the benefit of students where feasible.
> Principle 3: Students should not be wholly defined by their visible data or our interpretation of that data.
> Principle 4: The purpose and the boundaries regarding the use of learning analytics should be well defined and visible.
> Principle 5: The University is transparent regarding data collection, and will provide students with the opportunity to update their own data and consent agreements at regular intervals.
> Principle 6: Students should be engaged as active agents in the implementation of learning analytics (e.g. informed consent, personalised learning paths, interventions).
> Principle 7: Modelling and interventions based on analysis of data should be sound and free from bias.
> Principle 8: Adoption of learning analytics within the OU requires broad acceptance of the values and benefits (organisational culture) and the development of appropriate skills across the organisation (The Open University 2014, Sect. 4).

The policy was extensively publicized when The Open University instituted it, as it was held up as an innovative model for other universities.

The immediate virtue of such standards is that they make the realist view of learning analytics untenable. OU's insistence that learning analytics support the university's mission and that the data be free from bias draws attention to the possibility that some applications would not do so. They also nearly universally recognize students as moral agents, which works quite strongly against more manipulative applications of learning analytics. Especially encouraging is the promotion or adoption of such standards by vendors and interest groups. Jisc and New America could very easily have continued promoting a dangerously naïve view of learning analytics; instead, their participation in these discussions legitimizes the issues and compels a more critical viewpoint. And through their connection with the professional and educational organizations that train data scientists and operate data systems, codes of professional ethics institutionalize these principles in ways that can counter institutionalized data systems, creating logics of appropriateness that, for example, make the "Drown the Bunnies" model quite literally unthinkable. Such statements are important steps toward information justice.

But these professional codes do have noteworthy weaknesses. Most pay scant attention to the kinds of structural conditions and power relationships that Slade and Prinsloo and that I have emphasized. Boon's first principle, for example, is to "Determine the full range of data available." That is wholly inadequate; a key principle of information justice is that injustice is often caused by the way data is created to begin with. If we begin from the data we have, we are quite likely to miss the injustices present in that data and then institutionalize those injustices in student support programs. The emphasis on transparency in all of these models disregards the challenges that open data can create in securing justice, and are frequently posited along with privacy protections as if the two are entirely compatible. Indeed, privacy is often treated as if it is the only ethical concern in learning analytics (see, e.g., Pardo and Siemens 2014). And most such approaches understand ethics as a matter of ensuring good faith. New America, for instance, believes that ethics can be achieved through advice such as "convene key stakeholders to make important decisions" and "design predictive models and algorithms so that they produce desirable outcomes." As the "Drown the Bunnies" case illustrated, who counts as a key stakeholder and what constitutes a desirable outcome depends significantly on the organizational and power structures of the university. For these statements to fully achieve their potential in a praxis of information justice—and for them to avoid being mere paper declarations that do little to influence actual outcomes—they need to be informed by an overarching concept of information justice.

7.3 "A Data Justice Movement"

"If we accept that higher education is a 'moral and political' practice," Prinsloo and Slade argue, "information justice as praxis can act as a powerful counter-narrative to the current hegemony of 'techno-solutionism' and the discourses of 'techno-romanticism'" (2017, p. 121, citations omitted). Information justice can result when a coherent theory of information politics is both informed by information practices and shapes our choices in the design and use of information and information systems. Ultimately a praxis of information justice must work from four key principles:

1. Context. Data is a social and political practice, with associated consequences. This requires ongoing work with information ethicists and practitioners—going beyond just information technologists to include at the least activists, legal and policy specialists, and journalists. One of the key questions here concerns the ways that information functions as a public good.
2. Critique. The injustices present in existing information practices have both distributive and structural dimensions that must both be understood in order to address them. We live in a well-established information environment, and cannot

simply propose a new environment de novo. Critiques of that environment and the structures that create and sustain it are necessary for a theory of information justice that is not merely abstract utopianism.

3. Charge. Positive principles for justice in information and information systems can be based on ethics of care, capabilities, and restorative approaches to justice. It isn't enough to critique; negative guidance (i.e., "Don't do that!" whether in the form of a claim of injustice or an assertion of an inviolable right) only gets a data scientist so far. Those designing new information systems will need guidance in building systems that promote information justice. Justice frameworks that posit positive obligations and not just negative injunctions are most likely to develop principles that charge data scientists with promoting positive action.
4. Culture. Specific information practices that promote justice must be not only proposed but institutionalized. These practices can be reflected in formal standards such as codes of ethics and public policies, as standard elements of theoretical models of information systems, and in educational practices as model problems and solutions for aspiring data scientists.

Certainly, there is much more work to be done in building a praxis of information justice—and happily, there is a growing, multidisciplinary community of excellent researchers and practitioners working on the problems. I am very excited to see where the praxis of information justice goes from here.

But a theory, even one oriented toward praxis rather than abstraction, is not enough to make change of its own. Many people involved in data science are rightly convinced of their own good faith, and need only considerations like what this concluding chapter has suggested to do very good things with data. But others (a long list of Silicon Valley tycoons fits in here) are too impressed with themselves to see the harms they are doing, and a few—think of the "fake news" industry, for example—are actively using contemporary advances in information systems to intentionally do harm to others. These groups are less likely to be convinced by rational principles. Political contestation will be necessary for information justice to become a reality. Thus Saitta was, in the tweet that started this project, right to call not for a data justice theory but a data justice movement.

Organizations such as Data Justice and Cardiff University's Data Justice Lab are addressing information justice specifically. The Digital Justice Lab is an academic research center, but works with a strong focus on digital social movements. Data Justice is a US-based policy organization that challenges specifically the economic injustices of data practices through public outreach and policy campaigns. Other groups with wider focus are taking an interest in information justice as well. Color of Change, a US civil rights group, has actively led efforts to engage big data from a civil rights perspective. The Data & Society Research Institute is engaging in policy and information systems research on ethics and human rights in big data, and has formed the Council for Big Data, Ethics, and Society to engage in public activism and practitioner engagement. These organizations all actively promote political contestation of information systems and practices, with the result that the principles of information justice can influence outcomes and promote social and political

change. They also support the most promising political strategy for challenging existing information practices, exploiting gaps in information systems. Data politics is inevitable but not deterministic. Gaps in political and information systems are always present, and can be used to promote more virtuous data politics, developing counter-narratives and undermining seemingly hegemonic institutions.

I suspect one of the most vital roles for an information justice movement would be building the capabilities needed for participation in information systems. This would include both skills and technology. Donovan (2012) notes that the success of the Map Kibera project is connected both to the provision of GIS training to participants and users and to the development of local ownership and control. Stearns (2012) calls more broadly for data literacy campaigns modelled on anti-smoking campaigns "that can fundamentally shift people's understanding and relationship with their personal data." Organizations that are part of the information justice movement can provide this training, along with enterprise-level computing capacity and connections to social and political institutions. They can also provide alternatives to direct participation in the form of investigative and data journalism that may be more successful in some circumstances (Swartz 2009, 2012). Ultimately it is the organizations in civil society, not philosophers, that make it possible for marginalized groups to participate collaboratively or to challenge embedded power structures in information systems.

It remains vital that the praxis of information justice and social movements contesting information practices be understood as complementary; there is neither a hierarchy nor division of labor to information justice. An intellectual framework for understanding intellectual justice is, one hopes, indispensable for those who wish to bring it about. It can direct attention to possible causes and solutions, and provide paradigmatic cases that serve as starting points for action. The act of developing and maintaining such a theory also offers a critical perspective on the practice of an information justice movement. But, though each in their own ways, the scholar is as privileged as the programmer, the bureaucrat, or the activist. The critical perspective that the philosopher or the social scientist takes on an information system is applicable to academic work, and as difficult to execute from inside as any other. A close relationship between activists and theorists provides challenges to theory from practice that allow for theoretical growth.

A praxis of information justice is desperately needed today, not just in so-called "information societies" but globally, north and south. We can pursue data in good faith without any kind of ethical malice and, because of the structural injustices in data, still produce unjust outcomes. Exhortations to be more ethical as individuals are welcome but insufficient to make much headway toward a more just information environment. Thorny issues remain hidden in the details, to be sure. But as information becomes a primary, public good, we will have no choice but to understand information justice as an essential element of a just society.

References

Boon, R. (2016, November). Sharing data with students to inform decision making. *Ask eAIR*. https://www.airweb.org/eAIR/askeair/Pages/Rachel-Boon.aspx. Accessed 17 Nov 2016.

Cegłowski, M. (2016). *The moral economy of tech*. Presented at the 28th SASE annual meeting, Berkeley, CA. http://idlewords.com/talks/sase_panel.htm. Accessed 11 Nov 2016.

Clifford, D. (2013). Ethics, politics and the social professions: Reading Iris Marion Young. *Ethics and Social Welfare, 7*(1), 36–53. https://doi.org/10.1080/17496535.2012.667139.

Crawford, K. (2015). Can an algorithm be agonistic? Ten scenes from life in calculated publics. *Science, Technology & Human Values*. https://doi.org/10.1177/0162243915589635.

Dencik, L., Hintz, A., Cable, J. (2016). Towards data justice? The ambiguity of anti-surveillance resistance in political activism. *Big Data & Society*, *3*(2). doi:https://doi.org/10.1177/2053951716679678.

Dewey, J. (1954). *The public and its problems*. Athens: Swallow Press.

Donovan, K. (2012). *Seeing like a slum: Towards open, deliberative development*, *SSRN Scholarly Paper No. ID 2045556*. Rochester: Social Science Research Network. http://papers.ssrn.com/abstract=2045556. Accessed 5 Mar 2013.

Ekowo, M., & Palmer, I. (2017, March). *Predictive analytics in higher education: Five guiding principles for ethical use*. The New America Foundation. https://na-production.s3.amazonaws.com/documents/Predictive-Analytics-GuidingPractices.pdf. Accessed 18 May 2017.

Heeks, R., & Renken, J. (2016). Data justice for development: What would it mean? *Information Development*. https://doi.org/10.1177/0266666916678282.

Introna, L. D. (2015). Algorithms, governance, and governmentality: On governing academic writing. *Science, Technology & Human Values*. https://doi.org/10.1177/0162243915587360.

Kane, M. T. (2006). Validation. In R. L. Brennan (Ed.), *Educational measurement* (4th ed., pp. 17–64). Westport: American Council on Education/Praeger.

Kitchin, R., Lauriault, T. P., & McArdle, G. (2015). Knowing and governing cities through urban indicators, city benchmarking and real-time dashboards. *Regional Studies, Regional Science, 2*(1), 6–28. https://doi.org/10.1080/21681376.2014.983149.

McMillan Cottom, T., & Tuchman, G. (2015). Rationalization of higher education. In R. A. Scott, & S. M. Kosslyn (Eds.), *Emerging trends in the social and behavioral sciences: An interdisciplinary, searchable, and linkable resource* (pp. 1–17). http://onlinelibrary.wiley.com/book/10.1002/9781118900772. Accessed 22 Dec 2015.

O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy* (1st ed.). New York: Crown.

Pardo, A., & Siemens, G. (2014). Ethical and privacy principles for learning analytics. *British Journal of Educational Technology, 45*(3), 438–450. https://doi.org/10.1111/bjet.12152.

Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Cambridge: Harvard University Press.

Prinsloo, P., & Slade, S. (2017). Big data, higher education and learning analytics: Beyond justice, towards an ethics of care. In B. Kei Daniel (Ed.), *Big data and learning analytics in higher education* (pp. 109–124). Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-06520-5_8.

Rawls, J. (2005). *A theory of justice* (Original ed.). Cambridge, MA: Belknap Press.

Sclater, N., & Bailey, P. (2015, June). *Code of practice for learning analytics*. Jisc. https://www.jisc.ac.uk/sites/default/files/jd0040_code_of_practice_for_learning_analytics_190515_v1.pdf. Accessed 18 May 2017.

Scott, J. C. (1998). *Seeing like a state: How certain schemes to improve the human condition have failed*. New Haven: Yale University Press.

Slade, S., & Prinsloo, P. (2013). Learning analytics: Ethical issues and dilemmas. *American Behavioral Scientist, 57*(10), 1510–1529. https://doi.org/10.1177/0002764213479366.

Stearns, J. (2012, November 13). We need a "Truth" campaign for digital literacy and data tracking. MediaShift. http://www.pbs.org/mediashift/2012/11/we-need-a-truth-campaign-for-digital-literacy-and-data-tracking318.html. Accessed 14 Mar 2013.

Swartz, A. (2009, April 23). Transparency is bunk. Aaron Swartz's Raw Thought. http://www.aaronsw.com/weblog/transparencybunk. Accessed 3 Mar 2013.

Swartz, A. (2012, July 3). A database of folly. Crooked Timber. http://crookedtimber.org/2012/07/03/a-database-of-folly/. Accessed 3 Mar 2013.

Taylor, L. (2016). The ethics of big data as a public good: Which public? Whose good? *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 374*(2083), 20160126. https://doi.org/10.1098/rsta.2016.0126.

Taylor, L. (2017). What is data justice? The case for connecting digital rights and freedoms on the global level. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.2918779.

The Open University. (2014, September). *Policy on ethical use of student data for learning analytics*. http://www.open.ac.uk/students/charter/sites/www.open.ac.uk.students.charter/files/files/ecms/web-content/ethical-use-of-student-data-policy.pdf. Accessed 18 May 2017.

Young, I. M. (1990). *Justice and the politics of difference*. Princeton: Princeton University Press.