Public Administration and Information Technology 33

Jeffrey Alan Johnson

Toward Information Justice

Technology, Politics, and Policy for Data in Higher Education Administration

Springer

Public Administration and Information Technology

Volume 33

Series editor

Christopher G. Reddick, San Antonio, TX, USA

More information about this series at http://www.springer.com/series/10796

Jeffrey Alan Johnson

Toward Information Justice

Technology, Politics, and Policy for Data in Higher Education Administration

Springer

Jeffrey Alan Johnson
Institutional Effectiveness, Planning,
and Accreditation Support
Utah Valley University
Orem, UT, USA

ISSN 2512-1812 ISSN 2512-1839 (electronic)
Public Administration and Information Technology
ISBN 978-3-319-70892-8 ISBN 978-3-319-70894-2 (eBook)
https://doi.org/10.1007/978-3-319-70894-2

Library of Congress Control Number: 2017957849

Printed on acid-free paper

This Springer imprint is published by Springer Nature
The registered company is Springer International Publishing AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface

This project began with a tweet:

> Not only must data sovereignty trump open data, but we need active pro-social countermeasures—a data justice movement. (@Dymaxion [Eleanor Saitta], June 27, 2012)

It turns out to be a lot harder to get out of academia than we think. I left a tenure-track faculty position in 2008. I enjoyed teaching, whether introductory American government or a political theory curriculum I had been allowed to design myself. But it was all I enjoyed. Being a political theorist, I went where I had a job, in a city I had been stationed as a Marine and swore I would never return to. It was an institution that was struggling to do more than provide the basics to students who had already been left behind by the US educational system, and too often indifferent to those students. I had some wonderful colleagues, but also had too many colleagues who had given up, marking time until their age plus years of service equaled 80. As the only political theorist (and the entirety of the philosophy program as well), I was completely isolated from colleagues in my field. At the same time, research in political theory had become tedious for me. I was tired of reading the latest on someone who had not written in half a millennium. Jeffrey Issac's criticisms of political theory as having removed itself from politics (which, as I discuss in Sect. 1.4, play an important role in this book) had always impressed me, but they hit home as I attended conferences in which papers about such topics as "Arendt and the End of Investigative Journalism" turned out to be two pages on investigative journalism bookending 28 pages on Arendt.[1] Five years out of graduate school, the sacrifices we never realize we make to make a living in academia had already burned me out.

I packed up my apartment, my then-fiancé, my dog, and her cats and moved to someplace beautiful (it is rough commuting along the Wasatch Front every day, but someone has to) and close to family (Hi, Mom!), intending to go to work in public policy. Still wanting to be a good citizen of my university and knowing the challenges

[1] I have a real paper in mind as I write this. The theorist and the topic have been changed to protect the tenure-seeking assistant professor who, I have no doubt, needed lines on a CV more than interesting conversations. This was—and still is—a structural problem in political theory.

of getting another tenure-track search, I had submitted my resignation at the beginning of the academic year, effective at its end. But that was the 2007–2008 academic year, and the job market in May 2008 was decidedly not what it was in September 2007. Government jobs were scarce because of budget cuts, and I discovered that I was overqualified by education and technical skills and underqualified by experience. The few interviews were perfunctory. I met the minimum qualifications, but the Kafkaesque decision was already made: I would be bored doing this work so I was better off unemployed; I could not be asked to work for someone with only a bachelor's degree and a decade of experience; I did not have educational qualifications because teaching an environmental policy course did not count as taking an environmental policy course. So I returned to higher education, starting in the Institutional Research and Information office at Utah Valley University 18 months after I thought I was out of higher education.

For a political theorist, it turned out that I was a pretty damn good institutional researcher—pay attention in research methods, kids. The first thing one learns when starting in institutional research is that there is a thing called institutional research. Almost everyone I know in this field got here by accident, through a job announcement where the KSAs match the skill set of most graduate degrees and lead to someone saying, "Hey, I can do that." That makes the field interesting, though. Few jobs surround one with many very smart people from many very different backgrounds. The second thing one learns, however, is more circumspect. Institutional researchers have a lot of data. As one begins to see the amount of data we have from so many different sources, especially about students, it is hard not to worry about what we can do with that data. Trained to ask these questions, having worked mainly in the politics of science and technology when I was a graduate student and professor, I began to think about them in my work. They kept me up at night and kept me arguing about them around the water cooler. But it was Eleanor Saitta's tweet about data *justice* that provided a framework in which my initial questions coalesced into a project.

Toward Information Justice is the result of that project. Those who work with information, especially in higher education, struggle with the myriad ethical concerns our field presents. We can protect privacy, but sometimes that comes at the cost of denying students the opportunity to make informed decisions about their educational careers. We worry about disclosing data, but often do not consider the implications of creating data. We too easily assume that data is objective and apolitical rather than taking responsibility for (unavoidably, I will argue) practicing politics with our students. The aim of this book is to move us toward a coherent moral picture of information as it is used in public administration based on philosophical conceptions of justice. I do so in much the same way that environmental philosophers have understood environmental justice. That requires first understanding data as a political practice, and then rebuilding that practice so that it is consistent with what we consider to be right and/or good in social and political life. Hopefully, a theory of information justice can help practitioners resolve these challenges. To be sure, this book does not get there. But it does set the direction, and adds a valuable dimension to the array of incredible work being done in this field, by people I am proud to find are as insightful and stimulating as colleagues as I ever had as a professor.

Three months before I started this project, I had told a colleague and dear friend, Stephanie Mora Walls, that I was now an institutional researcher, out of political theory altogether, and would not continue to coauthor with her anymore. I owe her an apology.[2] And I owe aspiring theorists a warning. Once a political theorist, it turns out, always a political theorist.

Orem, UT, USA Jeffrey Alan Johnson

[2] Though not nearly as much as one might think, since her book, *Individualism in the United States: A Transformation in American Political Thought*, turned out to be a pretty damn good one without any further meddling from me.

Acknowledgments

There are many people whom I must thank for making this book happen. The project started with a tweet and came to a successful conclusion largely because of people that I know entirely or keep in touch with through Twitter. So many people have pushed me to keep asking these questions, encouraging along a project that met with indifference in some areas of my own fields. No one has been more encouraging than Tressie McMillan Cottom. Tressie, whose work has been amazing since long before she was a mononym, immediately believed in this project and talked me through some of the basic sociological concepts involved that, as a political scientist, I had no idea existed. (As an aside, no one who took a graduate course in contemporary political thought over the past two decades should not know what intersectionality is, but hey, white feminism FTW, right? And yes, that is a subtweet.) In a similar vein, I must thank Tod Massa for his contributions and support. Tod is doing some of the most important work in institutional research today. Chapter 3 of this book is nearly 15,000 words that Tod conceptualized in three: "counting to one." That should be a fundamental concept in data science, and is certainly central to the ideas in this book. To have their support, and the kind words that they expressed about this project, is far more than I have any business asking. To call them my friends is a deep honor.

So many other friends and colleagues have offered their support and insight. The broader Twitterati of this project includes Mark Carrigan, Karla Carter, Jessie Daniels, Melonie Fullick, Chris Gilliard, Karen Gregory, Erin Hansman, Paul Prinsloo, Lee Skallerup, Evan Selinger, Bonnie Stewart, Loralyn Taylor, and Audrey Watters. Whether through encouragement, inspiration, or both, they have kept this project alive during times when abandoning it sounded awfully appealing. Mike Krywy introduced me to Twitter, and for that alone he deserves better than the Winnipeg Jets did this season. Off Twitter, my friends Steve Vanderheiden and Christopher Robinson remind me that political theory is not just the navel-gazing exercise that I so narrowly saw when I left my faculty position. Every time a discussant asks whether I have considered using Foucault or confuses the lack of canonical political theorists for a lack of political theory,[1] I look to them. They do important work, and encourage me both directly and by example to continue working in the field I found so wonderfully challenging when I first discovered it. Evelyn Cruz and I discussed how these concepts apply to information questions in public health as I was developing the ideas in this book, which on many occasions helped me understand where I was mistaking a peculiarity for a principle and makes me look forward to making a similar analysis of information in public health, and continues to be as encouraging today as she was when I met her in graduate school. She also has not aged since then. Thank you all for your help on this.

If I have had any success in this book, it comes because I have been able to bring political theory and institutional research together. I may or may not be a particularly good theorist, but I am a damn good institutional researcher, and I cannot take the slightest credit for it myself. I have been blessed with two of the finest mentors for which one could ask. I joined IRI with the barest minimum understanding of databases. Michael Dover patiently talked me through data systems, data structures, and the many flavors of Structured Query Language used at Utah Valley University. In the process, he taught me the fine art of institutional research as well as its technical details. Every bit of technical knowledge in this can be traced to what Michael has taught me (or, because I am sure there are things that are not quite right in here despite my best efforts, to what I failed to learn properly). Just as Michael taught me how to practice institutional research, Marcus Jorgensen, under whom I worked as Assistant Director of Institutional Effectiveness and Planning, taught me what it meant to lead in higher education. His exceptional kindness and inclusive leadership are without question the reason behind any success I have had as a leader in institutional research and effectiveness.

Many other colleagues in institutional research have supported this work in various ways. At UVU, Linda Makin, Marc Jorgensen, Robert Loveridge, and Tim Stanley allowed me to pursue a far more active research agenda than most institutional researchers have the opportunity to, often finding conference support to make sure that the ideas in this book got before the many audiences that it needed to engage. They, along with David Knowlton, Quinn Koller, Michael Minch, and Angela Ward, allowed me to bounce ideas around to see that they made sense to both institutional researchers and social scientists. I have twice been honored with awards from my colleagues in the Rocky Mountain Association for Institutional Research for this research. I am deeply grateful for the support especially of Dawn Kenney and Nicolas Valcik among the many RMAIRians who continually reminded me that this work was worth pursuing, especially given the resistance that this work has seen from groups at the national level.

I also thank the many colleagues at UVU and other institutions whose work provided the material for much of this book. Much of it is based on my experiences at UVU, and so reflects the university's ongoing efforts to manage information. While

[1] This time, I am not protecting anyone by changing the facts. These have been staples of my experiences with this project at political science conferences and explain a good bit about why I have found so much more inspiration from digital sociology.

I have been appropriately critical at times, I am proud of the work UVU does for its students and the people of Utah. Those with whom I work share an uncommon commitment to the success of our students, many of whom are seeking a second—or perhaps seventh—chance at an education. There is no questioning their good faith even as challenges arise, a principle that drives the fundamental point of my argument that good faith alone is insufficient to achieve information justice. With the exception of UVU's president, Matthew S. Holland, I have declined to identify specific people in my analysis. In conducting my research, I have examined structures, policies and procedures, public actions, and publicly available (either through direct dissemination or that would be available through open records requests) documents of the university, its systems, and its officials rather than the officials themselves; it does not draw on my direct interaction with any faculty, staff, or students.[2] I should, of course, note that the analysis presented here is strictly my own, and should not be taken as representing in any way the policy of Utah Valley University or the views of its leadership.

A good bit of the content of this book has been previously published or presented, and benefited greatly from the review process and from comments from discussants and colleagues. A previous version of Sect. 2.1 and a portion of Chap. 7 was previously published as "From Open Data to Information Justice," *Ethics and Information Technology* 16:4 (December 2014), pp. 263–274, copyright Springer Science + Business Media Dordrecht; the final publication is available at Springer via https://doi.org/10.1007/s10676-014-9351-8. A previous version of Sect. 2.2 and the part of Sect. 7.2 on normative validity was included in "Ethics of Big Data in Higher Education," *International Review of Information Ethics* 21 (July 2014), copyright IRIE. Sections 3.1, 3.2, and 3.3 are an expanded versions of "Information Systems and the Translation of Transgender," *Transgender Studies Quarterly* 2.1 (February 2015), pp. 160–165, copyright Duke University Press. Most of Sect. 3.4 was published as "Representing 'Inforgs' in Data-Driven Decisions," in Karen Gregory, Tressie McMillan Cottom, and Jessie Daniels (eds.), *Digital Sociologies* (Bristol, UK: Policy Press, 2017), Chap. 11, copyright Policy Press of the University of Bristol. All of the above are published with the permission of the respective copyright holders under the terms of the publications agreements permitting author re-use. Much of this work was also presented in working form at several conferences. The participants in the Eastern Sociological Society Digital Sociology Mini-Conferences were exceptionally helpful across the board, and I am very grateful that they have allowed a narrow-minded political scientist into their fold. The participants in the first Power, Acceleration and Metrics in Academic Life Conference provided a wonderful environment for working through the issues presented in Chap. 4. I appreciate as well those who provided helpful suggestions at RMAIR conferences, Western Political Science Association Annual Meetings, and Midwest Political Science Association Annual Conferences. Finally, I need to thank the anonymous peer reviewers for their kind comments and suggestions on the final manuscript.

[2] This research thus does not constitute research involving human subjects within the meaning of federal regulations.

My deep regret with this project is that it took so long. Partially that is the downside of researching while an administrator, but most of it is my own laziness. My friend and editor, Lorraine Klimowich, should never have had to put up with that, and I cannot thank her enough for her patient encouragement. But the regret comes from someone to whom I cannot apologize: my father. In June 2016, he was diagnosed with advanced lung cancer and given anywhere from 6 months to 2 years. Lorraine, compassionate as always, said we could meet that deadline if I could complete the manuscript promptly. But that deadline was hopelessly optimistic; my father passed away only 3 weeks later.

In my eulogy for him, I praised him as a man who shared his life with his family. "As we all tend to," I wrote, thinking that we would mourn the loss of the pop culture world we shared before I mourned him, "I thought there would be more time. More time to share the world he created for our family." There was not, and now I regret the time I wasted while (not) trying to complete this project. My father never saw this because of the many days I just did not feel like writing. "I cannot imagine a man more genuinely excited for someone than my father has been for me, or a man who wanted more sincerely for his children to share their lives with him," I said then. I wish so deeply that he could have seen this come to fruition, so that I could have shared it with him.

Contents

Chapter 1
Introduction

Abstract This chapter situates questions of information within a broader critical-constructive theory of technology. I first define information justice as the fundamental ethical judgment of social arrangements for the distribution of information and its effects on self-determination and human development, a judgment that must be understood in both distributive and structural terms. Studying information from the perspective of justice is, however, complicated by the widely held but ultimately unsupportable claim that technology is morally neutral. Instead, I suggest a constrictive view of technologies in which values play a central role in their development, which raises the possibility of critically examining technologies in relation to alternatives that could have emerged. This critical-constructive view of technologies guides the rest of the book, serving as a foundation for theorizing the challenges that my work as an institutional researcher has presented. I conclude by examining the difficulties—but also the opportunities—of writing political theory from one's own experience.

In the 1990s, the government of India began a program to digitize and open land records. Digitizing the Record of Rights, Tenants, and Crops (RTC) along with demographic and spatial data was intended to empower citizens against state bureaucracies and corrupt officials through transparency and accountability. Sunshine would be the best disinfectant, securing citizens' land claims against conflicting records. In fact, what happened was anything but democratic. The claims of the lowest classes of Indian society were completely excluded from the records, leading to the loss of their historic land tenancies to groups better able to support their land claims within the process defined by the data systems. Far from empowering the least well off, the digitization program reinforced the power of bureaucracies, public officials, and developers (Donovan 2012).

Two decades later, student data management firm inBloom provided data storage and aggregation services to primary and secondary schools enabling them to track student progress and success using not only local data but data aggregated from schools nationwide. Many districts and some entire states adopted inBloom, which was backed by education reform giants such as the Bill and Melinda Gates Foundation and the Carnegie Foundation. inBloom promised that their system, by putting advanced data in the hands of teachers and administrators, would provide

J.A. Johnson, *Toward Information Justice*, Public Administration and Information Technology 33, https://doi.org/10.1007/978-3-319-70894-2_1

the infrastructure layer for a personalized learning ecosystem that would better meet students' needs while improving efficiency. But the aggregation of such data raised deep concerns about student privacy. After several states backed out of the arrangement because of these concerns, the company ceased operations in 2014. CEO Iwan Streichenberger attributed inBloom's failure to its "passion" and a need to build public acceptance of its practices—in essence rejecting the legitimacy of the ethical concerns its critics raised (Singer 2014). Whether one accepts the legitimacy of those claims or dismisses them as old-fashioned, there is no question that inBloom's failure was not one of inadequate technology but of inadequate ethical vision: inBloom failed to appreciate the moral risks of its technologies and business model, and failed to convince the public of new principles that would support them.

Why the technologies of open data or data-driven personalized learning—and so many other information technologies that claim to bring about not just efficiency and prosperity but fairness, democracy, and freedom—failed to live up to their promises is the question to which this book is devoted. Information, as a social practice and a social structure, raises the same kinds of questions that we might ask of any other practice or structure: What should we do with it? How should it—and control over it—be distributed? These are questions that information ethicists have explored in some detail, especially with regard to privacy. But there are other, deeper questions that information ethicists are just beginning to explore, ones that emerge at the challenging intersection between ethics and social science: What role does data play in the structure of society, and society in the structure of data? How does information shift distributions of goods (material or otherwise) or balances of social and political power, especially among social groups? What beliefs—beliefs about information, beliefs about politics and society, beliefs about people—are assumed by and embedded in information systems? These questions, in turn, assume answers to more deeply philosophical questions about society's relationship with information: How would a good society manage information, using it to further the best ends possible? What practices give us the fairest information processes and outcomes?

These are classical questions of justice. They give rise to a need for what I will call in this book "information justice."

1.1 Toward a Theory of Information Justice

Information justice refers to the fundamental ethical judgment of social arrangements for the distribution of information and its effects on self-determination and human development. It is a subset of the broader notion of political justice, applied to questions of information and information technologies. A theory of information justice helps us understand the conditions under which a society can be said to be securing political justice within the realm of information. It is a critical question to be asked of "the information society" whatever that may mean, as any vision of a society driven by information should be expected to achieve justice generally only to the extent that its fundamental social institution is itself just. Even in more modest visions of society in which information technology is not the definitive structuring force, one cannot deny that information and information technologies are as important to the functioning of contemporary societies—not only the post-industrial north but increasingly the global south as well—as the political economies of those societies (though of course, no more independent of the political economy than the political economy is of information).

A political philosophy that sees information as a socio-technical practice displays certain similarities to environmental philosophy. Environmental philosophers have long acknowledged questions of justice. David Schlosberg notes that the field is marred by a weak understanding of justice itself, making "most theories of environmental justice . . . incomplete theoretically" (2004, p. 517). Relying on the work of Iris Marion Young and of Nancy Fraser, he argues for a more expansive understanding of justice composed not only of the common distributive framework but also of a need to secure recognition of and participation by all groups in society. Using this framework, he is able to develop a framework of "critical pluralism" for environmental justice that makes sense not only philosophically but also in light of claims by social movements dedicated to securing environmental justice. Schlosberg's success in framing a more complex vision of justice in environmental issues suggests the value of a similar framing for information practices and technologies. Drawing on Schlosberg's conclusion that "justice itself is a concept with multiple, integrated meanings" (2004, p. 536), it may be possible to engage the challenges of information effectively from a justice-centered perspective.

That is to say, we can build a more effective social theory of information by addressing it in relation to, as Serge-Christopher Kolm defines justice, "the central ethical judgment regarding the effects of society on the situation of social entities" (1993, p. 438).[1] Justice is the primary standard by which social and political structures, actions, and practices are evaluated. Echoing Aristotle, John Rawls calls justice "the first virtue of social institutions, as truth is of systems of thought" (2005, p. 3); Young considers justice "the primary subject of political philosophy" (1990, p. 3). Information privacy can be understood as a specific kind of political situation or condition of a social entity, that regarding information about the entity, which is affected by the situations and actions of other social entities. A framework for information privacy can thus be developed that evaluates situations according to judgments about the rightness of that situation, and that (hopefully) promotes information practices that tend toward right situations.

[1] Kolm adds the caveat, "with respect to each entity's valuation of its own situation for its own purposes" (1993, p. 3). This is, to be sure, a common enough feature of most modern theories of justice, especially those rooted in classical liberal political thought. But it cannot be a definition of the problem of justice in itself, as many theories of justice explicitly reject the idea that the entity's own valuation is central. Plato's theory of justice is the paradigmatic case: in arguing that justice is "the minding of one's own business and not being a busybody" (Plato 1991, p. 111 [433a]), i.e., of fulfilling one's naturally ordained role in society, Plato explicitly rejects as unjust the pursuit of one's own purposes and holds instead that justice is to be judged with respect to nature's valuation of an entity's situation rather than the entity's valuation.

Broadly speaking, society might affect the situation of social entities in two ways. Distributive justice concerns the effects that occur “when the purposes of several such entities oppose each other, and the issue is how to arbitrate among their competing claims” (Kolm 1993, p. 438). Distributive justice might arbitrate among claims to material goods, but also claims to rights or political power. Arguably, most social effects on individuals can be understood as questions of the distribution of some good across social entities. However, not all social claims can be reduced to a distributive framework without doing significant violence to the claim itself, for example by ignoring the structural context that gives rise to the claim or taking as fixed social matters that are the product of relationships and processes (Young 1990, pp. 16–30). Hence we might also speak of structural justice, “the degree to which society contains and supports the institutional conditions necessary for the realization of…the values comprised in the good life,” values primarily concerned with self-development and self-determination (Young 1990, p. 37). Social choices such as the ones made in creating data, using it, and opening it for others to use will often have implications for both the distribution of material and social goods and for the social structures that shape individuals’ control over themselves.

From this perspective, the open data of the Indian land records system and the student data collected by inBloom are, in themselves, neither just nor unjust, nor do they inherently further justice or injustice. This is, as I will show in detail momentarily, not because open data is technologically neutral but because open data only exists in relation to a broader information system that gives it meaning: Open data as a-thing-in-itself does not exist in the real world. Moreover, openness is not the only value that ought to be pursued in an information system; data privacy, for example, is equally important (Nissenbaum 2010) and may often conflict with openness (Kaminski 2012). Whether we open or restrict data is thus best understood as one among many intermediate decisions in building an information system, decisions that should be made based on what will further justice given the nature of the data and circumstances. What is ultimately needed, then, is a way of understanding data in the context of an information system and in relation to justice directly: a framework for information justice. Such a framework would allow ethicists and practitioners to systematically identify the different ways in which data can present issues of justice, the relations among them, and the principles by which data can be made more just. Such a theory might pursue three parallel lines of inquiry: inquiries into moral principles, socio-technical practices, and institutions by which we might evaluate and govern data; practices that are conducive to achieving information justice; and the aims, capabilities, and conditions for the success of a social movement that aims to promote social justice.

The discussion in this book serves as a starting point to the study of information as a matter of justice. It should not be read as an indictment of any data practitioners. The problems identified herein are mostly structural in nature. If contemporary societies—affluent and otherwise—are to be as structured around data as many expect, we will need to know how existing social structures are perpetuated, exacerbated, and mitigated by information systems. We will need to know what the ideal information system looks like. Most important, we will need to know what can be done about it. These questions of justice presented by the information systems and practices now emerging in most societies—how the questions arise and what we might do about them—are the focus of this book.

1.2 The Myth of Technological Neutrality

Information justice differs from traditional notions of justice in that its object is explicitly technological. Understanding information justice demands understanding how justice can apply to technologies. That is challenging. With a culture of technological neutrality (Johnson 2006) and radical individualism (Walls and Johnson 2011) dominating the information technology industry, it is exceptionally easy for data scientists and users to accept current data practices and outcomes as natural or inevitable, and to make data use the only moral question of interest. More dangerously, one might take information technologies as instances of what Langdon Winner called "inherently political technology," which "unavoidably brings with it conditions for human relationships that have a distinctive political cast" (1980, p. 128). Advocates see technologies that make democratic politics or individual liberty inevitable through a "naïve technological determinism" in which technology "molds society to fit its patterns" (1980, p. 122). The fundamental rejection of this view, and the recognition that technologies are neither mere artifacts nor the outcome of a purely scientific, rational design process will prove to be a central premise of the later chapters.

Technological neutrality is present in fields far broader than just information technology. To say that a social conception of information raises questions of justice that until now have not been explored is not to say that moral questions about information have never been asked, nor that useful answers to such questions have not been offered. There is a well-established literature among scholars of philosophy, law, and information studies of information ethics. There is an even longer tradition in social and political thought of philosophical reflection on technology generally. The question of information justice sits between these perspectives, subsuming the specific questions asked by information ethics into a larger moral framework while maintaining connections to the complexities of individual technologies that are often lost in more general philosophies of "technology."

What, precisely, do we mean when we speak of technology and technologies?[2] Most views of technology focus on *technē*, paying little attention to *logos*. It is

[2] For my purposes here, we can conceive of technology as an intellectual structure that leads to the development of specific technologies as materialized ways of doing things. But generally speaking, I will use "technology" typically as shorthand for the complex of technology and the technologies it creates simply because "technology and technologies" gets cumbersome with repetition. To be sure, there is a continuum here: As one moves from a general idea of technology through information technology, then databases, then relational databases, then Oracle 12c, then Banner ODS, and finally the Banner installation at UVU, we move from more intellectual to more material. As will become clear in this chapter, both the physical and the material are ontologically essential to

somewhat surprising that "technology" generally does not refer to the study of something. To some extent this is a function of how "-ology" has come to be used to designate that which has been studied as much as the study: the biology of the mollusk, the ecology of Arctic, the methodology of a study. It is perhaps endemic among contemporary society (and perhaps even modern life in general) that we confuse *logos* and *episteme*, not only etymologically, but more importantly practically: In modern life, method is knowledge. No place is this more the case than with technology and technologies.

A useful place to begin understanding technology pragmatically is thus with the word itself. Technology, Larry Hickman (2001) notes, literally means inquiry into technique. But it is used more commonly to designate (a) techniques, tools, and artifacts; (b) systems of these; and (c) applied science. When techniques, technical systems, and applied science work well, there is no need for inquiry into them. It is when they fail in some sense that inquiry into them is necessary, i.e., that we need technology in the literal sense. Technology, in strict speech, is thus "invention, development, and cognitive deployment of [physical and intellectual] tools and other artifacts brought to bear on raw materials and intermediate stock parts, with a view to the resolution of perceived problems" (2001, p. 12). We can use as a convenient shorthand for this "systematic inquiry into technique."

But the problems of technology that we see are (or at least appear to be) found in areas defined by more conventional definitions of technology. They arise in techniques themselves. The problem is whether a particular technique should be used for a particular purpose, whether some people should be allowed to use a technique, whether a technique poses a threat to a particular social value. This, of course, raises the question of what constitutes a technique. In Hickman's interpretation of John Dewey we see him focusing on "tools and other artifacts brought to bear on raw materials and other intermediate stock parts," that is, on tools that we use to interact with the world, both as it is given by nature and created by us. The emphasis is on artifacts themselves. But we use these tools to carry out certain actions, to complete specific tasks. There is thus a *technē*, a craft or technique, to every artifact. It is when we conduct inquiry into these crafts that we engage in technology, that is, in the study of technical things.

Here we see a great divergence from conventional definitions of technology. In conventional definitions, as suggested above, technology is ultimately an artifact of some sort, usually a physical one but sometimes intellectual (a concept that we use to act, such as "markets") or manual (a specific method of manipulation, as a physical therapist might use to inflict useful pain on a patient). Even in manual technologies, the technique reduces the human to machine, carrying out tasks as if human practitioners are automata, reducing the human to an artifact.

---

any technology; we can thus conclude that there is, analytically, no way to understand technology without analyzing specific technologies, and that there are no technologies except built technologies: the size of a rocket depends principally on where it is expected to come down, so there is no actually existing rocket technology without a target in mind.

This artifact-driven view of technology leads to the thesis of technological neutrality. Technological neutrality is a vision of technology that begins with Bacon's *New Atlantis*[3] and continues to be reflected quite strongly in popular discourse about technology. The thesis starts from the observation, shared with many critical perspectives on technology, that technologies are in important ways morally ambiguous (Feenberg 1991). As Melvin Kranzberg famously put it in what has come to be known as Kranzberg's First Law, "Technology is neither good nor bad; nor is it neutral" (Kranzberg 1986). The thesis of technological neutrality is built on this ambiguity of technology, but takes the first clause of Kranzberg's Law to mean the opposite of the second.

The basic premise of technological neutrality is that technology is value-neutral. Technologies are simply physical and intellectual tools that have no intrinsic value. They can be used in different ways, some of which are good and some bad. It is human action that assigns value to a technology. Thus the normative evaluation of technologies focuses not on the technologies themselves but on what one does with them. Actions, not technologies, hold moral values (Tiles and Oberdiek 1995, pp. 13–17). The neutrality thesis can thus be stated as follows: *Technologies are value-neutral tools that are used to fulfill valued functions; therefore moral characteristics can be attributed only to uses of technologies and not to technologies themselves.* This view is seriously deficient, as I will show below; nonetheless, it remains the dominant view in contemporary western culture.

We can see this dominance most strongly in discussions of the responsibility of scientists and technologists for their creation. Two common (though ultimately flawed) arguments from neutrality identify a very limited scope for responsibility among scientists for their work. Both rely strongly on the ambiguity in use of technology. The first suggests that the fact that technologies have both good and bad uses shows that a technology is neither good nor bad; goodness and badness attach to use. Since, the argument seems to assume, science can only gain value through technology (in this case understood as "applied science") the neutrality of technology implies the neutrality of science and thus the freedom of the scientist from moral responsibility. Responsibility lies with those who use technologies, not those who create them. This is the view of Tom Lehrer's satirical version of German-American rocket scientist Wernher von Braun: "'Once the rockets go up, who cares where they come down/That's not my department,' says Wernher von Braun" (Lehrer 1965). A second argument suggests that the same body of scientific knowledge can lead to different technologies, some good and some bad. Since science can lead to both good and bad technologies, then it must be neutral itself. Again the scientist is exempt from moral responsibility by the neutrality of their work (Forge

[3] One might claim that Bacon's view is of technology as ameliorative rather than as neutral. This would be folly. The point of *New Atlantis* is not to say that technology is unambiguously good; if it were, we would have to take Bacon as both hopelessly naïve from the perspective of contemporary technological practice and also blind to the social problems that technology created in his own time. Surely this is not the case; Bacon most certainly would have recognized that technology could be harmful. We should thus take *New Atlantis* as positing the possibility of a technological utopia as the best of the possibilities that technology presents.

1998). In the first argument, the value-neutrality of technologies directly insulates scientists from responsibility because it places responsibility on those who use the technology. In the second, the value-neutrality is shifted from technology to science, but the ambiguity of technology remains.

To build an alternative view requires rejecting two key premises of the thesis of technological neutrality. If technology is more than just a tool to be used for whatever purpose one chooses, and if ends are part of the artifact, then its claim to value-neutrality becomes unsupportable. To assert that technology on the whole is permeated by embedded values and that using technologies embeds those values in society at large is a central claim of most critical theories of technology since the Second World War. Martin Heidegger (1993, pp. 307–341) argues that technology sees the world as standing reserve and ultimately leads to humans understanding other humans as such. Herbert Marcuse (1991) focuses our thinking on the role of technology in upholding bourgeois rule and encouraging commodification. Michel Foucault (1995) demonstrates the role of technology in imposing discipline and normalization. Richard Merelman (2000) shows that the political values implicit in modern technologies are fundamentally different from those in postmodern technologies. These all suggest that technology is itself value-laden, and that by implementing technology in any form one implements values.

The Social Construction of Technology (SCOT) approach is one of the more promising social science approaches to understanding technology as value-laden. SCOT agrees with the various critical perspectives on technology that values are inherent features of technologies. But it does so in a far more sophisticated way. The SCOT program treats the development of technology not as a processed fixed by nature (as the neutrality thesis assumes) or universal social forces (as Heidegger, Marcuse, and Foucault do in various ways). Technologies are created in a historically contingent process in which scientists and technologists make choices that are rooted, implicitly or sometimes explicitly, in non-scientific judgments.

Technological development, in the SCOT approach, is seen as a process of variation and selection that is guided by the meanings given to the artifacts by social groups. These meanings are historically contingent social factors at work in the development of the technology. Key to this process is the idea of the interpretive flexibility of a technological artifact. Relevant social groups, those who have some role in the process of development, hold competing social meanings of the artifact. The artifact is, in essence, underdetermined by its natural characteristics like its physical operation, use, or utility in ways very similar to how constructivist approaches to science see scientific theories and empirical observations as underdetermined by nature. As the technology develops to its final form—a process of closure—these contingent meanings are lost through a process of stabilization in which the interpretive flexibility is gradually reduced by social processes rather than natural characteristics as some form of the artifact becomes dominant. Closure of the technological development process results in a technology that appears to be fully natural and developed through a linear, teleological process. But the SCOT program shows that there is nothing inevitable in a technology: "'successful' stages in the development are not the only possible ones," and the selection of successful and unsuccessful stages are to be explained symmetrically by appealing to the social meanings at work in the choices that scientists and technologists make. Meanings, not nature, function, or utility, are the ultimate determinants of the form of a technology (Bijker 2001).

Both the critical political theories of technology and the SCOT empirical program lead to the same conclusion. Rather than being value-neutral, technologies embody and institutionalize certain values. Technologies are value-laden. The neutrality thesis cannot be maintained, and a fundamental contradiction in the superficial understanding of technology is exposed. This understanding of technology implies the precise opposite of the neutrality thesis. Technologies are shaped by normative social factors, not only by natural forces or a naturalized concept of utility. Ideas about the good, the beautiful, the healthy, the profitable are as much a part of technologies as the physics or chemistry of the artifact. Artifacts are designed and practices developed with these goals in mind, and these are ontologically part of the technologies as much as their physical characteristics. Far from technology being value-neutral, values are inherent in technologies.

What might these values look like? An analysis of the role of values in technology based on the constructivist framework leads to four ethical claims about the structure of technological values. The first is that values are embedded in technologies and thereby in society as a whole as well. A technology is not a value-neutral material tool because it is part of a structure of value-laden meanings. As Pinch and Bijker explain, "Obviously, the sociocultural and political situation of a social group shapes its norms and values, which in turn influence the meaning given to an artifact. . . . [D]ifferent meanings can constitute different lines of development" (Pinch and Bijker 2005). Constructivists hold that these meanings are ontologically part of the associated technologies (i.e., the technology cannot exist in its current state separately from these meanings) and embed the underlying values in technologies. If values are embedded in technologies, those values become embedded in society as well when the technology is implemented in society. As actors practice the technology, they bring about the consequences of the values embedded in it regardless of the values that the user holds. Implementing a technology is thus, Feenberg argues, the act of choosing "civilizational alternatives" (Feenberg 1991), different societies differentiated by the values embedded in them by technologies.

A second conclusion is the imposition of values comes with *each* technology, not just with technology in general. Technologies do have common features. If technologies are built by common social structures, the values of those structures should be embedded in the technologies that result. If technology itself has some common value—for example, understanding improvement as increased efficiency—that value should be present in all technologies. But the common features of technology do not exhaust the set of embedded values. Understanding the social place of a technology demands understanding it specifically, as each will be composed of different meanings and therefore embed different values than others. A concept of human psychology is at work in both medical testing and mass media, but it is a very different one: rational action is embedded in medicine, while unconscious motivation is embedded in television commercials. Each specific case demands its own analysis.

The third point is closely related. If specific technologies, and not just technology in general, can embed values, then each will embed somewhat different values based on the contingencies of the relevant social groups, the process of stabilization, and the contingencies of the experiences that underlie the key relationships in the technologies. The embedded values of specific technologies and of technology in general will thus be pluralistic rather than monistic. Technology in general can be standing reserve, commodified and bourgeois, and normalizing simultaneously. Online shopping may encourage normalization through advertising at the same time that it empowers consumers to express their individual sense of style by expanding their choices. It is possible to embed many different values in a technology, and even to embed conflicting ones. Understanding the social consequences of technology requires understanding the complex patterns of value in each specific technology rather than (or at least in addition to) a general monistic theory of technology.

The final point is the most consequential for political practice: the values embedded in *how* we do (i.e., the technology) can conflict with those of *what* we do (the action itself or its larger social context) when the neutrality thesis guides our understanding of the technology. The multiple values that could be embedded are now seen as either choices that individuals make in deciding how to use a "mature" technology or the natural (and therefore value-neutral) features of the technology itself. But if values are embedded in the technology, then the choice is made not in choosing how to use the technology but in the design process itself. In practice, the original values remain embedded in the technology, and implementing it remains an act that implements those values as well. By the time that the technology is ready for use (and thus ripe for the kinds of choices that the neutrality thesis focuses on), the values that it will embed in society will already be embedded in the technology by the process of constructive stabilization. Using the technology in any sense will embed those values whether we actually hold those values or not, choosing the resulting society whether we want it or not.

This leads to an important conclusion about normative problems associated with technologies. In a society dominated by technological neutrality, technologies will often pose irresolvable conflicts among the values embedded in and implemented through a technology and the values held by society more generally but not embedded in the technology. When we implement technologies, we assert their values as well, bringing about a particular society regardless of the values that we claim to hold. It is thus the former set of values, not the latter, that govern the social consequences of those technologies. The result is that the opportunity to choose among alternative directions for society is missed, hidden by the neutrality thesis.

1.3 A Critical-Constructive Alternative

A common thread in critical perspectives on technology is the rejection of realist or positivist views of data in favor of constructive views along the lines of the SCOT approach. Such views are deeply challenging to commonly held ideas about the moral status of data itself and the information technologies that manage data. As an alternative to technological neutrality, I present a critical-constructive view of technology that makes the details of technological development a central question. Technologies are formed through a process of selection in which alternative forms of the technology are winnowed into a final form by social and political forces as much as by scientific and engineering ones. These alternative forms allow one to explore the values of a technological system critically, opening technologies to examination as questions of justice. This philosophy of technology forms the basis for the analysis in the rest of the book.

Langdon Winner (1993) strongly criticizes the SCOT framework on several grounds related to its treatment of normative issues. He argues, in my view correctly, that SCOT is not generally concerned with the social consequences of technology and that it is generally ignorant of the larger moral and political questions that technology poses. A similar critique is offered by Hans Radder (1992), though his approach focuses primarily on normative implications of constructivist methodology rather than of the constructive nature of technologies themselves. One must certainly recognize the limits of Winner's critique: His claims do little to fundamentally challenge the SCOT program itself as these criticisms are less theoretical failures than consequences of the fact that the SCOT program is a program trying to empirically explain the development of a technology.[4] But in a broader sense the point is compelling: SCOT alone cannot be critical of technology in the way that other philosophers of technology have been.

What is necessary is a critical-constructive approach to the values in technologies. That possibility emerges in considering the alternative forms of the technology that could have been. Young's core premise of critical theory takes centrality here:

> Critical theory is a mode of discourse which projects normative possibilities unrealized but felt in a particular given social reality. Each social reality presents its own unrealized possibilities, experienced as lacks and desires. Norms and ideals arise from the yearning that is an expression of freedom: it does not have to be this way, it could be otherwise. (1990, p. 6)

Technologies offer the possibility of many possible ends-in-view, and a critical view of technology facilitates rather than restricts making effective, critically reasoned choices in these questions. As philosopher John Dewey argued,

> New technologies and techniques are multi-valent, that is, that they offer all sorts of new possibilities and that it is the obligation of those who use them to choose the best of those possibilities and then rework them in order to render them more valuable. (Hickman 2001, p. 59)

[4] In other words, Winner is Reviewer #2 complaining that SCOT scholars didn't write the paper he wanted to read.

If one replaces the word “multi-valent” in this passage with “interpretively flexible” and shifts the locus of responsibility from use to development, one has not only a position very similar to the SCOT approach but with the addition of an obligation on the part of those constructing the technology to do so responsibly and critically. Technology appears neutral in a sense *because* of its interpretive flexibility—because it is swimming in a sea of indeterminacy—in that it does not inherently entail any one set of values until closure is reached.[5] But it will ultimately be value-laden as closure is reached and possible forms of the technology are foreclosed. Those who move the technology toward closure are responsible for the values that are ultimately embedded in a technology because they are making the design choices that do so.

Dewey holds that the apparent neutrality of science and technology leaves society “forced to consider the relation of human ideas and ideals to social consequences which are produced by science as an instrument” (1981, p. 390). Science and technology have social responsibilities, he argues; they “must, in short, plan [their] social effects with the same care with which in the past we have planned [their] physical operation and consequences” (1981, p. 392). To leave the choice of these consequences to private interests is to abdicate the responsibility that technology has to society. It may appear problematic that Dewey sees that responsibility as control until one notes that for Dewey control means most fundamentally the ability to act in a *self*-controlled manner, that is, to act with knowledge and understanding that allows one to bring about in practice the consequences that one expects from one’s beliefs (1981, p. 395). If the closure of technology will result in some values being built into society, then it is indeed irresponsible not to inquire into whether those values *should* be built into society.

Embedded values are seen not as universal claims but as ends-in-view that are therefore subject to evaluation and revision as well. As Dewey puts it:

> Only recognition in both theory and practice that the ends to be attained (ends-in-view) are of the nature of hypotheses and that hypotheses have to be formed and tested in strict correlativity with existential conditions as means, can alter certain habits of dealing with social issues. (1981, p. 407)

At the very least, this critical-constructive philosophy of technology demands a kind of Weberian inquiry into technological values: we identify the values that are present, clarify the values by making them more logically coherent, draw out the implications of these values, and predict the consequences that one might expect from implementing technologies with particular values embedded in them (1949, pp. 20–21, 52–55). We likely will go at least as far as invoking the later ethics of Dewey’s predecessor in the development of pragmatism, Charles Sanders Peirce, who defines ethics as the “study of what ends we are deliberately prepared to adopt” (Peirce 1992, p. 200, vol. 2). The evaluation of norms under pragmatic

---

[5] I use the word “neutral” here for consistency with Dewey’s language. It would be more proper to say that technology is pluralistic in that even given interpretive flexibility, a technology will not permit all value possibilities equally. I believe that this position is more consistent with Dewey’s larger ideas regarding technology as well.

inquiry compels us to change our technologies if we are not prepared to deliberately adopt the ends that are embedded in technology because it reveals that we hold doubts about the rightness of those ends.

It might be possible to go a step further than this. Cheryl Misak (2000) holds that pragmatic inquiry is necessarily responsive to moral as well as observational experience. She argues that in truth-seeking inquiry, the assertion of a proposition entails that one believes that it is true, that one is committed to defending it, and that one is committed to abandon it in the face of compelling evidence and argument against it because one seeks truth in making a claim. This makes one sensitive to experience, which, Misak rightly shows, means more than just observational experience; a proof can be seen as an analytical experience. Misak shows that moral inquiry is subject to certain kinds of experience under conditions similar to those of the natural sciences. One's moral judgments, for example, are shaped by background beliefs which vary much more than those of scientists but operate in the same fashion. Thus she concludes that one's moral claims are sensitive to one's experience—and that of others—in precisely the same way that other kinds of inquiry require. So long as one maintains that one's moral belief is true, one is committed to respond to empirical and analytical experience just as with one's empirical beliefs. Critical-constructive technology should thus be able to criticize the beliefs that are inherent in technology much as it could criticize empirical beliefs, at least within a broad framework of moral pluralism.

In building a theory of information justice, this book challenges especially such ideas of technological neutrality and determinism in information technology. For all of the celebration of (and weeping and gnashing of teeth over) the purported ubiquity of data collection (e.g., Shilton 2009) and data as the "detritus" of human life (Learmonth 2009) in contemporary affluent societies, data—which we can understand preliminarily as systematically collected and stored information—does not, in fact, simply happen, nor is it a neutral, objective reflection of reality. Data exists only when information is transformed into data through a process of formatting, recording, making it retrievable and relatable, and communicating that information. It is, in an important sense, a form of communication between actors that embeds the assumptions and worldview of those actors in what is communicated. It is, like all technologies, a construct, an operationalization of an actor's concept and reality, interpreting between the physical world and the intellectual structures by which actors understand that world, and embedded in a set of social practices by which it is created, interpreted, and used. It exists as just one element of a technology of data analysis that also includes statistical methodologies, data management systems, and ends for which data can be used. Data systems are thus neither stores of objective information nor inherently democratic technologies but rather technological arrangements that serve as forms of order: "ongoing social process[es] in which scientific knowledge, technological invention, and corporate profit reinforce each other in deeply entrenched patterns that bear the unmistakable stamp of political and economic power" (Winner 1980, p. 126). Data systems should thus be viewed critically in the sense that Iris Young wrote of critical theory: "Each social reality presents its own unrealized possibilities… it does not have to be this way, it could be otherwise" (1990, p. 6). This makes data amenable to political analysis: Why should the data be the way it is rather than some other way? That is the fundamental question guiding the analysis in this book.

1.4 Theorizing from One's Own Experience

While this is primarily a work of social theory, it was spurred in part by questions arising in my own experience with information systems in higher education and is written in close conversation with socio-technical practices, especially in higher education administration. It thus requires some deep exploration of the actual structures and practices of information technologies, and a justification for relying on my experience in that exploration. I will frequently draw on the data system in place at Utah Valley University (UVU), where I worked as a Senior Research Analyst in its Institutional Research & Information (IRI) office from 2009 to 2013. That experience involved extensive work in data extraction and limited database design and administration, primarily in the Banner Operational Data Store (ODS) database. This is supplemented by narrative analysis of the Structured Query Language (SQL) implementing the data systems and the data standards established by the federal Integrated Postsecondary Education Data System (IPEDS) and the Utah System of Higher Education (USHE) reporting processes.

Since UVU's systems are a key touchstone for this work, it will be valuable to understand a bit about them. UVU's data backbone during this time was the Ellucian Banner relational database running on an Oracle 10 g database server.[6] Banner consists of a normalized set of several thousand data tables managing student and administrative data and optimized for Online Transactional Processing (OLTP)—entry and modification of individual data points to maintain records of transactions—locally referred to as "Prod" (a reference to it as the production database). The bulk of institutional data analysis is performed using the ODS, which consists of a denormalized set of fewer but much larger tables optimized for Online Analytical Processing (OLAP)—extracting full datasets for analysis. The data contained in the ODS is either identical to or derived from that in Prod but organized into a different structure of fields and tables. Both databases are extensively customized for UVU. Prod and the ODS also connect to several other data systems, including the advising information system, Ellucian Student Success CRM, and the learning management system.

[6] As a full review of database structure and operation is beyond the scope of practicality here, tedious for those already familiar with them, and redundant given the many excellent sources available, this discussion presumes a basic, non-technical understanding of databases. I have aimed to provide enough background to understand the points in my argument in ways that do not overly burden those unfamiliar with databases with technical knowledge but are still recognizable to technical specialists. I apologize to readers of both sorts to the extent that I haven't succeeded in that.

Most government reporting comes from three customized relational tables. One table, referred to locally as STUDENT,[7] contains information that is constant about individual students across courses within a term such as demographics, contact information, or overall academic characteristics. The second table, COURSE, contains information that is constant across all students in a section for a term. The final table, STUDENT_COURSE, contains information specific to a student within a specific course, such as course grade or (since some courses can award variable credit) credits attempted. Using appropriate joins, STUDENT, COURSE, and STUDENT_COURSE can provide most of the information that the institution would need to understand its students and academic offerings. For example, joining STUDENT and STUDENT_COURSE would allow the institution to determine the distribution of courses taken by major and gender. STUDENT_COURSE would identify the courses taken by each student; STUDENT would provide the major and gender information. Each table is a "live" data table, showing data as it exists currently for all terms (including any transactions that affect data for a term after the term has ended, such as retroactive withdrawals from courses). A set of "freeze tables" contain data snapshots allowing time-series analysis throughout a term, and include freezes for the official census and end-of-term reporting dates.

These frozen data from the official reporting dates is used principally for state and federal government reporting. But there is a strong expectation that data reported by the institution for non-government purposes, including that used to make and justify decisions, will be consistent with the government reporting data. For example, between 2010 and 2012, UVU created a web-based data dashboard to provide more specific information on retention and graduation rates than was reported to IPEDS. It nonetheless relied on IPEDS definitions of retention and graduation rates, demographic categories, and reporting cohorts. The cohort definition is especially important, as the IPEDS cohort includes only first-time, full-time, degree-seeking undergraduates entering in fall, a relatively small portion of UVU's students. Because of the expectation that locally used data will be consistent with government reporting data, the data processes in place at UVU are defined disproportionately by the rules that govern the three customized government reporting tables.

My work with UVU's data systems forms the basis for developing a political theory of information. Theorizing based on this experience raises two challenges of justification. The first is methodological. While certainly the experience with this system is less systematic as a data collection technique than a traditionally empirical study would demand, given that the objective of this book is to establish a theoretical framework for understanding data as a type of social artifact that influences the achievement of social justice, it does not seem unreasonable to interpret

[7] Table and field names will be indicated in capital letters, with TABLE_NAMES in Roman typeface and *FIELD_NAMES* in italics. Specific table names have been replaced with generic descriptive names to maintain data security and facilitate functional understanding. These descriptive names are often correspond with similar tables and fields included in a standard Banner installation that may exist but are generally not used at UVU. Field names have also been changed where the name in the table is sufficiently obscure to make understanding difficult for the reader.

that experience using frames and techniques common to emergent methods in social science. The approach used here shares some (but not all) features with constructivist grounded theory (Charmaz 2008). This approach is especially appropriate for the study of information systems on three grounds that are especially relevant to the study of information justice:

> [F]irst, it was useful for areas where no previous theory existed; second, it incorporated the complexities of the organizational context into the understanding of the phenomena; and third, that [grounded theory method] was uniquely fitted to studying process and change. (Urquhart 2007, p. 341)

There are clear parallels between grounded theory and the work presented here. As I move between experience and theory, I use an abductive approach to building theory from experience in which both methods of inquiry and substantive findings are emergent rather than predetermined, testing the concepts developed previously for consistency with further iterations of inquiry. My approach also works at a distance from existing literature on other problems in information systems and technological ethics in order to avoid artificially constraining the emergence of a broader theory of information justice. (Urquhart 2007, pp. 350–351)

However, I must stress that understanding the creation of data using grounded theory was not intent at the outset of this project; grounded theory is itself emergent in this research. It does not, for instance, rely on the formal data collection processes of open coding or memo writing. Kelle (2005) and Charmaz (2008) provide exceptional reviews of these specific techniques, defending respectively the two distinct methodological approaches created by the schism between Glaser and Strauss, the founders of grounded theory. But this may well be a virtue; at the least it is not as great a weakness as guides to grounded theory would imply. I would suggest that the focus on specific methods in that schism has missed the real strength of grounded theory: its reliance on abductively created theoretical concepts that are iteratively tested and refined. It is this aspect on which I draw in developing a theory of information justice.

This view of grounded theory would, especially, be more consistent with the approach's Peircean roots, in which, I have previously argued (Johnson 2000), the origin of theory is a creative act and science consists not in the body of knowledge but in subjecting claims abstracted from experience to the examination of further experience. Specific approaches to code development are not necessary for the success of grounded theory in the same way that, for example, successfully passing tests of statistical significance is for quantitative research using a hypothetical-deductive methodology. From this perspective the test of good grounded theory is its tendency to iteratively approach theoretical saturation rather than its compliance with any particular research procedure, and specific coding processes are evaluated from a purely instrumental perspective (i.e., is it helpful for moving toward theoretical saturation). The lack of compliance with such procedures in this book might thus argue for its inefficiency but not its inadequacy as a work of grounded theory.

That said, this work is not remotely intended to approach theoretical saturation, and its quite weak implementation of grounded theory methodology is merely an initial iteration of the process and thus valuable as a preliminary approach to the emerging question of information justice. Ultimately, while written in conversation with and as an interpretation of experience, this book is a work of normative social theory, the aims of which include making sense of the empirical and structural contexts of a set of normative questions and showing that understanding the former are essential to answering the latter. My methods are suitable for that context—given the importance of structure and practice in my argument, they are far more suitable than straightforward philosophical theorizing—and I make no further claim to any sort of methodological rigor appropriate to more strictly empirical research.

But while the methods may be sufficient for theorizing my own experience, this only heightens the second challenge: What makes my own experience, rather than claims to universal principles, worth theorizing? Political theory is not oriented toward theories of the particular. This was the heart of Jeffrey Issac's (1995) seminal—by which I mean widely read, widely praised, and in practice widely disregarded—article, "The Strange Silence of Political Theory." Isaac famously criticized political theory for its complete disregard of the collapse of communism as a topic for study—two of 384 articles in the major journals in the field published between 1989 and 1993 addressed the fall of the Iron Curtain. He argued that political theory had become too focused on the problems of "normal science" presented by the Western philosophical canon, which "engenders intellectual conformity and inhibits more engaged, colloquial, relevant kinds of inquiry." As enabling as the canon can be, it can also be "a cloak…that conceals and obstructs political reality and our ability to experience it and interrogate it." In consequence, political theory prefers abstract problems:

> It seems almost beneath us to examine mundane, practical political problems located in space and time, in particular places with particular histories. These inquiries, we apparently reason, can be safely left to historians and political scientists. How much more edifying, rigorous, hip, virtuous, it is to discuss the constitution of the self, the nature of community, the proper way to read an old book, or the epistemological foundations of lack thereof that are involved in examining mundane political problems. (1995, p. 643)

The problems of the real world, for Isaac's contemporaries (many of whom are still active two decades later), serve as examples of theory rather than objects for theory to engage and develop itself through. "Political theory," he writes, "fiddles while the fire of freedom spreads, and perhaps the world burns." (1995, p. 649)

Isaac's alternative motivates this book. Without rejecting the importance of the abstract or the exegetical, he called for political theory:

> …to acknowledge this world as a source of intellectual and practical problems, to engage it in all of its empirical and historical messiness, to demonstrate that our categories help to illuminate this political reality and, dare I say, to improve it…. Real political problems ought not be the pretext for scholarly investigations of other things; they should be what drives our inquiries. (1995, p. 646)

Academic conversations about the disciplinary canon (of authors and topics) cannot be the only form of political theory, in Isaac's view. Instead, political theory must embrace the kind of pragmatic political theory that was once characteristic of American political life, less concerned with ideological anchors and more concerned with living politics that makes major trends intelligible.

It has taken time—longer than it took me to move out of political theory and into administration because of the strange silence of political theory, not only on 1989 but on so many other political events—but we see much improvement today. As of this writing, the current volume of *Political Theory* includes essays on climate change and reinsurance (Lehtonen 2017) and on Nietzsche's place in ethnographic fieldwork (Ignatov 2017) along with ones on Plato (Valiquette Moreau 2017) and Adam Smith (Pitts 2017). But it still has not published an article on information technology. Indeed, for several years I had taken to referring to myself as the world's leading expert on information justice—by default.[8] That likely reflects in part Isaac's criticisms of a political theory that remains focused on the canon even if it is broadening its view somewhat. The focus on canonical thinkers and a standard syllabus of topics about which we can theorize makes new topics difficult to engage, even to find. This is where it becomes important to theorize one's experience. Especially as so many who are trained in political theory move into other walks of life amidst diminishing opportunities for the standard academic career of the late twentieth century—be they "alt-ac" or "post-ac"—there becomes the opportunity for so many to ask new questions simply by looking around at their work as I did, and asking the questions Young poses: Must it be this way? How could it be otherwise? Our answers to such questions will enrich both political theory and human practice.

1.5 Plan of Study

The aim of this book is to develop a political theory of information and its associated technologies in which justice serves as the primary consideration in normatively evaluating information practices. Chapter 2 examines two cases in which data presents questions of justice. Many argue as a philosophical principle that data sources should be available as widely as possible, the principle at the heart of the open data movement. But as I argue in that chapter, open data can just as easily lead to injustice: Like garbage in programming, "Injustice in, injustice out" ought to be a principle of data. In the second case, I consider what big data means for higher education. After discussing some recent examples, I identify two types of ethical challenges in the increasingly common use of predictive analytics at universities:

[8] Wrongly, it turns out. As far as I can tell, the term "information justice" was first used by Martha Smith to refer to a program which aimed "to conserve nature and to preserve humanity through creative uses of the technologies of information, knowledge, and memory . . . using the practices of rights, responsibilities, and caring connections" (2001, p. 520). There are significant differences between this approach and my own, largely in that Smith's concept of information justice is exclusively instrumental to other justice concerns; there is no argument that information may present questions of justice in itself. One might usefully characterize it as information-driven justice, in which practices of information are directed toward securing social justice broadly. But Smith does take quite seriously the preconditions for using information in the achievement of global justice, and certainly cannot be seen as anything less than a predecessor to the concepts of information justice currently in development.

challenges related to the direct consequences of the systems and those rooted in the ideology of scientism that inspire them. Both the open data and big data cases prove quite problematic if the aim is just data.

Chapters 3 and 4 establish the political processes and structures behind information systems. In Chap. 3, I show that data is not an objective representation of reality but rather a constructed translation of observations into legible elements designed to support, broadly speaking, governance (be it by the state or by private actors). Both technical and social structures influence this translation; the technical aspects of database architecture are insufficient by themselves to define this translation regime. Such regimes can contain three characteristic translations: normalizing translations that separate the normal from the deviant, atomizing translations that separate complexity into individual elements, and unifying translations that group diverse characteristics into categories. At the same time, these data systems translate their subjects into "inforgs," representations that consist of bundled information rather than actually existing subjects. These acts of translation, I conclude, are significant exercises in political power. Chapter 4 extends the analysis of the previous chapter to the role of metrics in political practice, using the U.S. standard graduation rate metric as a case. I argue that information is best understood as a process of communication in which observation is encoded into data through the translation regime and then decoded into metrics which are then institutionalized in political processes. In both processes, political factors are prominent, making metrics a political outcome at the least. I go further, however, showing that metrics play important distributive roles in politics, allocating material and moral goods as well as the conditions of political power. Metrics also exercise political control directly, working much like administrative procedures to select favored outcomes without direct legislative intervention and building the capacity of the state to exercise control over policy areas.

Chapters 5 and 6 examine two frameworks for justice in relation to information. In Chap. 5, I seek to go beyond contemporary theories of information privacy by subjecting the standard information flow models to analysis from the perspective of justice. I examine two perspectives on justice. At the least, one can see privacy as connected to justice instrumentally, that is, privacy is valuable not as a requirement of justice directly but because it is a useful means of achieving justice. This is, I argue, hardly adequate as an entire theory of information justice but it is too easily given short shrift in discussions of privacy (especially by the wealthiest Silicon Valley titans who can protect their interests directly). A more robust approach looks to theories of distributive justice. Theories of distribution that focus on the distributive process can address two significant weaknesses in information flow models of privacy, weak conceptions of informed consent and the inability to address the original acquisition of information. Pattern theories of distributive justice shift the focus from distributing information to distributing privacy rights, and provide significant insight into what it means to have rights to be left alone or forgotten. Each of these theories makes useful contributions to our understanding of privacy. But they are not wholly adequate to the task; for this, one needs to understand justice structurally as well as distributively.

Chapter 6 engages information from the perspective of structural justice using a case study of learning analytics in higher education, drawing heavily on the "Drown the Bunnies" case at Mount St. Mary's University in 2016. This case suggests the outlines of an increasingly common approach to promoting student "success" in higher education in which early academic and non-cognitive data, often from students at other universities, are used to build a student success prediction algorithm that uses a triage approach to intervention, targeting middling students while writing off those in most need of help as inefficient uses of resources. Most common ethics approaches—privacy, individualism, autonomy, and discrimination—capture at best only part of the issues in play here. Instead I show that a full analysis of the "Drown the Bunnies" model requires understanding the ways that social structures perpetuate oppression and domination. Attention to more just organizational, politico-economic, and intellectual structures would greatly attenuate the likelihood of cases such as the Mount St. Mary's University case, adding an important dimension to information justice. I conclude by contrasting the "Drown the Bunnies" model with an implementation of learning analytics at UVU, which did much better in part because of structural preconditions that support justice.

The concluding chapter (Chap. 7) summarizes the arguments of this book, situating them amidst the booming literature on information ethics that has emerged over the (too) long process of writing it. 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.

References

Bijker, W. E. (2001). Technology, social construction of. In N. J. Smelser & P. B. Baltes (Eds.), *International encyclopedia of the social & behavioral sciences* (1st ed., pp. 15522–15527). Amsterdam: Elsevier.

Charmaz, K. (2008). Grounded theory as an emergent method. In S. N. Hesse-Biber & P. Leavy (Eds.), *Handbook of emergent methods* (pp. 155–170). New York: The Guilford Press.

Dewey, J. (1981). In J. J. McDermott (Ed.), *The philosophy of John Dewey* (Phoenix ed.). Chicago: University of Chicago 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 March 2013.

Feenberg, A. (1991). *Critical theory of technology*. New York: Oxford University Press.

Forge, J. (1998). Responsibility and the scientist. In M. Bridgestock (Ed.), *Science technology and society: An introduction*. New York: Cambridge University Press.

Foucault, M. (1995). *Discipline and punish: The birth of the prison (Second Vin.)*. New York: Vintage Books.

Heidegger, M. (1993). In D. F. Krell (Ed.), *Basic writings from "Being and time" (1927) to "The Task of thinking" (1964)* (Rev. and expanded ed.). London: Routledge.

Hickman, L. A. (2001). *Philosophical tools for technological culture: Putting pragmatism to work*. Bloomington: Indiana University Press.

Ignatov, A. (2017). The earth as a gift-giving ancestor: Nietzsche's perspectivism and African animism. *Political Theory, 45*(1), 52–75. https://doi.org/10.1177/0090591716656461.

Isaac, J. C. (1995). The strange silence of political theory. *Political Theory, 23*(4), 636–652. https://doi.org/10.1177/0090591795023004005.

Johnson, J. A. (2000). *Abductive inference and the problem of explanation in social science.* Presented at the Midwest Political Science Association. Chicago.

Johnson, J. A. (2006). Technology and pragmatism: From value neutrality to value criticality. In *Western political science association annual meeting*. Albuquerque.

Kaminski, M. (2012). Reading over your shoulder: Social readers and privacy law. *Wake Forest Law Review, 2*(Online), 13–20.

Kelle, U. (2005). "Emergence" vs. "Forcing" of empirical data? A crucial problem of "Grounded Theory" reconsidered. *Forum Qualitative Sozialforschung/Forum: Qualitative Social Research*, *6*(2). http://www.qualitative-research.net/index.php/fqs/article/view/467/1000

Kolm, S. C. (1993). Distributive justice. In *A companion to contemporary political philosophy* (pp. 438–461). Oxford: Blackwell.

Kranzberg, M. (1986). Technology and history: "Kranzberg's Laws". *Technology and Culture, 27*(3), 544. https://doi.org/10.2307/3105385.

Learmonth, M. (2009). Next-gen creatives focus on Web's data detritus. *Advertising Age, 80*(21), 14.

Lehrer, T. (1965). Wernher Von Braun. On *That was the week that was*. Reprise/Warner Bros. Records.

Lehtonen, T.-K. (2017). Objectifying climate change: Weather-related catastrophes as risks and opportunities for reinsurance. *Political Theory, 45*(1), 32–51. https://doi.org/10.1177/0090591716680684.

Marcuse, H. (1991). *One-dimensional man: Studies in the ideology of advanced industrial society*. Boston: Beacon Press.

Merelman, R. M. (2000). Technological cultures and liberal democracy in the United States. *Science, Technology & Human Values, 25*(2), 167–194. https://doi.org/10.1177/016224390002500202.

Misak, C. J. (2000). Truth, politics, morality: Pragmatism and deliberation. London: Routledge.

Nissenbaum, H. (2010). *Privacy in context: Technology, policy, and the integrity of social life*. Stanford: Stanford Law Books.

Peirce, C. S. (1992). In N. Houser & C. J. W. Kloesel (Eds.), *The essential Peirce: Selected philosophical writings* (Vol. 1–2., Vol. 1, Peirce Edition Project ed.). Bloomington: Indiana University Press.

Pinch, T. J., & Bijker, W. E. (2005). The social construction of facts and artifacts: Or how the sociology of science and the sociology of technology might benefit each other. In W. E. Bijker, T. P. Hughes, & T. J. Pinch (Eds.), *The social construction of technological systems: New directions in the sociology and history of technology* (pp. 17–50). Cambridge, MA: MIT Press.

Pitts, J. (2017). Irony in Adam Smith's critical global history. *Political Theory, 45*(2), 141–163. https://doi.org/10.1177/0090591715588352.

Plato. (1991). *The Republic of Plato*. (2nd ed.) (trans: Bloom, A.). New York: Basic Books.

Radder, H. (1992). Normative reflexions on constructivist approaches to science and technology. *Social Studies of Science, 22*(1), 141–173. https://doi.org/10.1177/0306312792022001009.

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

Schlosberg, D. (2004). Reconceiving environmental justice: Global movements and political theories. *Environmental Politics, 13*(3), 517–540. https://doi.org/10.1080/0964401042000229025.

Shilton, K. (2009). Four billion little brothers?: Privacy, mobile phones, and ubiquitous data collection. *Communications of the ACM, 52*(11), 48–53. https://doi.org/10.1145/1592761.1592778.

Singer, N. (2014, April 22). InBloom student data repository to close. *The New York Times*, New York, p. B2.

Smith, M. (2001). Global informational justice: Rights, responsibilities, and caring connections. *LIbrary Trends, 49*(3), 519–537.

Tiles, M., & Oberdiek, H. (1995). *Living in a technological culture: Human tools and human values*. London: Routledge.

Urquhart, C. (2007). The evolving nature of grounded theory method: The case of the information systems discipline. In A. Bryant & K. Charmaz (Eds.), *The SAGE handbook of grounded theory* (pp. 339–360). Los Angeles: SAGE.

Valiquette Moreau, N. (2017). Musical mimesis and political ethos in Plato's Republic. *Political Theory, 45*(2), 192–215. https://doi.org/10.1177/0090591715591587.

Walls, S., & Johnson, J. A. (2011). *From beginning to end: The transformation of individualism in classical liberalism*. SSRN. http://ssrn.com/paper=1767067

Weber, M. (1949). *The methodology of the social sciences*. New York: The Free Press.

Winner, L. (1980). Do artifacts have politics? *Daedalus, 109*(1), 121–136.

Winner, L. (1993). Upon opening the black box and finding it empty: Social constructivism and the philosophy of technology. *Science, Technology & Human Values, 18*(3), 362–378. https://doi.org/10.1177/016224399301800306.

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

Chapter 2
Open Data, Big Data, and Just Data

Abstract This chapter examines two cases in which data presents questions of justice. Many argue as a philosophical principle that data sources should be available as widely as possible, the principle at the heart of the open data movement. But as I argue in that chapter, open data can just as easily lead to injustice: Like programming, "Injustice in, injustice out" ought to be a principle of data. Social privilege can color the data that is opened and create serious inequalities in who can access and use ostensibly open data. Open data can also establish standards that exclude knowledge that is not part of the data system. In the second case, I consider what big data means for higher education. After discussing some recent examples, I identify two types of ethical challenges in the increasingly common use of predictive analytics at universities: challenges related to the direct consequences of the systems and those rooted in the ideology of scientism that inspire them. Both the open data and big data cases prove quite problematic if the aim is just data.

> "Technology is neutral—it's what you do with it that turns it into a public good." Condoleezza Rice, at ASU-GSV Summit. (@DeanOlian [Judy Olian] 2016)

> "So for example, pollution in China, environmental degradation is a hot political topic in China, and people can walk outside of their apartments, or wherever, and they know they can't breathe in Shanghai or Chengdu, or whatever. And for a long time the government was giving a pollution index number that clearly didn't bear any resemblance to reality. The U.S. Embassy started publishing a number or putting a number up, but there's also now apparently an app that you can buy that will measure the pollutants. So just the provision of information challenges the monopoly on information that an authoritarian government depends on for control and acquiescence." Condoleezza Rice (Freedman 2014)

With due respect to the former Secretary of State, technology cannot be both neutral and deterministic. If it is neutral, then authoritarian regimes will be able to use it to further their monopoly on information. If it is inherently democratic, then it is a public good regardless of what one does with it. Rice's contradiction, like many who make such arguments, is in part rooted in an equivocation on the meaning of "technology." In the former, technology is a pure object, often even an abstract concept: technology is neutral in that one can build specific technologies to further (more or less) any end. In the latter, technology has become embodied and purposed:

J.A. Johnson, *Toward Information Justice*, Public Administration and Information Technology 33, https://doi.org/10.1007/978-3-319-70894-2_2

a cell phone is used to provide information that "get[s] a little doubt in." Google's Jared Cohen argues:

> But the one silver lining in all of this is, the totalitarian societies—the true cults of personality—have literally been eliminated by the Internet in the same way that scientists were able to get rid of smallpox. Once North Korea changes, you'll never see a cult of personality again, because the ability to create a society without doubt will no longer be possible. (Freedman 2014)

And yet Donald Trump's campaign and presidency—not to mention cults of personality venerating Silicon Valley elites—certainly suggest the internet can do exactly that. Neither neutrality nor determinism seem to be effective in understanding the social effects of information technology.

This chapter takes a more complex view of the social effects of information technology (both in the abstract and in the form of specific technologies). Instead of assuming that information technology will be inherently good for society, I explore the way that questions of justice arise when information technologies are implemented in a society. I study two cases in which information technology has been held to be inherently and deterministically good: the open data movement and the use of big data and learning analytics in higher education. Each case explores injustice along different registers, open justice showing the mechanics of injustice, and big data demonstrating the conceptual levels at which injustice emerges. In both cases, the initial claims of, essentially, technological determinism founder on the myriad connections between technology and society, such as the values and assumptions built into the technologies, the complex of problems and applications in which the technologies are used, and the social structures within which the technologies operate. The argument is not that open data and learning analytics are inherently bad; such would be every bit as deterministic as their advocates argue. But Kranzberg's (1986) famed formulation that "Technology is neither good nor bad; nor is it neutral" applies well here. The fact that information technologies can have good or bad (more realistically, good *and* bad) outcomes does not mean that they are neutral either. The values and structures of technology in society ensure that any information technology will raise complex questions of justice.

2.1 Opening Government Data

With the proliferation of data in contemporary information societies comes an increasingly common call for that data to be publically accessible: an open data movement. This movement claims that open data will support democratic politics and individual liberty, unequivocally allowing individuals to use the wealth of data produced by governments and enterprises greater control over their lives and improving both their material and social conditions. "Free-as-in-speech" software and the aphorism that "Information wants to be free" as well as a distrust of political authority and consequent belief that "sunlight is the best disinfectant" have led many to argue as a philosophical principle that data sources should be available as widely as possible:

> The Internet is the public space of the modern world, and through it governments now have the opportunity to better understand the needs of their citizens and citizens may participate more fully in their government. Information becomes more valuable as it is shared, less valuable as it is hoarded. Open data promotes increased civil discourse, improved public welfare, and a more efficient use of public resources. (Open Data Working Group 2007)

The movement has come to be reflected in public policy. The U.S. government implemented an open data policy through the Office of Management and Budget's Open Government Directive, which called for agencies throughout the executive branch to take steps promoting transparency, participation, and collaboration in the publication and use of government data (Orszag 2009). Whether public or private, open data generally consists of a commitment to make data available publically in non-proprietary, machine-readable formats at the lowest level of granularity possible. As expressed by the U.S. National Science Foundation (2012), "The key principle being applied in executing the elements of the NSF Open Government Plan is: *Unless shown otherwise, the default position shall be to make NSF data and information available in an open machine-readable format.*" Similar programs range from international organizations such as the EU INSPIRE directive to local governments (Rich 2012).

This view of open data as inherently democratic is problematic, as we shall see, rooted in both a naïve view of technology and a simplistic view of politics. Open data has the quite real potential to exacerbate as much as alleviate injustices. So it comes as no surprise that open data's track record does not match its promises. The digitization of land records in the Karnataka region of India is a widely discussed case in point (Donovan 2012; Gurstein 2011; Raman 2012; Slee 2012). Three programs digitized the Record of Rights, Tenancy, and Crops (one type of land title record among others); the age, caste, and religion of owners and tenants; and spatial data. The former programs (called *Bhoomi* and *Nemmadi*, respectively) were created by the state government; the latter was part of the National Urban Information Systems program developed by the Government of India. A public-private partnership made the information accessible through internet kiosks deployed throughout the state. The promise was a system that would increase transparency and secure the rights of land tenants. The reality was a system that shifted power—and land—from local landholders to real estate developers. This is, unfortunately, rather typical of open data projects that simply approach openness as a technical condition of access. Such approaches to openness present challenges to justice in a number of ways.

2.1.1 Injustice In, Injustice Out

The constructed nature of data makes it quite possible for injustices to be embedded in the data itself. Whether by design or as unintended consequences, the process of constructing data builds social values and patterns of privilege into the data. Where those values and privileges are unjust, the injustice is then a characteristic of the data itself; no amount of openness can remedy such injustices, just as no amount of statistical processing can undo inaccuracies in the original data. "Garbage in, garbage out" is a central concept in data ethics.

Data emerges often in the interaction of an individual with a bureaucratic organization such as the state or a business. But people and groups differ in their propensity to interact with such organizations. This difference provides an important point by which privilege can enter into data. Data over-represents some, and where those over-representations parallel existing structures of social privilege, it over-represents those already privileged and under-represents those less likely to be part of data producing interactions.

Interactions with the state are rife with disparities that reflect social privilege. One well-studied example is the undercount of the decennial U.S. census (Prewitt 2010). Since the problem of undercounting was first quantified in the mid-Twentieth Century, black and Hispanic households have been undercounted at higher rates than non-black households. The causes of this undercount are myriad:

> Households are not missed in the census because they are black or Hispanic. They are missed where the Census Bureau's address file has errors; where the household is made up of unrelated persons; where household members are seldom at home; where there is a low sense of civic responsibility and perhaps an active distrust of the government; where occupants have lived but a short time and will move again; where English is not spoken; where community ties are not strong. (Prewitt 2010, p. 245)

Two commonalities in these explanations are striking: the extent to which these causes are barriers to interaction with census takers and the extent to which they are correlated with racial and class privilege. The latter causes the undercount to disproportionately affect disadvantaged groups (hence, Prewitt argues, the focus on race in debates over census methodology between 1980 and 2000), while the former prevents those groups from being represented accurately in census data. Similar problems exist in collecting data on groups such as the homeless (Williams 2010).

Groups might also be disproportionately willing to participate in some interactions over others, such as differences in thresholds for reporting building code violations between the affluent and poor (Schönberger and Cukier 2013). This is an especially significant problem in the collection of public health data on minorities, where trust in government may be lagging. Migrant groups, especially indigenous groups, refugees, and undocumented workers, frequently fear that data collected by the state will be used to their disadvantage. In many cases, such communities maintain gatekeeper institutions through which outsiders must work in order to interact effectively with the community. These groups use such structures in part as protection from states and social actors that have histories of conflict with the group, or where the groups are accustomed to high-context institutions that provide a basis for trust. But the result is that even where such groups want the data being collected, the processes that generate trust in the data collectors exclude them from the datasets.[1] Since those groups tend to be those that lack privilege, this also embeds privilege in data.

[1] Evelyn Cruz, e-mail correspondence, March 29–31, 2013.

Such privileges are not confined to interactions with the state. Residential segregation especially is often tied to forms of institutional discrimination that would influence how often individuals interact with bureaucracies. Zenk et al. (2005) found that low-income, predominantly African-American neighborhoods in Detroit were, on average, 1.1 miles further from a supermarket than predominantly white neighborhoods with similar incomes, with consequently increased dependence on smaller food stores such as convenience stores or groceries. Similarly, Cohen-Cole (2011) argues that consumer credit discrimination based on the racial composition of applicants' neighborhoods is linked to increased use of payday loans. In both cases, the use of less bureaucratized businesses by groups already suffering from discrimination in the form of de facto residential segregation (either as the legacy of formal segregation or because of ongoing discrimination) results in data that is statistically biased against such populations and reinforces whites' privileged position. Businesses can analyze the needs of the (disproportionately white) customers with whom they interact and adapt accordingly; benefits thus accrue to the beneficiaries of social privilege.

Transforming information about a datized moment into data is equally problematic. Only some of the information about that moment will be datized. What information will be is not a natural consequence of the interaction but a design choice on the part of the data architects that reflects their purposes, resources, and values. An institutional survey director noted to me that survey data at the institution is subject to state open records laws and sometimes requested by the public and state legislators. As a result, the survey director encouraged the practice of not collecting data that the institution would not be comfortable making public.[2] In this case, the concern was privacy, but this reasoning is at least as likely when more self-interested motives are present. Regardless of the motivation, though, such decisions are value-laden; thus the data built on such decisions will embody those values and transmit them in the process of using the resulting data.

Less conscious assumptions such as those part of worldviews shaped by social privilege will also shape such decisions and likely be less amenable to challenge to the extent of their invisibility to lack of diversity among the data collectors. Higher education "net price calculators," which the U.S. government requires all institutions receiving Title IV aid to produce, are designed to help students and their families estimate the likely cost of attending an institution given the prevalence of "high-tuition, high-aid" business models. This assumes that the net price is what is important to students. But Sara Goldrick-Rab (2013) argues that the gap in applications to elite colleges between high-achieving, high-income and high-achieving, low-income students reported by Hoxby and Avery (2012) is rooted in "sticker shock" at the high gross price of such institutions among low-income families in spite of the institutions' often much lower net prices. Their disregard of net price is in part a lack of information, but more significantly a consequence of such families' lack of trust in institutions generally and substantially higher risk to such families if educational institutions fail to maintain the initial promises of aid, conditions that

[2] Jane Doe (pseudonym), personal communication, March 20, 2013.

make the net price of the institutions less credible: "Being told that a college is *likely* to give you aid is not the same thing as *getting the aid*, [emphasis in original]" Goldrick-Rab writes. Such students choose to apply at less expensive (and consequently less selective) institutions as they present less risk to themselves and their families.

If Goldrick-Rab is correct, the credibility that the middle class finds in state and social institutions that have generally protected their interests should be seen as underlying the decision to collect and report average aid amounts that do not vary by income: Middle class families can credibly take average aid as typical of people like them; low-income families cannot. One might expect the same to be true of first-generation students. With family members unfamiliar with the operations of universities, they will often be unaware of issues such as net price or even understand the financial aid process at all. Yet this background knowledge, like the credibility of a measure, is assumed in the selection of data to be collected. Those privileged with such knowledge find their privileges reinforced by this data; those who are not so privileged are further disadvantaged when they cannot see the data as meaningful.

Thus we find the outcome of the digitization of land records in India described in Sect. 2.1. The selection of the RTC as the definitive data form had consequences for the distribution of land ownership. Raman argues that the programs result in the exclusion of the claims of the Dalit caste (often referred to as "untouchables"), which are often not documented in the RTC records but are well supported in other sources. Adding to this the question of how that information is stored increases the complexity of the issue. Key features in the problematic *Bhoomi* experience with open data were not only the selection of only certain types of documentation for inclusion in the land title data but also the decision to store the resulting data in a relational database system (Raman 2012). These aspects of the system design effectively precluded informal and historical knowledge from being part of the open data system; such knowledge, which was the basis of the existing land claims of *Dalits*, cannot be queried by the systems used to store the data. The two features both inform and reinforce each other: excluding narratives and other unstructured data obviates the need for systems that can handle unstructured data such as those using text-analytics or Unstructured Information Management Architecture (UIMA), while the choice of a relational database precludes the use of narrative information.

The choice of the RTC and demographic data, and the decision to accord only the RTC legal status, is also a consequence of the programs' homes in the state department of revenue, as this data was already held by these departments and is needed by the department in the course of their responsibilities. But it also reflects a bureaucratic mindset:

> The architects of the *Bhoomi* and the *Nemmadi* projects viewed the prevalence of multiple records as a manifestation of "inefficient record keeping", "corruption of field bureaucrats" and the opacity of land records due to lack of modern systems of documentation …. They sought to resolve the conflicts by identifying a single owner to a single plot of land by according a legal status to the digital RTC. (Raman 2012)

This bureaucratic mindset builds data that reflects the bureaucratic values of efficiency and consistency, doing so at the cost of excluding data that cannot be accommodated to those values. Donovan (2012) cites this as an instance of Scott's (1998) "seeing like a state," in which the local government sought to simplify society by making it legible. The open data system incorporated this value in its choice of what to datize about the moment in which land was transferred. This incorporated a value structure into the data, one that is clearly not neutral in the competition for power.

Because of the myriad ways that social privilege can become embedded in datasets, open data cannot be expected to universally promote justice. It can just as easily marginalize groups that are not part of the data, people whose lack of privilege excludes them from the kinds of interactions that produce data and makes their viewpoints invisible to those who collect data. Opening datasets composed of such data simply propagates the injustices that came into the data as it was collected. Whatever steps are taken to promote fairness in using data that is at its root unjust, the result will almost inevitably be unjust as well. Data is very much a case of "Injustice in, injustice out."

2.1.2 Open to Whom?

Normatively "clean" data is a necessary starting point for the just use of data, but it is by no means sufficient to ensure just outcomes. While open data advocates assume that, once open, the use of data is entirely unproblematic, making data meaningful in fact requires turning raw information into "intelligence": conclusions that can inform actions or serve as the basis for evaluations. Data intelligence requires bringing many complementary structures to bear on the data itself, the absence of which can lead not to data equality but to "empowering the empowered" (Gurstein 2011). Gurstein posits a seven-layer model for promoting effective use of open data that identifies many of the most important complementary structures:

1. Sufficient internet access that data can be accessed by all users.
2. Computers and software that can read and analyze the data.
3. Computer skills sufficient to use them to read and analyze data.
4. Content and formatting that allows use at a variety of levels of computer skill and linguistic ability.
5. Interpretation and sense-making skills, including both data analysis knowledge and local knowledge that adds value and relevance.
6. Advocacy in order to translate knowledge into concrete benefits.
7. Governance that establishes a regime for the other characteristics.

In the absence of these conditions it is not likely that any open data will promote justice. Britz et al. (2012) argue that these conditions are required by Amartya Sen's capabilities approach to justice; in the absence of these conditions, diverse individuals are not able to use information to act on or become something that they value.

The *Bhoomi* program described in the previous section illustrates the problems that can arise in the absence of these conditions. Raman (2012) describes real estate developers as the main beneficiaries of the *Bhoomi* program. They are better positioned to gain access to and use the digital RTC records both because they have greater computational capabilities and interpretative skills in relation to the political and legal practices governing land tenure under the program. At the same time, they also have greater social and political power with which they can assert their interpretation of the data, increasing the probability that it will be the accepted interpretation. Open data under conditions of unequal capabilities—what Raman refers to as the "capture of information"—led to frequent mass evictions of residents of slums from "productive" (i.e., desirable to developers) parts of cities where previously their ability to present conflicting claims could at least stall such processes (Raman 2012).

This problem is likely to be exacerbated by the emergence of "big" data. While the term has come to mean virtually all things to all people, four key threads emerge. The first is size: big data is often the result of device use or transactions, and so is much larger than an ordinary dataset. A common way of expressing the size is to say that "Your data might not fit easily on an Excel spreadsheet. Big Data doesn't fit on your laptop" (Charles 2013). Big data is frequently measured in petabytes, more than one million times larger than the gigabytes that measure memory in a desktop computer. But the role of size in big data is controversial; to a very important extent "big" data is Yodian: size matters not. Big data is as much about integrating multiple data sources, sources that lack common structure and in many cases lack structure at all (Craig and Ludloff 2011). The combination of size, multiple sources, and unstructured data then presents the problem of having sufficient computing power to process the data as well as the methodological skills needed to extract useful information from the data, advantages that played important roles in the re-election of Barack Obama in the 2012 U.S. presidential election campaign (Scherer 2012). Often these methods are rooted in artificial intelligence and machine learning, and the resulting output of big data analysis is more often not simply descriptive or even explanatory but in fact predictive (Baepler and Murdoch 2010).

The emergence of big data is driven largely by dramatic reductions in the cost of computing power and storage, which have made it possible for data administrators to produce data characterized by all three key values in data administration: velocity, volume, and variance.

> The advent of clouds, platforms like Hadoop, and the inexorable march of Moore's Law means that now, analyzing data is trivially inexpensive. And when things become so cheap that they're practically free, big changes happen — just look at the advent of steam power, or the copying of digital music, or the rise of home printing. Abundance replaces scarcity, and we invent new business models. (Croll 2012)

The temptation is thus to think that the intersection of big and open data, and especially of those with open-source software capable of managing and analyzing it such as Linux, MySQL, R, QGIS, and Hadoop, should minimize the capabilities differences that plagued the *Bhoomi* program.

But these tools also have capabilities requirements that often go far beyond those of ordinary citizens. Hadoop supports distributed computing and the management of unstructured data, but setting up and maintaining a Hadoop system is by no means an ordinary user skill. R and QGIS are free, but developing the skills needed to conduct advanced statistical or GIS analysis takes time and money. Petabytes of storage and teraFLOPS of processing power are "trivially inexpensive" to a large organization but not something readily available to the non-professional. In January 2014, the largest external hard drive available on Amazon.com was a mere 32 terabytes and cost $4,461.This likely explains why open data projects remain dominated by state and business users: Enterprises have the capacity to take advantage of big, open data, a capacity that citizens lack. A data store developed in Manchester, England, pooled content from ten local authorities but resulted in little citizen use beyond proofs of concept such as a bus timetable. Uses have emerged where compelling business cases can be made, and the state itself—police in particular—has proved to be an important user of open government data (Archer 2012).

The result is that big data is not, in practice, open to citizens. Opening data may allow citizens to analyze individual datasets, producing useful descriptive statistics. The empowering potential of even this should not be dismissed. But "citizen-open" pales in comparison to what might be called "enterprise-open" data. Enterprises will have the resources to get the most out of open data as they will be able to apply the full range of big data capabilities to it. They will be able to join multiple datasets together even where the data lacks structure using non-relational databases. They will be able to use proprietary business intelligence software to develop predictive models based on the data, and employ staff with the skills to both build such models and use their results. Such data is open in the sense that there are minimal restrictions on access. Insofar as it can be managed and analyzed using tools that are, to an enterprise, cheap, simple, and widely available, it is fully open to enterprises. But to the extent that such data requires capabilities that are beyond those of ordinary citizens, it cannot be understood as open to them.

2.1.3 The Normalizing Database

Injustice can emerge in systems of data as much as in any particular parts of such systems. Many of the systems of data collection to which open data advocates seek access can be usefully understood as disciplinary in nature (Adams 2013). As developed by Foucault (1995), disciplinary systems exist when individuals, regulated by a combination of detailed control and constant surveillance, self-discipline their behavior to reflect "normalizing judgment": an evaluation not of obedience to a command but of conformity to a standard of normalcy. This normativity can both impose itself on those who might wish to deviate from it and marginalize those who actually do so. Thus, to the extent that disciplinary systems take advantage of open data to impose unjust normalizing judgments or impose normalizing judgments unjustly, open data presents the possibility of undermining social justice.

This is astonishingly common in educational data, and usually deliberately and explicitly so. The U.S. Department of Education's Gainful Employment regulations required institutions to both disclose to potential students and report to the federal government information about program completion, employment of graduates, and student loan repayment. The regulations were a response to concerns about whether for-profit educational institutions were taking advantage of student aid programs to support programs that would not lead to "gainful employment" and thus expose students to excessive debt burdens and waste taxpayers' money. Preliminary data indicated that approximately 5% of programs covered under the regulations would not have met any of the benchmarks for employment and debt, jeopardizing their eligibility to offer aid. A Department of Education spokesperson stated that the regulations had led institutions "to think about what they were doing" and cut underperforming programs, a conclusion echoed by a spokesperson for Corinthian Colleges, a parent company for several for-profit colleges. The Gainful Employment regulations are a classic disciplinary program: hierarchical observation in the form of reporting requirements that are examined by an authority leads actors to adhere to an imposed norm on their own without direct coercion from the authority.

The Integrated Postsecondary Educational Data System (IPEDS) is the major postsecondary education data reporting process used in the United States. IPEDS requires educational institutions that offer Title IV financial aid to provide an extensive list of information about the institution to the National Center for Education Statistics (NCES), which then makes the data available publicly via the internet. Institutions that fail to comply risk losing their eligibility to award federal financial aid. While most of the data submitted is either demographic or input-driven (e.g., number of students enrolled or amount of state funding received), nearly all output measures IPEDS requires institutions to report concern retention and graduation. Institutions must report a first-year retention rate and graduation rates within specified percentages of normal program time. IPEDS does not require any measures of student performance, such as grade point averages, standardized test scores for postgraduate admissions, or licensing exam statistics.

These items establish the norm to which judgment is oriented: universities exist not in order to increase students' intellectual capabilities but in order to award degrees within the amount of time a normal person takes to get through the program. It must be stressed as well that "normal" most certainly does not mean "average." In practice, no disciplinary system can provide the kind of universal surveillance that Foucault describes, in which the universal possibility of observation is sufficient to ensure the self-discipline of the systems' objects. IPEDS limits the scope of surveillance by directing institutions to report graduation and retention rates on a specific subset of students, those who had first enrolled at the institution with no previous postsecondary education during a fall term intending to pursue the highest undergraduate degree offered by the institution on a full-time basis. This, too, is thus part of the norm: The "normal" student that postsecondary institutions exist to serve is the classic college student, going off to college immediately following high school graduation, studying full-time with minimal outside commitments.

IPEDS normalizes the 4-year residential university. Colleges and universities self-discipline themselves to conform to this normalizing judgment.

Educational institutions, in turn, are relying on big data techniques to create disciplinary systems that control their students. Austin Peay State University has developed an electronic advising system that suggests courses based on students' degree requirements, the extent to which courses can meet requirements for several degrees should students change their majors, and the likelihood of success in the course. Students must work through the system at registration, though they may disregard the recommendations after reviewing them. The system is a response to the problem of maintaining student aid and graduation rates:

> [Austin Peay Provost Tristan] Denley points to a spate of recent books by behavioral economists, all with a common theme: When presented with many options and little information, people find it difficult to make wise choices. The same goes for college students trying to construct a schedule, he says. They know they must take a social-science class, but they don't know the implications of taking political science versus psychology versus economics. They choose on the basis of course descriptions or to avoid having to wake up for an 8 a.m. class on Monday. Every year, students in Tennessee lose their state scholarships because they fall a hair short of the GPA cutoff, Mr. Denley says, a financial swing that 'massively changes their likelihood of graduating. … When students do indeed take the courses that are recommended to them, they actually do substantially better,' he says. (Parry 2012)

Certainly the institutional worldview that understands student success as simply completing a degree and its interest in maintaining financial aid should be apparent here. But this system, like similar systems at Arizona State University and Rio Salado College, goes a step further, using hierarchical observation and examination to promote student self-compliance with "wise choices" as the institution understands them. Here the tools of analysis and the construction of the data combine to create a data system that, open or closed, is about the institution imposing its values on students who may not share them; the data collected and analyzed is data that is relevant to a particular vision of education (credentialing) and of student success (completion). Opening the data (for instance, by allowing students to understand how the recommendations are made) does not change that in the slightest.

Hence the opening of data can function as a tool of disciplinary power. Open data enhances the capacity of disciplinary systems to observe and evaluate institutions' and individuals' conformity to norms that become the core values and assumptions of the institutional system whether or not they reflect the circumstances of those institutions and individuals. Both individuals who deviate from these norms and the institutions that specialize in serving them are marginalized in policy debates; the surveillers evaluate all institutions according to the norm (and indeed data may only exist regarding it), and the institutions internalize the norms and orient their actions to them. With the norms reflecting the power structure of the society in which they developed, they reiterate the patterns of justice and injustice that open data set out to ameliorate.

2.2 Big Data in Higher Education

Data mining and predictive analytics are increasingly used in higher education to classify students and predict student behavior. Institutions of higher education, in some cases working with commercial providers, have begun to use these methods to recommend courses, monitor student progress, individualize curriculum, and even build personal networks among students. Institutional researcher E. Rob Stirton argues that data mining, as a major part of business intelligence, is part of a radically different future for higher education in general and institutional research in particular:

> Preparing predictive models through data mining changes the focus from trends and past performance to future-oriented projections, thereby allowing planning strategies to be based on leading indicators and scenarios, which further leverage our investment in people and computers. The story changes from describing what happened to foretelling what will likely occur. Providing statistically significant predictive analytics would alter every institution's approach to Strategic Enrollment Management. (Stirton 2012)

But while the potential benefits of such techniques are significant, realizing them presents a range of ethical and social challenges. Those who implement these techniques in higher education will thus be called on to not only build the technical processes but also to protect students, institutions, and society from their side effects.

One might consider two kinds of challenges that data mining poses for institutional researchers. The immediate challenge considers the extent to which data mining's outcomes are themselves ethical. Individually, those subject to data mining—primarily but by no means exclusively students—must be respected as human beings when data mining is used to understand their characteristics and guide their actions. This means protecting both their privacy and their individuality. Institutionally, data mining may undermine the purposes of higher education in a democratic society or the missions of individual institutions. A deeper challenge, one not readily apparent to institutional researchers or administrators, considers the implications of uncritical understanding of the scientific basis of data mining. Excessively scientistic views neglect the problems of acting on conclusions that are erroneously perceived to be scientifically justified and of the meanings, assumptions, and values that are embedded in data mining applications.

2.2.1 *Data Mining and Predictive Analytics*

Data mining and predictive analytics[3] encompass practices and methods that vary greatly in familiarity to those with quantitative backgrounds typical of researchers in education and the social sciences. Some techniques, such as various regression

[3] For the purpose of this chapter, I will use the terms *data mining* to refer to the general task of identifying relationships in large datasets without *a priori* theoretical bases and *predictive analyt-*

methods, are familiar but used in different ways. Other techniques, such as k-means clustering and decision tree algorithms, have been used extensively in business—the oft repeated examples of Netflix, Amazon, and Target are now clichés in data mining—but are only recently coming to the attention of institutional researchers and educational administrators.

Data mining presents different challenges to its users than do inferential research methods—often called "academic analytics" (Baepler and Murdoch 2010). At the outset of research, academic analytics, like inferential approaches to both social scientific research and business analytics, begin from a model developed *a priori* by the researcher. The purpose of data analysis is to test the hypothesized relationships predicted by the model. Data mining, however, eschews the hypothetico-deductive process, relying instead on a strictly inductive process in which the model is developed *a posteriori* from the data itself. The model does not need to be tested against the dataset from which it is derived, as the algorithm ensures an accurate fit to that data (Baepler and Murdoch 2010; Two Crows Corporation 2005).[4]

Operating without theory requires much different mathematical techniques than academic analytics. The inferential statistics used in academic analytics work from mathematical theory and include in most cases quite specific assumptions about the underlying data (e.g., that it is normally distributed or homoskedastic); techniques are designed to minimalize computational requirements and rely on detailed specification of model form at the outset. Predictive analytics relies heavily on machine learning and artificial intelligence approaches. These take advantage of vastly increased computing power to use brute-force methods to evaluate possible solutions. Detailed model specifications are not necessary at the outset, as the process is said to "learn" the best model form over multiple iterations of the algorithm (Two Crows Corporation 2005).

The results of the two approaches are also significantly different. Academic analytics produces models whose main goal is to characterize the general tendencies in a dataset. This is most clearly the case for descriptive statistics, but measures of association and hypothesis testing statistics also have the same goal of explaining, in a single value, the general relationship between variables or the degree to which distributions would be expected by typical random variation. Even regression models, which do in principle yield predicted values for individual cases, are most typi-

ics to refer to the mathematical and computational techniques used in the practice of data mining. Readers are advised, however, that this distinction is introduced in the paper for clarity and is not based on more broadly accepted convention in the field; the two terms are often used interchangeably in the broader literature.

[4] It is, of course, advisable that the model be tested against a new dataset, often a portion of the original dataset reserved for that purpose. With some predictive analytic techniques this is necessary, as it is possible for the model to over-fit the data. Neural nets, for example, will inevitably produce a model that exactly matches the dataset on which the net is trained if allowed sufficient iterations and hidden layers, but once the model begins to incorporate stochastic variation, it will show increasing error when applied to data on which the model was not trained (Two Crows Corporation 2005).

cally used to evaluate general relationships: β is interpreted as the effect on the dependent variable of a one standard-deviation change in an independent variable, r^2 is the proportion of variance explained, and p is the likelihood that the general relationship is attributable to random variation (King 1986). Predictive analytics, on the other hand, is designed to characterize specific cases, generating a predicted value or classification of each case without regard to the utility of the model for understanding the underlying structure of the data. Many predictive analytic techniques, in fact, do not yield models capable of generalized interpretation at all (Two Crows Corporation 2005).

The result of these three procedural differences is the key practical difference between academic analytics and data mining. Under the right circumstances and with appropriate limitations, the results of an inferential test are intended to be interpreted causally. Inferential research in retention can thus be said to aim at explaining why retention occurs, and relationships between variables that cannot be understood causally—ones displaying multicollinearity or that are likely to be spurious, for instance—are of no use (Pollack 2012). Data mining, however, aims strictly at identifying data relationships. Models such as Neural Nets or Classification and Regression Trees (CART) are difficult or impossible to interpret generally; the lack of theoretical guidance in the machine learning process makes even interpretable models such as Decision Trees or Multivariate Adaptive Regression Splines (MARS) as likely to include spurious as causal variables, especially when such variables display significant multicollinearity. Such variables are valuable in data mining because they may be more effective indicators of the response variable than an ultimately causal variable that is obscured by interactions.

This is the key—perhaps the sole—reason that a strictly inductive, non-hypothesis driven approach is of value: Data mining works for the quite different purposes for which it was designed, purposes which do not include ascribing causality (Baepler and Murdoch 2010). The aim of data mining is to identify relationships among variables that may not be immediately apparent using hypothesis-driven methods. Having identified those relationships it is possible to take action based on the fact that the relationships predict a given outcome. For example, retailer Target is able to identify pregnant customers based on changes in their habitual purchasing patterns. Target mined purchasing data from customers who had signed on to the company's baby registry and was able not only to identify pregnant customers who were not in the registry, but were able to determine their approximate due date. Using this data, Target was able to tailor advertising to those customers, with the aim of changing their overall shopping habits, an opportunity that coincides with major life changes (Duhigg 2012). Clearly nothing identified by Target's efforts to data mine purchases was causal. But for Target's purposes, cause was not relevant; the company simply sought cues that would predict when a customer would be inclined to purchase particular items. Data mining is the ideal tool for such situations.

2.2.2 *Higher Education Applications of Data Mining*

The use of data mining has attracted increasing attention in higher education over the past decade. Educational data mining aims at "making discoveries within the unique kinds of data that come from educational settings, and using those methods to better understand students and the settings which they learn in" (Baker 2010). As Delavari, Phon-Amnuaisuk, and Beizadeh argue:

> The hidden patterns, associations, and anomalies that are discovered by data mining techniques can help bridge this knowledge gap [between what those carrying out educational processes know and what they need to know] in higher learning institutions. The knowledge discovered by data mining techniques would enable the higher learning institutions in making better decisions, having more advanced planning in directing students, predicting individual behaviors with higher accuracy, and enabling the institution to allocate resources and staff more effectively. (Delavari et al. 2008)

The growing interest in data mining is spurred, in part, by the increasing quantity of data available to institutional researchers from transactional databases, online operations, and data warehousing (Baepler and Murdoch 2010).

Initial research projects using data mining approaches studied several different types of outcomes such as student satisfaction (Thomas and Galambos 2004) and student assessment (Delavari et al. 2005). Based on this initial research, Delavari et al. (2008) suggested a wide range of potential applications, including predicting alumni contributions, predicting standardized test scores, creating learning outcome and institutional typologies, predicting outcomes and intervention success, predicting student performance and identifying at-risk students, and identifying appropriate academic programs for each student. Similarly, Baker (2010) suggests four areas of application: building student models to individualize instruction, mapping learning domains, evaluating the pedagogical support from learning management systems, and scientific discovery about learners. Kumar and Chadha (2011) suggest using data mining in organizing curriculum, predicting registration, predicting student performance, detecting cheating in online exams, and identifying abnormal or erroneous data. More recent applications have embraced such suggestions, exploring course recommendation systems (Vialardi et al. 2009), retention (Zhang et al. 2010), student performance (Ayesha et al. 2010; Baradwaj and Pal 2011), and assessment (Llorente and Morant 2011).

Unfortunately, these studies do not make for a promising foundation for the practice of educational data mining, because they suffer, on the whole, from major methodological flaws. None of the predictive efforts provide control data, for instance, commonly reporting a generic "accuracy rate" that is not even clearly described. For example, the course recommendation system designed by Vialardi and colleagues aimed to predict success in the recommended course and steer students away from courses in which they were likely to be unsuccessful. They reported a 73.9% accuracy rate with 80.2% of errors being false negatives. While this sounds impressive, the absence of any sort of proportional reduction in error measure of model accuracy prevents evaluation. If the pass rate for the course was 50%, the model would be impressive indeed. But if the pass rate is 90%, the model is less accurate than simply predicting that all students would pass, and thus offers no improvement on existing methods. This problem is also present in the study by Zhang and colleagues. Llorente and Morant provide a crosstabulation of results with only column percentages and do not even provide the sizes of their treatment and control groups, making it impossible to determine the statistical significance of their findings. Barawaj and Pal provide no evidence at all that their decision tree in fact has any predictive value.

There is also an exceptionally casual attitude toward attributing causation. Delavari and colleagues, for example, report a strong relationship between instructors' performance and their marital status among those with weaker academic qualifications (2008). This relationship is almost certainly spurious, probably epiphenomenal with age and experience. Similarly, Thomas and Galambos hold, "studying a single student body begins to identify aspects of the college experience *that most affect* student satisfaction" (2004, p. 265, emphasis added), without any effort to describe a causal relationship between satisfaction and the variables identified by their CHAID method. Given that data mining was not designed to support causal inferences and provides no means for identifying potentially spurious relationships, such claims are not supportable. Both of these problems will prove problematic when considering the ethics of their use.

The cases of data mining reported in the scholarly literature above have been primarily pilot projects, limited to predictions for individual courses or academic programs. In spite of this and their methodological problems, however, data mining is gaining hold operationally at the institutional level. The most common applications are within courses. Rio Salado College has developed a system that predicts student success in online courses based on early performance in the course. The system provides information to instructors about predicted student performance so that instructors can intervene to promote success. The system's developer claims to be able to predict course success on the eighth day of class with 70% accuracy. Success with intervention, such as using welcome emails to encourage students to log in on the first day of class, has been mixed, however, according to descriptions in the media. A system in use at Arizona State University uses data mining to personalize content in online courses by adapting the course content to each student. The system, developed by educational software company Knewton, provides content in online and hybrid math courses based on student behavior and past performance, focusing students on the concepts they need help with, sequencing lessons based on individual needs, and presenting content in formats suited to their learning style. Data mining has filtered into traditional classrooms as well with systems such as Harvard's Learning Catalytics. The system matches students for in-class discussions based on answers to practice problems with the aim of stimulating discussion. Students with differing answers to the practice problem are matched together in real time to debate their answers (Parry 2011, 2012). A similar system is in place at the University of Texas (Deliso 2012).

The other major application of data mining has been in advising. Course recommendation systems are in place at several universities, including Arizona State University, the University of Florida, and Austin Peay State University. Austin Peay's "robot adviser" is a response to the findings of behavioral economics that show the difficulty of making good choices when confronted with an overwhelming number of options and the often substantial consequences of marginal differences in performance. It uses recommendation algorithms similar to those used by Netflix to suggest courses based on major, degree requirements, student performance, and the past performance of similar students. Its grade predictions are accurate to one-half letter grade, administrators report, and they believe that students perform better when following the recommendations. ASU and Florida go a step further, monitoring student progress through their academic programs and sometimes intervening to force student action. ASU's eAdvising system requires students to choose a major and develops a plan for when to take courses. The plans front-load key courses so that students who aren't suited to the major are identified early. Students are marked "off-track" based on enrollment and performance, and may be forced to change majors after two such semesters. Austin Peay is implementing a similar system (Parry 2011, 2012).

Other applications of data mining are less common but may indicate where universities are taking data mining in the future. ASU mines campus identification card swipes at campus facilities to model campus social networks and student behavior, with an eye toward identifying lack of social integration or changes in behavior that suggest a student may withdraw. It can combine that data with other administrative data, for instance, requests for transcripts, to identify students to whom advisors should reach out (Parry 2012). ASU also mines Facebook data from students who have installed the university's Facebook app and recommends other students with similar interests (Deliso 2012). Admissions and recruiting are also growth areas for data mining. ConnectEDU, an online social network platform, operates as an eHarmony-like matching site for colleges. It matches students to colleges where they will fit well and allows colleges, indirectly, to contact students whose ConnectEDU profiles fit the institution's recruiting program (Parry 2011).

2.2.3 *Consequentialism: The Immediate Challenge*

Nearly from its inception, data mining has raised ethical concerns. Once implemented, a series of challenges for both the individuals who are the subjects of data mining and the institution that bases policy on it arise as consequences.[5] The most prominent of these are the related problems of privacy and individuality. The privacy of subjects in a data mining process is primarily a factor of information control: a subject's privacy has been violated to the extent that the opportunity for consent to

[5] In this section, I use "consequences" and related terms strictly in a non-technical sense, referring to moral conditions that arise consequent to the implementation of a data mining process. At this point, I take no position on the relative merits of formally consequentialist or deontological ethical theories in evaluating those circumstances, though it will become clear to readers familiar with the distinction through the examples that follow that I believe that both kinds of ethical theory at least raise questions that data miners should address.

collection or use of information is absent or in which personal information flows are used in ways that are incompatible with their social context (Nissenbaum 2010; van Wel and Royakkers 2004). The potential of data mining to violate personal privacy spans a range of applications. At its least intrusive, the data collection and storage capabilities that make data mining possible allow those with whom one interacts to develop a complete dossier about those interactions. This leaves one's privacy unprotected by the failures of human memory. Mining that data allows one to infer information about the data subject that some would not be comfortable divulging themselves (as in the Target example described above). At its worst, privacy violations allow for the manipulation of or discrimination against the subject, for example, by price discrimination and restrictive marketing (Danna and Gandy 2002).

These risks are very much present in higher education applications of data mining. Course recommendation or advising systems that consider student performance are a way of developing a comprehensive picture of student performance, in essence, an electronic reputation that the institution maintains and makes available to faculty and staff through dashboard and stoplight processes and administrative rules. Unlike a personal reputation among faculty in one's major, an electronic reputation seems more difficult to escape. It seems unreasonable to expect that Rio Salado College's students universally want a system to identify them as more likely to fail; even if the intent is to encourage faculty to reach out to those students, undoubtedly many students would feel stigmatized instead. Arizona State University's effort to identify students who intend to transfer is clearly not information that students would consistently want to divulge, as one ASU student reported (Parry 2012).

Privacy concerns can easily give way to challenges to individuality. To be sure, such challenges are not new; older techniques that describe central tendencies and typical relationships can easily be seen as contributing to a collectivization of subject, where all are treated identically based on the assumption that they are all "typical" students. Data mining can go far toward overcoming this because it recognizes and models diversity among subjects. Thomas and Galambos, for instance, used the CHAID decision tree method to find "a significant dimension of diversity among the undergraduates in a public research university … identifying different satisfaction predictors for different types of students" (2004, p. 259). To the extent that the model is reasonably comprehensive and causally supportable (necessarily by other means) and that the data mining technique does, in fact, aggregate characteristics to something that represents the whole person, this individualization is to be preferred over collectivization.

But while academic analytics tends to collectivize the students by treating them all identically to the central tendency case, data mining has a tendency to disaggregate the whole individual into nothing more than the sum of a specified set of characteristics. Data mining can create group profiles that become the persons represented:

> Profiling through web-data mining can, however, lead to de-individualisation, which can be defined as a tendency of judging and treating people on the basis of group characteristics instead of on their own individual characteristics and merits. … In non-distributive group profiles, personal data are framed in terms of probabilities, averages and so on. (van Wel and Royakkers 2004, p. 133)

These profiles treat the subject as simply a collection of attributes rather than a whole individual, and interfere with treating the subject as more than a final predictive value or category. Course recommendation systems are just such a case; students are encouraged to do what students like them have done before. Austin Peay's system does not consider student interests, while Arizona State's eAdvising system is built specifically to identify students whose "ambitions bear no relation to their skills" (Parry 2012). This suggests that the students, far from being understood as individuals, are simply bundles of skills that need to be matched to an outcome.

At its extreme, data mining can undermine individuals' autonomy. Broadly speaking, autonomy can be understood as the ability to critically reflect on and act so as to realize or modify one's preferences, particularly preferences among conceptions of the good. This raises the questions of whether coercion and paternalism are ever justified, questions that are often addressed on the basis of a principle of preventing harm to others, furthering ends that the objects of the paternalism values themselves, or addressing a limited capacity for autonomy on the part of the object (Dworkin 1995). An especially complicated form of interference is the creation of disciplinary systems, wherein the control of minutiae and constant surveillance lead subjects to choose the institutionally preferred action rather than their own preference, a system that generally disregards autonomy (Foucault 1995).

Data mining can easily be coercive, paternalist, or disciplinary. ASU's system of compelling students making insufficient academic progress to change their major is very much coercive. The sense of promoting "wise" choices in Austin Peay's course recommendation system is a classic example of paternalism. Classifying students and communicating the classification to the professor used at Rio Salado College is virtually identical to Foucault's example of the Nineteenth Century classroom (1995, pp. 146–149) and could be expected to have similar effects: encouraging conformity to a set of behaviors that the institution has conceived of as successful. One might justify these interferences with student autonomy as preventing waste of taxpayers' money (a harm to the taxpayer, arguably), as furthering the educational ends that students presumably have when they enroll, or as guidance for those who are still not fully mature or lacking information about the consequences of a decision. But it remains necessary to provide such a justification in each case, as violations of the principle of autonomy are generally justified only as exceptions to the broad aim of allowing each person the maximum autonomy consistent with all others also having such autonomy.

It is not only the individuals whose data is mined, however, whose moral status comes into question when institutions use data mining. Many of the applications of data mining discussed above present moral concerns regarding the institution not as an actor but as one affected by the action. These chiefly concern the relation of data mining to the purpose of higher education, especially in a liberal democratic society. There are, of course, many such purposes. Peters argues that education is a process that leads "to the development of an educated man in the full sense of a man whose knowledge and understanding is not confined to one form of thought or awareness" (2010, p. 14), a perspective that one might call critical education. Flathman (1996) goes further, arguing that education ought to enable the individual to make critically informed choices among conceptions of the good life, which he sees as the essence of liberal education. University of Pennsylvania president Amy Gutmann argues that democratic education ought to prepare students to participate in the processes of public deliberation over policy that guide representative government; higher education, especially, has an important place as a refuge for unpopular ideas, promoting values for professions that are not promoted by market forces such as professional virtue, and promoting communities that share intellectual and educational values (1999, pp. 172–193).

At the same time, higher education also has more practical purposes. It smacks of elitism to deny that students should pursue higher education for vocational purposes. It is as naïve to disregard higher education's role in establishing and maintaining social classes as it is cynical to disregard its role in promoting class mobility. Moreover, discussion of the purpose of "higher education" in general ignores the fact that each university may have its own specific purposes as well, deriving from its history, community, and governance. The practical and unique purposes are as important to a university's moral circumstances as are general views of what higher education should be.

Data mining can both contribute to and undermine these purposes. Mining data to find courses and majors in which students will be successful, like Arizona State, Florida, and Austin Peay do, may contribute to the vocational goals that many students have when they enroll in higher education. Students who find fields in which they will be academically successful are, it stands to reason, more likely to be professionally successful as well. But at the same time, those may be courses and majors in which students are successful because they are not challenged; likewise, personalized curriculum may provide the easiest path to course completion but not the surest path to learning. Where they are not challenged academically, they may not ever be critically educated in Peters' meaning. Where they are not challenged by divergent ideas, they may not ever be liberally educated in Flathman's sense or able to deliberate rationally as Gutmann would have them. ASU's social data mining is especially problematic for both democratic education and the status ambitions of many students in that it will almost certainly tend to reinforce the class relationships that students have when they enroll, preventing them from deliberating with a view toward the perspectives of others and from forming networks with others that would aid their social mobility.

2.2.4 Scientism: The Deep Challenge

The consequential challenges of data mining are the most prominent ones, but they are not the only ones. In fact, the most difficult challenges may be ones of which institutional researchers are least aware. In the process of designing a data mining process, institutional researchers build both empirical and normative assumptions, meanings, and values into the data mining process. These choices are often obscured by a strong tendency toward scientism among data scientists. For philosophers of science and technology, the term refers (almost always critically) either to the claim that the natural sciences present both epistemologically and substantively the only legitimate way of understanding reality or to instances of scientific claims being extended beyond the disciplinary bounds in which the claim can be supported (Peterson 2003).[6] Such perspectives introduce the temptation to uncritically accept claims that purport to have scientific backing. This was a recurring theme in Twentieth Century political philosophy, one reflected in Dewey's (1954) critique of expertise, Arendt's (1973) analysis of Hitler's racial theories, and Habermas' (1990) communicative ethics. Given the mathematical precision and rigor of data mining, the temptation to accept the results as scientifically established and thus an unequivocal representation of reality is strong.

Scientism has a long tradition in the social sciences, and especially in the study of education (Hyslop-Margison and Naseem 2007). Critics of scientism in education see a fetishization of the scientific method, which manifests itself in contemporary policies such as *No Child Left Behind* and mandates "scientific" evidence of effectiveness as an authoritative practice of politics (Baez 2009). The preponderance of such methods in education research—and especially in the kinds of studies produced by institutional research offices—point to the assumption that traditional scientific methods are the ideal approach to understanding contemporary higher education. Indeed, one need look no further than the AIR standards for designation of a presentation as a "scholarly" paper: "Scholarly papers must include research questions, methodologies, literature reviews, and findings" (Association for Institutional Research 2012). Surely one would not dismiss disciplines such as philosophy or literature as non-scholarly for not being organized as AIR suggests; that is not the appropriate organization for scholarly work in those disciplines. The AIR standards confuse "scholarly" with "empirical," a confusion rooted in the positivist dismissal of the non-observable as unknowable "metaphysics."

Scientism is a trap that, if not avoided, can do substantial harm to students. But unfortunately, current examples of data mining in higher education have embraced, rather than rejected, scientism. The lack of attention paid to the major methodological flaws described in the previous section is a good example of scientism at work. A non-scientistic perspective critically evaluates methods and evidence before taking action upon it. But the casual attitudes toward causality and the ignorance of even statistical uncertainty in the studies of data mining in higher education suggest that the authors have taken an uncritical attitude toward the underlying science of data mining. Assuming that the relationships uncovered by data mining are inherently causal and reasonably certain can lead to ineffective actions and actions that reinforce rather than interdict causal mechanisms. Similar problems can occur when uses of data mining are insufficiently appreciative of the uncertainty present in the models; especially among users who only see the predictions and are unfamiliar with the model itself, predictions of a marginal change in likelihood can easily be implemented as predestined certainty.

[6] See, for example, Haack's (1993) critique of Quine's naturalism for a technical treatment. A useful non-technical perspective on scientism can be found in Kitcher (2012).

Rio Salado College's lack of success with intervention is telling. The welcome emails assumed that the relationship between first-day login and course success was causal; encouraging students to log in on the first day would thus increase their likelihood of success. But if both course success and first-day login are caused by students' self-motivation, a single email is unlikely to affect course success even if it does result in a first-day login; a sustained effort rather than a one-time intervention is needed. While this intervention is unlikely to harm, at the least an opportunity has been missed to make an effective intervention. The same cannot be said of potential actions stemming from the findings about lecturer marital status by Delavari et al. (2008). If the relationship between lecturer marital status and student performance is epiphenomenal to that between lecturer experience and student performance but is nonetheless used in hiring practices,[7] the university will certainly have harmed the subject of the model.

The problem of scientism in data mining goes deeper than just poor methodology. Part of the scientist epistemology is the claim that science is objective, and thus it—and its products—is value-neutral. But one of the key recent findings in both the philosophy and the sociology of science is the value-ladenness of science and technology. This is more than just claims of biases in scientific inquiry that deviate from the norms of such inquiry; it is an inherent feature of science and technology that they embody and embed values as they are created within a complex web of technical and social interdependencies (Nissenbaum 2010, pp. 4–6). Contingent meanings are as important as evidence and function in their development, as scientists and technologists make choices among equally likely possibilities or equally useful practices. Design intent and assumptions about user behavior are especially significant sources of embedded values in technologies. As technologies are themselves embedded broader structures when implemented, the values embedded in the technologies become embedded in the social context in which the technologies are used. The iteration of the technology development cycle reinforces this relationship: social values are embedded in technologies, and technologies reinforce those values (Johnson 2006).

The connection between technological artifact and social purpose suggests that data mining applications in higher education are best understood as part of a problem-model-intervention nexus. In developing models data miners link their own meanings, values, and assumptions to similar ones taken from the problem and the intended intervention. Richard Clark points in this direction when he criticizes personalized learning for its assumption that students' performance is rooted in different learning styles (a pedagogical theory that has seen its support eroded by recent research) and for questionable interpretations of data points, such as what

[7] Delavari and colleagues do not identify the university in which their study is conducted or its location, thus whether there are legal constraints that would prevent such a policy is unknown. Even if there are such constraints, however, such constraints are external to the criticism being made here; the finding and the failure of the authors to address the question of its spuriousness suggest that such conclusions are likely in areas in which the law presents no such constraint to designing an intervention around a spurious relationship that would harm the subjects of the model.

differences in time spent on a topic or learning method indicate (Parry 2012). When used properly—that is, predictively rather than causally—these criticisms lose some of their effect; if students who spend more time with video than text in one lesson are more successful when presented with video in the next lesson, it does not matter whether the relationship is epiphenomenal to an underlying motivational effect, personal preference, or difficulty with the material. The students' past behavior is sufficient to predict the success of an intervention regardless of a causal relationship. Of course, susceptibility to scientism in this respect is also likely to make one susceptible to the previous respect as well; when (mis)interpreted causally, the embedded values and assumptions of interventions based on data mining can easily become self-fulfilling prophesies.

Even when properly used, the values embedded in a model nexus become part of the institutional context. Vialardi and colleagues note that predictive analytic models "are based on the idea that individuals with approximately the same profile generally select and/or prefer the same things" (2009, p. 191). This very behaviorist model of human nature is at the foundation of every data model. While it is generally reasonable, one should note that it directly contradicts the rational utility maximizer model of human nature used in microeconomics or the habitual perspective of behavioral economics, and has very different implications for interventions. This is especially problematic in that interventions often incentivize behavior, a prescription best suited for rational utility maximizers. Similar processes embed more specific values in specific models. Most models are developed with both a problem and an intervention in mind, as can be seen in Austin Peay Provost Tristan Denley's description of the university's course recommendation system:

> Denley points to a spate of recent books by behavioral economists, all with a common theme: When presented with many options and little information, people find it difficult to make wise choices. The same goes for college students trying to construct a schedule, he says. They know they must take a social-science class, but they don't know the implications of taking political science versus psychology versus economics. They choose on the basis of course descriptions or to avoid having to wake up for an 8 a.m. class on Monday. Every year, students in Tennessee lose their state scholarships because they fall a hair short of the GPA cutoff, Mr. Denley says, a financial swing that "massively changes their likelihood of graduating." (Parry 2012)

The wisdom of a student's choice and the difficulty of making such a choice under these circumstances is part of the model; what it is to predict is not just a choice that the student will like but one which will be, from the institution's perspective, wise. And the model is specific about what constitutes wisdom: conformity to a utility function that values high grades and rapid progress toward graduation.

The question that arises here, then, is threefold: are users aware of the assumptions, meanings, and values embedded in the data model; are they consistent throughout the problem-model-intervention nexus; and is the inclusion of them justifiable? This is a question that is specific to each application of data mining in higher education, because the question is not whether values should be included in the application *per se*. There will be values in any technology; ethical applications of data mining are not value-free. They can, however, be value-conscious, even value-critical (Johnson 2006). They ask whether likelihood of success in a course is a good standard to use for recommending that a student take the course; perhaps a wise choice is one that gives opportunities to develop wisdom through struggle rather than to maintain the highest GPA possible. They ask whether a behavioral prediction regarding academic progress makes sense as the basis for a utilitarian intervention; perhaps habitual behavior needs more impetus for change than a changing utility function. Often, one hopes, the values included are entirely reasonable, and perhaps even necessary. But ethical data mining can't happen if the ethical and philosophical assumptions behind the models are not considered.

2.3 Conclusion

In Tom Lehrer's (1965) song "Wernher von Braun," the titular hypocritical/apolitical rocket scientist denies responsibility for his creations: "'Once the rockets go up / Who cares where they come down / That's not my department' / says Wernher von Braun." Data systems, similar to von Braun's rockets, are too often assumed to be value-neutral representations of fact that produce justice and social welfare as an inevitable by-product of efficiency and openness. Rarely are questions raised about how they affect the position of individuals and groups in society. But data systems both arbitrate among competing claims to material and moral goods and shape how much control one has over one's life. These are the two classic questions of philosophical justice, raising the question of information justice. Information presents questions of justice as data is created, as it is used, even by its mere existence in a data system. It presents immediate questions about the consequences of information and deeper questions about the ideology of information technology itself.

Data scientists cannot be content to say that the use of their systems is someone else's problem: where the rockets are meant to come down determines the design of the system. Understanding information as a social product requires that information scientists work with an eye toward the social, asking critical questions about the goals, assumptions, and values behind decisions that are too easily—but mistakenly—seen as merely technical. Information science requires an understanding of information justice, which requires an understanding of justice itself.

References

@DeanOlian [Judy Olian]. (2016, April 18). Technology is neutral – it's what you do with it that turns it into a public good. Condoleezza Rice, at ASU-GSV Summit. *Twitter*. https://twitter.com/DeanOlian/status/722251515629441024. Accessed 22 Apr 2016.

Adams, S. (2013). Post-panoptic surveillance through healthcare rating sites. *Information, Communication & Society, 16*(2), 215–235. https://doi.org/10.1080/1369118X.2012.701657.

Archer, P. (2012). *Report on using open data: Policy modeling, citizen empowerment, data journalism*. Brussels: W3C. http://www.w3.org/2012/06/pmod/report. Accessed 3 Mar 2013.

Arendt, H. (1973). *The origins of totalitarianism*. New York: Harcourt Brace Jovanovich.

Association for Institutional Research. (2012). 2013 forum presenter information. *Association for Institutional Research Annual Forum*. http://forum.airweb.org/2013/Pages/Information/Presenting.aspx. Accessed 1 Apr 2013.

Ayesha, S., Mustafa, T., Sattar, A., & Khan, M. (2010). Data mining model for higher education system. *European Journal of Scientific Research, 43*(1), 24–29.

Baepler, P., & Murdoch, C. J. (2010). Academic analytics and data mining in higher education. *International Journal for the Scholarship of Teaching and Learning, 4*(2), 17.

Baez, B. (2009). *The politics of inquiry: Education research and the "culture of science"*. Albany: State University of New York Press.

Baker, R. S. J. D. (2010). Data mining for education. In B. McGaw, P. Peterson, & E. Baker (Eds.), *International encyclopedia of education* (3rd ed.). Oxford: Elsevier.

Baradwaj, B. K., & Pal, S. (2011). Mining educational data to analyze students' performance. *International Journal of Advanced Computer Science and Applications, 2*(6), 63–69.

Britz, J., Hoffmann, A., Ponelis, S., Zimmer, M., & Lor, P. (2012). On considering the application of Amartya Sen's capability approach to an information-based rights framework. *Information Development*. https://doi.org/10.1177/0266666912454025.

Charles, N. (2013, March 4). Big data madness and my football prediction model. *Wallpapering Fog*. http://www.wallpaperingfog.co.uk/2013/03/big-data-madness-and-my-football.html. Accessed 24 May 2017.

Cohen-Cole, E. (2011). Credit card redlining. *Review of Economics and Statistics, 93*(2), 700–713. https://doi.org/10.1162/REST_a_00052.

Craig, T., & Ludloff, M. E. (2011). *Privacy and big data*. Sebastopol: O'Reilly. http://proquest.safaribooksonline.com/9781449314842.

Croll, A. (2012, August 2). Big data is our generation's civil rights issue, and we don't know it. *O'Reilly Radar*. http://radar.oreilly.com/2012/08/big-data-is-our-generations-civil-rights-issue-and-we-dont-know-it.html. Accessed 12 Mar 2013.

Danna, A., & Gandy, O. (2002). All that glitters is not gold: Digging beneath the surface of data mining. *Journal of Business Ethics, 40*(4), 373–386.

Delavari, N., Beizadeh, M. R., & Phon-Amnuaisuk, S. (2005). Application of enhanced analysis model for data mining processes in higher educational system. In *2005 6th international conference on information technology based higher education and training*, F4B–1–F4B–6. doi:https://doi.org/10.1109/ITHET.2005.1560303.

Delavari, N., Phon-Amnuaisuk, S., & Beizadeh, M. R. (2008). Data mining application in higher learning institutions. *Informatics in Education, 7*(1), 31–54.

Deliso, M. (2012). How big data is changing the college experience. *OnlineDegrees.org*. http://www.onlinedegrees.org/how-big-data-is-changing-the-college-experience/. Accessed 12 Sept 2012.

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.

Duhigg, C. (2012). How companies learn your secrets. *The New York Times Magazine*. https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html?pagewanted=1&_r=2. Accessed 16 Feb 2012.

Dworkin, G. (1995). Autonomy. In R. E. Goodin & P. Pettit, A. (Eds.), *Companion to contemporary political philosophy* (pp. 359–365). Cambridge, MA: Blackwell.

Flathman, R. E. (1996). Liberal versus civic, republican, democratic, and other vocational educations: Liberalism and institutionalized education. *Political Theory, 24*(1), 4–32.

Foucault, M. (1995). *Discipline and punish: The birth of the prison* (2nd ed.). New York: Vintage Books.

Freedman, M. (2014, March 26). What is the relationship between technology and democracy? *Insights by Stanford Business*. https://www.gsb.stanford.edu/insights/what-relationship-between-technology-democracy. Accessed 22 Apr 2016.

Goldrick-Rab, S. (2013, March 20). What have we done to the talented poor? *The EduOptimists*. http://theeduoptimists.com/2013/03/what-have-we-done-to-the-talented-poor.html. Accessed 24 May 2017.

Gurstein, M. (2011). Open data: Empowering the empowered or effective data use for everyone? *First Monday*, *16*(2). http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/3316/2764. Accessed 5 Mar 2013.

Gutmann, A. (1999). *Democratic education*. Princeton: Princeton University Press.

Haack, S. (1993). *Evidence and inquiry: Towards reconstruction in epistemology*. Oxford: Blackwell.

Habermas, J. (1990). *The philosophical discourse of modernity*. Cambridge, MA: MIT Press.

Hoxby, C., & Avery, C. (2012). *The missing "One-Offs": The hidden supply of high-achieving, low income students*. (No. w18586. Cambridge, MA: National Bureau of Economic Research. https://doi.org/10.3386/w18586.

Hyslop-Margison, E. J., & Naseem, M. A. (2007). *Scientism and education empirical research as neo-liberal ideology*. Dordrecht: Springer. http://public.eblib.com/EBLPublic/PublicView.do?ptiID=337528. Accessed 1 Apr 2013.

Johnson, J. A. (2006). Technology and pragmatism: From value neutrality to value criticality. In *Western political science association annual meeting*. Albuquerque. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2154654.

King, G. (1986). How not to lie with statistics: Avoiding common mistakes in quantitative political science. *American Journal of Political Science, 30*(3), 666–687.

Kitcher, P. (2012, May 3). The trouble with scientism why history and the humanities are also a form of knowledge. *The New Republic*. http://www.tnr.com/article/books-and-arts/magazine/103086/scientism-humanities-knowledge-theory-everything-arts-science. Accessed 1 Jan 2013.

Kranzberg, M. (1986). Technology and history: "Kranzberg's Laws". *Technology and Culture, 27*(3), 544. https://doi.org/10.2307/3105385.

Kumar, V., & Chadha, A. (2011). An empirical study of the applications of data mining techniques in higher education. *International Journal of Advanced Computer Science and Applications, 2*(3), 80–84.

Lehrer, T. (1965). In W. Von Braun (Ed.), *On That was the week that was*. Reprise/Warner Bros Records.

Llorente, R., & Morant, M. (2011). Data mining in higher education. In K. Funatsu (Ed.), *New fundamental technologies in data mining* (pp. 201–220). New York: InTech. http://www.intechopen.com/books/new-fundamental-technologies-in-data-mining/data-mining-in-higher-education.

National Science Foundation. (2012). *The national science foundation open government Plan 2.0*. http://www.nsf.gov/pubs/2012/nsf12066/nsf12066.pdf. Accessed 12 Mar 2013.

Nissenbaum, H. (2010). *Privacy in context: Technology, policy, and the integrity of social life*. Stanford: Stanford Law Books.

Open Data Working Group. (2007). 8 Principles of Open Government. https://public.resource.org/8_principles.html. Accessed 7 July 2015.

Orszag, P. R. (2009, December 8). Open government directive. *Office of Management and Budget*. https://obamawhitehouse.archives.gov/open/documents/open-government-directive. Accessed 24 May 2017.

Parry, M. (2011, December 11). Colleges mine data to tailor students' experience. *The Chronicle of Higher Education*. https://chronicle.com/article/A-Moneyball-Approach-to/130062/.

Parry, M. (2012, July 18). College degrees, designed by the numbers. *The Chronicle of Higher Education*. https://chronicle.com/article/College-Degrees-Designed-by/132945/

Peters, R. S. (2010). What is an educational process? In R. S. Peters (Ed.), *The concept of education* (pp. 1–16). Oxford: Routledge.

Peterson, G. R. (2003). Demarcation and the scientistic fallacy. *Zygon, 38*(4), 751–761. https://doi.org/10.1111/j.1467-9744.2003.00536.x.

Pollack, P. H. I. (2012). *The essentials of political analysis* (4th ed.). Washington, DC: CQ Press.

Prewitt, K. (2010). The U.S. decennial census: Politics and political science. *Annual Review of Political Science, 13*(1), 237–254. https://doi.org/10.1146/annurev.polisci.031108.095600.

Raman, B. (2012). The rhetoric of transparency and its reality: Transparent territories, opaque power and empowerment. *The Journal of Community Informatics*, *8*(2). http://ci-journal.net/index.php/ciej/article/view/866/909. Accessed 5 Mar 2013.

Rich, S. (2012, July 20). Palo Alto, Calif., to launch open data initiative. *Government Technology*. http://www.govtech.com/policy-management/Palo-Alto-Calif-Open-Data-Initiative.html. Accessed 12 Mar 2013.

Scherer, M. (2012, November 7). Obama wins: How Chicago's data-driven campaign triumphed. *Time Swampland*. http://swampland.time.com/2012/11/07/inside-the-secret-world-of-quants-and-data-crunchers-who-helped-obama-win/print/. Accessed 13 Mar 2013.

Schönberger, V., Cukier, K. (2013, March 6). Big data excerpt: How Mike flowers revolutionized New York's building inspections. *Slate Magazine*. http://www.slate.com/articles/technology/future_tense/2013/03/big_data_excerpt_how_mike_flowers_revolutionized_new_york_s_building_inspections.single.html. Accessed 8 Mar 2013.

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

Slee, T. (2012, June 25). Seeing like a geek. *Crooked Timber*. http://crookedtimber.org/2012/06/25/seeing-like-a-geek/. Accessed 5 Mar 2013.

Stirton, E. R. (2012). The future of institutional research – business intelligence. *eAIR*. https://www.airweb.org/eAIR/specialfeatures/Pages/default.aspx. Accessed 10 Sept 2012.

Thomas, E., & Galambos, N. (2004). What satisfies students? Mining student-opinion data with regression and decision tree analysis. *Research in Higher Education, 45*(3), 251–269.

Two Crows Corporation. (2005). *Introduction to data mining and knowledge discovery* (3rd ed.). Potomac: Two Crows Corporation. http://www.twocrows.com/intro-dm.pdf.

Vialardi, C., Bravo, J., Shafti, L., & Ortigosa, A. (2009). Recommendation in higher education using data mining techniques. In T. Barnes, M. Desmarais, C. Romero, & S. Ventura (Eds.), *Educational data mining 2009: 2nd international conference on educational data mining, proceedings* (pp. 190–199). Cordoba: International Working Group on Educational Data Mining. http://www.educationaldatamining.org/EDM2009/uploads/proceedings/vialardi.pdf.

van Wel, L., & Royakkers, L. (2004). Ethical issues in web data mining. *Ethics and Information Technology, 6*(2), 129–140. https://doi.org/10.1023/B:ETIN.0000047476.05912.3d.

Williams, M. (2010). Can we measure homelessness? A critical evaluation of "Capture-Recapture". *Methodological Innovations Online, 5*(2), 49.1–4959. https://doi.org/10.4256/mio.2010.0018.

Zenk, S. N., Schulz, A. J., Israel, B. A., James, S. A., Bao, S., & Wilson, M. L. (2005). Neighborhood racial composition, neighborhood poverty, and the spatial accessibility of supermarkets in metropolitan Detroit. *American Journal of Public Health, 95*(4), 660–667. https://doi.org/10.2105/AJPH.2004.042150.

Zhang, Y., Oussena, S., Clark, T., & Kim, H. (2010). Use data mining to improve student retention in higher education: A case study. In J. Filippe & J. Cordiero (Eds.), *Proceedings of the 12th international conference on enterprise information systems* (Vol. 1, pp. 190–197). Funchal: SciTePress.

Chapter 3
The Construction of Data

Abstract In this chapter, I show that data is not an objective representation of reality but rather a constructed translation of observations into legible elements designed to support governance (be it by the state or by private actors). Both technical and social structures influence this translation; the technical aspects of database architecture are insufficient by themselves to define this translation regime. Such regimes can contain three characteristic translations: normalizing translations that separate the normal from the deviant, atomizing translations that separate complexity into individual elements, and unifying translations that group diverse characteristics into categories. At the same time, these data systems translate their subjects into "inforgs," representations that consist of bundled information rather than actually existing subjects. These acts of translation, I conclude, are significant exercises in political power.

Whether in business, government, or higher education, pressures toward "data-driven" or "evidence-based" decisions are ubiquitous, promising more insight, more efficiency, and better outcomes than was previously possible. Implicit in this view, however, is a scientifically realist view of data: Data can save us because it is an objective representation of observed reality that can thus transcend politics to bring organizations to the correct decision. But if data is a social construct requiring acts of choice and interpretation in its creation, then it becomes political, its power masked behind its false realism. The structures that shape these choices are thus central to understanding information justice.

This chapter establishes the translation regime as a mechanism by which the social construction of data takes place, and suggests that translation regimes should be viewed as political structures rather than technical ones. Data exists because organizations such as universities or states have a need to make the domains in which they act legible. Doing so, however, requires some process that narrows the many possible representations of a given state of the world to a single data state. This process is carried out within translation regimes: systems of technical rules and social practices that establish a one-to-one correspondence between a given state of the world and a data state. The technical structures of a relational database, such as tables, functions, business rules, and queries, translate states of the world into data states based on standards established by social structures such as cultures, states,

J.A. Johnson, *Toward Information Justice*, Public Administration and Information Technology 33, https://doi.org/10.1007/978-3-319-70894-2_3

and organizations. These regimes operate in a non-neutral fashion, carrying out a set of characteristic translations that favor certain groups over others. As such, information systems design is a political act, among other things shaping representation, asserting and protecting interests, and constructing normalized and deviant identities. Because these political acts are carried out through the technical structure of the translation regime, they appear as technical outcomes, making it more difficult to challenge them.

3.1 Data as Reality Made Legible

The ubiquity of data in contemporary society hides its peculiarity. Data is a very specific form of information, one in which the subject is broken down atomistically, measured precisely (in the sense of being measured to quite specific standards that may or may not involve a high level of quantitative precision), and represented consistently so that it can be compared to and aggregated with other cases. That this form of knowledge is more common in highly structured institutions and rose to ubiquity with the modern, bureaucratic state and the capitalist enterprise should surprise no one. Creating data should be regarded as a social process in which reality is made legible to the authorities of an institutional structure.

Scott (1998) argues that the driving force behind the creation of data is the need to make the subjects governed by an institution legible.

> Certain forms of knowledge and control require a narrowing of vision. The great advantage of such tunnel vision is that it brings into sharp focus certain limited aspects of an otherwise far more complex and unwieldy reality. This very simplification, in turn, makes the phenomenon at the center of the field of vision more legible and hence more susceptible to careful measurement and calculation. Combined with similar observations, an overall, aggregate, synoptic view of a selective reality is achieved, making possible a high degree of schematic knowledge, control, and manipulation. (Scott 1998, p. 11)

Legible knowledge transforms reality into standardized, aggregated, static facts that are capable of consistent documentation and limited to the matters in which there is official interest. Such facts emerge from a process in which common representations are created into which cases are classified and which can then be aggregated to create new facts on which the state will rely in making decisions (1998, pp. 80–81).

The importance of legibility for governance can be seen most clearly when Scott contrasts legible knowledge with local knowledge. The latter, with all of the specific practices, details, and dynamics of reality, is impossible to use for the kind of broad governance characteristic of the modern state; it lacks commonality with other localities and is not objective to outsiders. This obstructs governance in two ways, first by preventing synoptic understanding by authorities and then by denying the governing algorithms of the bureaucracy the standardized inputs they need to produce a standardized output. "A legible, bureaucratic formula which a new official can quickly grasp and administer from the documents in his office" (1998, p. 45) is a necessity for modern governance in both the state and the enterprise.

The need for legibility defines not only the form but also the substantive nature of data. It is common to regard data from a scientific realist perspective in which data is a technical artifact, a representation of information about some subject that is stored such that it can be related to other such representations. This is, for example, the approach used in the United Kingdom's Data Protection Act 1998. The act defines data as a qualified synonym for information: "Data is information which …" followed by a list of technical conditions relating to storage and processing; personal data is defined by the data's relation to an individual identifiable either in the data itself or in relation to other data, and sensitive personal data includes information about a specific list of personal characteristics (Information Commissioner's Office n.d., pp. 19–23).

This is a quite problematic view of data, however, as it suggests that the process of representing reality[1] is an automatic, even algorithmic process. Such a view is naïve, however; like virtually all technologies (Johnson 2006), data is a socio-technical construct in which human agency and social structure is central (Nissenbaum 2010, pp. 4–6) and the path from reality to data is contingent rather than determined (Seaver 2014). Rather than being an automatic process with a one-to-one relationship between reality and data, data states are underdetermined with a one-to-many relationship between reality and data: one state of the world can give rise to many possible data states, some of which are incommensurable with others. In order to make the world legible to human authorities and algorithmic bureaucracies, one data state must be chosen to represent a state of the world from among many possibilities. Reality constrains those possibilities but it does not, by itself, fully reduce the state of the world to a single data state.

Netflix provides an exceptionally valuable case, as it explicitly attempts to datize a cultural product and thus makes the socio-technical nature of the company's recommendation engine clear. Netflix's process is at its heart a form of structuralism, disassembling films into their smallest component parts and reassembling the "altgenres" (hyperspecific categories that Netflix users see organizing films that Netflix recommends, such as "visually appealing intellectual action-thrillers") that describe the common structures across films. Structuralism reveals the contingency of Netflix's purportedly objective recommendations:

> When you break an object down into its parts and put it back together again, you have not simply copied it—you've made something new. A movie's set of microtags, no matter how fine-grained, is not the same thing as the movie. It is, as Barthes writes, a "directed, *interested* simulacrum" of the movie, a re-creation made with particular goals in mind. If you had different goals—different ideas about what the significant parts of movies were, different imagined use-cases—you might decompose differently. There is more than one way to tear apart content. …Netflix's altgenres are in no way the final statement on the movies.

[1] For the purpose of this paper, I take "reality" to mean the physically existing world as interpreted by actors within it. Here I follow Charles Sanders Peirce in his seminal essay "The Fixation of Belief" in which there is an underlying reality that cannot itself be perceived but that can be asymptotically approached through repeated observation (Peirce 1992). This leaves open an interpretive space but does not deprive the concept of reality of meaning, allowing it to be bracketed as a distinct but related problem from that addressed in this paper.

> They are, rather, one statement among many—a cultural production in their own right, influenced by local assumptions about meaning, relevance, and taste. "Reverse engineering" seems a poor name for this creative practice, because it implies a singular right answer—a fact of the matter that merely needs to be retrieved from the insides of the movies. We might instead, more accurately, call this work "interpretation." (Seaver 2014)

Broadening Seaver's analysis to data generally, each possible data state can be regarded as a potential interpretation of the underlying state of the world to be datized. The contingency of the final forms of data requires some external source of stability in order for data to bring legibility to the world (Mitev 2005). What is needed is a process of translation from reality to data that constructs a single representation by serving as the external source of stability for representation. Such a process is inherently endogenous to the creation of the data as long as multiple interpretations are possible. In a realist view of data, all but one of these states must be regarded as errors or biases in the data, which can be corrected by validating the data against itself or the reality it purports to represent until a single data state that is fully consistent with reality remains. But the self-correcting process of scientific realism cannot do this; rules for interpretation are required in order to have data in the first place. All possible final data states will appear consistent with reality because they follow the rules of the specific interpretive process that leads to them. These processes have legitimized the data states resulting from them as the only acceptable representation of reality, all else—local knowledge in particular—being dismissed as anecdotal evidence.

Classifying individuals within a system of gender relations is a good paradigmatic case that demonstrates the operation of a constructivist understanding of data beyond the domain of cultural production. Simply within the binary gender system common in western cultures, people might be represented within a data system either by sex or by gender. These categories are not reducible to each other; the existence of transgendered and intersexed people is sufficient to make sex and gender incommensurable within such binary systems. Moreover, there is no inherent reason that a data system needs to be limited to a gender binary, even in predominantly Western contexts: Facebook recently introduction of more than 50 custom gender descriptions from which its members can choose (Facebook n.d.). The intellectual construct "gender" is thus insufficient to determine how data systems will represent a specific person; the reduction of gender realities to specific categories cannot be an objective, value-free, observational process. In spite of this, most data systems rely on the same binary coding frame, one in which gender is taken to have a one-to-one correspondence with biological sex. The representation of individuals' place in the system of gender relations is thus determined by neither reality nor by the technical requirements of the data system. It is a choice on the part of developers to reduce an exceptionally complex reality to a specific legible form.

3.2 The Translation Regime

In order for the process of selecting a data state from among the many possible ones to be, in fact, legible, the process must be a rule-governed one. Creating data from reality is not simply an interpretation but a translation (or, more precisely, a series of translations) in which substantive content embedded in a set of technical rules determines how reality will be represented in the data system. For a relational database,[2] those rules are largely, but not entirely, contained within the data system itself, expressed as technical specifications within the database. The construction of data in relational databases consists mainly in the design and selection of rules such that they implement the demands of the content sources and only secondarily, when the rules and content sources are insufficiently precise, in the direct interpretation of reality by those entering data into the data system. Collectively, one might refer to these structures as the translation regime for a data system.

3.2.1 The Technical Structure of the Translation Regime

The most basic technical element of the relational database translation regime is the structure of individual data tables. The fields selected for inclusion in an individual table do much more than selecting which aspects of reality will be stored (though they most certainly do that as well). Those fields break down that reality into component parts. This is, of course, only a selection of the parts of the reality, and recombining these parts creates only an interpretation of reality rather than an objective and complete representation of it. A simple case is found in the fields representing the name of a person who is represented in a table row. The STUDENT table uses *FIRST_NAME*, *MIDDLE_NAME*, and *LAST_NAME*. These fields cannot be recombined to generate the formal name of everyone the data purports to represent. Truncating *MIDDLE_NAME* to a middle initial translates a student who goes by "G. Gordon Liddy" into "George G. Liddy"; a student from a country where family names precede given names named "Mao Zedong" is translated into "Zedong Mao."

A more complex example is seen in the information kept on a student's academic program. This data is kept in a hierarchy of fields within STUDENT: *DEGREE*,

---

[2] In a more general theory of data, the choice of database type would itself be understood as part of the translation regime. Raman (2012) shows that the choice of a relational database rather than one based on Unstructured Information Management Architecture (UIMA) to maintain land claims in India prevented the storage of knowledge held in narrative form, as was common among *Dalits* in the region. Narrative knowledge would have to be translated into an atomic structure in order to be stored in a relational database; in this case, such knowledge was simply excluded in favor of that contained in state-produced documents that could be stored in a relational system. Since all of the data currently used by UVU is contained in relational databases, the influence of database type must be investigated in another context.

COLLEGE, *DEPARTMENT*, *MAJOR*, and *CONCENTRATION*. This hierarchy standardizes the grouping of students by program in ways that may or may not reflect the actual operation of the program. The Behavioral Science major includes concentrations such as Psychology and Family Studies with such significant overlap in coursework, administration, and faculty that distinguishing students by concentration introduces an institutional separation of students to the data that is absent in reality. The major's Anthropology concentration, however, has much less overlap with the other concentrations. There is, in addition, a separate Social Work major that has stronger connections with the Psychology and Family Studies concentrations than the Anthropology concentration. The data fields, however, translate these varied conditions into a single, hierarchical set of student groups.

Moreover, each field in a table includes a definition restricting the type of data that can be entered into the field. These definitions define at the least the type of characters that can be put into the field and the maximum length of the field content; often field definitions might also include number formats, specialized formats such as times, or more precise tests of valid data. As such, they define what form the resulting data must take and proscribe the use of other forms. For a data element that represents a well-defined condition this is straightforward. But for conditions with more variability it is not at all so. Field definitions may thus permit or prohibit the entry of data that is valid in relation to reality but not in relation to the field definition. *TAX_ID*, for example, is defined as a variable character text field (to preserve the leading zero in some Social Security Numbers) of up to 63 characters. A more strictly defined field (for example, a fixed-width text field limited to nine characters) would prevent the entry of Federal Taxpayer Identification Numbers, which some students may have instead of a Social Security Number. The more flexible field definition of *TAX_ID* thus supports the translation of a wider range of conditions.

Commonly, some fields within a table will be indexed. Indexing a field stores information about the content of the field separately from the table itself, allowing the field to be searched rapidly. Typically, a table would index fields on which records would be selected, and then other data in those records could be returned promptly. In a small table, the difference in response time and server load between an indexed and non-indexed field may be minimal, but in a very large system might be the difference between practical and impossible searches. Indexing thus creates privileged translations of data, in extreme cases making fields that are conceptually equivalent incommensurable where one is indexed and the other is not: The indexed field is, effectively, the only field that can be used to represent the data in practice, and thus the only translation available for use. In COURSE, descriptive course information fields such as *SUBJECT*, *COURSE_NUMBER*, and *SECTION* are not indexed, but *COURSE_REFERENCE_NUMBER* is. This makes it quite practical to refer to courses by reference number and to identify descriptive course information given a reference number, but somewhat more difficult to do starting with the descriptive information, especially in the absence of other limits on the data needed.

Beyond the structure of individual tables, one might also look to the structures of a database that validate data across tables. Validation tables function in ways similar to field definitions. A validation table contains a list of values that are acceptable for use in a field, used commonly in fields that contain categorical data with a limited number of possible values. The validation table for *COUNTY* contains a list of all counties in Utah, along with three residual values for all other cases. This prevents the entry of invalid county names. In the process, however, the validation table also determines the conceptual framework for the field itself. In this case, *COUNTY* becomes a characteristic held only by people from within that state. This is even clearer in the example of gender. The validation table for *GENDER* includes only the values "Male," "Female," and "Unspecified," imposing a binary gender schema on the people represented in the field. The "Unspecified" value as a residual is an especially strong reinforcement of the gender binary in this common validation frame: if one is not either male or female (whether because the translation regime insists on correspondence between sex and gender thus denying the existence of transgender identities, or because the person identifies as some form of non-binary identity), one is not even a residual "Other." One is presumed to, in reality, identify with one of the binary values and simply did not communicate that identification to those collecting data. These examples make clear the special importance of residual representation, often an afterthought, in validation tables' role in the translation regime.

A more complex validation structure is a business rule. Business rules place conditional requirements or constraints on the data in one or more fields based on the content of other fields, within or across tables. A common use might be to either require or proscribe certain external actions, for example, preventing a contract from being issued before a credit check has been performed by requiring that a row exists for a customer's credit check in a table of credit check data before a row can be created for that customer's contract in a table of contracts. Business rules can also be used to validate data across fields, preventing the entry of a state other than Utah in *STATE_ADMIT* (the state in which the student resided at admission) and a Utah county in *COUNTY_ADMIT*. In much the same way as validation tables, business rules impose a conceptual framework on the fields that they govern by limiting the data that can be entered to data that is consistent with the underlying concept. The central concept underlying a hierarchy of state and county of admission is the authority of a unitary state over its citizens at the local level. A business rule upholding that hierarchy would thus reinforce the structure of authority within state government in the United States. UVU's lack of such a rule has led to an exceptional amount of inconsistent data and thus inhibits the translation of a geographic location such as a street address into a political one such as a legislative district.

The relationships among data in different tables further shape the translation regime. In a relational database, data tables are structured so that tables can be joined to each other on common elements to allow cases in one table to be matched to cases in another. In the absence of appropriate commons field on which to join, however, data in different tables cannot be related to each other. The UVU reporting tables are designed expressly to facilitate this: COURSE and STUDENT_COURSE can be joined on the combination of *COURSE_REFERENCE_NUMBER* and *TERM*; STUDENT and STUDENT_COURSE can be joined on *STUDENT_ID* and *TERM*. Joining STUDENT to COURSE requires joining all three tables. As a result, the translation of a particular characteristic of reality into an individual data field is also a translation of it into a context created by an extensive set of other data fields. A student is not simply a Botany major; joining STUDENT and STUDENT_COURSE makes the student a female Botany major who has not taken a course in the major in three semesters. This translation is much more interesting to those responsible for increasing retention of women in STEM degree programs.

All of the structures discussed above involve primary translation: the translation of a state of the world into data. But translation regimes include as well secondary translation processes, translating not reality into data but rather existing data into new data. Functions are a common structure that performs secondary translation. Fields can be defined with functions. Functions calculate a value for a field based on the content of other fields rather than being populated through direct entry of data. The function that calculates STUDENT.*INSTITUTIONAL_GPA* combines *CREDITS_ATTEMPTED* and *CREDITS_EARNED* across all rows in STUDENT_COURSE for a student to create a representation of that student's academic performance that does not exist in the absence of the function: an aggregate performance indicator. Functions thus widen dramatically the range of data contained in the data system and produced by the translation regime, illustrating the extent to which translation is not solely about creating equivalent representations of existing data but also about creating new data through the combination of existing data.

The data stored in a database is not necessarily the data that will be used in the final representation of reality. Data from relational databases is extracted through queries that specify precisely what data will be extracted, how it will be combined in new fields, and how it will be aggregated. A query will, at the least, specify which fields to retrieve for a record, and will usually specify which records to retrieve as well. The query thus selects, for example, whether students' academic performance will be represented by *INSTITUTIONAL_GPA* or *OVERALL_GPA*. But queries can also use the same set of functions that are used to define fields. A particularly common translation used at UVU when extracting demographic data for survey samples[3] is the creation of a binary ethnicity field. Given that minorities make up less than 20% of the UVU student body (Institutional Research & Information 2012a, p. 18), surveys are rarely large enough to provide reliable data when broken down by individual ethnic categories. Survey sample queries thus categorize student ethnicity using a function that parses *PRIMARY_ETHNICITY* and *RACE_COUNT* to identify students as either White or minority. This function is written into the standard queries that are used to generate samples, and often included in ad hoc queries for particular projects.

Queries are the final point in a relational database where translations take place. That does not mean, however, that the technical structures of the translation regime are limited to those processes that take place between data entry and data extraction. Applications, whether software systems or analytical processes that connect to a

[3] UVU commonly refers to the group invited to participate in the survey as a "sample" even when the invited group is in fact a census of a sub-population of students such as graduates in a term. I use "sample" in this sense here as well.

data system, can further translate the data extracted. UVU's "Stoplight" risk warning system translates 20 possible trigger conditions that the institution identified as characteristic of students at risk of failing courses into a color-coded risk rating that is shown to advisors and on class rosters (UVU Student Retention 2013). Stoplight operates as a new application built within the ODS, with a custom table carrying out this secondary translation and feeding data from it to advisors and instructors. UVU also maintains a website presenting data on mission fulfillment that is built, in part, on data that is extracted from the data system then aggregated and represented graphically using business intelligence software, translating individual data points into aggregated visual data. These applications are the point at which data finally meets a human who must act on the data, and thus mark the boundary of the technical structure of the translation regime.

3.2.2 The Social Sources of the Translation Regime

While the substantive content of the translations is inscribed in, and to an extent constrained by, the technical structure of the database, the bulk of the substance comes from sources external to the database itself. Culture, the state, the institution, and private sector actors all provide content for the translations that is then built into the technical structures.

As much as the language of conscious design and engineering permeates both the theory and practice of information systems, their conformity with their origin communities' cultural structures suggests that sociological institutions are at least as important. Like many organizational forms, data systems include "not just formal rules, procedures, or norms, but the symbol systems, cognitive scripts, and moral templates that provide the 'frames of meaning' guiding human action" as a mutually-constituting element of social action (Hall and Taylor 1996): Data systems are both composed of and instantiate cultural institutions. Some of these are relatively straightforward, such as *GENDER* including only "Male" and "Female" as valid values. A binary frame of meaning shared by most people in the institution defines which values are built into the validation table, for the most part without a conscious decision to do so.

A more interesting example is *STUDENT_CLASSIFICATION*. This field, in a technical sense, divides undergraduates into four classes based on the number of credits completed. Ostensibly, this indicates progress toward degree. But the values in *STUDENT_CLASSIFICATION* are more than categorizations or translations of a number. "Freshman," "Sophomore," "Junior," and "Senior" is a cultural script for understanding the social relations of a traditional residential institution. At an institution where many students are part-time, married, or returning adults and there is no campus housing, neither a 4-year academic career nor a distinct cohort are relevant concepts to most students.

STUDENT_CLASSIFICATION thus operates not as a reflection of student behavior or program structure but rather as a script for (mis)understanding and relating students to each other and the institution. “Freshman” is a frame of meaning that assigns attributes to a student; the First Year Experience program acts toward such students as that frame says is appropriate, stressing that “College is different from high school,” that independence is exciting but can be overwhelming, and that participation is a good way to make friends: messages appropriate to a traditional freshman on a residential campus but not to a recently retired Marine pursuing her bachelor’s degree as a start to a second career. Similarly, the National Survey of Student Engagement samples all first-year students on the basis of institutional classification and asks a series of questions about differences between high school and college engagement experiences even though many freshmen at UVU are closer their children’s graduation from high school than their own.

Political influences operate in a much more clearly conscious fashion, usually being deliberately designed into the data structures. The state shapes translations primarily by establishing formal data standards. Data standards define substantively and sometimes technically the content of a data field or record. UVU’s main data standards, as is true for most public higher education institutions, are found in two sources: the federal IPEDS Glossary (National Center for Education Statistics n.d.) and a series of data dictionary files from USHE (Utah System of Higher Education 2013a, b, c, d). The USHE standard for *CITIZENSHIP_STATUS* in the student table, for example, translates the many categories of rights to presence in the United States under US law to five categories: US citizen, US national, resident alien (which includes all documented immigrants entitled to stay indefinitely in the United States), non-resident alien, and “non-immigrant undocumented students” (Utah System of Higher Education 2013c). This last category is particularly interesting, as it marks a quite significant departure from the typical discourse of undocumented *immigrants*, translating a person’s intentions as well as their position within the immigration regime. The USHE data standard for *GENDER* became a stricter one with the inactivation of the “Unspecified” value in the USHE standards in 2012 (Utah System of Higher Education 2013c, p. S-13). This prohibited missing data in *GENDER*. As a result, gender nonconformity is no longer even translated as missing data; all students are translated into one of the binary gender categories. With such cases being typical examples of how data standards translate reality, it is clear that they should be viewed as substantive translations, not simply as technical coding procedures.

The translations created by data standards can be quite complex, especially when multiple data standards can apply to the same set of data. STUDENT supports three distinct data standards for ethnicity data to support competing, and in some cases conflicting, data standards. The USHE standard for *ETHNICITY* defines an eight-character field in which each character position represents an ethnicity with which a student might identify, with multiple identifications allowed, chosen among Hispanic or Latino, Asian, Black or African-American, American Indian or Alaska Native, Native Hawaiian or Pacific Islander, White, Non Resident [sic] Alien, or Unspecified (Utah System of Higher Education 2013c, p. S-14). IPEDS currently used the Office of Management and Budget standards, in which students select all groups with which they identify among American Indians or Alaska Natives, Asians, Blacks or African-American, Native Hawaiians or Other Pacific Islanders, or Whites and then identify whether or not they are Hispanic or Latino (National Center for Education Statistics n.d., p. R). UVU also supports older IPEDS standards that define students by a single ethnicity.

To do this, STUDENT includes one binary field for each possible ethnicity that it might report, a count of the total ethnicities selected by the student, and a primary ethnicity to be used with standards that do not support multiple ethnic identities. Reality having been translated into these data fields, a further translation of the data fields into the reporting identities takes place in querying and extracting data for reporting. This creates an extensive complex of translations that are not entirely consistent. The same student may be "White" in *PRIMARY_ETHNICITY,* multiracial in *IPEDS_ETHNICITY,* "Minority" in a query dividing students into "White" and "Minority," and "Non Resident Alien" in *USHE_ETHNICITY*. While inconsistent, none of these is fundamentally incorrect either as a translation of the ethnicity fields in STUDENT nor as translations of reality.

Political systems have more subtle means at their disposal to influence the translation regime as well. Especially for public institutions but, given the public mission of higher education generally, to some extent for all higher education institutions there exists a principal-agent relationship between the polity and those institutions similar to that between legislatures and bureaucracies. That relationship subjects the translation regime to many of the same oversight pressures as any regulatory regime. One of the most common responses to such pressure is bureaucratic anticipation: agencies, seeing signals from legislators about their desired outcomes, anticipate direction from the legislature and move to secure those outcomes without waiting for that direction to be made explicit (which, in many cases, never happens because the need for direction has been met) (Weingast and Moran 1983). This is not simply having the foresight to see a new formal requirement coming and implement it in advance; it is an act of anticipating the demands of political actors and meeting them as a means of satisfying those actors whether or not the demands are formalized.

Anticipation was a key factor in designing one of UVU's signature data applications, its Student Success and Retention dashboard (Institutional Research & Information 2012b). The dashboard was designed to assess efforts to improve the first-year retention rates and graduation rates reported to IPEDS. The appropriate federal data standard is thus the rates for the IPEDS cohorts: first-time, full-time, bachelor's degree-seeking students entering in the fall term. At the time this was being developed, however, the US Department of Education had begun public discussion of revised data standards to take effect in the 2014–2015 data collections, and constituencies and their legislators in both Utah and the federal government had raised significant concerns about whether higher education was meeting the needs of non-traditional students. Those involved in designing the dashboard recognized that significant political pressure was building to demand student success data for part-time and transfer students. The dashboard as completed in 2010, well before NCES made decisions on the new standards, was thus based on a fall new student cohort with both full-time and part-time students, and designed in a way that would facilitate the addition of transfer students by creating a transfer student cohort. UVU was able to provide part-time data to the institutional administration, the community, and political actors well before it faced a formal requirement to do so. Neither standard was implemented by NCES until 2013, taking effect with the 2014–2015 IPEDS data collection, well after UVU had begun tracking the success of part-time students.

The private sector, both for-profit and non-profit, is an important source of content as well. Because UVU's data system is a customized version of a widely used commercial higher education data system, much of the translation regime's content comes from Ellucian, the makers of Banner. When UVU adopted Banner in 2005 and implemented the ODS in 2009, the institution started from a standard Banner database schema and then customized it to meet specific needs on the UVU campus (such as USHE reporting). This requires a notional higher education institution whose needs are representative of most institutions around which the out-of-the-box version of Banner can be designed; elements of the schema that were left unmodified thus reflect Ellucian's conception of what the content of fields should be based on that notional institution. UVU's class rosters, produced by an application within the Ellucian Luminis web services platform connecting to Banner data, provide students' formal names even with a preferred name field available. The University of California, Davis, has in fact implemented an option within Banner that allows students to use preferred names rather than formal names on many university documents including class rosters (Easley 2014), demonstrating that the standard form for class rosters in Banner is not a technical or legal constraint but an assumption on the part of the software designers.

The non-profit sector's contributions to the content domain of higher education translation regimes should not be discounted. Institutions, in a bid to increase transparency (or at least the appearance thereof), are frequently participants in voluntary data sharing processes, each of which comes with their own data standards that may or may not be coordinated with others. The Voluntary System of Accountability (VSA) is one of the largest among public institutions, providing both input and outcome information with the aim of demonstrating the value of an institution's programs. UVU also participates in the Consortium for Student Retention Data Exchange, a program that facilitates peer benchmarking of multi-year retention and graduation rates. In both cases, these organizations' data standards exist alongside government standards. The tables supporting UVU's retention and graduation rate application, described above, translate student enrollment data into both IPEDS and CSRDE standards.

Despite these many external pressures, institutions themselves are important influences on the content of the translation regime. Data standards do not always offer precise operational definitions and logics to determine the data value; they often couple conceptual definitions with a set of valid end states, leaving institutions considerable leeway in the translation process itself. Institutions nearly always control the technical implementation of data standards. Under different alternatives, a particular state of the world can be translated into different values within a data standard depending on how the translation is performed. UVU can thus choose whether to use the state from an applicant's current or permanent addresses when selecting the value of *STATE_ORIGIN*, which USHE simply defines as "The state code indicating the student's state of origin as described at the time of their first application to the institution, if one is available" (Utah System of Higher Education 2013c, p. S-11). That decision is embedded in functions and validation procedures, but the data standards do nothing to specify how the function should evaluate a student living in Twentynine Palms, California, who considers her permanent home to be Moab, Utah, so long as the function returns a valid value. That a function for doing so is provided in the base Banner package does not prevent the institution from changing that. The decision of how the function evaluates the primary data fields to create a secondary translation is a design choice rather than a predetermined outcome, one made and implemented at least in part by the institution.

The institution is also the data collection point, giving it the power to choose both what data to collect and what interactions to translate into data. This is a powerful tool in shaping data: interactions and characteristics that are not turned in to data are not simply missing; they are untranslatable and hence illegible. This prevents them from being considered in decisions. The standard Banner package includes a field for students' religious preferences. UVU does not collect that data from its students, however. Ostensibly, this is because of a concern that asking students to identify religious preferences would create the impression that UVU was supporting the dominant religion of its community, The Church of Jesus Christ of Latter Day Saints. This has not prevented UVU's Institutional Research & Information office from including that question on its student opinion surveys, the most recent of which that asked the religion question found that 77.2% of students identify with some form of the LDS faith (Institutional Research & Information 2013, p. 45). That data is not included in Banner, however; more than 75% of students' data records have a null value for *RELIGION*. As a result, the institution does not routinely consider religion in its decision-making, even though such a large number of students sharing a common worldview present many of the classic problems of in-group/out-group dynamics.

Religion is, to UVU, illegible. This is not at all to say that the institution is hostile to either LDS Church members or non-members; its President, Matthew S. Holland, is the son of one of the highest authorities in the LDS Church and an active church leader in his own right and yet has consistently promoted religious inclusivity toward those outside the LDS Church as an important element of UVU's Core Themes. But the decision not to collect data with which to populate *RELIGION* does leave religious preferences opaque to the institution. The institution cannot ask questions about the role of religion, either as a belief system or as a social institution, in the operation of educational programs. It cannot consider whether students who are not LDS Church members have lower retention rates, a possible sign that they feel excluded from the social life of the campus. It cannot consider whether LDS Church members are less likely to complete the FAFSA and thus to receive Pell Grants, a possible consequence of a strong ethos of self-sufficiency and financial conservatism within LDS theology and culture. UVU is quite effective addressing these questions within the limits of survey research methods, but a full canvas of students over time is impossible. This leaves UVU unable to "read" a characteristic that is central to many students' identities.

3.3 Characteristic Translations Within the Regime

The data translation regime is not substantively neutral; it favors certain types of outcomes over others. In a relational database such as that used at UVU, one can identify at least three characteristic types of translations in the data (as well as, of course, numerous translations that are relatively unique and not analyzed here). These characteristics translations describe how the ontological character and meaning of states of the world commonly change over the course of the translation process. The result is that translations are most often analytically incommensurable with the reality they purport to express: the words attached to the conditions may be similar, but they are embedded in an entirely new structure.

3.3.1 Normalizing Translations

One type of translation establishes certain states of the world as part of the realm of normally existing conditions, thus implicitly establishing all other states as deviations from normalcy in some sense. Such translations typically have the effect of reducing the states of the world to only those within the realm of the normal data states. Those represented in the database are thus represented only to the extent that they are capable of being represented within that normal realm; to the extent that they deviate from the normal world as it exists within the database they cease to exist analytically.

The simplest normalizing translation is from relevance to existence. Data is collected based on what the collectors find relevant to their interests: it may shed light on a question they need answered or a decision they may make, or it may be needed to comply with requirements of an external authority. Data is not collected, however, on matters that are not of interest to the institution, nor on matters for which the existence of data is counter to the institution's interests. One common objection to data collection and analysis within IRI was that UVU could be forced to make the data or subsequent analyses of it public under Utah's open records laws. Most frequently this objection was used with projects that might collect data that subjects might consider sensitive but that was not protected by privacy laws, a not unreasonable protection but nonetheless one that is driven by a specific interest on the part of the university. Those characteristics or states of the world were considered irrelevant to decisions, and thus not collected.

But when questions arise about such characteristics, irrelevance turns into nonexistence. The characteristics about which there is no data frequently function not as unknowns which need to be estimated or otherwise accounted for in analysis, but are rather ignored, treated analytically as if they do not exist or, at best, subsumed into platitudes about "context" that fade into the background when the data is available. This is more than just saying that nonexistent data does not exist: it is not data about a given characteristic that is translated into nonexistence but the characteristic itself. Having determined that religion is irrelevant to decision-making and not collected information about it, UVU's students cease, analytically, to have religious preferences.

A similar process takes place with regard to the conditions that a characteristic might take on. Translation regimes transform the diversity of possible conditions of a characteristic into a set of acceptable data values. Those conditions that cannot be represented by a valid data state become represented not as themselves but as deviance: the data is missing; it is given a residual category value such as "other," "not applicable," or "not available"; it is forced into one of the valid data states even if that does not actually represent the state of the world. So the diversity of gender identifications are translated into categories of normalcy that are represented by the values "Male" and "Female," and invalid data that exists in a state of deviance from normalcy, first as "Unknown" and then, with the deprecation of that value in the USHE data standards, into a forced choice of a valid but untrue data state. Transgender identities are not simply statistically rare; they are abnormal. And as in the case of irrelevant characteristics of the world, deviant conditions of the world become analytically nonexistent, assumed to be trivial exceptions to a meaningful interpretation of reality.

It is important especially to understand what it means to say that states of the world *analytically* cease to exist. The qualifier is an important limitation. No one at UVU would deny that many students are religious; the lack of data does not preclude thinking about the characteristic. In a culture where decisions are legitimated in part based on the ability to support them with data analysis, a characteristic that is not datized cannot be analyzed, and so decisions about it cannot be legitimated and are unlikely to be built into policy. Nor can assumptions about the characteristic be questioned. This is perhaps the most pernicious aspect of the translation of relevance. While a characteristic for which data is unavailable may not exist analytically, it may be very prominent culturally, in many cases functioning as part of an ideal type representation and assumed to be true of all cases. The culture of the region carries with it a strong religious identity. The result is the assumption that any one student is a member of the LDS Church until they are known to be otherwise.

It is also important to recognize that analytical normalcy is different from social normalcy, by which I mean the existence of certain conditions as the normal or typical condition from which other conditions vary. Self/other distinctions are a form of social normalcy: whites or men represent the normal or typical, while people of color or women are an "other" defined in relation to the norm. The analytical normalcy that I posit here includes both the typical and other categories in normalcy; deviance constitutes existence outside of the set of recognized categories rather that existence within one of the atypical categories. Analytical normalcy does not imply social normalcy. "White" is, analytically, merely one category of *PRIMARY_ETHNICITY*, not different from other values within the translation regime despite being the only socially normal value. Nor, however, does analytical normalcy challenge social normalcy: the equation of "Male" with normal takes place outside of the translation regime, so that when the translation regime categorizes someone as male or female it does nothing to prevent the substitution of typical and atypical.

The translation of irrelevance to nonexistence played a significant role in the creation of UVU's "15 to Finish" program (First Year Experience and Student Retention 2014), which encouraged students to take 15 credits per semester in order to graduate in 4 years. The assumption behind the program is that students who attend full-time are not only more likely to graduate on time; they are more likely to graduate at all. One of the core messages is of the program that it is better to reduce or eliminate outside work in order to attend full-time, even if that requires students to take out loans, because they will be more likely to finish, finish faster (especially within the limits of Pell Grant eligibility), and spend more years earning an income commensurate with their completed degrees. The analysis performed in support of the program did indeed show that this was the case. But it did not consider whether this was practical for all students. UVU does not collect effective data regarding the family status or family income of its students; the only systematic effort at data collection regarding the number of children students have or parents' income that is integrated into Banner is the FAFSA, but institutional privacy protections limit the transfer of FAFSA data outside of the financial aid office and low rates of FAFSA completion make such data unrepresentative in any case. Students with high family incomes might find it much easier to attend school without working, while those with families might find it especially difficult to reduce or eliminate outside work. Yet neither group exists at an analytical level. The program does have a strong ethos that 15 credit hours may not be appropriate for all students because of their family status or availability of parental support. But that is not implemented formally in the way that, for instance, the various triggers are built into the Stoplight program. "15 to Finish" is for all students, "with exceptions, of course." These characteristics are irrelevant to the institution's data collection efforts, become illegible because they are not collected, and ultimately cease to exist as part of the "normal" world in which administrators operate.

3.3.2 Atomizing Translations

One of the generally accepted best practices of relational database design is that data fields should be atomic, representing one and only one value for one and only one characteristic. To the extent that this is practiced (and it usually is), the result is that translation regimes will represent the world in atomistic terms, fragmenting characteristics that are defined as much by their relationship to other characteristics as by their specific conditions into distinct fields that are not connected to each other. These fields are then analyzed in isolation from each other rather than in the contexts that make them meaningful to the people represented in the database.

Individual identity is highly susceptible to atomization. Complex, intersectional identities frequently bring together different categories of identity into a coherent whole that does not exist within a database: "Jewyoricans" are fragmented into atomistic categories of religion, residence, and ethnicity without the relationships among them that are central to the identity of Jewish New York residents of Puerto

Rican descent. These categories reflect both the principle of atomicity—separate fields for separate characteristics—and the data standards to which the institution must conform. The USHE reporting standards for STUDENT maintain separate fields for ethnicity and state of origin (and, of course, do not collect information about religion) (Utah System of Higher Education 2013c). This makes representing complex identities that reflect not just one or another aspect of one's identity but the intersection of or relationships among multiple aspects of one's identity quite rare; data is often analyzed along ethnicity or gender, often sequentially but rarely both at once. There are people represented in UVU's data who are Black, and people who are female; there are some who are both. But there are no Black women in the data.

Atomizing translations can be especially complex when trying to translate a narrative into data. In such cases, it is necessary not only to separate characteristics but also to reduce complex conditions into nominal data states that conform to the validation rules and data standards. One might consider the case of students who transfer in large numbers of credits that reflect their prior educational experiences that are at best tangentially related to their current educational ambitions but don't meet the requirements of their current degree program. All of these characteristics are included in STUDENT: *PREVIOUS_EDUCATION* captures whether the student was enrolled at another university in the past, *TRANSFER_CREDITS* reflects the number of hours brought in, *TOTAL_CREDITS* identifies the number of credits earned at all institutions, and *STUDENT_CLASSIFICATION* performs a secondary translation that characterizes overall progress toward the degree. But the fields don't reflect the narrative of a student entering, leaving, and returning with different educational goals and having far more credits than are actually needed to graduate while not being anywhere near completing the current program. A student may be classified as a "Senior" but have perhaps 2 or 3 years of additional coursework to complete in order to graduate. The narrative that provides meaning to the value in *TOTAL_CREDITS* is lost; it is reduced to a name: 142.

In most cases, atomizing translations are driven by the content domain rather than the technical domain; the latter merely implements atomicities that are already practiced in other contexts. Technical limitations do not force atomicity on those using extracted data. The different characteristics can be quite easily brought together through simple concatenations of fields or crosstabulations of extracted data. The IPEDS data standards in fact do exactly this. Institutions are required to report enrollment by ethnicity separately for men and women, allowing the federal government to see the intersectionality of the two conditions. UVU's Student Success and Retention Dashboard allows analysis by two characteristics at once, making it possible to see the effects of a wide range of two-dimensional intersectionalities, and with some rather awkward technical gymnastics a very narrow set of three-dimensional ones, on graduation and retention rates. The multi-character *ETHNICITY* field in the USHE data standards shows how a secondary translation of atomic data can create a field that captures the complexity of multiracial identities.

Narratives, too, can be stored in data systems. Banner includes a data table in which comments can be stored. These can provide the narratives that are stripped away by atomizing translations—if users actually use them. Extracting data from comments is notoriously difficult, requiring complex expressions, tedious analysis, and careful interpretation to make them legible to the institution. Suggesting that the best source for a particular data point is found in COMMENT is universally reviled with UVU's IRI office, but it can be, and sometimes is, done. But this is rarely the case, and even when it is the narrative structure of the comment is rarely used fully, the preference being to identify a nominal value that may be extracted from a comment rather than a more structured field. To be sure, many of these cases involve a re-translation of atomic data. But they do show that other possibilities exist and thus make clear the nature of translation as a choice rather than an inherent technical limitation.

3.3.3 Unifying Translations

In spite of the imperative toward atomicity, there is a counter-tendency toward unity in translation regimes as well. The detailed and diverse conditions of reality frequently exceed the capability of data systems to store them or the ability of analysts to manage them. A characteristic that can have thousands of potential values, especially when those values are expressed in a nominal level of measurement, does little to bring legibility to the state of the world. Diverse states of the world must often be translated into a small number of values that bring many different conditions together into a common data state.

One translation process that unifies disparate conditions is grouping a large number of possible conditions into a small number of data values. This creates a unified group that may not, in fact, exist in reality or that is at least far more complex than is expressed in a single value label. The USHE ethnicity categories are an example. The standard defines "Asian" as "A person having origins in any of the original peoples of the Far East, Southeast Asian, or the Indian subcontinent including for example Cambodia, China, India, Japan, Korea, Malaysia, Pakistan, the Philippine Islands, Thailand, and Vietnam" (Utah System of Higher Education 2013c, p. S-14). The common label "Asian" hides a wide range of diversity within the definition; UVU had Asian students admitted from 27 countries other than the United States among its Fall 2013 students. It seems reasonable to expect that they would have considerable differences among them, and in many cases might find more in common with other racial groups. The borders Pakistan shares with Iran, Afghanistan, and Tajikistan defines Pakistanis as Asians and thus unifies them with students from Japan or Indonesia while separating them from citizens of the surrounding countries who are defined as "White," a category that includes those "having origins in any of the original peoples of Europe, the Middle East, or North Africa." Similar differences in racial identity exist between African and Black immigrants from Africa or the Caribbean and among immigrant groups themselves (Benson 2006), who are nonetheless unified into a single category of "Black or African-American."

Data systems may also unify characteristics temporally. Characteristics that vary over time become essential and fixed in data systems, stripping away the contingency that is often at work in them. Again, the USHE ethnicity standards are instructive here. The USHE standards make reference consistently to the origins of the student, suggesting that ethnic identity is a fixed part of a person's overall identity. As a result, it can be stored in the data systems and reported consistently over the course of a student's academic career. But there is considerable evidence that ethnic identity is not essential; rather it is a characteristic that is situated in particular circumstances and can change with them, such as when the student moves from a public to a private space or into and out of spaces dominated by heritage identities (Zhang and Noels 2013). One might expect this to be especially strong among students who identify with multiple ethnic or racial groups. This situational variability is not captured by the data system, however; the permanence of the data state implies a permanence to the state of the world it purports to represent that may be accurate on average but may not be so at any given moment.

3.4 Translating the Subjects of Data Systems

The technical structures of a relational database, such as tables, functions, business rules, and queries, translate states of the world into data states based on standards established by social structures such as cultures, states, and organizations. These regimes also translate the entities about which data is collected into "inforgs," entities that exist solely as bundles of information. Within many of the structures that guide data use and data-driven decision-making inforgs behave quite differently than people, fundamentally changing the power dynamics of representation in decision process. I explore two structures related to representation in this section. First, inforgs significantly complicate the way that data-driven decision processes can be considered representative of students. While a less data-driven process emphasizes a trustee model of representation in which the decisionmaker is seen as acting in the best interest of the student, a data-driven process that translates students as inforgs requires decisionmakers to create constructs that ultimately represent themselves rather than students. Standard approaches to protecting student privacy are also considerably more problematic in translated data processes. Approaches to privacy typically rely on restricting the flow of information. A traditional approach views this as a protection of an individual. But when the individuals exist solely as inforgs as in a data-driven decision process, restrictions on the flow of information destroy or at least degrade the inforg itself, excluding the associated person from the process.

3.4.1 Inforgs in Data-Driven Decision Processes

In recent decades, higher education in the United States has seen dramatically increasing corporatization, bureaucratization, and rationalization of higher education derived from the for-profit sector but increasingly common in the public and private non-profit sectors as well. A central feature of this has been the emergence of accountability regimes, in which

> a politics of surveillance, control, and market management disguise[es] itself as the value-neutral and scientific administration of individuals and organizations (Tuchman 2009). Related to strategic planning, this accountability regime supposedly minimizes risks for an organization (or corporation) by imposing rules about how work will be done and evaluated. (McMillan Cottom and Tuchman 2015, p. 8)

The scope of such regimes goes far beyond traditional notions of legal and financial risk, reaching into the realm of operational control through data-driven decision-making processes. Accrediting bodies demand that mission fulfillment and student learning be demonstrated through "meaningful, assessable, and verifiable data—quantitative and/or qualitative, as appropriate to its indicators of achievement" (Northwest Commission on Colleges and Universities 2010, Sect. 4.A.1) and that institutions practice "regular, systematic, participatory, self-reflective, and evidence-based assessment of its accomplishments" (Northwest Commission on Colleges and Universities 2010, Sect. 5.A.1). The results of these data-driven analyses are "used for improvement by informing planning, decision-making, and allocation of resources and capacity" (Northwest Commission on Colleges and Universities 2010, Sect. 4.B.1). Institutions that fail to use appropriate data-driven processes to evaluate mission fulfillment and student learning risk punitive actions by accreditors. For example, in June 2013, the Middle States Commission on Higher Education, the largest of the regional accrediting bodies in the US higher education system, issued warnings that the accreditation of ten schools was in jeopardy; nine of these institutions had failed to demonstrate compliance with standards relating to planning, effectiveness, and learning assessment (Middle States Commission on Higher Education 2014).

The reliance on data in assessment, evaluation, and planning—arguably the most important decision processes in a university—is a paradigmatic case of the broader model of data-driven decision-making. Mandated at the primary and secondary levels in the United States by the now superseded No Child Left Behind Act of 2001, data-driven decision-making compels institutions to use data "to stimulate and inform continuous improvement, providing a foundation for educators to examine multiple sources of data and align appropriate instructional strategies with the needs of individual students" (Mandinach 2012, p. 72). The model is based on business management theories (especially those derived from manufacturing), including Total Quality Management and Continuous Improvement. The model organizes and interprets multiple types of data into information that is meaningful to the users. This then becomes actionable knowledge when users evaluate and synthesize the available information, ultimately using the information to either inform discussion or to choose actions. This process is cyclical and takes place within a range of varying organizational contexts (Marsh et al. 2006). The result is held to be a more rigorous and informed decision process that allows educators to teach more effectively and administrators to operate more efficiently and reliably (Mandinach 2012).

Unexamined in this model is the nature of the data that is driving decision-making. Data is, from the perspective of data-driven decisions, seen as an objective representation of a real world. This realist view is fundamentally flawed, however. In order to understand what a data point means, it must be understood as a representation of something within a nexus of problems, models, and interventions rather than as an abstracted object. The process of making reality legible reflects a fundamental problem: the relationship between that which is to be represented and the data state ultimately representing it is one-to-many; therefore data systems must select a single data state from among the many possible in order to produce legible knowledge. Hence the second key element: that data is itself constructed by social processes. I have elsewhere (Johnson 2015) called this process the *translation regime*, which one might define as the set of implicit or explicit principles, norms, rules, and decision-making procedures through which single, commensurable data states are selected to represent states of the world[4] that provides an external source of stability for the data system and allows it to bring legibility to the represented conditions (Mitev 2005). One could look to gender as a paradigmatic case of translation, with myriad possible gender expressions reduced to a small number of values, most commonly "male" or "female," by data standards and validation tables that reflect social norms, in particular those at work in the accountability regime of the institution.

From this perspective, data-driven decision-making takes place within an abstracted model world that resembles any reality external to it in one of many possible ways selected by the translation regime. In a data table, data exists in columns where the data has a common framework, but it also exists in rows that relate data points in different columns to each other through association with some sort of entity: data is information *about some things*, students and courses in the case of UVU's core institutional research data systems. These things in the database can have no more objective existence than the characteristics that the database attributes to them. The translation regime does not simply translate the characteristics of objectively existing entities into the columns of a database; those entities that make up the rows are also translations, whose existence is defined strictly by the information with which they can be associated.

These data entities are best described as what philosopher of information Luciano Floridi terms "inforgs":

> In many respects we are not stand-alone entities but rather interconnected informational organisms or *inforgs*, sharing with biological agents and engineered artefacts a global environment ultimately made of information, the infosphere. This is the informational environment constituted by all informational processes, services, and entities thus including informational agents as well as their properties, interactions, and mutual relations. (Floridi 2010, p. 9, emphasis in original)

[4] This definition follows that of Krasner's (1982, p. 186), used to define regimes in international relations.

An inforg is characterized as an entity that is de-physicalized, typified (represented as an instance of a class of identical objects), perfectly clonable, and existing only through its interactions with other inforgs. While the extent to which this ontology, which Floridi calls "informational structural realism," is an adequate description of being more broadly remains controversial, the sense of inforgs inhabiting an infosphere captures well the ontology of the model world in which a data-driven decision process takes place. In such a model world, data consists of signifiers of states that attach to inforgs. In a star schema, for instance, data is divided into fact tables that describe entities and dimension tables that describe conditions that those entities can take on. Each row in the fact table represents one entity, named by the data table's primary key, and that entity has no characteristics other than the facts stored in the row, that can be joined to the row, or that are stored in the related dimension tables. These inforgs are thus the only kind of entity that can exist within a data-driven decision process.

3.4.2 Informational Representation

Decisions in higher education are political decisions in the most basic sense: they are decisions made to govern a collective entity, in this case a postsecondary educational institution. As such, those that are affected by this decision, as in all political decisions, have a legitimate claim that they ought to have meaningful input into it in some fashion. This is the origin of the problem of representation, a problem not challenged by the fact that the decision takes place in a bureaucratic rather than legislative institution. Presumably, then, decisionmakers in higher education intend for their decisions to represent, in some form and among other considerations, the students about whom they are making decisions.

One might analyze different modes of representation along two dimensions. The first concerns the level of participation. Participatory models involve all those who have a claim to input in the process of making the decision; representative models vest that power in a relatively small group of individuals who act for the group as a whole. A second dimension considers the relationship between the decisionmakers and the group. Promissory models view the decisionmaker as an agent who acts on behalf of those they represent as principals, while autonomous models allow the decisionmakers the freedom to act on their own. The most common models fall into either the autonomous/participatory or the promissory/representative quadrants. Direct democracy, in which all members of the polity participate directly in policy-making, is the standard case of the former; the trustee-delegate dichotomy, in which representatives act respectively in the best interests of the represented or as the represented themselves would, is the basis of the latter.

This is not to say that the only coherent models of representation fit into one of these two quadrants. Frameworks of representation in the two other quadrants are less commonly observed but nonetheless important. In descriptive representation, representatives act without any moral obligation toward the positions of the represented but “in their own backgrounds mirror some of the more frequent experiences and outward manifestations of belonging to the group” (Mansbridge 1999, p. 628). This correspondence of backgrounds acts as a mechanism to ensure correspondence between the interests of the representative and the represented so that a representative acting in their own self-interest is coincidentally acting in that of the represented as well rather than acting out of an obligation to do so. Descriptive representation is an important case of representation that is both autonomous and representative used especially to study representation in bureaucracies (see, for example, Wilkins and Keiser 2004). Jean-Jacques Rousseau’s *On the Social Contract* proposes a system in which citizens participate directly in government but represent not their particular individual wills but the “will that one has as a citizen,” which he terms “the general will,” thus directly participating in government but as an agent of the collective body of citizens that serves as principal. However, neither of these models is of practical value in higher education decision processes. In the case of descriptive representation, decisions are made by actors who cannot resemble the key characteristic of those they might be taken to represent: administrators are not students. Concepts related to the general will have never been shown to be sufficiently clear in any applied context to be of use in making a specific decision. Analysis of representation will thus focus on the direct and promissory models of representation.

In a personalized decision-making context, which we might define in contrast to a data-driven process as one in which either single or multiple decisionmakers use their personal judgment to make what they consider the best decision given the available information under some degree of uncertainty, higher education tends toward a trusteeship model of representation. Even at the smallest of institutions, direct participation in all decisions is impractical because of the number of students and of decisions involved in governing the institution. But there is also a strong strain of paternalism in decision-making at colleges and universities. Students, it is frequently held, cannot be counted on to do what is best for them. Consider, for instance, Austin Peay State University’s use of predictive analytics in student advising:

> [Provost Tristan] Denley points to a spate of recent books by behavioral economists, all with a common theme: When presented with many options and little information, people find it difficult to make wise choices. The same goes for college students trying to construct a schedule, he says. They know they must take a social-science class, but they don’t know the implications of taking political science versus psychology versus economics. They choose on the basis of course descriptions or to avoid having to wake up for an 8 a.m. class on Monday. Every year, students in Tennessee lose their state scholarships because they fall a hair short of the GPA cutoff, Mr. Denley says, a financial swing that ‘massively changes their likelihood of graduating. (Parry 2012)

Such students would, if they chose themselves, make choices that run counter to their true interests (presumably, in receiving a generic college degree at minimum cost); decisionmakers must therefore choose not what the students *would* choose but what they *should* choose. Such a model of representation is defensible only to the extent that the decisionmakers do, in fact, have an adequate view of that interest.

This model of representation breaks down when students are translated into inforgs. Initially, one is tempted to see the translation of students (or of anyone with a claim to voice in a political process) as a gain for direct participation. The promissory models both break down when applied to inforgs. The trustee and delegate approaches both require a unifying concept that acts as the wholeness of the represented (interest or will, respectively) that guides how the agent acts on behalf of the principal, one that is lacking when the principal is no more than a bundle of information: which piece of information defines that unifying concept? But while a personalized process of direct participation requires some complex structure that allows universal participation in the process of developing policy alternatives, manages extensive deliberation among those alternatives, and aggregates preferences into a decision, a data-driven process can bring the participants in as inforgs and then aggregate their informational characteristics. The capacity for participation in data-driven decision-making is apparently limited only by the power to collect and process the information that constitutes the inforgs.

This understanding of representation assumes that inforgs have an objective or realist ontological status, existing in their own right rather than being constituted by actors outside of themselves: the data row represents a physically existing student as they are in the "real" world rather than existing as an inforg that has been created by someone other than the represented. The analysis of the data structures above shows that this is not the case. Inforgs are themselves social constructs, and both their existence and their characteristics reflect the same social pressures and structures that data fields do. As such, the idea that inforgs are capable of being independently represented in a data-driven decision process is fundamentally unsound; what is represented is the constructive activity of those creating the inforgs. There is the appearance of direct participation, but the participants are not representations of students but actants created through the translation regime. What is represented is as much the constructors' understanding of students that is built into the data driving the decision process.

Data-driven decision processes thus present a fundamental contradiction. While they are instituted as objective processes, it is clear that no process of representing students can take place within them without the process of data creation also being a process of imposing external values and assumptions. The inforgs are created by those who create the data system, and decisions about them can only be made if decisionmakers supply their own concepts of interests of will to guide the application of promissory models of representation. This is, to be sure, true of personal decision models as well, but in those models there is a clear connection to individuals against which those assumptions can be checked. In a data-driven model, there is nothing to check against beyond the data; the students exist solely as data. The objectivity of the process, its supposed virtue, is thus a fiction needed to make the process work.

3.4.3 Destructive Privacy Among Inforgs

Representing inforgs becomes more seriously compromised when considered in relation to information privacy. In the United States, students are protected first and foremost by federal laws including but not limited to the Federal Education Rights and Privacy Act (FERPA), but also by a range of state laws, institutional policies, and data handling standards. All of this is meant to ensure that students are able to maintain a sphere of personal identity and activity safe from intrusion by others, including others' knowledge about the student. Most commonly this is protected by the twin principles of consent and anonymity: personal information may only be used or transferred with the consent of the subject; all other information must be stripped of personally identifying characteristics before use or transfer (van Wel and Royakkers 2004). Certainly these opt-in or opt-out procedures are the bedrock of most institutions' privacy policies, with the latter likely far more common than the former.

Increasing pressures on personal privacy have given rise to more complex perspectives on privacy. It is increasingly common to interpret privacy as a property right in information about one's self. Subjects hold initial ownership rights in information about them, and can exchange that information contractually in information markets, receiving appropriate compensation—or they can refuse to permit the use of such information in the absence of sufficient compensation to encourage the transaction (Solove 2004, pp. 76–81). This approach makes sense, for example, of the willingness of so many to give access to their email to Google: in exchange for an outstanding product, consumers are willing to allow Google to use the information captured to generate profit for itself. Alternatively, Helen Nissenbaum (2010) argues for a reliance on social context to protect privacy. As technosocial systems, the context of information flows is as much a defining feature of data exchange and use as the content of that information flow. The combination of situation, actors, information attributes, and practices of transmission for accepted information exchanges constitute an existing norm of practice that may be violated in the case of new uses of information, such as a data mining practice. Changes in this context that are not supported by its underlying norms are violations of the contextual integrity of the information flows, and in the absence of separate justification violate one's privacy rights. More recently, the European Court of Justice has embraced a "right to be forgotten" under which individuals are entitled to have information about them essentially destroyed, in the instant case by having Google remove links to information about them from search results (Costeja González 2014).

The common thread of each of these approaches to privacy is that they aim to restrict flows of information across parties, transactions, or both. This restriction is frequently considered the essence of data privacy. The centrality of collection (the flow of information from a subject to a data system) and dissemination (the flow of information across data systems or from a data system to subjects) in common definitions of information privacy makes restrictions on flow the *sine qua non* of data privacy. Such a model of privacy is at least plausibly appropriate for the governance of subjects who are persons; preventing the transfer of information will, presumably, prevent those receiving information from using it to do harm to the subjects of that information. This meets the fundamental criteria of a wide range of ethical frameworks, such as Mill's harm principle, which permits the infringement of one's liberty in order to prevent harm to others, or the more recent proposal of a Hippocratic Oath making "do no harm" the first principle in the use of information and communication technology for development (Mill 2011, p. 17; Rodrik 2012).

Restricting the flow of information fundamentally fails, however, when the subjects are constructive inforgs. The flow of information is what translates subjects (in this case, students) into inforgs in the first place. To restrict that flow is to change the inforg itself. Such restrictions might, for instance, limit the data known about an inforg in absolute terms as privacy restrictions prevent the transfer of certain types of information (when, for example, the subject opts out of sharing of internet use information). Or it might do so in relative terms as it prevents the transfer of information from one source (when the subject installs a privacy plug-in in Chrome) but allows that same transfer from another source (when the subject doesn't bother reading the 31-page terms and conditions for the latest iOS update). Since an inforg is nothing more than a typified and clonable bundle of information, a difference in the information constituting the inforg violates the principles of typification (the difference resulting in inforgs that are instances of two different types) and clonability (the difference distinguishing two instances as different rather than as clones), and is thus the creation of a different inforg.

This becomes even more problematic when a subject opts out of a data system altogether. For a constructive inforg, a complete data opt-out is not simply a withholding of information; it is a complete destruction of itself as an inforg. Prohibition of data flows prevents the inforg from being constructed in the first place. It is perhaps only slightly overdramatic to characterize complete restriction of the flow of data as information suicide for a constructive inforg, as the inforg that protects its privacy ceases to exist in the model world of the data-driven decision process. The physical entity corresponding to the inforg (in this case, the actual student) is at best reduced to context—that there are some students who are excluded by privacy protections. But context, again, exists only in relation to data, which is to say in relation to inforgs. Students who opt to protect their privacy thus exist only as others to the inforgs' selves, defined not individually as entities in themselves but collectively as a typified characteristic of the inforgs (i.e., as a group of identical entities of which the inforgs are not members). Reduced to context that is meaningful only in relation to entities that have corresponding inforgs, those students cease to exist analytically and instead are subsumed as information into inforgs corresponding to other students.

That further complicates the problem of representation as well. Partial restrictions change how subjects are represented; complete prohibitions exclude subjects from being represented entirely. Students are faced with a difficult choice: they can be represented (with varying levels of adequacy given the process of constructing inforgs) in the data-driven decision processes that run the institution that shapes a significant part of their lives both now and long into the future, or they can choose to minimize the extent to which that institution is allowed into the student's sphere of private activity and identity. To exactly the extent that students choose one good, they undermine the other. In personalized decision processes, the unifying concepts of principal-agent representation can moderate this, with decisionmakers taking into account expressions of students' best interests and wills regardless of—and perhaps taking into consideration—the privacy status of individual students, as these are not data-dependent. In data-driven decision processes, however, with those unifying concepts absent and decisions formally constrained by the available data, representation and privacy are fundamentally irreconcilable.

3.5 Conclusion

These transformations are political acts. The actors that design translation regimes are building structures that embed values and relationships within them that can advantage certain groups over others as the data rather than the actors it represents comes to play a defining role in decision processes. The translation regime begins by representing some groups and excluding other groups, representing some characteristics of individuals but not other characteristics of those same individuals, and representing the data subjects as the data system's designers would represent them rather than as the subjects would. In UVU's data systems, non-credit students and non-degree seeking students do not exist under most circumstances; nearly all queries are designed to filter such students out unless information about them is needed specifically. English as a Second Language students were until recently treated as non-degree seeking and thus left unrepresented in most data-driven decisions. Students' ethnicity is represented but their religion, the most commonly discussed aspect of identity in the student interviews, is not. White students are represented as an ethnicity rather than seeing themselves as ordinary people (who seem to lack ethnicity), as one White student described himself. These translations are not necessarily hostile to the students' representation, but they do quite clearly shape it.

Just as there are many characteristic translations, there are many political acts that take place through them. The creation of data systems is an assertion of self-interest on the part of the designers; the data system embeds their interests in the decision process but not those who have no influence on the design processes; the latter have no way to make themselves and their interests legible even to institutions that might want to take them into account in good faith, let alone those who might deliberately seek to exclude them. The categorization of characteristics creates and fragments groups that could assert their aims to the institution: Black women are forced to choose to work within the defined fields of *GENDER* and *ETHNICITY* to meet their needs and thus to accept racial inequality within the feminist movement or gender inequality within Black culture rather than identifying as Black women specifically and pursuing an intersectional strategy (Hill Collins 2009). Defining states of the world as valid or invalid (e.g., transgender identities) is at the least an imposition of a normalizing judgment through a means other than surveillance, one that has the same kind of potential to construct individuals and groups as hate speech (Butler 1997).

Similarly, data-driven decision-making becomes much more problematic when we recognize that data is made, not collected. As decision-making increasingly takes place within model worlds created by the process of collecting, managing, and analyzing data, it increasingly transforms people into inforgs and marginalizes considerations not rooted in data as mere context.[5] Data-driven decision-making is part of a larger ethos, one connecting managerialism, technocratic government, and neoliberal politics, that increasingly pervades higher education. The problems of representation and privacy, and especially the tension between the two, stem from the very core of this ethos.

Much of the politics that one would typically expect as groups compete is present in the translation regime. The politics of the translation regime is different, however, in that it is hidden behind a facade of technical specifications. The translations are, superficially, not exercises of power but simply functions and validation tables that store ostensibly objective information about reality. The scientistic ontology and ideology (Haack 1993; Peterson 2003) embedded in information systems creates the appearance of an apolitical process that is not open to contestation. It thus becomes quite difficult to engage from a political perspective. It cannot be challenged technically, as the translation regime is internally valid and self-legitimating. Any test against reality will confirm the validity of the regime so long as the rules are complied with, because the rules include what data can be considered. Data from within the regime will be correct, and there is no such thing as "data" from outside of the regime. The translation regime creates data; all else is anecdote and thus illegitimate. Challenges to the politics of the translation regime must, then, overcome the issue of legitimacy before the regime can be questioned.

The translation regime is thus a significant and problematic form of political power. Integrating both the technical and the social to render its subjects legible to the exercise of power, the characteristic translations that it produces also exercise power in their own right. As such, the fact that data is constructed through translation, among other processes, presents the need for a theory of information justice. Such a theory must rely on neither controlling the possession of information nor its use. If information is simply representational, these would be adequate safeguards. Privacy rights could protect transfer of information, and substantive regimes similar to human subjects protections might prevent against harmful uses. But the constructive nature of data makes these inadequate. Neither privacy use ethics addresses the content of information that is, within the internal framework of the translation regime, accurate. These approaches cannot address the questions that arise in build-

[5] To be sure, one might argue that the portrayal of data-driven decision-making presented here is something of a straw-man argument that neglects the subtleties of and importance of context in the models advocated in higher education. I would argue to the contrary that those models themselves only pay lip service to context; the more context can be used to override data and the more that conflicting data points are to be considered in the decision process, the less data-driven decision-making is distinct from personalized decision-making. If there is something distinct about data-driven decision-making, it is that data must take priority over context.

ing data systems in ways that their translations further rather than undermine the individuals represented in them. Instead, a theory of information justice should be oriented toward understanding data as a socio-technical system, promoting design practices that minimize their potential for domination and oppression.

References

Benson, J. E. (2006). Exploring the racial identities of black immigrants in the United States. *Sociological Forum, 21*(2), 219–247. https://doi.org/10.1007/s11206-006-9013-7.

Butler, J. (1997). *Excitable speech: A politics of the performative*. New York: Routledge.

Costeja González, M. (2014). Google Spain SL, Google Inc. v Agencia Española de Protección de Datos (es), EU:C:2014:31.

Easley, J. A. (2014, March 18). If it's your preferred name, then we prefer it, too. *Dateline: News for Faculty and Staff*. http://dateline.ucdavis.edu/dl_detail.lasso?id=14756&dn=031814. Accessed 19 Mar 2014.

Facebook. (n.d.). *How do I edit basic info on my Timeline and choose who can see it?* https://www.facebook.com/help/276177272409629. Accessed 25 Feb 2014.

First Year Experience and Student Retention. (2014). *15 to finish*. https://www.uvu.edu/success/15tofinish/. Accessed 24 May 2017.

Floridi, L. (2010). *Information: A very short introduction*. Oxford: Oxford University Press.

Haack, S. (1993). *Evidence and inquiry: Towards reconstruction in epistemology*. Oxford: Blackwell.

Hall, P. A., & Taylor, R. C. R. (1996). Political science and the three new institutionalisms. *Political Studies, 44*(5), 936–957.

Hill Collins, P. (2009). *Black feminist thought: knowledge, consciousness, and the politics of empowerment* (2nd ed.). New York: Routledge.

Information Commissioner's Office. (n.d.). *The guide to data protection*. http://ico.org.uk/for_organisations/data_protection/~/media/documents/library/Data_Protection/Practical_application/the_guide_to_data_protection.pdf. Accessed 25 Feb 2014.

Institutional Research & Information. (2012a). *Fact book 2012*. https://www.uvu.edu/iri/documents/additional_resources/factbook2012.pdf. Accessed 12 Mar 2014.

Institutional Research & Information. (2012b). *Student success/retention.* http://www.uvu.edu/iri/indicators/. Accessed 12 Mar 2014.

Institutional Research & Information. (2013). *Student omnibus survey – Fall 2012 results*. http://www.uvu.edu/iri/documents/surveys_and_studies/Omnibus%20Student%20Survey%20-%20Fall%202012%20Results.pdf. Accessed 20 Mar 2014.

Johnson, J. A. (2006). Technology and pragmatism: From value neutrality to value criticality. In *Western political science association annual meeting*. Albuquerque. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2154654.

Johnson, J. A. (2015). Information systems and the translation of transgender. *TSQ: Transgender Studies Quarterly, 2*(1), 160–165. https://doi.org/10.1215/23289252-2848940.

Krasner, S. D. (1982). Structural causes and regime consequences: Regimes as intervening variables. *International Organization, 36*(2), 185–205.

Mandinach, E. B. (2012). A perfect time for data use: Using data-driven decision making to inform practice. *Educational Psychologist, 47*(2), 71–85. https://doi.org/10.1080/00461520.2012.667064.

Mansbridge, J. (1999). Should blacks represent blacks and women represent women? A contingent "Yes". *The Journal of Politics, 61*(3), 628–657.

Marsh, J. A., Pane, J. F., & Hamilton, L. S. (2006). *Making sense of data-driven decision making in education: Evidence from recent RAND research*. Santa Monica: RAND Corporation.

http://www.rand.org/content/dam/rand/pubs/occasional_papers/2006/RAND_OP170.pdf. Accessed 16 Sept 2014.
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.
Middle States Commission on Higher Education. (2014, June 26). *Summary of commission actions on institutions*. http://www.msche.org/institutions_recentactions_view.asp?dteStart=4/29/2014&dteEnd=6/26/2014&idCommitteeType=0&txtMeeting=Commission. Accessed 22 Sept 2014.
Mill, J. S. (2011). *On liberty (project Gutenberg eBook.)*. London: The Walter Scott Publishing, Ltd.. http://www.gutenberg.org/files/34901/34901-h/34901-h.htm. Accessed 30 Sept 2014.
Mitev, N. N. (2005). Are social constructivist approaches critical? The case of IS failure. In *Handbook of critical information systems research: Theory and application* (pp. 70–103). Northampton: E. Elgar Pub.
National Center for Education Statistics. (n.d.). *The integrated postsecondary education data system – Glossary*. http://nces.ed.gov/ipeds/glossary/. Accessed 11 Nov 2015.
Nissenbaum, H. (2010). *Privacy in context: Technology, policy, and the integrity of social life*. Stanford: Stanford Law Books.
Northwest Commission on Colleges and Universities. (2010). *Accreditation standards*. http://www.nwccu.org/Standards%20and%20Policies/Accreditation%20Standards/Accreditation%20Standards.htm. Accessed 22 Sept 2014.
Parry, M. (2012, July 18). College degrees, designed by the numbers. *The Chronicle of Higher Education*. https://chronicle.com/article/College-Degrees-Designed-by/132945/.
Peirce, C. S. (1992). In N. Houser, C. J. W. Kloesel, & Peirce Edition Project (Eds.), *The essential Peirce: Selected philosophical writings* (Vol. 1–2., Vol. 1). Bloomington: Indiana University Press.
Peterson, G. R. (2003). Demarcation and the scientistic fallacy. *Zygon, 38*(4), 751–761. https://doi.org/10.1111/j.1467-9744.2003.00536.x.
Raman, B. (2012). The rhetoric of transparency and its reality: Transparent territories, opaque power and empowerment. *The Journal of Community Informatics*, *8*(2). http://ci-journal.net/index.php/ciej/article/view/866/909. Accessed 5 Mar 2013.
Rodrik, D. (2012). *A hippocratic oath for future development policy a hippocratic oath for future development policy*. http://www.policyinnovations.org/ideas/commentary/data/000244
Scott, J. C. (1998). *Seeing like a state: How certain schemes to improve the human condition have failed*. New Haven: Yale University Press.
Seaver, N. (2014, January 30). On reverse engineering: Looking for the cultural work of engineers. *Medium.com*. https://medium.com/anthropology-and-algorithms/d9f5bae87812. Accessed 28 Feb 2014.
Solove, D. J. (2004). *The digital person: Technology and privacy in the information age*. New York: New York University Press.
Utah System of Higher Education. (2013a, July 1). *Course data submission file, 2013–2014 submission year*. http://higheredutah.org/wp-content/uploads/2013/09/rd_2013DataDict_Course.pdf. Accessed 10 Mar 2014.
Utah System of Higher Education. (2013b, July 1). *Graduation data submission file, 2013–2014 submission year*. http://higheredutah.org/wp-content/uploads/2013/09/rd_2013DataDict_Graduation.pdf. Accessed 10 Mar 2014.
Utah System of Higher Education. (2013c, July 1). *Student data submission file, 2013–2014 submission year*. http://higheredutah.org/wp-content/uploads/2013/09/rd_2013DataDict_Students.pdf. Accessed 10 Mar 2014.
Utah System of Higher Education. (2013d, July 1). *Student_Course data submission file, 2013–2014 submission year*. http://higheredutah.org/wp-content/uploads/2013/09/rd_2013DataDict_Student_Courses.pdf. Accessed 10 Mar 2014.

UVU Student Retention. (2013). *Stoplight report*. http://www.uvu.edu/retention/advisors/stoplight.html. Accessed 13 Mar 2014.

Weingast, B. R., & Moran, M. J. (1983). Bureaucratic discretion or congressional control? Regulatory policymaking by the federal trade commission. *Journal of Political Economy, 91*(5), 765–800.

van Wel, L., & Royakkers, L. (2004). Ethical issues in web data mining. *Ethics and Information Technology, 6*(2), 129–140. https://doi.org/10.1023/B:ETIN.0000047476.05912.3d.

Wilkins, V. M., & Keiser, L. R. (2004). Linking passive and active representation by gender: The case of child support agencies. *Journal of Public Administration Research and Theory, 16*(1), 87–102. https://doi.org/10.1093/jopart/mui023.

Zhang, R., & Noels, K. A. (2013). When ethnic identities vary: Cross-situation and within-situation variation, authenticity, and well-being. *Journal of Cross-Cultural Psychology, 44*(4), 552–573. https://doi.org/10.1177/0022022112463604.

Chapter 4
The Political Life of Metrics

Abstract This chapter extends the analysis of the previous chapter to the role of metrics in political practice, using the U.S. standard graduation rate metric as a case. I argue that information is best understood as a process of communication in which observation is encoded into data through the translation regime and then decoded into metrics which are then institutionalized in political processes. In both processes, political factors are prominent, making metrics a political outcome at the least. I go further, however, showing that metrics play important distributive roles in politics, allocating material and moral goods as well as the conditions of political power. Metrics also exercise political control directly, working much like administrative procedures to select favored outcomes without direct legislative intervention and building the capacity of the state to exercise control over policy areas.

To recognize metrics as political is, in some senses, trivial. Metrics, and the data underlying them,[1] are quite often the locus of "office" politics as competing interests in an organization seek to shape them to their own advantage. Consider, for instance, academic departments within a university that are allowed to count majors themselves before making funding requests: one department uses declared majors knowing that a large number of early students will change to other majors after the introductory courses, while another counts all students who have taken several courses in the department so as to include those students who enjoy the subject but want a major with better job prospects.

In this sense, though, calling data political in this way is to call it pathological; this kind of data is contrary to the apolitical standard that data should provide.

[1] For the purpose of this chapter, I will use "data" to refer to a representation of a purportedly unified construct and "metric" to refer to combinations of data points that provide a standard for evaluation. Paradigmatically, the number of students in the IPEDS Graduation Rate Survey cohort is data. The GRS150 graduation rate is a metric composed of three data points (the number of students in the GRS cohort, the number of graduates from that cohort, and the number of students excluded from the cohort) and a mathematical transformation of those data. Metrics might also include growth rates in a single type of data over time, transformations of other metrics, or comparisons to other data or to a benchmark value. They would rarely, if ever, be single data points themselves, as such provide no basis for evaluation; likely, there is an implicit relationship to other data in metrics that are so defined.

J.A. Johnson, *Toward Information Justice*, Public Administration and Information Technology 33, https://doi.org/10.1007/978-3-319-70894-2_4

Hence there is also an alternative sense of the (a)politics of data that sees it as overcoming office politics, transcending political divides with an unequivocal, observable truth: "You can't argue with the data," proponents of apolitical data are apt to exclaim. Departments cannot game the numbers to their advantage when the data come predefined for them. The office politics of data cannot happen, in this view, unless there is poor data, and the solution is to create sound data, data that is reliable and objective, that provides a true view of the world as it really is. The manager who can rely on data can thus overcome politics with fact.

Of course, such a view can only hold if data can be, in fact, objective. If it cannot—as the previous chapters have shown—then politics cannot be overcome. The position that metrics are apolitical tools for management is thus unsupportable. This chapter presents an alternative view, showing instead that metrics are inherently political by a variety of definitions of politics, using a case study of the Graduation Rate Survey (GRS) cohort used in the United States' Integrated Postsecondary Education Data System (IPEDS) and its use in decision-making, primarily at Utah Valley University (UVU). I explore three political aspects of the GRS cohort definition that allow it to function as a normalizing translation regime: its development by the United States Congress and the Department of Education (USED) during the 1990s as a reporting element, its transformation into a key point in policy debates surrounding the "completion agenda," and its effects on current campus decision-making.

This critical political theory of information situates metrics such as the IPEDS graduation rate within a broader process of encoding and decoding data, with metrics as one mechanism for the decoding phase of data communication but one that strongly influences the encoding phase as well. Metrics, as the outcome of a process of communication and translation, are first and foremost a product of political institutions and processes. Metrics and the data that informs them then enter other political processes by being transformed into socio-political institutions, doing much more than simply measuring and informing. Such institutions operate both as structures embedded in social and political practice that shape individual behavior and as culturally-specific practices that shape ways of understanding the world and responding to it. This influences political outcomes as distributive processes and as establishing relationships among people and groups.

4.1 Encoding Data as a Political Process

In the United States, authority over higher education is divided between the federal and state governments. The states' role is primarily operating the public university[2] systems that educate the majority of American university students and secondarily

[2] In U.S. usage there is no precise, formal distinction between colleges and universities. Institutions that use "college" in their names are more typically either smaller "liberal arts" colleges that originated as institutions to train primary and secondary teachers or 2-year community colleges

regulating the private higher education sector. The Utah System of Higher Education (USHE) supports two research universities, four regional universities offering primarily bachelors's degrees (including Utah Valley University), a community college, and a 2-year rural residential college. While the universities have significant autonomy, USHE is primarily responsible for system-wide policy and, of particular importance to the politics of metrics, sets statewide data reporting standards and processes that are commonly used in campus decision-making as well.

The federal government plays some regulatory role, especially with regard to the private accreditation bodies that play the most direct role in regulation, but its major role is funding higher education through research grants and student financial aid programs (known commonly as "Title IV" programs after the section of the Higher Education Act of 1965 that created the major programs in use today). The universities eligible to participate in federal financial aid programs for higher education bear significantly increased regulatory burdens consequent to receiving federal money. This chiefly takes the form of reporting requirements, most notably reporting to IPEDS. IPEDS consists of a series of surveys collecting a wide range of aggregate[3] data on institutional characteristics, admissions, enrollment, completion, financial aid, human resources, and institutional finance.

One of the more controversial data elements has been graduation rates, which are based on the graduation of students included in the GRS cohort. While the collection of education data dates to 1867, graduation rates for post-secondary universities in the United States have only been collected since the 1997–1998 academic year (Fuller 2011, pp. 5–6). Between 1966 and 1987, the Higher Education General Information Survey (HEGIS), IPEDS' immediate predecessor, collected data on the number of degrees awarded; the HEGIS completions survey was replaced by the IPEDS Completions Component in the 1987–1988 academic year, collecting data for the number of degrees and other formal awards conferred annually. While additional information has been added to the Completions Component, especially as required under the Higher Education Act (HEA) of 1998 and the Higher Education Opportunity Act (HEOA) of 2008, the basis of the Completions Component remains reporting on the number of awards and the number of students who receive them (Fuller 2011, pp. C1–C3). Neither HEGIS nor the Completions Component by themselves collect graduation rates (i.e., the percentage of students graduating), nor could the Completions Component be combined with other IPEDS components to calculate such a rate prior to the creation of the GRS.

---

oriented toward vocational training and programs that transfer to institutions offering bachelor's degrees. They are almost—but not quite—always purely undergraduate institutions. Universities are typically larger than liberal arts colleges (though community colleges can span the entire range of institutional size) and usually (but, again, not always) offer graduate programs of widely varying scope. The distinctions, however, are primarily nominal and not analytically useful due to the vast overlap in institutional characteristics. This chapter will thus conform to the more common international usage, using "university" to refer to all U.S. institutions offering post-secondary degrees.

[3] The Department of Education is barred from collecting student unit record data under sec. 113 of the Higher Education Opportunity Act of 2008. While efforts to change this are very nearly constant, none have yet come close to success.

Both HEGIS and the original IPEDS components operated in an era when the federal collected data primarily for research purposes. One of the Congressional purposes in the Department of Education Organization Act of 1979 was "to promote improvements in the quality and usefulness of education through federally supported research, evaluation, and sharing of information" (sec. 102), a purpose furthered with the creation of the Office of Educational Research and Improvement (sec. 209) that had responsibility for HEGIS and subsequently, until a statutory reorganization of the department in 2002, IPEDS. The orientation toward research-driven data collection changed significantly with the Student Right-to-Know Act of 1990, shifting emphasis from supporting research to providing "consumer" information. The Congressional findings in the act showed concern with student performance and the universities' required educational commitments, in particular finding that "knowledge of graduation rates would help prospective students and prospective student athletes make an informed judgment about the educational benefits available at a given institution of higher education" (sec. 102). The language of data submission also changes, from a language of collection by USED to one of disclosure by universities.

This change in emphasis gave rise to the collection specifically of graduation rates rather than just counts of degrees awarded. Section 103(a) of the Student Right-to-Know Act requires universities to make a "Disclosure of Completion or Graduation Rates" as an amendment to section 485(a)(i) of the HEA (codified in 20 U.S.C. 1092), which notably concerns "information dissemination activities for prospective and enrolled students" about academic programs and financial aid, a quite different language than that of disclosure. The calculation of graduation *rates* is significantly more complex than that of completions, however, as it includes both time and base population dimensions. At least two questions immediately arise: which students to include in the denominator of the rate, and when rates will be calculated relative to the student. Not all student groups are equally useful in making an informed judgment about educational benefits. Students not seeking degrees, for example, do not complete a program in any traditional sense, and prior credits earned make it difficult to treat first-time and transfer-in students as part of the same statistical population. The time constraint is equally complicated; part-time students are likely to take much more time than full-time students, while students who transfer into the institution will take less time than first-time students.

These questions bring one back to the key principle of a critical-constructive theory of information technology, Young's recognition that current practice "does not have to be this way, it could be otherwise" (1990, p. 6). It is worth considering that Congress and USED had several alternatives available, and chose these particular data definitions over those alternatives for reasons. Including all students is a viable standard if the aim is to assess the likelihood that an entering student will complete their degree; the transfer advantage is not particularly relevant in that case. Nor is it entirely necessary to report at a threshold time; graduation rates for a cohort could be updated annually for an (in principle) indefinite period. There are thus many possibilities for an operational graduation rate measure—but a need to have only one. A critical-constructive theory of information technology asks why this data standard should have prevailed rather than the others that are equally representative of reality.

If one accepts a realist account of data, this is exceptionally problematic. Data cannot present an objective representation of an underlying reality if there are many data states that can represent a given state of the world. This is entirely unproblematic, however, if data is understood as a form of communication (though, to be sure, it makes much social practice around data exceptionally problematic). Communication involves producing and consuming a message in discursive form—a form that precludes communication of a "raw" event—such that "the broadcasting structures must yield encoded messages in the form of a meaningful discourse. The institution-societal relations of production must pass under the discursive rules of language for its product to be 'realized" (Hall 2006, p. 165). As we saw in the previous chapter, in the case of data it is not simply the rules of language that encode meaning: the meaning of "19457033" is not purely found in rules of language (whether natural or programming) but in semantically and pragmatically meaningful field names (e.g., "STUDENT_ID" or "PHONE_NUMBER") that produce not just contextually different but fundamentally incommensurable meanings. Without a translation regime, "graduation rate" remains a concept to be argued over rather than a fixed data point.

This issue of encoding data through a translation regime may explain why creating an operational definition of the graduation rate took 18 years. Student Right-to-Know initially defined the base population in section 1092(a)(1)(L) as "certificate- or degree-seeking, full-time students entering such institutions" and, in section 1092(a)(3), measures completion as graduation from the program or transfer to another institution for which the program provides preparation "within 150% of the normal time for completion of or graduation from the program." It also defined, in section 1092(a)(4), three classes of students who could be excluded from the graduation rate cohort upon leaving without graduating: those who left to serve in the military, in a religious mission, or in the Peace Corps. USED published final regulations implementing the Student Right-to-Know requirements in 1999 (34 C.F.R. 668), which further specified the graduation rate cohort to include, at most universities, only those students who entered the institution during the fall term (with first-time students who entered in the summer and continued to fall considered to have entered in fall), and added two more categories of exclusions: students who are deceased or who are totally and permanently disabled. Neither the restriction to fall entry nor the additional exclusions are explicitly authorized beyond the general authorization in the Educational Sciences Reform Act of 2002 to collect data that is "useful for policymaking at the Federal, State, and local levels" (20 U.S.C. 9547) that is the current statutory basis for IPEDS. This was further codified by HEOA in 2008, which added a 200% of program time graduation rate in section 1015a(i)(1)(J), added a recalculation provision for schools with large numbers of exclusions in section 1092(a)(4)(B), and gave responsibility to NCES for collection through IPEDS in section 1015a(i)(4).

IPEDS implements these requirements by defining the GRS cohort, currently, as "all students who enter an institution as full-time, first-time degree or certificate-seeking undergraduate students during the fall term of a given year" (National Center for Education Statistics n.d.-a, "Fall Cohort"), and the graduation rate as the 150% program time rate under Student Right-to-Know. IPEDS collects the initial cohort size during the year students enter the institution. It then collects the number of graduates and the number of authorized exclusions at 100%, 150%, and 200% of program time to calculate the graduation rates (colloquially, the "GRS100," "GRS150," and "GRS200" rates), calculating them as the number of completers divided by the adjusted (initial less exclusions) cohort (National Center for Education Statistics n.d.-a, "Graduation Rate") for the highest undergraduate degree offered by the institution as well as for all undergraduate degrees and certificates. The rate explicitly excludes students entering in fall terms as part-time or transfer-in students and all students entering in other terms. IPEDS has recently added, without specific statutory authorization, several Outcome Measures that include separate cohorts of part-time and transfer-in fall students, but this is explicitly not described as a completion or graduation rate (National Center for Education Statistics n.d.-b).

Even the IPEDS definition is insufficient to fully operationalize the GRS cohort at the level of individual universities; each institution's data systems are locally developed and implemented (though frequently using data architecture built on commercial products). Hence the translation regime includes elements from other sources as well. At most universities, students may add or drop classes or even enroll or withdraw from attendance altogether at many points in the semester. The IPEDS standard simply requires universities to report data as of their census date, the date on which the institution must report data to external authorities or on which the institution adopts data as official. UVU offers not only full semester but also half-semester classes that begin after the census date defined by USHE. As the census date is the 21st day of classes for the full semester, students who enroll for a second half-semester class before the census date can be included in the GRS cohort, but those who enroll after the census date cannot even though they take the same classes. In this way, state data standards and institutional practices are as determinative of the encoding as are federal regulations.

This definition, with emphasis on the GRS150 rate, serves as the standard definition of graduation rates in US higher education. But local data systems are also important in determining the encoding. UVU operates two distinct sets of data: a near real-time system and a set of data freezes that reflect the real-time data system as of the date of the freeze. The latter is used especially to provide data as of the census date or end-of-term data. Data changes in the former are not always corrected in the latter. A common case of such change occurs when the university receives a transcript indicating a student has transferred from another institution well after the student has been reported as a first-time student. Such errors are not simply oversights by students; it is possible to check students against the National Student Clearinghouse (NSCH) database that includes most but not all students who have attended a Title IV-eligible institution. It is UVU's data collection process, justified by a need for transcript information for transfer students and perceived weaknesses in NSCH data, that results in these data states. Some data processes, moreover, are built around primarily archival data that produces cohorts as they were reported at the time and not as they would be reported based on subsequently cleaned data. It is thus possible to be a first-time student in one data system and not in another. And while it would be possible to, through technical means, arrive at a consistent data state, it is not possible to choose one without considering non-technical considerations such as whether to consider the cohort a group defined at the outset of a program (thus not subject to correction after its establishment) or a status associated with individual students (which would need to be updated as new information was obtained). These are questions of policy, not of technology.

The encoding of the translation regime is thus only partially technical, and this hybridity is built into the GRS definition at its core. The selection of which students are included or excluded in the cohort and how long they have to graduate are not driven by any technical standard. They reflect, rather, a particular social construct of universities and of their students: a normalizing translation of all students into traditional college students that renders all others invisible or nonexistent. In this paradigmatic type of student, an American university student goes away to a residential university the fall after graduating from high school, chooses a degree program upon enrollment, attends for 4 years of full-time study, and then receives their bachelor's degree. Only within this normalizing translation does the cohort definition make sense. Many students will begin enrollment in a winter or spring term, yet they are not thought of as "typical" and will not be included in an institution's graduation rate. The assumption that first-semester enrollment is representative of enrollment throughout a student's academic career is also dubious for anyone but traditional students from traditional families at traditional institutions. These assumptions are justified normatively, as representatives of an ideal type around which the U.S. system of higher education is designed, rather than as empirically adequate representations of actual students.

The 150% time restriction is especially illustrative of the GRS150's normalization of the traditional student. At a typical U.S. university, students taking the minimum credit hours required to be classified as a full-time student (typically 12) would nominally graduate from a bachelor's degree program in 5 years, 125% of program time. But given the incremental number of credit hours for which students typically enroll (in blocks of three or four credit courses in most cases), the maximum practical number of credit hours for a part-time student is nine rather than 11; such students would graduate in 14 semesters, or 175% of program time. And in 1991–1992, as these standards were being developed, the average student at public and private non-profit degree granting universities took only 16.95 credit hours per *year*, taking 15 semesters—188% of program time—to graduate from a typical bachelor's degree program. As the Student Right-to-Know requirements were being implemented, it was already clear that the majority of students were not included in the standard and if they were they would not meet it. The 150% time standard thus cannot be understood as representing the majority of students or as accounting for part-time attendance without also understanding those who selected and implemented it as at best spectacularly uninformed. Rather the 150% time standard should be seen as the time by which all "normal" students would have graduated—and thus as defining what constitutes normalcy itself. This norm is the key function of the translation regime toward which the technical definitions in law, regulation, and IPEDS rules are directed.

4.2 Decoded Metrics as Political Institutions

There are several hundred data points in the IPEDS data system. Often data is used in policy or academic research on higher education as envisioned by the original legal authority granted to USED. But very few of the IPEDS data points have taken on the significance of the GRS rates, and especially the GRS150. This, too, makes little sense from a realist perspective: the GRS rates are often poor operationalizations of the construct that they are purported to represent. But it again makes sense if one views data as a process of communication. Once encoded by the producer, communication must then be decoded by its audience:

> It is this set of decoded meanings which "have an effect", influence, entertain, instruct or persuade, with very complex perceptual, cognitive, emotional, ideological or behavioural consequences. In a "determinate" moment the structure employs a code and yields a "message": at another determinate moment the "message", via its decodings, issues into the structure of social practices. (Hall 2006, p. 165)

In this frame, the users of data, analogous to the audience of a television program, are active participants in communication rather than passive consumers of predefined information. Only once the data is decoded by its users does it have any effect on social practice. But just as in the case of encoding, decoding data is not solely governed by the rules of language. To the extent that users are aware of it, data can be decoded with reference to the data structures of the translation regime, but by no means does the translation regime determine the decoding. Data users also bring their own structures and meanings to the decoding processes in the form of the nexus of problems, models, and interventions into which the data is incorporated. These translate the data point or aggregations into a representation of some consideration relevant to the users' circumstances and intended courses of action.

These decoding frames, once well established and widespread, can then act as political institutions: "the symbol systems, cognitive scripts, and moral templates that provide the 'frames of meaning' guiding human action" (Hall and Taylor 1996, p. 947). Institutions operate cognitively rather than rationally, "providing the cognitive scripts, categories, and models that are indispensable for action" in that they allow understanding the world and interpreting the behavior of others; they "influence behaviour not simply by specifying what one should do but also by specifying what one can imagine oneself doing in a given context" (Hall and Taylor 1996, p. 947). As a result, action in institutionalized contexts is driven by a logic of appropriateness, "more on identifying the normatively appropriate behavior than on calculating the return expected from alternative choices" (March and Olsen 1989, p. 22).

This is not to say that institutions operate universally without challenge or change. There are many complexities to metrics as institutions. March and Olsen identify two key complexities:

1. Institutions "appear to be bureaucratic, rigid, insensitive, or stupid" and "imperfection is often manifest, especially after the fact," yet they persist because routinized rather than individually autonomous behavior is necessary for widespread coordination of social activity.

2. Institutions are neither internally nor externally inherently consistent or monolithic, allowing actors to choose which routines to follow, a choice still based on the logic of appropriateness amidst conflict and ambiguity and that follows a reasoning process akin to common law legal reasoning (March and Olsen 1989, pp. 24–26).

To these, one might add that as the circumstances in which an institution is relied on to determine an appropriate action becomes less frequent or less salient, institutions may well break down, losing their ability to compel action even where they did formally apply. Through these processes, political institutions may grow, evolve, interact, and eventually die, not because of rational calculus or functional suitability—institutions often persist long after their function is gone or there is no calculus of utility to sustain them—but as organic elements of social structure.

The GRS150 is an institutionalized metric. While USED established the GRS150 metric in the 1990s, policy organizations have played the foremost role in translating the GRS150 from statistic to institution in the early Twenty-first Century. Organizations such as the Lumina Foundation and the Bill and Melinda Gates Foundation, in cooperation with various associations of universities, have aggressively supported a "completion agenda" that called for universities to double student completion rates by 2020. The American Association of Community Colleges made a "Sample Completion Commitment Statement" available to its members in 2011, in which "[INSERT NAME OF YOUR INSTITUTION]" takes responsibility for completion:

> With the "completion agenda" as a national imperative, [INSERT NAME OF YOUR INSTITUTION] has an obligation to meet the challenge while holding firmly to traditional values of access, opportunity, and quality.
>
> …We believe the "open door" must not be a "revolving door," and that [INSERT NAME OF YOUR INSTITUTION] must take responsibility for student success.
>
> … We believe to change [*sic*] in institutional culture, from emphasis on access only to emphasis on access and success.
>
> … We commit to acting on facts to make positive changes in the interest of student success and college completion.
>
> We commit to promoting faculty and staff development focused on evidence based educational practice. (American Association of Community Colleges 2010)

In the national discourse on higher education, in academic research, and especially as this larger discourse is engaged in institutional program management, the development of the completion agenda acts as a cognitive script for policy entrepreneurs, government agencies, the media, and institutions. That script idealizes completion of a post-secondary credential as the path to economic success (McMillan Cottom 2017), identifies completion of a (job-qualifying) credential as the overriding goal of higher education both individually and systemically, and characterizes students who leave before completing their degrees as having been failed completely (i.e., those with some college but no degree are in no better position to qualify for a job than those with no higher education, and are likely saddled with substantial debt) by the institutions, which bear primary responsibility for student success.

The completion agenda is the dominant frame for decoding the GRS150. This has, in turn, made the GRS150 the *de facto* national standard metric of an institution's success in promoting completion. The GRS150 is the touchstone for much higher education policy, which in recent years has focused on degree completion. Graduation is, for example, one of the chief concerns of university leaders; a 2005 survey of university presidents found that only budget, institutional reputation, and a shared vision for the institution were more important measures of success, and it was far more important than the other major student outcome in contemporary higher education discourse, job placement (Selingo 2013). This reflects persistent criticism of U.S. higher education for failing its students, not enough of whom receive degrees once they enroll in academic programs: Contemporary American universities have been dubbed "failure factories" (Schneider 2008) as the university jeremiad genre has exploded.

Completion is thus a high priority, and the GRS150 is the key metric toward which UVU (like most institutions) has directed its efforts toward improving completion rates. Initial work focused on collecting more effective data on completion. This led to a 3-year dashboarding project to collect individual-level completion data for all students included in a GRS cohort since 1998, allowing the data to be cross-tabulated by a wide range of demographic, academic, and institutional characteristics (Institutional Research & Information 2012). The dashboard now serves as the primary metric for one of the standards the institution uses for evaluation by its accrediting body, the Northwest Commission on Colleges and Universities, and for evaluating its student success efforts. It is also the key parameter supporting the academic success initiatives that the university is using to improve student success.

But the GRS is problematic for these purposes. Evaluation of UVU's graduation and completion metrics shows consensus on three principles: that the GRS does a poor job of representing UVU, that UVU should provide more representative measures and interpretive tools to support them, and that locally developed measures will not be accepted as supporting accountability and must be complemented by measures that allow for national comparison (University Planning Advisory Committee 2017, p. 7). It is this last consideration that is determinative. This seems quite surprising in some important ways. UVU's students are not the typical students envisioned in the GRS definition. Fewer than half of UVU's fall-entry students come directly from high school, the majority of its students enter part-time, and many transfer from other universities. Many are returning adults. Many enter in the spring term rather than in fall. As a result, only 19.9% of UVU's fall 2016 student body are part of an active GRS cohort. By the time allowable exclusions are considered, a university of more than 35,000 students that awards more than 5000 degrees and certificates annually is evaluated on the success or failure of fewer than 1000 students.

UVU is in no way unique here. The weaknesses of the GRS component have long been recognized throughout U.S. higher education. The Council of Regional Accrediting Commissions (C-RAC, presumably because someone wisely decided "crack" might not project the most favorable image), faced with increased federal pressure to scrutinize institutions with low graduation rates, announced in 2016 that regional accreditors would heighten oversight of institutions with "graduation rates" (which has been universally interpreted to mean the GRS150) below threshold values. But the council also recognized the weaknesses of the GRS, especially its exclusivity:

> Recognizing that one or two data points are insufficient to make a qualified judgment as to the educational quality of an institution, accreditors will also review additional information. This will include the number and percentage of students counted and transfer rates, in order to provide valuable and thorough context to the Federal data used for graduation rates, which sometimes reflect a very small fraction of students at an institution. (Council of Regional Accrediting Commissions 2016)

C-RAC's concern reflects a much longer history of criticism. As required by the Higher Education Opportunity Act of 2008, USED's Committee on Measures of Student Success developed a series of recommendations to support more accurate description of student success for 2-year institutions (Committee on Measures of Student Success 2011). NCES subsequently convened several technical review panels on the addition of a broader set of Outcome Measures (TRP #37, February 2012; TRP #40, October 2012; TRP #45, September 2014; TRP #50, August 2016) leading to the implementation of the Outcome Measures component to IPEDS in 2015–2016 alongside the existing GRS component.

It is surprising that UVU should rely on a metric that is widely accepted as flawed. But it is to be expected if we think of metrics as political institutions. Institutionalized metrics serve many of the fundamental purposes of political institutions:

> Routines make it possible to coordinate many simultaneous activities in a way that makes them mutually consistent. Routines help avoid conflicts; they provide codes of meaning that facilitate interpretation of ambiguous worlds; they constrain bargaining within comprehensible terms and enforce agreements; they help mitigate the unpredictability created by open structures and garbage can processes by regulating the access of participants, problems, and solutions to choice opportunities. Routines embody collective and individual identities, interests, values, and worldviews, thus constraining the allocation of attention, standards of evaluation, priorities, perceptions, and resources. (March and Olsen 1989, p. 24)

In spite of appearing "bureaucratic, rigid, insensitive, and stupid" (March and Olsen 1989, p. 24), the GRS150 serves as a common ground for understanding student success. The UVU dashboard provided a consistent methodology for measuring retention (also defined within the GRS framework) and graduation across the university by relying on a common metric that, critically, had external validation, supporting coordinated effort to improve completion. It is difficult to see how that would be possible without some kind of institutionalized metric. The difficulties of the federal Postsecondary Institution Ratings System initiative, which ultimately foundered on the difficulty of evaluating institutions' success, can be understood as an attempt to deinstitutionalize the GRS150: Institutions might agree that the GRS150 is poor, but without an institutionalized metric there was no legitimate alternative. Hence, as is often the case, actors accept the institution as the only viable course of action—the appropriate way to act under the circumstances. We can think of many alternative statistics that we might wish were institutions, but we cannot think of an alternative institution. Thus the GRS150 stands.

The use of the GRS150 in performance funding formulae is especially instructive in understanding the limits of an institution. As described below in more detail, 16 states use GRS-based metrics, predominantly the GRS150 and first-year retention rates based on the GRS cohort, in their funding formulae in spite of the well-established weaknesses in the GRS methodology described above. But those weaknesses push the boundaries of appropriateness when expanded from baseline measures of student success to allocative tools for budgeting. The logic of appropriateness that supports the GRS150 as a measure of, for example, mission fulfillment comes up against other logics of appropriateness that demand recognizing institutional diversity in higher education system management. In Utah, for instance, Utah State Board of Regents Policy R312 recognizes four different types of university roles, most of which are incompatible with a focus purely on first-time, first-year students. This institutional conflict was resolved by moving from GRS-based completion metrics in the initial performance funding formula to metrics based on completions per FTE in 2016–2017, reflecting performance funding practices in 31 states. That has not (yet) changed the GRS150's status as an institution in higher education governance, but it does establish limits to it.

The completion agenda institutionalizes a problem-model-intervention nexus in which the GRS150 is a normalizing representation of completion, completion is framed as a problem of institutional accountability, and the completion problem is solved by "acting on facts" and "evidence based educational practice" that changes institutional cultures. The interventions proposed are directed toward improving completion as measured by the "facts," i.e., GRS rates: programs that would assist primarily full-time, traditional students who have academic weaknesses. The GRS rates are thus best understood as representations not of the students but of problems defined in relation to conditions that may extend well beyond that which are of immediate concern, conceptual and empirical understandings, and the set of possible or intended interventions. At UVU, like many universities, completion is certainly a matter of student success, but it cannot be understood apart from the need to *demonstrate* that success and the expectation that it will be held accountable for it in some fashion.

This conception of the problem connects to models of student success that understand the university, rather than the student, as the chief determinant of that success though its academic and student support programming. Models focused on student attributes beyond the university's control are inconsistent with the accountability narrative. Such models do not go unacknowledged: they are frequently discussed in informal conversations, and UVU's President, Matthew Holland, notes frequently that solutions to completion that focus on recruiting better students—solutions offered frequently by vendors of learning analytics software—are the easiest way to improve completion rates but are inconsistent with the university's mission (Utah Valley University 2015), a conclusion widely supported by the campus community (University Planning Advisory Committee 2016). Thus UVU's student success and retention programs are directed toward changing both institutional practices and student characteristics. The University Planning Advisory Committee, a campus-wide planning body designed to communicate with senior leadership, has called for improved advising and changing mathematics general education requirements (University Planning Advisory Committee 2014) to improve completion. The Student Success and Retention office's Completion Plan includes an Early Alert program that allows faculty to direct students toward academic support as well as a Stoplight program to identify students at risk of withdrawing from the institution (Student Success and Retention 2017). University College offers study skills courses and a course designed to improve students' resilience. Behind each of these programs is the belief that the university has the ability to influence student behavior commensurate with its accountability for it. The problem-model-intervention nexus exists as an institution, one in which UVU adopts models of the problem and potential interventions that are consistent with the completion agenda because they are the only institutions available to decode the data into a metric.

4.3 Metrics as Determinants of Political Outcomes

The encoding and decoding processes establish metrics as both political outcomes and political institutions. But once established, metrics also take on political functions, by which I mean that they do things to carry out the routine processes of politics: They distribute material and moral goods, and they structure relationships among political actors.

4.3.1 Metrics as Distribution Mechanisms

The dominant vision of politics among Twentieth Century U.S. political scientists (and, consequentially, the most common if not quite dominant view globally) sees it as a framework for distribution of material and moral goods in society. The classic formulations are distributive: Lasswell's (1950) "who gets what, when, and how" and Easton's (1953) "authoritative allocation of values" represent a fundamental continuity across the behavioral and post-behavioral eras in political science in the United States. From these perspectives, distribution is "the major, if not sole, function of the polity," and there is special focus on the distribution of political power (Mitchell 1961) in its many forms—whether formal authority, individual rights, or "soft" forms of power. All of these, along with more mundane goods and services distributed through taxation, appropriation, and operation of government programs (or the refusal to do so), can be understood from a perspective in which the political system allocates or controls their distribution across political actors in society.

The politics of metrics is by no means exceptional in this respect. Once data is encoded and then decoded as a metric, the metric can take on many distributive roles, authoritatively allocating many different values. The simplest form of this is when metrics allocate material goods. As of 2015, 38 states used or were implementing some form of performance-based funding to allocate resources to higher education institutions. In 16 of these, GRS rates, nearly always including the GRS150, are used as part of performance funding formulae. Utah used the GRS100, GRS150, and GRS200 for completion and the GRS cohort for retention rates initially before later moving to a completions and enrollment metric. Florida also used all three GRS rates for completion and the GRS cohort for retention. Illinois, Kansas, and Michigan included the GRS150 in its formula for 4-year institutions. Indiana and Missouri used the GRS100 for completion. Montana used the GRS cohort for retention; North Carolina used it for a 12-credit hour progress rate and a success measure similar to the federal Outcome Measures. Many other states used retention and completion rates without publically specifying the GRS cohort or a GRS rate, but due to the need for national comparison it is all but certain that these states were using GRS-based metrics (National Conference of State Legislators 2015).

A change in the GRS150 definitions, then, has the potential to reallocate millions of dollars in higher education funding. One can usefully, if not perfectly, quantify this using the Utah performance funding formula and IPEDS data. In the 2016–2017 Utah performance funding cycle, UVU's target graduation rate was 39.4% (based on the 66th percentile of universities in the Carnegie Public 4-year and above, Baccalaureate Colleges—Diverse Fields classification that admitted 90% or more of applicants), and its graduation rate for the students formally part of the performance funding process (basically, the GRS150 for the cohort that entered in Fall 2008 and graduated by Summer 2014) was 27.8%. As a result, UVU forfeited $191,145 in performance funding for graduation efficiency (Buhler 2015). Based on data for the Outcome Measures, using the 8-year graduation rate for all students rather than the GRS150 would have increased UVU's combined graduation rate for all award levels from 32.0% to 35.7%.

Directly comparable data for the peer universities used in the performance funding formula is not available, because Outcome Measures are not yet publically available. So it is impossible to know what the effect of this change on the target graduation rate—and thus on funding awards—would be precisely. But with 10 of UVU's peer universities having nominally selective admissions processes and five not offering associate's degrees, it seems reasonable to conclude that UVU would gain significantly compared to its peers in the performance funding formula. It has a slightly higher percentage of transfer students among its new student population than the average of its peers (35.0% versus 32.0%, respectively, in Fall 2013, the data year used in awarding 2015–2016 performance funding) and a substantially higher percentage of part-time students (38.2% compared to 19.5%). If UVU closes 25% of the gap between it and its peers by a change in the metric to an all-student Outcome Measure (equivalent to a 2.9 percentage point improvement in the GRS150), it gains $48,500—about half the cost of a new faculty line.

A more straightforward switch from the GRS150 to the GRS200—a change of a single element in the definition of the graduation rate metric—would still favor UVU given its large numbers of part-time students (many of whom entered as full-time students and are thus included in the GRS cohorts). Sixth-year completers make up a noticeably higher proportion of UVU's GRS150 completers than among the performance funding peers (22.2% against 18.6%), so one expects that the GRS200 for UVU's 2007 cohort will continue to improve relative to its peers, on average as the time threshold increases. GRS200 data for the Fall 2007 cohort was only reported in winter of 2016–2017, and is not yet available publically through IPEDS, so one cannot make a direct comparison to actual FY 2015–2016 funding. But using data for the 2006 cohort, which was available for 2015–2016 performance funding awards, rather than the 2008 cohort tells a dramatic story. Based on the GRS150 rates for the 2006 cohort, UVU would have only received 43.0% of its potential award, amounting to $283,524, a loss of $181,751 compared to the actual award based on the 2008 cohort. This reflects a very poor GRS150 rate of 16.9% for what was then Utah Valley State College. The GRS200 is another matter. While UVU's peers gained on average 2.8 percentage points moving from the GRS150 to the GRS200 for that cohort, UVU gained 11.1 points. Switching to the GRS200 nets UVU an additional $147,366 compared to the actual award. A hypothetical such gain for the 2008 cohort—3 points for peers and 11 points for UVU—would net an additional $138,105.

Multiplied by dozens of institutions in 16 states, the implications of defining timely graduation as 150% or 200% of program time constitute a significant reallocation of higher education funding across institutions. All of these speculations became moot, of course, when USHE changed its performance funding metric for graduation efficiency from graduation rates to degrees and certificates awarded per FTE. The new metric is favorable to UVU, which awarded 26.4 degrees or certificates per 100 FTE in 2016–2017, thanks to serving large numbers of transfer-in and part-time students who aren't counted in GRS-based graduation rates, and will thus likely result in UVU making good on a larger portion of its potential performance funding awards. One might even suspect that awards per FTE is well on its way to institutionalization itself based on its more favorable evaluation across higher education and frequent use in performance funding formulae. But in either case, decisions about metrics are decisions about policy. These speculations show that making decisions about metrics, whether big differences between metrics measuring significantly different constructs or small differences in the selection and design of metrics for an established construct, is a way of performing one of the most fundamental policy actions: allocating resources to government functions.

Metrics often allocate more than material goods. They often allocate moral goods as well, for example, rights or recognition. By institutionalizing operational definitions of the groups of concern in the problem-model-observation nexus, metrics allocate recognition, legitimacy, and participation rights. Those that are within the operational definition are recognized as having a legitimate place in the nexus and can make a claim to participate in the process or to receive a benefit. Those who are outside of that definition are not necessarily consciously excluded from participation but the metric guides those within the nexus to each other, and to fail—or sometimes refuse—to recognize those outside of the nexus as being important to the process. This is not a consequence of the constructs themselves. Often a conceptual definition would include a far broader range of participants than are included when the constructs are operationalized in metrics.

UVU offers a number of scholarships to support completion. The general principles behind the programs are to provide scholarships to students within 1 year of completing their degrees in order to ensure that students do not fail to graduate when they have completed most of their programs. The completion scholarships are, at the heart, a low-hanging fruit strategy: Improve completion rates by intervening where it can most readily make a difference in outcomes. By intervening in the final year, the completion scholarships act when there are fewer possibilities for extraneous factors that prevent the intervention from producing the expected outcome.

Not all students are eligible for these scholarships. There are three main completion scholarship programs at UVU: first-generation completion grants, summer completion grants, and Wolverine Completion grants. The first two reflect aspects of the problem and intervention; first-generation students have a lower than average graduation rate, and low summer utilization provides an opportunity for students to take more courses before the GRS150 deadline. This is, of course, a textbook case where the metric defines the problem (graduation in less than 150% of program time) and intervention (accumulate more credits before time expires). The general structure of the completion grants programs is built around a GRS-driven operationalization of completion found in the UVU Completion Plan (Student Success and Retention 2017) and its associated implementations such as the Student Success and Retention data dashboard (Institutional Research & Information 2012). The GRS150 also defines eligibility directly for the Wolverine Completion grants, which UVU describes as a "[f]inancial aid program created specifically for students in IPEDS [GRS] cohorts who have completed 90 credits or more toward a bachelor's degree" (Taylor 2016, p. 61).

Certainly, as scholarships, these are additional cases in which the GRS150 allocates material resources. But they also reflect ways in which the GRS150 allocates recognition and legitimacy. Those who are included in the GRS-based measures have a priority claim to participate in conversations about completion. Eligibility for the GRS cohort—and thus the capability of contributing to UVU's GRS150 rate—confers recognition that a student is or is not part of the completion problem. Students whose completion is consistent with the GRS framework—those for whom "on-time completion" is a meaningful goal as opposed to ongoing progress that will lead to a degree eventually—have standing to be represented in the completion dialogue. They are incorporated into solutions to completion, while students who face completion challenges that are not reflected in the GRS framework—students who entered UVU in the spring or as part-time students—stand as also-rans, students whose completion is a good thing but not among the university's priorities.

Metrics also shape political power through the control of information. Metrics are, in essence, institutionalized information. This makes them distributors of political power by institutionalizing and giving consequence to information asymmetries. Those with access to information institutionalized within a logic of appropriateness may have greater independence of action: Some have authority to produce information, and their information is considered appropriate and included in the problem-model-intervention nexus. Others' information is not: it is at best seen as contextualizing, often dismissed as anecdotal, sometimes considered false simply because it conflicts with the "official" metric. As a result, those favored by the information asymmetries have a greater range of options available to them within the logic of appropriateness and have a greater role in the process. One cannot deny those who hold the key information a place in the process of intervention, and those with the information can withhold or provide it to their own advantage.

This was very much the case in the controversy over Mount St. Mary's University's plan to dismiss some students before the IPEDS reporting deadline in 2016, made infamous by then-President Simon Newman's comment about the students that "sometimes you have to drown the bunnies" (Svrluga 2016).[4] This was a project to manipulate the GRS-based retention rate by dismissing students who were unlikely to be retained to the following fall before the reporting date so they will not be included in the GRS cohort (and thus in the denominator of the GRS retention rate). Newman intended to use a survey administered at the orientation for entering students to identify those to be dismissed. Strong opposition from the faculty delayed analysis of the survey and implementation of the plan to dismiss students beyond the IPEDS reporting date, however. This ultimately empowered the university's institutional researchers to submit a cohort that had not been affected by the survey process in order to comply with the GRS' reporting date definition. That date, an element of the metric's definition, shifted decision-making power at Mount St. Mary's University from the president to the office controlling the information and reporting process: A major initiative was thwarted not because of authority relationships in the university's formal hierarchy but by the definition of the metric.

4.3.2 Metrics in the Politics of Control

While distribution is a central theme in contemporary politics, it is not the only way in which metrics influence political outcomes. Metrics are elements of political control, shaping processes in ways that are at best partially understood (and often deeply misunderstood) as simple allocative measures. Metrics often function as part of quasi-algorithmic procedures designed to control decision-making by a political actor. One of the fundamental insights that political science made into the policy process during the late Twentieth Century was that such procedures are very much instruments of political control (McCubbins et al. 1987). Famously enough to become a mononym, McNollgast[5] showed that Congress ensures that the bureaucracy executes the law within the scope of legislative intent less through the textbook means of punitive budgeting, advice and consent to appointments, and

[4] This case is discussed in more detail in Sect. 6.1.1.

[5] And thus to treat themselves as a non-gendered single author on its web site, e.g., "McNollgast is most well-known for its early articles that helped introduce positive political theory (PPT) into the study of administrative law" (Weingast 2013).

statutory revision than by designing administrative procedures that overcome bureaucrats' advantages in information asymmetry and allow favored groups to assert their interests directly to bureaucrats. This achieves the most common form of legislative intent, which is less to produce legislatively favored policy outcomes than to produce policy outcomes preferred by legislatively favored constituents.

While McNollgast focused largely on the administrative procedures used by bureaucrats to make implementing regulations (e.g., the procedural requirements of rulemaking under the Administrative Procedures Act of 1946) and information disclosure requirements such as the Freedom of Information Act, it is not at all difficult to understand required metrics as having the same effect. No metric is universally useful, and (in principle) no rational representative chooses a metric without considering how their favored constituencies will fare under such a metric. Metrics should be seen, then, as *prima facie* supporting the needs of specific constituencies. By designing metrics into legislation, as Congress did in the Student Right-to-Know Act, legislators constrain the behavior of agencies such that an agency will, on its own, secure the favored outcome.

The Congressional intent of the GRS150 is not hard to understand. The language of informed choice, consumer information, and on-time graduation combines with the specification of first-time, full-time students in the Student Right-to-Know Act to make clear that required reporting—and policy analysis by the Department of Education based on the data reported—is intended to favor traditional students. By setting the GRS150 as an essentially national standard, Congress can be seen as requiring that institutions pursue policies that favor (or at least meet the needs of) those traditional students regardless of their effects on other students, about whom neither USED nor Congress know much. The GRS150 remedies an information asymmetry by eliminating the additional information available to institutions (about students not in the GRS cohort) from consideration in the policy sphere. Graduation rates for part-time transfer students are (at least until the Outcome Measures gain currency) not standardized; an institution's report of them is anecdote rather than national data and thus not part of the graduation rate.

Universities must respond to the needs of the educated, middle-class voters favored by Congress because universities must report a graduation rate specific to their needs. No matter what an institution can show about completions per FTE or extended graduation rates (for example, the Consortium for Student Retention Data Exchange collects data out to 12 years), the traditional families of traditional students who expect a traditional university experience are in a position to ask, "But what about your graduation rate?," confident (or even without considering to the contrary) that "graduation rate for people like my child" is implicit in the question. Hence comes the widely held belief that UVU must respond to the GRS150 however unrepresentative it is (reinforcing its status as a social institution, one notes), and the prevalence of news stories about returning adult graduates—"man bites dog," one suspects—but scholarship programs for students in the GRS cohorts.

When viewing this from the relationship between legislatures and (potentially) favored constituencies, this remains a distributive form of politics: Favored groups are given opportunities to intervene in the political process that other groups lack.

The key here is that this distributive perspective fails to understand the relationship between the bureaucracy and the legislature. This is neither an allocation of power across agencies nor between the legislature and an agency, though certainly those dynamics can be present as well. This is a relationship between master and servant, in which the master claims that all actions of the servant take place with the master's authority and acts to constrain the servant not by limiting its power but by executing that power in specific ways. UVU is not given a range of powers by Congress, NCES, NWCCU, or USHE beyond which it cannot go. It retains the full authority to make policies prioritizing whatever students it wants to prioritize. But it is told by an agency with the power to effectively terminate the institution's operations (by denying it eligibility to offer federal financial aid) that it will be subject to "special attention" if the graduation rate for traditional students falls below 25%; whether the GRS150 is an appropriate measure for the institution or not, the university will be required to provide "information about the conditions that may have led to low graduation rates and how the institution is working to improve" (Council of Regional Accrediting Commissions 2016). This is a question about outcomes, not authority. It should not be understood from a purely distributive perspective.

The political power of metrics can go well beyond their immediate use. One important determinant of policy success is the capacity of the political system to make and implement policy:

> Decisions made by governments cannot always be carried through; there is no law guaranteeing that government authorities will attempt only those interventions that they really can execute. The administrative organization of government is crucial, especially when policies calling for increased government intervention are to be implemented. Governments that have, or can quickly assemble, their own knowledgeable administrative organizations are better able to carry through interventionist policies than are governments that must rely on extragovernmental experts and organizations. (Skocpol and Finegold 1982, pp. 260–261)

This state capacity enables governments to act effectively in areas where capacity already exists, but seriously handicaps governments where they need to build new capacity. Hence, Skocpol and Finegold demonstrate the success of agricultural policy and the failure of industrial policy in the United States during the New Deal. A long history of agricultural policy tied especially to the place of land grant colleges gave the federal government substantial capacity in agriculture that it lacked in industry, where the government was handicapped by lack of a skilled bureaucracy and had to draw on businesses themselves for expertise.

State capacity does not come solely from an organization chart, however. Much of the advantage the United States had in agricultural policy in the 1930s came from an established knowledge base developed through schools of agriculture (and especially agricultural economics, which arose as a separate discipline from general economics and was more strongly rooted in institutionalist research traditions). A political system's success relies not just on administrative organization but on administrative knowledge. An agency (or constellation of agencies) that can act coherently and effectively but does not know what to do and how to do it is no more likely to successfully make and implement policy than one that is being built from nothing. This is why making the world legible to policy, as discussed in the previous chapter, is so central to the entire project of data-driven management: Metrics allow political systems to intervene knowledgably and to overcome resistance from civil society, thereby intervening more effectively and accomplishing its policy goals. Metrics create state capacity and thus shift power to the state.

By establishing the GRS150, USED was able to control the political landscape of completion. It is able to determine which institutions are performing “well” or “poorly” and to act accordingly, as it has done with the College Scorecard, for example. The department is first able to identify a concern with completion that it could not have seen (at least not in a purportedly objective way) without a graduation rate metric. That data allowed the department to establish completion as a policy problem—it allowed USED to set the political agenda—and to extend the terms of that problem beyond being one of consumer information and choice (as described in the Student Right-to-Know Act) into what would become the completion agenda, connecting with actors in civil society and in Congress that would support such an agenda in a classic “iron triangle” relationship. It could then bring pressure on groups such as C-RAC to enforce a graduation rate standard that C-RAC itself acknowledges is not a sound basis for evaluating institutional effectiveness.

USED quite simply could not do this without the GRS150; the concept “graduation rate” without the associated metric is too nebulous terrain for USED to act effectively. Its success when using the GRS150 is usefully contrasted with the effort to build the Postsecondary Institution Rating System (PIRS). Following the decision that the GRS150 would not be the exclusive metric for completion (U.S. Department of Education 2014), the completion metric became one of the most contested issues in the system, ultimately leading the Obama administration to drop the initiative entirely. PIRS moved outside of the state capacity that had been built with the GRS150, ceding ground to actors in civil society (especially the for-profit higher education sector) and fragmenting policymaking capacity by opening rifts between Congress and the executive branch.

4.4 Conclusion

This study of the GRS150 shows that information exists both as the object of politics and as a force in politics. The encoding processes by which raw information about the environment becomes data and the decoding processes by which data becomes a metric suited for use in a nexus of problem, model, and intervention mean that we cannot consider data simply an objective, apolitical solution to politics. Just as the choice to use data over other approaches—anecdote or interpretation or pragmatism or revelation—is a political act, the choices to encode observations in one data frame rather than another and to decode data through one statistical methodology rather than another is an act of politics as well. And those political acts have political consequences, distributing resources, allocating legitimacy, controlling decisions, and building capacity. This is not because data is politicized by actors who seek advantage from it, but because data is inherently[6] political. There simply is no such thing as apolitical information.

And if data is political, if it is not objective, then it is no longer a purely technical question. Isaiah Berlin (1979) distinguishes between questions that can be resolved through observation and formal reasoning—technical questions—from those philosophical questions that cannot. But unlike the logical positivists who held the latter to be meaningless, Berlin shows that the absence of answers does nothing to delegitimize the questions. "Which students should be included in the graduation rate?" is not a trivial or meaningless question for its lack of a technical answer. While there may not be a demonstrably correct answer to the question, there are surely better and worse answers, and we can give reasons for choosing one answer over another and reasonably hold ourselves and others accountable for choosing an answer. We can rely on a preference for reliability over validity, for example, arguing that a consistent standard for graduation rates is fairer to institutions and students than arbitrary decisions, and we can criticize the use of the GRS150 for its concern for the already privileged over the least-well-off in society. Dismissing these questions as if they don't exist, as if they are "nonsense upon stilts" or, as is more common, as extraneous to the model, does not make them go away. Our information choices are answers to these questions whether we ask them or not.

But dismissing the political questions about data does obscure them, and prevents us from seeing the political and ethical consequences of choices we pretend not to make. Nothing is less ethical than to pretend there are no ethical questions to be answered. We must, therefore, open the consideration of what the data should be to a political perspective. That means that we must subject information to analysis as a matter of justice.

References

American Association of Community Colleges. (2010, December 1). *Sample completion commitment statement*. https://www.aacc.nche.edu/About/completionchallenge/Documents/CalltoAction_Writable.docx. Accessed 3 Dec 2015.

Berlin, I. (1979). Does political theory still exist? In H. Hardy (Ed.), *Concepts and categories: Philosophical essays of Isaiah Berlin* (pp. 143–172). Oxford: Penguin Books.

[6] One is tempted to say "always-already," as is the current fashion. But that is not quite right. The central point of the continental formulation is to suggest a point after which it is impossible to conceive of the time before: Humans are always-already linguistic because the only way we can frame a time before language is to use language. We are thus already linguistic in the present (having acquired language at some point in the past and not needing to do so now), but that acquisition must appear to have always been the case because there is no possibility of understanding what was before. That is not how data is political. Data has not become politicized in such a way that we can never again understand its pre-political state. Data is political from the moment it comes into existence. Data is thus always-but-not-already political because it is impossible to create data *at any time* without engaging in politics.

Buhler, D. L. (2015, July 22). *2015–2016 USHE performance funding model and allocations*. https://higheredutah.org/pdf/agendas/201507/TabR.pdf. Accessed 9 May 2017.

Committee on Measures of Student Success. (2011, December). *Committee on measures of student success: A report to secretary of education Arne Duncan*. United States Department of Education. https://www2.ed.gov/about/bdscomm/list/cmss-committee-report-final.pdf. Accessed 9 May 2017.

Council of Regional Accrediting Commissions. (2016, September 21). *Regional accreditors announce expanded review of institutions with low graduation rates*. http://www.nwccu.org/Pubs%20Forms%20and%20Updates/Forms/PDF%20Files/Press%20Release%20FINAL%209-16-161.pdf. Accessed 9 May 2017.

Easton, D. (1953). *The political system: An inquiry into the state of political science*. New York: Knopf.

Fuller, C. (2011). *The history and origins of survey items for the integrated postsecondary education data system (No. NPEC 2012–833)*. Washington, DC: National Postsecondary Education Cooperative. http://nces.ed.gov/pubs2012/2012833.pdf. Accessed 28 Oct 2015.

Hall, P. A., & Taylor, R. C. R. (1996). Political science and the three new institutionalisms. *Political Studies, 44*(5), 936–957.

Hall, S. (2006). Encoding/decoding. In M. G. Durham & D. M. Kellner (Eds.), *Media and cultural studies: Keyworks* (Revised ed., pp. 163–173). Oxford: Blackwell.

Institutional Research & Information. (2012). *Student success/retention*. http://www.uvu.edu/iri/indicators/. Accessed 12 Mar 2014.

Lasswell, H. D. (1950). *Politics: Who gets what, when, how*. New York: P. Smith.

March, J. G., & Olsen, J. P. (1989). *Rediscovering institutions: The organizational basis of politics*. New York: The Free Press.

McCubbins, M. D., Noll, R. G., & Weingast, B. R. (1987). Administrative procedures as instruments of political control. *Journal of Law, Economics, and Organization, 3*(2), 243–277.

McMillan Cottom, T. (2017). *Lower ed: The troubling rise of for-profit colleges in the new economy*. New York: The New Press.

Mitchell, W. C. (1961). Politics as the allocation of values: a critique. *Ethics, 71*(2.) https://www.jstor.org/stable/2379509. Accessed 7 May 2017.

National Center for Education Statistics. (n.d.-a). *The integrated postsecondary education data system—Glossary*. http://nces.ed.gov/ipeds/glossary/. Accessed 11 Nov 2015.

National Center for Education Statistics. (n.d.-b). Outcome measures. *IPEDS 2015–16 data collection system*. https://surveys.nces.ed.gov/ipeds/VisNextYearForms.aspx?year=3&survey=13&form=101&nextYearForm=101&index=0&ri=0&show=all&instid=30101. Accessed 11 Nov 2015.

National Conference of State Legislators. (2015, July 31). *Performance-based funding for higher education*. http://www.ncsl.org/research/education/performance-funding.aspx. Accessed 7 May 2017.

Schneider, M. (2008). *The costs of failure factories in American higher education*. (No. Education Outlook Number 6. Washington, DC: American Enterprise Institute. https://www.aei.org/publication/the-costs-of-failure-factories-in-american-higher-education/. Accessed 26 Nov 2015.

Selingo, J.L. (2013). *What presidents think: A 2013 survey of four-year college presidents*. *The chronicle of higher education*. http://www.maguireassoc.com/wp-content/uploads/2015/08/Chronicle-Presidents-Survey-for-Education-Counsel-2.pdf. Accessed 11 Nov 2015.

Skocpol, T., & Finegold, K. (1982). State capacity and economic intervention in the early New Deal. *Political Science Quarterly, 97*(2), 255–278.

Student Success and Retention. (2017, March). *UVU completion plan*. Utah Valley University. https://www.uvu.edu/retention/docs/uvu_completion_plan_march2017.pdf. Accessed 10 May 2017.

Svrluga, S. (2016, January 19). University president allegedly says struggling freshmen are bunnies that should be drowned. *The Washington Post: Grade Point*. https://www.washingtonpost.com/news/grade-point/wp/2016/01/19/university-president-allegedly-says-struggling-freshmen-are-bunnies-that-should-be-drowned-that-a-glock-should-be-put-to-their-heads/. Accessed 17 Nov 2016.

Taylor, M. (2016, November 7). *Student affairs PBA*. https://www.uvu.edu/pba/docs/2016sa.pdf. Accessed 9 May 2017.

University Planning Advisory Committee. (2014, October 9). *Meeting minutes, October 9, 2014*. http://www.uvu.edu/insteffect/docs/upac_min_2.25.16.pdf. Accessed 25 May 2017.

University Planning Advisory Committee. (2016, February 25). *Meeting minutes, February 25, 2016*. http://www.uvu.edu/insteffect/docs/upac_min_2.25.16.pdf. Accessed 25 May 2017.

University Planning Advisory Committee. (2017, February 23). *2016–2017 Mission fulfillment self-evaluation*. Utah Valley University. http://www.uvu.edu/insteffect/docs/2017_upac_self_evaluation.pdf. Accessed 25 May 2017.

U.S. Department of Education. (2014, December 19). *For public feedback: A college ratings framework*. https://www2.ed.gov/documents/college-affordability/college-ratings-fact-sheet.pdf. Accessed 12 May 2017.

Utah Valley University. (2015). *UVU: State of the University—February 10, 2015*. Orem, Utah. https://www.youtube.com/watch?v=uZrp_iUNVkY. Accessed 25 May 2017.

Weingast, B.R. (2013, October 3). *McNollgast*. https://web.stanford.edu/group/mcnollgast/cgi-bin/wordpress/mcnollgast-2/. Accessed 11 May 2017.

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

Chapter 5
Distributive Information Justice (And Its Limits)

Abstract In this chapter, I seek to go beyond contemporary theories of information privacy by subjective the standard information flow models to analysis from the perspective of justice. I examine two perspectives. At the least, one can see privacy as connected to justice instrumentally, that is, privacy is valuable not as a requirement of justice directly but because it is a useful means of justice. This is, I argue, hardly adequate as an entire theory of information justice but it is too easily given short shrift in discussions of privacy (especially by the wealthiest Silicon Valley titans who can protect their interests more directly). A more robust approach looks to theories of distributive justice. Theories of distribution that focus on the distributive process can address two significant weaknesses in information flow models of privacy, weak conceptions of informed consent and the inability to address the original acquisition of information. Pattern theories of distributive justice shift the focus from distributing information to distributing privacy rights, and provide significant insight into what it means to have rights to be left alone or forgotten. Each of these theories makes useful contributions to our understanding or privacy. But they are not wholly adequate to the task; for this, one needs to understand justice structurally as well as distributively.

One of the few strains of coherence in the study of information privacy[1] is acceptance of its incoherence. Woodrow Hartzog and Evan Selinger argue that in both academic theory and U.S. law, "despite the widespread concern and extensive academic treatment of surveillance issues, the language and framing used in surveillance debate is diverse, inconsistent, and over-generalized" (2015, p. 1344); H. Jeff Smith, Tamara Dinev, and Heng Xu suggest that in information systems research "the findings and the theories that emerged have often relied on overlapping

[1] For the purpose of this paper, I will adhere, with some modification as this paper develops, to the distinctions among general privacy, physical privacy, and information privacy offered by Smith et al. (2011, pp. 990–991). *Physical privacy* refers to the privacy of "an individual and/or the individual's surroundings and private space" where *information privacy* refers to the privacy of information about one or more individuals, whether as individuals or as a group. General privacy refers to both. Unless otherwise specified, I will use *privacy* as a shorthand for information privacy, and specify when I refer to physical or general privacy. However, I depart from the authors above in confining privacy in any form to issues of access, for reasons that will be developed below.

J.A. Johnson, *Toward Information Justice*, Public Administration and Information Technology 33, https://doi.org/10.1007/978-3-319-70894-2_5

constructs nestled within loosely bounded nomological networks. This has resulted in a suboptimal cumulative contribution to knowledge" (2011, p. 990). Daniel Solove's pronouncement has an air of definitiveness[2] about it: "Privacy, however, is a concept in disarray. Nobody can articulate what it means" (2008, p. 1).

Information privacy has yet to be investigated explicitly as a question of justice. To do so, of course, requires one or more underlying theories of justice itself and thus provides a focus for developing a more general theory of information justice. That is, having developed a theory of justice that can engage privacy effectively, we will have (hopefully) found key insights to the problem of information more generally. In particular, we can understand the relationship between information and the major forms of justice. Indeed, Schlosberg's success in environmental justice comes by recognizing that justice exists in multiple forms and that our knowledge of justice flows from both abstract theory and social practice. I suggest three main ideas of justice in information privacy: an instrumental view in which information privacy is valuable for its contribution to justice in other spheres, the common distributive paradigm that forms the basis for much contemporary work on information privacy, and an enhanced structural form composed of fully integrated recognition and participatory dimensions. The instrumental and distributive approaches will prove to have significant value in understanding privacy but also significant limitations, especially in light of criticisms offered by the third approach, that suggest that information privacy is most likely to be engaged effectively by a structural approach to justice.

5.1 Instrumental Justice in Information Privacy

Information privacy may be related to justice only tangentially. One might ask what harm there is simply in having information; from a philosophical perspective, this is tantamount to asking if information privacy is a requirement of justice in itself or simply useful in the pursuit of justice in other forms. The latter is certainly a common attitude within the information technology industry, where many leaders have happily discarded information privacy as a concern. Facebook founder and chief executive Mark Zuckerberg has observed that information privacy is no longer a "social norm" as expansion of the internet has made people more comfortable with sharing information and more open about themselves: "That social norm is just something that has evolved over time," Zuckerberg believes (Johnson 2010). At the 2015 World Economic Forum in Davos, Switzerland, computer science professor Margo Seltzer stated, "How we conventionally think of privacy is dead" because "Privacy as we knew it in the past is no longer feasible"; tech entrepreneur Anthony Goldbloom observed, "I trade my privacy for the convenience. Privacy is not something that worries me " (Carter 2015). All of these suggest that privacy is not morally valuable in itself: it can be traded for convenience, abandoned because it is impractical, evolve out of existence.

[2] A false air, or course, coming at the beginning of a book Solove begins by saying "I am ready to set forth my theory of privacy" (2008, p. ix).

That is not to say, however, that privacy is irrelevant to justice. One could argue that, even if information privacy can be abandoned without inherent injustice, it should be maintained because it is a valuable tool for protecting justice. The harm caused by violations of one's privacy, from such a perspective, is not in the information transfer itself but in the (actual or potential) effects of the information transfer on the achievement of justice; holding private information about another is not unjust, but what one does with the information may be exceptionally so. Solove, for instance, finds arguments about the intrinsic value of privacy problematic, but concludes, "The value of ameliorating privacy problems lies in the activities that privacy protections enable" (Solove 2008, p. 85).[3] Certain kinds of injustice, one might with good reason argue, require information about the objects of action. Privacy is to be maintained as a way of preventing those who would use information unjustly from having one of the necessary conditions for the implementation of injustice. Privacy is thus, from this perspective, of solely—but by no means inconsiderable—instrumental value.

This is the approach to information privacy at work in employment discrimination law in the United States. It is generally considered a best practice for potential employers not to solicit information from job candidates regarding membership in "protected classes," those groups protected by the body of employment discrimination law. Those include "race, color, religion, sex, or national origin" under Title VII of the Civil Rights Act of 1964, persons over 40 years old under the Age Discrimination in Employment Act of 1967, those with disabilities under the Americans with Disabilities Act of 1990 (ADA), and genetic information under the Genetic Information Nondiscrimination Act of 2008 (GINA). The practice of restricting the flow of information about membership in protected classes can be seen as a fairly typical information privacy protection.

Federal employment discrimination law in the United States varies widely in its restrictions on collecting information. The Civil Rights Act does not, in fact, prohibit transfer of information about protected classes at all. The essential restrictions on employers under the Civil Rights Act are:

> (a)(1) to fail or refuse to hire or to discharge any individual, or otherwise to discriminate against any individual with respect to his compensation, terms, conditions, or privileges of employment, because of such individual's race, color, religion, sex, or national origin; or
>
> (2) to limit, segregate, or classify his employees or applicants for employment in any way which would deprive or tend to deprive any individual of employment opportunities or otherwise adversely affect his status as an employee, because of such individual's race, color, religion, sex, or national origin. (42 U.S.C. § 2000e-2)

This restriction is repeated nearly verbatim with exception of the description of the protected classes in question for the Age Discrimination in Employment Act (29 U.S.C. § 623), the ADA (42 U.S.C. § 12112), and GINA (42 U.S.C. § 2000ff-1). Indeed, many employers routinely collect information about membership in

[3] To be clear, Solove does not support a strictly instrumental theory of privacy, arguing rather than privacy is a *sine qua non* for certain kinds of activities that are of the essence of justice. He does nonetheless endorse the idea that privacy is of value even if it is not a question of justice itself.

protected classes from job applicants but isolate that information from hiring decisions. The advice against soliciting such information with other hiring information, and thus to further the protection of job applicants' privacy, should be seen as a protection of justice for both the applicant (from unlawful discrimination) and the employer (from unfounded claims of unlawful discrimination): employers cannot discriminate on the basis of a condition about which they have no knowledge.

The ADA and GINA do include provisions barring the collection of information itself. The ADA prohibits employers from:

> (6) using qualification standards, employment tests or other selection criteria that screen out or tend to screen out an individual with a disability or a class of individuals with disabilities…; and
>
> (7) failing to select and administer tests concerning employment in the most effective manner to ensure that, when such test is administered to a job applicant or employee who has a disability that impairs sensory, manual, or speaking skills, such test results accurately reflect the skills, aptitude, or whatever other factor of such applicant or employee that such test purports to measure, rather than reflecting the impaired sensory, manual, or speaking skills of such employee or applicant…. (42 U.S.C. 12112(a)(6)–(7))

Moreover, subsection (d) extends the anti-discrimination provisions of subsection (a), which beyond the paragraphs quoted above generally mirror those of the Civil Rights Act, to include collecting information through medical examinations. Employers are specifically barred from "conduct[ing] a medical examination or mak[ing] inquiries of a job applicant as to whether such applicant is an individual with a disability or as to the nature or severity of such disability" (subparagraph (d)(2)(A)); a similar provision (subparagraph (d)(4)(A)) protects employees.

All of these provisions, however, include exemptions allowing the collection of information about the ability of a person to perform job-related functions. In the case of employment discrimination, the use of tests is permitted where "the standard, test or other selection criteria, as used by the covered entity, is shown to be job-related for the position in question and is consistent with business necessity" (paragraph (a)(6)) and where impaired skills "are the factors that the test purports to measure" (paragraph (a)(7)). Similarly, within the restrictions of section (d), an employer "may make preemployment inquiries into the ability of an applicant to perform job-related functions"; may require exams before an employee begins work and make offers condition on the results of those exams; may require medical examinations for current employees where "such examination or inquiry is shown to be job-related and consistent with business necessity"; and may conduct voluntary exams as part of benefits programs or that evaluate employees' abilities to perform their jobs. All of these exceptions are subject to the requirement that such exams are administered regardless of disability, maintained as a separate and confidential medical record, and used only as permitted elsewhere under the ADA.

GINA similarly makes it unlawful to "request, require, or purchase genetic information with respect to an employee or a family member of the employee" (42 U.S.C. § 2000ff-1(b)). But it, too, contains a lengthy list of exceptions that allow employers to collect such information for a number of non-discriminatory circumstances: "inadvertent" requests for family history, the operation of wellness programs (with protections for consent, confidentiality, and anonymity), certifying family leave requests, publicly available information that includes family history, monitoring workplace safety (again with protections for purposes, consent, confidentiality, and anonymity), and the unique needs of forensic DNA analysis labs. The information collected under these exceptions is considered a confidential medial record, protected from disclosure and to be kept separate from employment records (42 U.S. Code § 2000ff–5).

The most stringent protections of information privacy in the employment discrimination regime are thus shown to be the most clearly instrumental: information privacy is valuable not in itself but as a protection from unjust discrimination. Information privacy is a question related to justice in employment discrimination; the latter is a matter of justice and so practices that further or inhibit discrimination raise questions of justice. But information privacy is instrumental to justice in employment and not a requirement of it. A government of Madisonian angels could be trusted with such information without undermining justice, a question that is moot if mere possession of the information is unjust. Among mere mortals, medical and disability information can be used where it is relevant to the ability to perform job duties given reasonable accommodations, and genetic information can be used for the administration of benefits or protecting workers from environmental hazards. Employment discrimination law does not see the possession or transfer of information as the harm, but as the means to harm. The remedy is to limit the possession and use of information such that it can only be used in the interests of justice.

This would explain the comfort of many information technology evangelists with the loss of privacy. It is not that they necessarily see privacy as inherently without value. But it is not an inherent right or essential feature of justice. It is valuable just to the extent that it is useful to protect those social features that are seen as inherent rights or essential features of justice. If information privacy is no longer feasible or has evolved out of existence, then one must look elsewhere for such protections—a belief that their own virtue upholds a system of perfect meritocracy, perhaps. If a calculus of utility overwhelms the protective value of information privacy with conveniences gained by giving it up, then it can be cast aside without moral loss—at least for those with the power to influence decisions about a society's information privacy practices. It is easy to be comfortable with the idea that privacy is dead even as one bemoans "the dawn of the age of genetic McCarthyism," as historian of science Sophia Roosth told the Davos panel, when one takes an instrumentalist view in which the central moral consideration behind privacy is that "By and large, tech has done more good than harm," as Seltzer argued (Carter 2015).

To be sure, it is hard to argue that information privacy does not have instrumental value aside from any claims about its intrinsic value. To take but one recent example, Terrell et al. (2016) found that women contributing to GitHub, an online software repository that provides code management and version control, are less likely to have their code accepted in software development projects than men when the contributors are outside of the project and their gender is known, but more likely to have code accepted than men when their gender is not identifiable. The privacy of the code contributors in this case clearly supports more equitable outcomes in the open source software development community. One might further argue that the instrumental dimension of justice in information privacy is underappreciated by the theories of privacy that are analyzed below, as Schlosberg argues in his criticism of theories of justice generally. An understanding of the necessary, sufficient, and contributory social conditions for justice should be as much a part of its study as the abstract principles of it; information privacy is not trivialized by suggesting that it has value for more than its own sake. But this is not at all to say, as those like Zuckerberg seem to, that privacy is valuable only instrumentally, that it can be disregarded if justice can be achieved by other means or if does not further justice. A justice-centered view of information privacy—especially one that takes the practice of privacy as seriously as the philosophy of it—certainly must consider the instrumental value of privacy but, as the following two sections demonstrate, this view is inadequate in itself.

5.2 Distributing Information Privacy

In philosophy, law, technology, and business practice, contemporary approaches to privacy are predominantly built around controlling access to information. Smith et al. characterize privacy concerns as:

> grounded in the growing "art of the possible" in the technological realm. The spread of ubiquitous computing and the seemingly unbounded options for *collecting, processing, distributing, and using* personal information trigger consumer worries,

and define information privacy strictly in terms of "*access* to individually identifiable personal information" (2011, p. 990, emphasis added). Similarly, Lita van Wel and Lambèr Royakkers argue, "Informational privacy mainly concerns the control of information about oneself. It refers to the ability of the individual to protect information about himself" (2004, p. 130); implicitly, one controls others' access to information about one's self and protects information about one's self from being disclosed to others. To frame information privacy as a matter of access is thus to make it a matter of the distribution of information, and thus justice in information privacy is a matter of distributive justice.

Distributive justice, the dominant philosophical framework for understanding justice, is focused on the just distribution of material and social goods. It is telling that the two major collections reviewing the state of social and political philosophy over the past 25 years (Gaus and D'Agostino 2013; Goodin and Pettit 1993) both include chapters or sections on specifically distributive justice but not on other forms, as if this is the only form that justice takes. Kolm (1993) holds that questions of distributive justice arise when the issue is how to arbitrate among competing claims by opposing groups; while this most often concerns scarce material goods, theorists such as Rawls (2005) have applied distributive frameworks to primary social goods such as liberty, political rights, and social positions as well. One can approach distributive justice from two perspectives; the most common in the study of justice holds that a distribution is just if the pattern of distribution that results from the distributive processes is just (regardless, for the most part, of what the distributive process looks like). This approach is less consistent with most theories of privacy however. A minority of distributive theorists focus on the distribution arising historically through a just process and hold that any distribution actually resulting from a just process is thereby just regardless of what that distribution looks like, an approach more consistent with the dominant view of privacy.

5.2.1 Process Distributions

In process theories of justice, the process by which a distribution actually arises historically is an end in itself; the actual pattern of distribution of goods resulting from it is not relevant to evaluation. Robert Nozick's theory of justice is exemplary of process theories of distributive justice.[4] For Nozick, "a distribution is just if it arises from another just distribution by legitimate means. . . . Whatever arises from a just situation by means of justice is itself just" (2013, p. 151). Contra Rawls, the end state is irrelevant; a distribution that could be reached by just means but was not is not just (the thief, who could have received goods from the victim by gift rather than theft but did not, is the paradigmatic case). Nozick posits two dimensions of justice in holdings. A principle of justice in original acquisition holds that "A person who acquires a holding in accordance with the principle of justice in acquisition is entitled to that holding"; a principle of justice in transfer of holdings holds that "A person who acquires a holding in accordance with the principle of justice in transfer, from someone else entitled to the holding, is entitled to the holding." By iterative application of these two principles, goods can be continuously redistributed in a just society without need for authoritative reallocation to remedy deviations from the just end state: "The complete principle of distributive justice would say simply that a distribution is just if everyone is entitled to the holdings they possess under the distribution" (Nozick 2013, pp. 150–151).

Nozick's principles of justice in original acquisition and transfer are primarily rooted in property rights and market mechanisms. Liberty is a common basis for principles of just transfer, with such theories returning often to Locke's claim that

[4] "Process theories of distributive justice" are distinct from distributive theories of procedural justice. The former determine the distribution of material and social goods by repeated application of a just process over the history of a good or society, while the latter determine entitlements to political, legal, or social process by according to the principles of just distribution of such entitlements. In general, it is initially sufficient for the purposes of this section to regard common uses of the term *procedural justice* as cases of the latter rather than as a distinct form of justice. But while this is a common interpretation of the idea, the implication of sect. 5.3 is clearly that procedural justice is better understood as a set of social structures, relations, and processes that guarantee one's ability to participate in determining one's actions and their circumstances, and thus as a species of structural justice as described in that section.

the state of nature is “a state of perfect freedom to order their actions, and dispose of their possessions and persons, as they think fit, within the bounds of the law of nature, without asking leave, or depending upon the will of any other man” (1980, sec. 4). This leads to the purely economic theories of Hayek (1994) or Friedman (2002) and is most easily applied in the distribution of material goods than of social goods. But the claim that justice in transfer is rooted in liberty of action is also important to many areas of social life in which consent is central; one could interpret the crime of rape, for example, as a failure to justly transfer a right to sexual interaction under a consent-based theory of justice in transfer.

Many theories and practices of privacy can be understood from a process-distributive perspective. In this view, privacy is protected to the extent that transfers of information are limited to those permitted under some principle of justice in transfer. The simplest form of this is seen in corporate privacy policies, in which consent is the implied principle of justice in transfer. The principle behind a publically available privacy policy is that consumers can understand how the receiver of the information intends to collect and use information about the consumer, allowing the consumer to choose whether to provide the information either directly or by undertaking actions such as making a purchase or visiting a web site that will allow the receiver to collect the information in the course of the action. The information collected in accordance with, for example, Google’s privacy policy is not a privacy violation because it was collected with the informed consent of the consumer.

This approach to privacy is present in another set of federal laws regarding information privacy in the United States, one that developed somewhat later than employment discrimination law but that nonetheless had enough overlap to suggest the complexities of information privacy in U.S. legislation. The Family Educational Rights and Privacy Act of 1974 (FERPA) and the Health Insurance Portability and Accountability Act of 1996 (HIPAA), together with their associated implementing regulations, are philosophically different than the employment protections described in sect. 5.1. Here, we see a series of information privacy acts in which privacy is an end in itself rather than an instrumental protection of some other justice concern.

FERPA bars, in principle, funding to “any educational agency or institution which has a policy or practice of permitting the release of education records (or personally identifiable information contained therein…) of students without the written consent of their parents” (20 U.S. Code § 1232g(b)(1)), with the consent rights shifting to the students themselves if they are at least 18 years old or attend a post-secondary institution (subsection (d)). Paragraph (1), of course, includes the usual lengthy list of exceptions for school and government program administration. Its implementing regulations’ purpose is explicitly the protection of privacy (34 C.F.R. 99.2), and the largest category of exception, the disclosure of “directory information,” is defined as “information contained in an education record of a student that would not generally be considered harmful or an invasion of privacy if disclosed” (34 C.F.R. 99.3). HIPAA does not specifically define standards but rather required the Secretary of Health and Human Services (HHS) to develop such standards, which led to the HHS Privacy Rule. The rule permits insurers and health care providers “to use and disclose protected health information” only for specific purposes related to medical practice and payment, prohibits the sale of such information, and requires that such discloses be limited to the minimum necessary information for the purpose of the disclosure. It exempts de-identified information from the rule entirely, but protects the information for 50 years after a person's death (45 CFR 164.502).

These are significantly different approaches to what is seen in employment discrimination law. FERPA does not at all restrict the collection of information about students by educational agencies or institutions, either formally as does the ADA or GINA or by presenting litigation risks as under the Civil Rights Act. Nor does it contain any restrictions on who can hold such information or under what circumstances they can hold it, as the ADA and GINA do. FERPA simply requires that any transfer of information outside of the exceptions be done with the student's or parent's consent. The same is true of HIPPA's relationship with insurers and health care providers, who can collect information virtually at will but operate under even more significant restrictions on the transfer of individually identifiable information. HIPAA demonstrates even more clearly its purpose in protecting privacy as an end in itself with its wide scope for the use of de-identified information and protections of privacy even after death. FERPA and HIPPA are not instrumental protections of some other goal. They act to create an operational principle of justice in transfer for educational and medical information.

A more complex version of the consent principle holds information privacy to be a commodity that consumers exchange for services. Individuals determine the value of services received—for example, free email from Google—and choose whether the value of the services exceeds the value of the information that they will provide to the service provider. If the individual determines that Gmail is a good enough product to justify allowing Google to read one's email and use that to target advertisements, they "purchase" the service with information rather than money (Campbell and Carlson 2002). The service provider thus gains a right to the consumer's information through a just transfer (i.e., a consensual and mutually beneficial exchange of considerations) of information from the consumer, a right constrained by the specific terms of the exchange. In either the pure consent or the commodity exchange framework, information privacy constitutes an injustice to the extent that information is collected or used without or in ways contrary to the consent of the individual from whom it is collected.

Process theories of distributive justice add considerable depth to both consent and commodity theories of privacy. Common dissatisfactions with privacy policies often focus on violations of an implicit theory of just transfer. Various failures of informed consent such as the length, complexity, accessibility, and incomprehensibility of the policies; the inability of consumers to negotiate the terms of the policy with receivers; or the necessity of the transactions governed by the policy entail that the distribution of information has arisen from a (presumably) just initial acquisition of information (i.e., the individual has all of the information about them and others have none) but undermined by unjust transfers. The Apple iTunes Terms of Service (2015) agreement, for example, runs nearly 21,000 words, suggesting a reading time of between 98 and 136 min without accounting for the difficulty of the text (Trauzettel-Klosinski and Dietz 2012). Analyzing it in Microsoft Word 2013 shows it has a Flesch-Kincaid Grade Level score of 16—a college graduate—and a Flesch Reading Ease score of 31.4, putting it above the reading comprehension level of the more than 70% of Americans without a bachelor's degree. Considering consent explicitly as a theory of just transfer calls attention to the need to articulate specifically the conditions under which consent to information transfer is meaningful at a level of specificity similar to that developed to govern informed consent in medical practice. A legalistic notion of consent might dismiss these concerns with a simple "caveat emptor," but a justice-centered view of privacy must take seriously whether the capacity to consent is absent from such a case and whether that violates a principle of just transfer of information.

The process approach to justice also exposes a serious weakness in consent-based and especially commodity models of privacy. Few such models offer a serious attempt to explain the principle of justice in original acquisition at work. FERPA and HIPAA assume that the information held by educational institutions and health care providers is held legitimately according to some unstated principle of justice in original acquisition.[5] It is only in the transfer (and in some cases under HIPAA, use) of information that privacy protections come into play under these privacy laws. Campbell and Carlson make a compelling and critical argument that the transfer of information beyond that immediately necessary for a transaction is rooted in a kind of panoptic self-surveillance in which the threat of exclusion from market benefits entices cooperation, but cite those who hold that "users should willingly allow for information collection in return for economic benefits" (2002, p. 593). Both sides, however, start from the premise that information beyond minimum transactional data is originally the property of the individual. It is not at all clear that this is so, or that it is adequate as a principle of just original acquisition. Some kinds of inferences about individuals' personal information are central to most social interactions; traditionally at least, in formal English it is necessary to know the gender of someone one wishes to refer to in the third person even where their[6] gender is

[5] To be sure, this does not imply that the laws assume it is held accurately. FERPA grants extensive access rights to educational information to the student or their parents, and requires institutions to have formal processes in place that allow students to challenge and correct inaccurate or misleading information. One can consider this at best a minor element of an otherwise absent theory of justice in original acquisition, to the effect that inaccurate information is not held justly. It is surely inadequate in itself as a principle of just acquisition.

[6] The use of *they* as a gender neutral, indeterminate gender, or gender diverse singular pronoun is increasingly common as a response to exactly this problem. The singular they has a long history in English usage; it is nearly ubiquitous in spoken English, quite common in the writings of prominent authors, and came to be stigmatized by prescriptive grammarians only in the eighteenth century. As Jürgen Gerner notes, "As there are no third-person personal, possessive, or reflexive pronouns in English which are both singular and gender-neutral, complete grammatical agreement is not possible [when the antecedent is both singular and not marked for gender, e.g., *anyone*]. The choice is between a violation of gender concord or a violation of number concord" (Gerner 2000). That number agreement is the obvious answer is now increasingly questioned.

That this is (and, with the exception of the two-century interlude of prescriptive grammarians insisting that the only objective solution was number agreement, always has been) acceptable, unknown or they may not see the gender binary as meaningful to them. It seems strained to say that information about one's gender is the "property" of that individual and cannot be transferred without that individual's consent.

At the same time, a theory of privacy based on procedural distributive justice seems to give the individual no rights with regard to such information as is minimally necessary for a transaction and blurs the line between transfer and acquisition. No reasonable theory of just original acquisition could conceivably argue that an online store has not justly acquired records of the items purchased through their site, the dates of purchase, and the means by which the purchase was made. But using only this data, without any transfer of data to another party, retailer Target was able, infamously, to predict which customers were pregnant, a finding that has stirred great controversy among privacy advocates (Hill 2012). I have argued elsewhere (Johnson 2014a) that so-called "big data" methods can result in privacy violations by creating new data about an individual that the individual would not have revealed themselves, whether or not they actually knew the information. A student might not be willing to disclose their probability of failing a course to their professor at the beginning of the semester; inferring information about membership in protected social categories (e.g., by identifying proxies for race) gets around laws barring the direct collection of such data. But the common sense perspective on acquisition and transfer would see this as unquestionably a matter of acquisition: the data collector has created the data themselves from other data that they held legitimately. Hence what would be a violation of privacy if it arose through transfer is achieved in a technically legitimate way by the creation of new data from data previously transferred, which is to say, by original acquisition. Transfer becomes acquisition, further separating the subject of the data from their rights regarding it.

Theories of consent, driven implicitly by a process approach to justice, only cover justice in transfer; justice in original acquisition is simply assumed. But a privacy framework that derives more explicitly from a process theory of justice will require not principles justifying transfer but justifying acquisition as well. Helen Nissenbaum's (2010) idea of contextual integrity can be interpreted as a partial solution here. Nissenbaum argues that the context of information flows is as important to privacy as the content of them; privacy violations occur when information flows beyond the "values, goals, and ends of the context" in which it exists. One could read the context as an initial limit on justice in original acquisition: the information was acquired not universally but for use in a specific context. To use it beyond that context, whether by transfer to another party or by the initial possessor, is to hold the information contrary to the principle of justice in original acquisition, a violation of the principles behind process theories of justice, since a just distribution can only arise by repeated operation of the principles of justice in original

however, does not solve the problem of ownership of one's gender information; it merely expands the range of options that a process theory of justice allows one to own. That one prefers *they*, or even *ze*, remains within the scope of owned information that then requires a theory of just transfer before one can speak of the person. This seems an exceptionally complicated way of understanding the matter.

acquisition and of justice in transfer. Target would thus be justified in using transactional data to execute the transaction and to manage its ability to repeat the transaction (on the assumption that the context of a transaction at Target includes some notion of iteration) but not to generate personal information about the consumer. Certainly, a principle of contextual integrity would be part of a principle of just transfer as well.

5.2.2 Pattern Distributions

While it is undoubtedly the most common way of engaging information privacy from the perspective of distribution, not all distributive approaches to privacy are oriented toward process. Hartzog and Selinger argue:

> Ideas about preventing the surveillance society from going too far usually focus on three desirable outcomes: (1) prevent certain groups from ever having access to certain types of information; (2) prevent certain groups from being able to use certain types of information in select contexts or in certain ways; and (3) make it harder for certain groups to be able to access or interpret information. (2015, p. 1345)

In each of these outcomes, information flow is restricted not by stating terms under which it can be (un)justly acquired or transferred but by describing a pattern of distribution following the distributive process that is (un)just: one in which certain groups have inappropriate access to information or use it inappropriately. The result is the ability to call a particular pattern of distribution that is the end state of a distributive process just regardless of the conditions or processes under which it arose historically, and to argue for information privacy practices as pragmatic policy solutions to bring about that distribution.

A theory of distributive pattern justice must generally consider two dimensions of distribution. Consider, for example, the basic utilitarian justice principle of maximizing utility. Kolm suggests that the basic framework for theories of distributive justice is in which material and social goods are distributed according to one or more "directly relevant ethical variables" (1993, p. 438) that may range from a basic political equality to individuals' varying needs for material goods. In the utilitarian framework, the directly relevant ethical variable is utility. However, in spite of his framing it is clear from Kolm's presentation (as well as the wide range of theories of distributive justice reviewed in Gaus and D'Agostino's (2013) eight chapters on distributive justice) that one must consider not simply the variable on which the distribution is based but also the distributive model itself. *Distribute goods according to utility* is itself inadequate, as there are many possible distributions in which utility is the only variable. Hence the utilitarian principle includes a specific distribution as well, that of maximizing utility for the greatest number of individuals. One could equally offer the argument that goods ought to be distributed such that utility was equal across individuals or such that total utility is maximized without violating the principle that utility is the only directly relevant ethical variable. Rawls offers a more complex theory that nonetheless addresses both dimensions, distributing material goods on the basis of a maximin function (one that maximizes the minimum outcome) applied to shares in social surplus as a directly relevant ethical variable, but primary social goods on a greatest equality function (one that maximizes outcomes subject to the constraint that all receive the same outcome) applied to natural liberty.

While process is important to some pattern theories of distributive justice, in those theories the process is usually instrumental to either reach or justify a particular pattern of distribution. Rawls' original position and veil of ignorance are well-known procedural solutions to the problem of how to ensure that individuals choosing principles of justice will do so without regard for their own interests. But Rawls is clear that the original position is intended to justify a set of distributive principles:

> The original position is not, of course, thought of as an actual historical state of affairs, much less a primitive condition of culture. It is understood as a purely hypothetical situation characterized so as to lead to a certain conception of justice. (Rawls 2005, p. 12)

It is a distribution of goods that conforms to the final distributive principles that Rawls identifies through the process of reasoning from the original position, not the process of choosing them, that define the just society. Indeed, the process determines only the principles of justice to which the end distribution must conform and not the actual distribution of goods.

The most straightforward pattern distributive framework for information privacy attempts to define a right to information privacy analogous to existing rights to personal privacy framed, initially, as a "right to be left alone" (Smith et al. 2011, p. 994). Such frameworks essentially posit a right to privacy as a primary social good, delineate the basis of the right as the directly ethically relevant variable, and use, presumably, equality as the principle of distribution. Rebecca Greene, for example, argues for a right to obscurity in political information about an individual:

> Political obscurity therefore describes a broader right than anonymity: it is the fundamental right to exist without one's political preferences being continuously recorded and, consistent with the right articulated in [*United States Department of Justice v. Reporters Committee for Freedom of the Press* (489 U.S. 749 (1989))], a right against state-facilitated cataloguing of one's political preferences. (Greene 2013, pp. 373–374)

This is a limited right that is equally distributed across all persons by virtue of being inherent in individual autonomy; in its absence, deliberate self-governance becomes exceptionally problematic. By distributing political obscurity as a primary social good individual autonomy and citizen participation is reinforced.

Often apparently rights-based approaches create only legal rights that put into effect an access control regime rather than creating a right in itself. The Canadian Personal Information Protection and Electronic Documents Act (S.C. 2000, c. 5) establishes ten principles for the protection of personal information to which organizations must comply. Those principles give individuals a range of legal rights against organizations holding information about them but do not create a general right to which one can appeal beyond the specific situations governed by the act. Similarly, the U.S. Department of Health and Human Services states that HIPAA "gives you rights over your health information, including the right to get a copy of your information, make sure it is correct, and know who has seen it" (U.S. Department of Health & Human Services Office for Civil Rights 2013). But these rights largely protect flows of information: they require the disclosure of purposes, the consent of the individual, limited collection and use, and individual access to the information among others. The acts do not create a general right to privacy that can be distributed; they create procedural protections against specific information flows, and so are best understood as either a means of implementing a process approach to justice in information privacy or as an attempt to delineate and operationalize an implicit right to information privacy.

The latter is probably the best interpretation of European Union law on information privacy. The European Data Protection Directive (Council Directive 95/46/EC) is among the earliest and most extensive legal regimes for protecting information privacy specifically as a right. The directive was specifically created to "protect the fundamental rights and freedoms of natural persons, and in particular their right to privacy with respect to the processing of personal data" (art. I, para. 1), in particular aiming to harmonize information privacy protections across member states in the face of increasing flows of personal information, in both government and commerce, across member states (recitals 5, 7, 10, and 11). The right to privacy is not, in fact, articulated substantively, but recital 10 references everyone's "right to respect for his private and family life, his home and his correspondence" under Article 8 of the Convention for the Protection of Human Rights and Fundamental Freedoms, tying the substance of the Directive to the existing body of privacy rights recognized within the EU. Recital 10 states that the Directive gives substance to and amplifies previous protections of privacy: Data collection and processing under the Directive must ensure both data quality and legitimacy, the latter secure by either consent or limited notions of transactional necessity; in general, processing of sensitive personal data is barred by Article 8 of the Directive; extensive information and access rights to an individual's own data are specified in Articles 10 through 12; and the controllers of data are subject to substantial regulation under Articles 16 through 20. This structure, of regulations that operationalize a principled right to general privacy as it applies to issues of personal information, is preserved in the proposed General Data Protection Regulation that took effect in 2016.

Claims of a fundamental right to specifically information privacy in the United States are less well developed. Many organizations concerned with information privacy have asserted that information privacy protected as a legal or moral right. The American Library Association's interpretation of its "Library Bill of Rights" argues that privacy is implicit in the bill's Article IV on resisting "abridgment of free expression and free access to ideas" and cites a chilling effect on those principles from breaches of privacy. But the basis claimed for this right relies primarily on legal precedents related to either receiving information in a library or to general privacy cases (American Library Association 2002) without a clear argument to that effect, which is complicated by the diverse legal bases for privacy claims under U.S. law. It may well be that one's choices to access information is included within the zone of personal behavior free from unreasonable state intrusion, implicit in ordered liberty, or basic to a free society (*Mapp v. Ohio* [1961], 367 U.S. 643); is within "penumbras, formed by emanations" from specific guarantees within the U.S. Bill of Rights (*Griswold v. Connecticut* [1965] 381 U.S. 479, 484); or are made with a reasonable expectation of privacy (*Katz v. United States*, 389 U.S. 347 (1967) Harlan, J., concurring). But which doctrine forms the basis of information privacy and how the connection is made will certainly matter for the practical contours of a right to information privacy.

U.S. courts appear increasingly willing to accept such claims, but a definitive doctrine has yet to emerge. The U.S. Supreme Court's recent decision in *Riley V. California* (573 U.S. ____ (2014); Docket No. 13–132) held that the information contained on a mobile phone enjoyed the protection of the Fourth Amendment's requirements for reasonableness in searches and thus could not be searched incident to the arrest of a person in possession of the phone. Observing that mobile phones "are now such a pervasive and insistent part of daily life that the proverbial visitor from Mars might conclude they were an important feature of human anatomy," the Court relied on the extensive information contained in smartphones to conclude that the digital content of such devices is categorically different from treating them as the kinds of physical objects assumed by the search incident to arrest doctrine articulated in *Chimel v. California* (395 U.S. 752 (1969)), *United States v. Robinson* (414 U.S. 218 (1973)), and *Arizona v. Gant* (556 U.S. 332 (2009)). With mobile phones, the Court held, "The sum of an individual's private life can be reconstructed through a thousand photographs labeled with dates, locations, and descriptions" as "the more than 90% of American adults who own a cell phone keep on their person a digital record of nearly every aspect of their lives—from the mundane to the intimate," a situation aggravated by the fact that a mobile phone is connected to data stored elsewhere through cloud applications. The court here clearly presents a defense not of personal or general privacy but of information privacy specifically, distinguishing between the permissible physical search of a mobile phone to determine if it poses a physical threat (e.g., if the arrestee concealed a weapon such as a razor blade in it) and a search of the information contained on it.

This remains, however, a weak theory of information privacy. It springs from general privacy; the question before the court is how one specific legal doctrine, the search incident to arrest rule, applies to the information carried on a device. The court links the issue to the traditional defense of Fourth Amendment protections, that of an overreaching exercise of police powers by the state. The court's ruling does not lead to any definitive implications beyond criminal process. For example, the court remains silent on whether there is a reasonable expectation of privacy in the contents of a mobile phone or simply whether the privacy implications are sufficient to overcome a weak case on behalf of the government's interest in the traditional justifications for searches incident to arrest, concluding that a search supported by a warrant is the proper course because that is the general principle of reasonableness in searches.[7] That would be important in determining, for example, whether the state can demand the right to search one's mobile phone in the course of a voluntary administrative process where significant discretion is allowed, such as approving membership in any one of the U.S. government's Trusted Traveler programs. And it of course says nothing about intrusions on privacy from commercial actors. Google, among many others, already maintains "the sum of an individual's private life" for every individual who uses a Gmail account or an Android phone. Amazon's examination of the "digital record of nearly every aspect of their lives" may be less problematic than the state's, but it is by no means obviously unproblematic.

I will argue below that viewing rights from a distributive perspective is especially problematic. But it seems puzzling that other distributive approaches to privacy are nearly nonexistent. This is certainly not because such claims cannot be made. Rawls' first principle of justice is that "each person is to have an equal right to the most extensive basic liberty compatible with a similar liberty for others" (2005, p. 60). Rawls here treats rights and liberties as essentially synonymous, so we could conceive of privacy as part of a system of basic rights and liberties that ought to be maximized subject to the constraint that all have the same liberty. This may seem like a trivial or even substanceless requirement, but it actually illuminates a quite serious problem in information technology, that of online harassment and domestic violence, especially toward women.

In 2013, after feminist critics of video game culture raised concerns about the depiction of women (and lack thereof) in games, gamers responded with an astounding level of online harassment of their critics. The harassment, almost entirely either anonymous or pseudonymous, included quite specific rape and death threats (for example, including critics' addresses and the times they would be assaulted), driving at least three women from their homes and prompting FBI investigations. An especially serious tactic used to silence critics was that of "doxing," or publishing personal information about the critics that would make them vulnerable to violence or harassment in the physical world (Dewey 2014). One critic, Anita Sarkeesian, had to cancel a speech at Utah State University after the university, constrained by state law prohibiting it from barring legally carried concealed weapons on campus, could not respond to an anonymous threat to carry out "the deadliest school shooting

[7] There is good reason to believe a general right to privacy with regard to information might flow from this decision. The court's recognition of the scope of information available through a mobile phone, including information held by third parties, was a major part of its reasoning that a substantial privacy interest was implicated. This would challenge the major obstacle to a broad doctrine of information privacy in existing case law, the third-party doctrine (*Smith v. Maryland*, 442 U.S. 735 [1979]), by recognizing that individuals retain a privacy interest in data shared with a third-party service provider. The U.S. Court of Appeals for the Second Circuit, in considering the legality of bulk telephone metadata collection in *ACLU v. Clapper* (No. 14-42-cv, slip op. at 83–90), suggested (in a judicial dictum, as it had already disposed of the case on non-constitutional grounds) that the third-party doctrine is brought under new scrutiny by the technology of bulk collection, relying in part on the concurring opinions of Justices Alito and Sotomayor in *U.S. v. Jones* (132 S. Ct. 945 [2012]) arguing that Global Positioning System tracking on a vehicle exceeds the reasonable expectation of privacy.

in American history" targeting both the lecture and the university women's center if she spoke (McDonald 2014). The threat read, in part:

> Anita Sarkeesian is everything wrong with the feminist woman, and she is going to die screaming like the craven little whore that she is if you let her come to USU. I will write my manifesto in her spilled blood, and you will all bear witness to what feminist lies and poison have done to the men of America. (Neugebauer 2014)

Supporters of Gamergate counter that the movement supports ethics in video game journalism.

To be sure, there are many deep problems with Gamergate—some of which I address in sect. 5.3—that contributed to this harassment. Among them is an asymmetry of privacy. The critics of games could not make their critiques from a position of anonymity or pseudonymity; their critiques are part of their professional identities and were unavoidably done from a public stance. In order for Zoe Quinn, Brianna Wu, and Sarkeesian to offer their criticisms publicly as professionals in the game development industry, they must have adopted a public identity with all of the constraints and risks that entails. The lack of such an identity is equally essential to their harassers. Protected by anonymity, both personal and technical, they could level threats that were of an unquestionably criminal nature (indeed, as the example above demonstrates, an unquestionably inhuman nature) without substantial risk. The Gamergate harassers thus exercise a degree of privacy not available to their critics. A Rawlsian principle of "greatest equal privacy" could provide guidance to those developing systems that could mitigate the asymmetry of privacy, taking away the protection that anonymity provides online harassment and domestic violence. Facebook, for example, maintains a real name policy that requires users to register and use their real names rather than pseudonyms in part to cut down on online harassment. One might make similar arguments about privacy from the perspective of need or capabilities (Brock 2013; Robeyns 2013).

Of course, such a principle would be open to challenge as well, with anonymity being a core principle of privacy, especially in technical solutions. If the greatest possible amount of privacy compatible with all having a similar privacy prohibits anonymity entirely, then the critics who argue that privacy is a dated concept have a strong argument. The Electronic Frontier Foundation considers the TOR web browser, which routes web browsing through multiple relays making users anonymous to the sites they visit, an essential privacy tool. TOR might allow a user to gather information about protest movements without being subject to state surveillance, but it would also allow anonymous access to an email account from which a threat like those made to Sarkeesian could be sent without fear of legal consequences. These kinds of tradeoffs are part of why EFF considers threat modeling an essential practice for surveillance self-defense. That, however, further undermines a greatest *equal* privacy principle, since different threats will entail different levels of privacy, not only for end users but for system designers. Perhaps a needs or capabilities approach would be an improvement here, though it is not entirely clear that this is so. Ultimately this may suggest that there are serious limits to what pattern approaches to distributive justice can offer a theory of privacy.

Finally, one notices a significant shift in the object of distribution in moving from process to pattern concepts of justice. In the former, the concern is with the distribution of information, while in the latter the concern is the distribution of privacy rights. It is not at all clear that they are the same thing. Distributions of information present an intersection of concerns with privacy and with open data (Johnson 2014b) in that aims of controlling data flows in the name of privacy directly conflict with commitments to expanding them in the name of transparency, a problem that suggests a solution of balancing the competing concerns. Distributing privacy rights, however, delineates a clear obligation that trumps the practicality arguments that prevail among advocates of open data. Sunshine may be the best disinfectant, but it certainly harms the patient. The two approaches to justice give very different answers to how tolerant of such harms we should be.

The process framework addresses information rather than privacy in part because it must. This is more than a definitional issue in which privacy is defined as a question of distributing information. The process approach simply does not work effectively when the issue at question is a right held to be universal and inalienable. In such cases, the principles of justice in acquisition and in transfer become trivial: the right is originally acquired when one achieves personhood (in whatever capacity one wishes to use, but in any case well before one can make decisions about one's privacy related to one's original acquisition) and the right cannot be transferred under any circumstances. In order for a process approach to distribute privacy rights as such rather than information, it must first show that privacy rights are somehow alienable. To say that they are wholly so does great violence to the concept of rights generally, so the question then becomes what aspects of one's privacy rights can be transferred, a question that is, in practice, no different than stating the kinds of information that can be justly transferred and the circumstances under which it can be transferred.

It is less clear that one could not develop a pattern theory that distributes information rather than privacy rights, but it is no easy task. Though not focused on justice in any strong sense and taking a wide range of positions on privacy, the open data movement is very much one positing an ideal pattern of information distribution. This might be an especially good focus for distributive theories rooted in need of capabilities (see, e.g., Britz et al. 2012). But information privacy poses an unusual challenge to pattern theories of justice. In most cases, pattern theories of distributive justice assume that the goods being distributed are in principle either universal (in the case of primary social goods) or scarce (in the case of material goods). In these cases, the problem is to ensure that everyone receives their fair share of the goods in question. To be sure, this is a problem in information justice more broadly, as those questioning the justice of open data argue (Donovan 2012; Gurstein 2011; Raman 2012; Slee 2012). But information privacy poses the opposite problem: restricting the distribution of a good that can be, in the age of electronic reproduction, initially produced and then reproduced infinitely at near-zero cost. A pattern distribution of information that protects privacy would thus need to both create a solution to a novel type of issue that runs counter to existing theories and also support such theories in other areas of information practice. It is not at all clear how this might be done other than by restricting the flow of information.

5.3 Beyond Distributive Privacy

Concepts of distributive justice have revealed some useful insights into specific aspects of information privacy. In summary:

1. Information privacy is of at least instrumental value in pursuing distributive justice generally.
2. The information flow paradigm must pay attention to the principles for justice in information transfer, as they typically do; for justice in original acquisition, which they typically do not; and for distinguishing between the two.
3. Robust conditions for consent are critical to making the information flow paradigm produce meaningful principles of justice in transfer.
4. Rights-based distributive approaches to information privacy are easily confused with legally established information flow approaches, and are most coherent as systems of information justice when they are used to delineate and operationalize an existing framework of privacy rights.
5. Theories of distributive justice distinct from distributions of rights may be able to provide some useful, practical guidance for information privacy.
6. The justice implications of distributing information and of distributing privacy rights are significantly different.

Nonetheless, distributive justice has not proven itself capable of providing a strong framework for a general theory of information privacy. None of these considerations fundamentally remake our understanding of privacy, nor do they do much by themselves to bring coherence to the field. And they have raised as many problems as they have solved.

The distributive paradigm is not, however, the only way of understanding justice. The same alternatives that Schlosberg found so useful in reconstructing environmental justice can serve as the basis for a useful justice-driven conception of privacy. Young's critique of distributive frameworks of justice provides some insight into why such frameworks seem to offer relatively little to theories of privacy as well as a starting point for a more effective justice-driven approach to privacy. Young's critique focuses on how distributive justice "regards persons as primarily possessors and consumers of goods" rather than considering "action, decisions about action, and provision of the means to develop and exercise capacities" (1990, p. 16). This leads to two related failures on the part of the distributive paradigm connected to social structure.

The first is that it obscures the structural conditions that underlie the distribution of material goods. The distribution of employment, for example, is a common object for the study of distributive justice, as seen in the analysis of employment discrimination law above. Young argues, however, that asking about the just distribution of jobs tends to assume rather than examine structures like the hierarchical division of labor, social stratification, and commodification that tend to determine the distribution of jobs. In many cases, Young argues, controversies over the distribution of goods are, in fact, controversies over these structures themselves: citizens oppose a hazardous waste treatment plant or a plant closing not because they see the distribution of environmental economic burdens per se but because they lack a voice in decisions that affect their lives. Distributive justice fails to capture these aspects of justice in the distribution of material goods, thus restricting the scope of claims to justice (1990, pp. 18–24).

This is a central issue in the question of information privacy, and one that receives scant attention. The distribution of information is a problem, but the problem is not simply a maldistribution of information but also, perhaps even more so, the structures that make it so. Ownership and commodification of information cannot be a solution to information flows because it is at the heart of the problem: Target can identify sensitive medical information that it has no business knowing—there would be no question of it being a privacy violation under HIPAA for Target to have determined which customers were pregnant by buying their medical records—because of the structure of economic activity in a free market system. The unequal positions of enterprise and consumer drive the latter's willingness to be surveilled by the former:

> Clearly, the exchange of information between consumers and suppliers is not equitable, as large corporations do not in the same transaction generally reveal to customers detailed information regarding their internal structure or operations…. [I]t is this very inequality in the relationship between consumers and suppliers of goods and services in the marketplace that compels individuals to provide personal information. The ability of the producer or supplier to set the terms of the contract that the consumer can only accept or decline defines the transaction as inherently inequitable…. The consumer is ultimately a "contract taker, rather than a contract maker," and thus provides the information in the belief that it represents a reasonable transaction cost…. [I]ndividuals are not necessarily aware of the degree of inequalities in their relationship with suppliers because marketers and advertisers have effectively concealed the consumerist Panopticon. (Campbell and Carlson 2002, pp. 591–592)

Similarly, state surveillance is unlikely to be seriously addressed with the kinds of legalistic concepts of criminal process rights and limited government seen in *Riley* and in *ALCU v. Clapper* when terrorism, the surveillance state's *raison d'etre*, exists as a state of exception to law (Agamben 2005): the U.S. prison at Guantánamo Bay, Cuba, exists solely because it is considered outside of U.S. legal jurisdiction (see *Rasul v. Bush*, 542 U.S. 466 (2004)).

These structural conditions run much deeper than the social contexts in which information is distributed. I have previously argued that information is substantively influenced by a translation regime that transforms underdetermined observations into a single data state (Johnson 2015). Information exists as a form of communication in which the translation regime encodes the information and a nexus of problems, models, and interventions decodes it. Both the translation regime and the problem-model-intervention nexus are constructed by social actors with social interests in mind, whether deliberately or as unconscious assumptions. This is not just to say that there are errors in information that can be remedied by the kinds of protections seen in FERPA or the European Data Protection Directive; the data is directly constructed by standards that are inevitably biased but, because the standards at least appear to be mechanical and objective, cannot be challenged as erroneous. The result is that information privacy cannot be about the distribution of information without also being about the sources and construction of that information.

This failure to highlight structural conditions is compounded, Young argues, when distributive justice is extended beyond material goods. The kinds of moral goods that distributive theorists—and privacy rights theorists—address distributively are "better understood as functions of rules and relations than as things," with the result being that distributive justice "tends to preclude thinking about what people are doing, according to what institutionalized rules, how their doings are structured by institutionalized relations that constitute their positions, and how the combined effect of their doings has recursive effects on their lives." Distributive justice cannot conceive of the way that both distributions and structures shape actions except by supposing that actions are constrained only by distributions of goods. Young finds this especially problematic when thinking of power as capable of distribution, as "power is a relation rather than a thing" mediated by structure of agency and actions that obscures especially the ways in which social structures and systems "exclude people from participating in determining their actions or the conditions of their actions" through processes and relationships rather than possession of some abstract concept of power (1990, pp. 24–33).

One sees this most clearly in Greene's otherwise strong argument for political obscurity. Generally, Greene's chief concern is the ability to hold and act on dissenting political views:

> Political obscurity refers to the state of one's political preferences being shrouded or otherwise difficult to discern or distinguish by others. A person enjoys political obscurity when she can go about her day as she so chooses without others perceiving or otherwise determining the nature of her political views. The politically obscure person is *able to control and manage the extent of disassociatio*n from the political views she holds (or once held) or political actions taken in the present and in the past. (2013, p. 373, emphasis added)

But when Greene translates this concept into a right that can be distributed (as quoted in the previous section, "the fundamental right to exist without one's political preferences being continuously recorded"), the language of action, and in fact the actors involved, disappears. Rather than controlling and managing, the person simply "exist[s]"; the others actively trying to know and influence her political views become a passive, impersonal occurrence that happens to a person.

This, too, is a fundamental inadequacy of a distributive justice-based theory of privacy. A right that is possessed gives no consideration to what one might do with such a right; it is merely distributed and its recipients wished the best of luck with it. Its moral status is not affected if the right should prove inadequate to the purposes for which its recipients intend to use it. In the case of information privacy, consumers' concerns about privacy are not simply about what data enterprises hold about them, but about how it will be used to further interests that may well conflict with the intentions and plans of the consumers. The memorable vignette from the Target case is of the father who first confronts the store manager about marketing goods for pregnant women and newborns to his teenaged daughter, only to later apologize, saying, "It turns out there's been some activities in my house I haven't been completely aware of. She's due in August. I owe you an apology" (Hill 2012). Target excluded both the father and the daughter from participating in decisions about how they would to respond to a major life event. The wrong is not (or at least not simply) knowing that the daughter was pregnant, or even knowing that she was sexually active, but rather that Target, not the woman, determined the circumstances under which her family would find out about it. That is not simply a question of a right to prevent certain flows of information, but of the basic justice of manipulative marketing from a company who had determined that childbirth was an excellent opportunity to shift someone's buying habits.

These two criticisms come together to understand why Gamergate cannot be understood as simply a matter of asymmetric privacy. Gamergate was not simply a heated dispute over any kind of principled matter (especially not "ethics in journalism"); the threats directed toward Quinn, Wu, and Sarkeesian were not just insults meant to hurt their feelings. Gamergate depends critically on the structure of gender relations in video game culture. It began when Quinn's ex-boyfriend, in a lengthy diatribe, accused her of being sexually unfaithful, it peaked when someone still unidentified threatened to massacre "the craven little whore" who has poisoned the *men* of America. Doxing them was a threat intended to silence them and maintain a system of structural power that favors men, one that was, due to the legally enshrined hyper-masculine culture of Utah that treats carrying a gun as the *sine qua non* of manhood, successful in at least Sarkeesian's case. These three women's information privacy was violated not simply in that personal information was distributed improperly but because information about them was used as a tool of domination and oppression. A justice-driven theory of information privacy will be inadequate if it cannot engage such issues.

There are promising approaches to privacy that are compatible with ideas of structural justice. Approaches treating privacy as a form of obscurity have been quite sensitive to structural issues. Greene's analysis does not rely on her rights-driven formulation, and in fact offers excellent analysis of the structural features of both information technology and petitioning processes to understand the injustices that she identifies:

> What if the real (and much more difficult to document) harm befell those who did not—or would not—sign the petition? What if the harm in releasing petition names is not to activists being mooned or shouted at as they advocate publicly for their cause? What if the real privacy victim is a mother of two, passing a petition circulator entering the grocery store, fearful that signing a petition—even for a cause in which she very much believes—might create a lifelong indelible association with that cause on her Internet record? (2013, p. 370)

She points, for example, to the publication of signers of a Maryland petition to bring an anti-same sex marriage measure to the ballot by an LGBT newspaper that resulted in one university terminating its chief diversity officer, who had signed the petition, as an example of how a lack of political obscurity undermines the capacity to act politically. Information technologies that undermine political obscurity by making possible frictionless gathering of information about people's political beliefs have a chilling effect on political participation and action.

Similarly, the broader idea of obscurity as a protection against state surveillance that Hartzog and Selinger (2015) argue for is rooted very much in an analysis of processes and relationships of power rather than distributive concepts. They define obscurity as making information hard, but not impossible, to find or interpret, arguing that this quality makes information "safe." Explicitly, they seem to mean safe from exposure: "when information is difficult to acquire or burdensome to interpret, the only people who will be inclined to do the detective work are those who deem the expense an acceptable cost." But implicitly, obscurity preserves the safety of its subject by "mak[ing] it hard (or harder), but not impossible, to discover irrelevant, inadequate, and embarrassing details" that would limit one's ability "to manage the accessibility and comprehension of social exchanges by outsiders, the loss of which can be quite harmful." It makes the actions of citizens less intelligible, which one might interpret as acting as a counter to the state's interests in legibility (Scott 1998). Obscurity thus functions as a structural rather than a legal right, protecting interests because social structures prevent their violation.

Neither Hartzog and Selinger nor Greene defend obscurity as a distributive right, nor do they frame it as a question of justice in information transfer. Indeed, Hartzog and Selinger argue quite explicitly that they do not aim to bar such transfers but simply to make them more difficult in order to shape relationships among actors and facilitate certain kinds of action for certain actors. And they can show how the structural conditions in which information is made and distributed affect one's ability to both develop one's capacities as a person and to participate in determining one's actions. One might draw a similar conclusion about Daniel Solove's problems-based approach to privacy: "A privacy invasion interferes with the integrity of certain activities and even destroys or inhibits some activities" (2008, p. 8). His imagery of the contrast between Orwell and Kafka reminds one that the harm of surveillance is not just the threat of discipline from Big Brother but also the feeling of being powerless against systems that shape one's life.

5.4 Conclusion

Young's critique, and the compatibility of some of the more successful approaches to privacy, suggests the virtue of a structural justice approach to information privacy. Information privacy is rife with unanswered questions of distributive justice. Information flow models of privacy assume answers to questions about justice in acquisition and transfer that may be indefensible on their own or incompatible with each other. Rights approaches often assume rather than articulate a justification of privacy itself. And there are considerable implicit differences between the two in what is to be distributed. It should not be seen as practical to address questions of information privacy without considering these questions. Young's critique of the distributive paradigm reveals deeper problems with understanding the question of justice in information privacy. Information privacy is as much a matter of social structure as it is of distributing material or moral goods, and the focus on distribution obscures the ways in which information privacy violations challenge the ability to participate in determining one's actions. These critiques suggest that a more productive line of inquiry would be to pursue information justice as a matter of primarily structural rather than distributive justice.

References

Agamben, G. (2005). *State of exception*. Chicago: University of Chicago Press.

American Library Association. (2002, June 19). *Privacy: An interpretation of the library bill of rights*. HTML file. http://www.ala.org/Template.cfm?Section=interpretations&Template=/ContentManagement/ContentDisplay.cfm&ContentID=132904. Accessed 8 Mar 2016.

Apple Inc. (2015, October 21). *iTunes store – Terms and conditions*. http://www.apple.com/legal/internet-services/itunes/us/terms.html. Accessed 4 Mar 2016.

Britz, J., Hoffmann, A., Ponelis, S., Zimmer, M., & Lor, P. (2012). On considering the application of Amartya Sen's capability approach to an information-based rights framework. *Information Development*. https://doi.org/10.1177/0266666912454025.

Brock, G. (2013). Needs and distributive justice. In *The Routledge companion to social and political philosophy* (pp. 444–455). New York: Routledge.

Campbell, J. E., & Carlson, M. (2002). Panopticon.com: Online surveillance and the commodification of privacy. *Journal of Broadcasting & Electronic Media, 46*(4), 586–606. https://doi.org/10.1207/s15506878jobem4604_6.

Carter, R. (2015, January 22). Privacy is dead, Harvard professors tell Davos forum. *Yahoo! Tech*. https://www.yahoo.com/tech/privacy-dead-harvard-professors-tell-davos-forum-144634491.html. Accessed 16 Feb 2016.

Dewey, C. (2014, October 14). The only guide to Gamergate you will ever need to read. *The Washington Post*. https://www.washingtonpost.com/news/the-intersect/wp/2014/10/14/the-only-guide-to-gamergate-you-will-ever-need-to-read/. Accessed 9 Mar 2016.

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.

Friedman, M. (2002). *Capitalism and freedom (40th anniversary ed.)*. Chicago: University of Chicago Press.

Gaus, G. F., & D'Agostino, F. (Eds.). (2013). *The Routledge companion to social and political philosophy*. New York: Routledge.

Gerner, J. (2000). Singular and plural anaphors of indefinite personal pronouns in spoken British English. In J. M. Kirk (Ed.), *Corpora galore: Analyses and techniques in describing English: Papers from the nineteenth international conference on English language research on computerised corpora (ICAME 1998)*. Amsterdam: Rodopi.

Goodin, R. E., & Pettit, P. (Eds.). (1993). *A Companion to contemporary political philosophy*. Oxford: Blackwell.

Greene, R. (2013). Petitions, privacy, and political obscurity. *Temple Law Review, 85*, 367–411.

Gurstein, M. (2011). Open data: Empowering the empowered or effective data use for everyone? *First Monday*, *16*(2). http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/3316/2764. Accessed 5 Mar 2013.

Hartzog, W., & Selinger, E. (2015). Surveillance as loss of obscurity. *Washington and Lee Law Review, 73*(3), 1343–1387.

Hill, K. (2012). *How target figured out a teen girl was pregnant before her father did*. http://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/

Johnson, B. (2010, January 10). Privacy no longer a social norm, says Facebook founder. *The Guardian*. http://www.theguardian.com/technology/2010/jan/11/facebook-privacy. Accessed 16 Feb 2016.
Johnson, J. A. (2014a). The ethics of big data in higher education. *International Review of Information Ethics, 21*, 3–10.
Johnson, J. A. (2014b). From open data to information justice. *Ethics and Information Technology, 16*(4), 263–274. https://doi.org/10.1007/s10676-014-9351-8.
Johnson, J. A. (2015). Information systems and the translation of transgender. *TSQ: Transgender Studies Quarterly, 2*(1), 160–165. https://doi.org/10.1215/23289252-2848940.
Kolm, S.-C. (1993). Distributive justice. In *A companion to contemporary political philosophy* (pp. 438–461). Oxford: Blackwell.
Locke, J. (1980). In C. B. Macpherson (Ed.), *Second treatise of government* (1st ed.). Indianapolis: Hackett Pub..
McDonald, S. N. (2014, October 15). "Gamergate": Feminist video game critic Anita Sarkeesian cancels Utah lecture after threat. *The Washington Post*. https://www.washingtonpost.com/news/morning-mix/wp/2014/10/15/gamergate-feminist-video-game-critic-anita-sarkeesian-cancels-utah-lecture-after-threat-citing-police-inability-to-prevent-concealed-weapons-at-e-vent/. Accessed 9 Mar 2016.
Neugebauer, C. (2014, October 15). Terror threat against feminist Anita Sarkeesian at USU. *Standard Examiner*. Ogden, Utah. http://www.standard.net/Police/2014/10/14/Utah-State-University-student-threatens-act-of-terror-if-feminist. Accessed 9 Mar 2016.
Nissenbaum, H. (2010). *Privacy in context: Technology, policy, and the integrity of social life*. Stanford: Stanford Law Books.
Nozick, R. (2013). *Anarchy, state, and utopia*. New York: Basic Books, a member of the Perseus Books Group.
Raman, B. (2012). The rhetoric of transparency and its reality: Transparent territories, opaque power and empowerment. *The Journal of Community Informatics*, *8*(2). http://ci-journal.net/index.php/ciej/article/view/866/909. Accessed 5 Mar 2013.
Rawls, J. (2005). *A theory of justice (Original ed.)*. Cambridge, MA: Belknap Press.
Robeyns, I. (2013). The capability approach (and social justice). In *The Routledge companion to social and political philosophy* (pp. 446–466). New York: Routledge.
Scott, J. C. (1998). *Seeing like a state: How certain schemes to improve the human condition have failed*. New Haven: Yale University Press.
Slee, T. (2012, June 25). Seeing like a geek. *Crooked Timber*. http://crookedtimber.org/2012/06/25/seeing-like-a-geek/. Accessed 5 Mar 2013.
Smith, H. J., Dinev, T., & Xu, H. (2011). Information privacy research: An interdisciplinary review. *MIS Quarterly, 35*(4), 989–1015.
Solove, D. J. (2008). *Understanding privacy*. Cambridge, MA: Harvard University Press.
Terrell, J., Kofink, A., Middleton, J., Rainear, C., Murphy-Hill, E., & Parnin, C. (2016). Gender bias in open source: Pull request acceptance of women versus men. doi:https://doi.org/10.7287/peerj.preprints.1733v1
Trauzettel-Klosinski, S., & Dietz, K. (2012). Standardized assessment of reading performance: The new international reading speed texts IReST. *Investigative Ophthalmology & Visual Science, 53*(9), 5452. https://doi.org/10.1167/iovs.11-8284.
U.S. Department of Health & Human Services Office for Civil Rights. (2013, February 7). *Your health information privacy rights*. Portable Document Format file. http://www.hhs.gov/sites/default/files/ocr/privacy/hipaa/understanding/consumers/consumer_rights.pdf. Accessed 8 Mar 2016.
van Wel, L., & Royakkers, L. (2004). Ethical issues in web data mining. *Ethics and Information Technology, 6*(2), 129–140. https://doi.org/10.1023/B:ETIN.0000047476.05912.3d.
von Hayek, F. A. (1994). *The road to serfdom* (50th anniversary ed./with a new introd. by Milton Friedman.). Chicago: University of Chicago Press.
Young, I. M. (1990). *Justice and the politics of difference*. Princeton: Princeton University Press.

Chapter 6
Structural Information Justice

Abstract This chapter engages information from the perspective of structural justice using a case study of learning analytics in higher education, drawing heavily on the "Drown the Bunnies" case at Mount St. Mary's University in 2016. This case suggests the outlines of an increasingly common approach to promoting student "success" in higher education in which early academic and non-cognitive data, often from students at other universities, are used to build a student success prediction algorithm that uses a triage approach to intervention, targeting middling students while writing off those in most need of help as inefficient uses of resources. Most common ethics approaches—privacy, individualism, autonomy, and discrimination—capture at best only part of the issues in play here. Instead, I show that a full analysis of the "Drown the Bunnies" model requires understanding the ways that social structures perpetuate oppression and domination. Attention to more just organizational, politico-economic, and intellectual structures would greatly attenuate the likelihood of cases such as the Mount St. Mary's University case, adding an important dimension to information justice. I conclude by contrasting the "Drown the Bunnies" model with an implementation of learning analytics at UVU, which did much better in part because of structural preconditions that support justice.

The explosion of student learning and behavioral analytics raises deep questions about whether it can be done within a meaningful frame of information justice. These questions that came to the forefront of public discourse in 2016 when Mount St. Mary's University President Simon Newman described using predictive student analytics to weed out students unlikely to be retained as a way to "drown the bunnies . . . put a Glock to their heads" (Svrluga 2016a). Using the Mount St. Mary's University incident as a touchstone case, this paper suggests that these concerns can best be understood within a framework of structural justice, which focuses on the ways in which the structures of predictive student analytics influence students' capacities for self-development and self-determination. I first examine four ethical concerns that arise in student analytics: privacy, individuality, autonomy, and discrimination, showing that these concepts offer some marginal critiques of "Drown the Bunnies" models but that they do little to understand the issues presented by the approach in general. I then turn to the structural concept of justice articulated by Young (1990), which finds justice rooted in aspects of social structure that promote

J.A. Johnson, *Toward Information Justice*, Public Administration and Information Technology 33, https://doi.org/10.1007/978-3-319-70894-2_6

or impede self-development and self-determination. I examine three types of structures: showing that organizational, politico-economic, and knowledge structures all pose significant challenges to justice in predictive student analytics. This approach is able to determine that "Drown the Bunnies" models are categorically disrespectful of self-development and in most cases permit little self-determination for students, demonstrating that information justice is at least equally, if not primarily, a question of structural justice.

6.1 Bunnies, Glocks, and Analytics

Predictive student analytics are algorithmic systems that use data from student behavior and performance to generate individual predictions for student outcomes.[1] Nominally, the purpose of student academics is to promote student success. But the case of Mount St. Mary's University[2] shows just how contentious definitions of student success can be and how such definitions are critical to understanding and using student analytics. Mount St. Mary's University thus provides a valuable case study of how predictive student analytics can lead to fundamental information injustices in higher education.

6.1.1 The Mount St. Mary's University Case

Mount St. Mary's University attempted to use predictive student analytics to improve the university's first-to-second year retention rates (Schisler and Golden 2016). The university had a 66% graduation rate for its 2009 cohort of first-time, full-time entering students, and 78% of its 2014 cohort returned in 2015. Both rates are well above the national average for 4-year institutions but commonly exceeded by private liberal arts colleges (Jaschnik 2016a). Those rates are based on standards used in the federal Integrated Postsecondary Education Data System (IPEDS), which identifies cohorts based on the institution's fall census date. The cohorts, with some generally minor adjustments, are the denominator for graduation and retention rates. Students who enroll as first-time, full-time students but withdraw from an

---

[1] For the purpose of this paper I will use the terms "predictive analytics" to refer to predictive methods generally and "predictive student analytics" or simply "student analytics" to refer to such methods used in educational contexts. The term includes but is not limited to the narrower concept of "learning analytics," which describes the use of such methods based on primarily academic behavioral and performance data to predict future academic performance and individuate course and program content.

[2] There are several higher education institutions in the United States with similar names. Mount St. Mary's University, the object of this case study, is in Emmitsburg, Maryland. It should not be confused with Mount St. Mary's College in Los Angeles, or Mount St. Mary College in Newburgh, New York.

institution before the census date are not included in the cohort. In Mount St. Mary's University's case, the census date was initially September 25, approximately 1 month after orientation for entering students (Schisler and Golden 2016).

Emails obtained by *The Mountain Echo*, Mount St. Mary's University's student newspaper, show that, at the suggestion of then-president and private equity investor Simon Newman, Mount St. Mary's University instituted a locally developed survey intended, ostensibly, to develop better metrics for student analytics to be administered at the student orientation. While the instructions for the survey suggested it was purely informational and designed to "help you [the student] discover more about yourself," the emails revealed that the intent of the administration was to dismiss 20–25 students before the IPEDS reporting date, based in part on the survey results:

> Newman's email continued: "My short term goal is to have 20–25 people leave by the 25th [of Sep.]. This one thing will boost our retention 4–5%. A larger committee or group needs to work on the details but I think you get the objective."

Emails from other campus leaders make clear that the use of the survey to dismiss students was made unilaterally by Newman in the face of strong opposition from those leaders, characterizing the risk that some of those dismissed might succeed as "some collateral damage." Associate Provost Leona Sevick stated that the plan would contradict catalog standards for dismissing students. In response to *The Mountain Echo's* investigation, chairman of the university's Board of Trustees John E. Coyne III characterized the program as part of "the Mount's efforts to improve student retention and to intervene early on to assure that incoming students have every opportunity to succeed at our university," and in December Newman stated in an email to the faculty, "It has never been a goal to 'kick out' first year students because they were not doing well." The plan failed to come to fruition as some campus leaders stalled the decision of whom to dismiss until after the IPEDS deadline, which the university extended by 1 week in an effort to identify students to dismiss (Schisler and Golden 2016).

The headline-grabbing aspect of the controversy was neither the intent nor the process; rather, it was a conversation between Newman and other campus officials that elevated the case to international news. Newman requested that the director of the university's first-year student symposium, Dr. Greg Murry, provide Newman with a list of students who were unlikely to return. In response to Murry's objections and in the presence of another faculty member, Newman said, "This is hard for you because you think of the students as cuddly bunnies, but you can't. You just have to drown the bunnies…put a Glock to their heads" (Schisler and Golden 2016). The exceptionally violent metaphor shocked the higher education community, generating extensive coverage based on *The Mountain Echo's* reporting. Both Newman and the Board of Trustees confirmed Newman's statement. The board characterized Newman's "Drown the Bunnies" comment as an "unfortunate metaphor" (Lee 2016), and Newman stated that he regretted the language but not the intent of his statement (Svrluga 2016a).

Table 6.1 Mount St. Mary's University Class of 2019 survey structure

Section	Content	Items
1	Introduction	0
2	New courses and programs	12
3	Internal-external locus of control scale (Rotter 1966)	5
4	Life events	8
5	Academic habits	12
6	Center for epidemiologic studies depression scale (Radloff 1977)	10
7	Connor-Davidson resilience scale (Connor and Davidson 2003)	5
8	Goz lab ten item personality measure (Gosling et al. 2003)	10
9	Delaying gratification inventory (Hoerger et al. 2011)	10
10	College students beliefs and values survey (Astin et al. 2010)	3
11	Interpersonal reactivity index (Davis 1983)	5
12	Short grit scale (Duckworth and Quinn 2009)	8
13	Miscellaneous items	11
14	Extracurricular activities	12

The controversy was aggravated by university efforts to blame the student journalists for revealing the information and then punish faculty for their opposition. Coyne initially dismissed the story as "the product of a disgruntled employee and the creative and destructive imagination of a student being spoon fed his information" (Schisler and Golden 2016). Backed by the Trustees, Newman terminated the paper's faculty advisor and a tenured philosophy professor and demoted the university's provost and a dean for "disloyalty" (though they were reinstated 3 days later). The controversy ultimately provoked an unsustainable backlash, as parents and alumni threatened to pull students and donations from the university and the university's accreditor, the Middle States Commission on Higher Education, requested an ad hoc report on implications of the process for accreditation (Jaschnik 2016b). In the face of growing opposition, Newman eventually resigned (Joseph and McPhate 2016).

Sensationalism aside, however, it is the process by which Mount St. Mary's College intended to dismiss students that raises questions of information justice in this case. The survey that Mount St. Mary's University used (see Table 6.1) was obtained and made public by *The Washington Post* (Svrluga 2016b). It consists of approximately 14 sections and 110 individual responses. It was developed locally but consisted of both locally developed individual items and items taken from a hodgepodge of psychometric instruments. Broadly, the survey addresses three topics: courses and programs, non-cognitive student characteristics, and student activities. Following the introduction, there is a section discussing possible new courses and programs that the university might offer. There are a wide range of programs listed, from philosophy, politics, and economics to civil engineering. Students were asked to rate their interest in the course or program and were offered an option expressing a willingness to pay "a small premium" to enroll in it. The bulk of the survey focuses on non-cognitive characteristics of students. This includes sections on both resilience and grit, personality inventories, sections about religious beliefs, and a section evaluating students for clinical depression. Most of these sections are abridged versions of established—in some cases, dated—psychometric instruments. The final section requests information on preferred activities (Mount St. Mary's University 2016).

It is clear that the use of the survey to identify students to be dismissed was either contrary to its designed intent or to deliberately deceptive. The introduction states:

> We firmly believe that the SAT and ACT exams, and even a GPA score, do not effectively value the potential of anyone. At The Mount, we look beyond these simple numbers and seek to understand what motivates each student, as well as understand what factors may be holding each student back from performing at his or her best. . . . We will ask you some questions about yourself that we would like you to answer as honestly as possible. There are no wrong answers. . . . We believe everyone here has potential to become someone way beyond what you may think possible right now. (Mount St. Mary's University 2016)

The appeal to self-development and support, and the claim that there are no wrong answers, is inconsistent with the aim of using the survey to identify students who will be discouraged from continuing their educations at the university. Far from becoming "someone way beyond what you may think possible right now," the survey—and much of its underlying psychometrics, designed to measure personality traits—assumes that the information on the survey represents consistent and stable aspects of students' character that are unlikely to change over the course of an academic career and are thus determinative of their success at Mount St. Mary's University.

Many in the administration at Mount St. Mary's University were well aware of and deeply concerned about the ethics of the "Drown the Bunnies" approach, and especially the survey. Dean Josh Hochschild, one of those involved in the email exchange obtained by *The Mountain Echo*, called the process "deeply disturbing," "highly intrusive and misleadingly framed," and "unethical," asking, "How can we in good conscience administer this?" Associate Provost Leona Sevick shared his concerns, noting that the dismissal process likely violated university policy stated in the catalog. Other university leaders were also cited in the article as voicing strong ethical objections. Nonetheless, Newman moved forward with the intent to dismiss students in spite of opposition, even pushing back the IPEDS reporting date (a move of dubious legality itself). The ethical objections were ineffective; ultimately the process was thwarted not by a conscious decision to do what was right but rather by stalling beyond the IPEDS deadline: Murry, the director of the first-year symposium, failed to provide the names of students to be dismissed, saying "We simply ran out the clock" (Schisler and Golden 2016).

6.1.2 A Generalized "Drown the Bunnies" Model

How the survey was to be used to identify the students who were not likely to be retained to their second year is unclear, and Newman's comment that "A larger committee or group needs to work on the details" at least suggests that there was never a concrete methodology leading from data to intervention. While a few of the source instruments were designed specifically for higher education (Astin et al. 2010) or measure constructs that have recognized applications in education (Connor and Davidson 2003; Duckworth and Quinn 2009; Hoerger et al. 2011), many of the instruments were designed for clinical use (Connor and Davidson 2003; Davis 1983; Gosling et al. 2003; Radloff 1977; Rotter 1966) and do not appear to have any established connection to educational settings. The use of the individual items frequently invalidates the source instruments, which are designed to operate as multiple-item indices (Astin et al. 2010; Connor and Davidson 2003; Davis 1983; Hoerger et al. 2011; Rotter 1966). There is no evidence of any effort by the university to test validity or reliability, either of the abridged instruments as measures of the constructs the source instruments intended to measure, or of their association with each other as an overall measure of the likelihood of academic success. And the use of questions about mental health and learning disabilities suggests serious civil rights concerns as well as responsibilities to ensure that students whose responses indicate, for instance, major depressive disorder receive appropriate mental health care. Reporting does suggest that the intent was to use the survey results in conjunction with academic performance information from the first month of classes (Schisler and Golden 2016). Absent a model that could predict how responses are associated with outcomes, the administration would be making marginally educated guesses about which students to dismiss.

While Mount St. Mary's University appeared to lack such a model, such models are increasingly available through predictive student analytics. Austin Peay State University pioneered the use of predictive student analytics through Degree Compass, a course recommendation system. Degree Compass uses student performance data from previous courses and recommends courses that would maximize students' GPA and thus their likelihood of maintaining scholarships and of graduating (Parry 2012):

> Inspired by recommendation systems implemented by companies such as Netflix, Amazon, and Pandora, Degree Compass successfully pairs current students with the courses that best fit their talents and program of study for upcoming semesters. The model combines hundreds of thousands of past students' grades with each particular student's transcript to make individualized recommendations. (Denley 2013)

The system works in part by predicting student performance in courses that the student might take, based on the performance of similar students in the past in addition to—or in the absence of—data on the student's own performance. While its stated purpose is to help students select courses that will lead to timely completion of the programs, it touts a 92% accuracy rate in predicting whether a student will receive a C grade or better in a course; Denley notes that completion success is a "hope" rather than an evidentiary claim. Following development support from the Bill and Melinda Gates Foundation, Degree Compass was acquired by learning management system vendor Desire2Learn, commercializing the product for use at any higher education institution (Denley 2013). It thus has the capability of predicting student success even for entering students who have no track record of performance, the kind of prediction that Newman intended to join to the survey results.

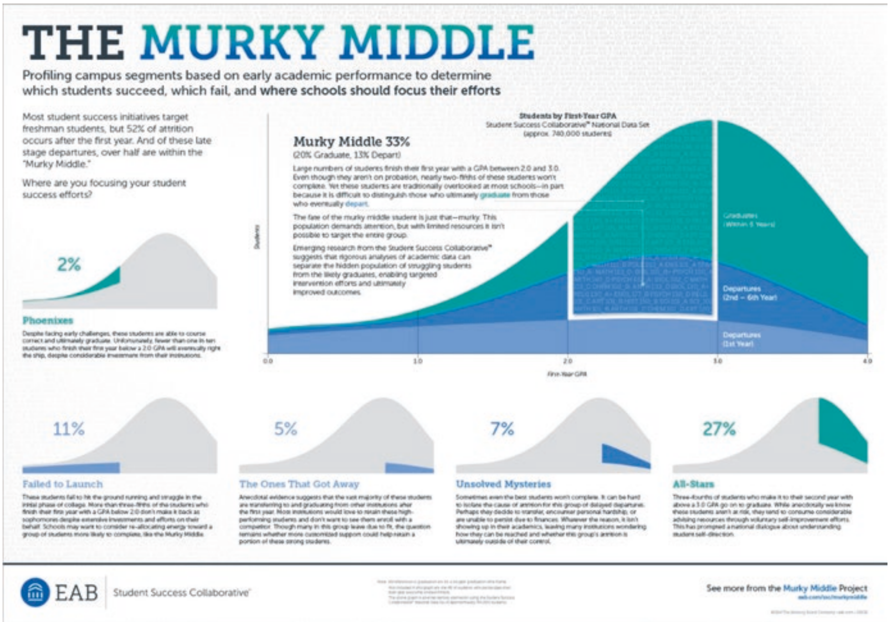

Fig. 6.1 A triage approach to student support (EAB Student Success Collaborative 2015)

The survey and academic data used by Mount St. Mary's University are also representative of recent approaches using predictive student analytics to "triage" students for support resources. Essential to nearly all triage models, whether in their origin in battlefield medicine or as incarnated in predictive student analytics for higher education, is the existence of three basic categories: those needing immediate care for a positive outcome, those who will see a positive outcome even without immediate care, and those who are beyond care. Commercial providers such as EAB (EAB Student Success Collaborative 2015) argue that predictive analytics can support student success by allowing institutions to focus their efforts on the students for whom support will make a difference, noting that the top third of students are likely to succeed without support, and the bottom third are unlikely to succeed even with support. For the latter group, EAB suggests that schools "consider re-allocating energy toward a group of students more likely to complete" (see Fig. 6.1). To do this organically, Mouth St. Mary's University would need to not only survey the class at entry but wait until students returned in their second year to build a model of retention. The survey could not be used to predict retention of the class of 2019, the first to take the survey, but it would allow the university to predict the outcomes of subsequent classes, especially combining the pre-enrollment non-cognitive data with past student performance data along the lines of the methods used by Degree Compass.

While the specifics of Mount St. Mary's University are certainly problematic, the importance of the case for information justice in higher education comes in the context of initiatives like the EAB Student Success Collaborative and Degree Compass. This combination of broad collection of academic and non-cognitive data, predictive analytics methods, and triage-driven intervention intended to directly influence key quantitative indicators of success is now the state of the art in student services for retention and completion. Many of its characteristics can be seen in applications of predictive analytics in other areas such as criminal justice, where modelling identifies defendants who have the highest likelihood of failing to appear in court or of offending again and sets higher bail amounts, longer sentences, or recommends against parole, in some cases effectively denying bail or imposing life sentences for relatively minor crimes (Christin et al. 2015). This represents a significant new challenge in higher education management, especially from the perspective of meeting universities' responsibilities to students, and is a major area of application for information justice. In honor of former President Newman, I call this combination the "Drown the Bunnies" model.

6.2 Student Analytics from an Ethics Perspective

Student analytics presents a complex array of challenges to information justice, many of which fall within existing perspectives on information ethics. It is not at all immediately obvious that higher education institutions should not select those students most likely to continue through their programs provided the institutions do not do so for their own benefit rather than that of the students. Especially at private institutions, higher education is a costly undertaking in money and time for even the most prepared students and their families. Beginning a program and failing to complete it is a significant use of resources to often negligible benefit. This is especially so for students who receive federal financial aid and then withdraw early in the term. These students are often obligated to repay the aid they received before they can receive further aid; one such semester can amount to a de facto lifetime bar from further financial aid where the repayment obligation exceeds the students' earnings capacity. Especially where this includes repayment of Pell Grants, such students will often be the least able to pay even for less expensive institutions on their own. This is the justification for the requirement that students receiving federal financial aid either have a high school diploma or equivalent, or demonstrate their ability to benefit from the institution's programs under the Higher Education Act of 1965 (as amended; 20 U.S.C. 1091[d]). While most evidence in the Mount St. Mary's University case points to a desire to improve a high-profile statistic rather than to genuinely meet the needs of students, a less cynical and more sophisticated approach could not be so easily dismissed. Under what circumstances would such a process be acceptable?

The array of challenges that inform this question are easy to subsume under a simple notion of ethical action, but complexity emerges quickly. In part, it emerges simply because the challenges are often problems of unintended consequences and from the fact that the challenges can weigh against each other. Complexity also emerges because these challenges are often driven not by the kinds of choices about action usually addressed by professional ethics codes but by unfounded beliefs about data science itself that hide the ethical choices in learning analytics behind a belief in technological neutrality and information objectivity. I have argued previously (Johnson 2014a) that there are at least four major ethical concerns in student analytics: privacy, individuality, autonomy, and discrimination. But here I wish to argue that these are of limited value for understanding a "Drown the Bunnies" student analytics model.

6.2.1 Privacy

The most visible challenge related to student analytics is privacy. Most commonly, privacy is understood (to the extent that it is not, as Solove (2008, p. 1) puts it, "a concept in disarray") in terms of information flows (Nissenbaum 2010): privacy rights are protected by preventing information from moving from those who hold it legitimately to those who have no right to such information. This, it would seem at first glance, is a minor consideration in student analytics, as information flows voluntarily from students to institutions (as in the survey), does not flow at all (institutionally generated information such as enrollments and grades never leaves the institution), or flows from the institution to a vendor acting as an agent of the institution and for which the institution remains ethically (and to varying extents legally) responsible. The strongest concern has been the flow of information to vendors who combine institutions' data for analysis, such as the controversial InBloom project in secondary education (Singer 2014). This is characteristic of a number of applications in higher education as well; the EAB Student Success Collaborative is an example of a process that explicitly touts the benefits of cross-institutional data collection to provide more accurate analytics (EAB Student Success Collaborative 2015). This view of privacy would unequivocally hold that the "Drown the Bunnies" model presents no concerns, since information is voluntarily provided and does not flow to third parties except as agents of the institutions themselves.

But this is a quite narrow view of both information flows and of privacy. In some cases, student analytics are inferring information about students that the students themselves would be reluctant to divulge to either the institution or an individual faculty member, essentially creating an electronic reputation that a student may not want to precede them into the classroom. Arizona State University, for example, developed a system to predict students who intended to transfer that included tracking card swipes used to access campus facilities (Parry 2012). Since students who transfer out have the same effect on institutions' retention and graduation rates as they would if they simply dropped out altogether, one might imagine such information feeding a triage process that effectively restricts potential transfer students' access to support services. To the extent that consent is an essential feature of privacy, these would be privacy violations even though information flowed only from one campus data system to another. Nissenbaum's solution to this is a well-founded preference for controlling the flow of information across contexts rather than merely across actors; privacy is violated not only when information that is provided to one actor flows to another without the consent of the subject but also when it flows from one use to another by the same actor without at least justification if not consent.

The contextual integrity perspective does raise some concerns about a "Drown the Bunnies" model of student analytics. The introduction to the survey establishes the context as self-development: manifesting the students' "potential to become someone way beyond what you may think possible right now." For this information to then be used to identify students to be dismissed (whether formally or through encouragement to pursue other educational options) is a transfer of information across contexts, as the initial justification for collecting the information no longer applies. But as I have argued previously (Johnson 2016a), the contextual integrity is maintained as long as a new and sound justification for the use of the information is available; a renewal of consent is not itself required. We have seen that such a justification can be provided. The question thus becomes whether the initial context represented a good-faith description of the intent or whether it was a deception of the students. In the Mount St. Mary's University case, it was at best a knowingly disingenuous administration of an instrument that may, at one point, have been intended to be used as described; the administration's intent in using the survey to dismiss students was well-established by the time the survey was administered. As such, it should be considered unacceptable. But this, too, is a very narrow ground for analysis, as it tells us nothing about the justice of a "Drown the Bunnies" approach generally, only that it is wrong to deceive students about it.

6.2.2 Individuality

The recognition of student individuality is also problematic in predictive student analytics. Here, in one sense, is one of the great advantages of predictive student analytics: while traditional statistics reduces all students to the central tendency (whether of a set of variables individually as in descriptive statistics or of a relationship among variables as in inferential statistics), the big data techniques behind student analytics can treat each case individually, recognizing its uniqueness and allowing analysis of the diversity of cases (Two Crows Corporation 2005). But student analytics presents the opposite challenge, disaggregating students into bundles of unconnected information meant to correspond to human beings, bundles of information that Floridi calls "interconnected information organisms, or *inforgs*" (Floridi 2010, p. 9, emphasis in original). In this case, students' individuality remains compromised by denying the existence of the narratives and meanings that make them a coherent being.

For example, a decision tree model used by one large state university to predict retention classified students first by first-term grade point average and then by whether they had visited their advisor. The prediction was useful, leading to a simple retention message: "Go to class. Do the Work. See your advisor. Graduate." But it is also problematic for many students. In all likelihood these characteristics are not distinct; a discursive or ethnographic lens might reveal them as part of an identity in which the students see themselves as students rather than, for example, as workers who are taking classes to advance on the job. Those who see themselves as students take continued enrollment for granted (rather than consciously deciding whether to enroll for each semester), visit their advisor regularly, and attend class and do their assignments (which is usually sufficient to earn at least a C grade) because students do each of these things: the determinative factor is not a causal relationship among the behaviors but rather their presence in the role of "student," part of a cognitive script that the successful students follow. The dehumanization of the student analytics approach is then imposed on the students as institutional decisions are made based not on humanistic complexes of individual and social meaning but on mechanical processes of measurement, classification, and response. The institution may require that all students must see their advisors before enrolling in order to move students to an apparently more successful classification. The analytics process that seemed to point to concrete, actionable information that can help students succeed may have the opposite effect, reinforcing a non-student identity that is an obstacle to success for those who struggle to find extra time in their schedules to see their advisors regularly.

This role-driven interpretation is a significant issue for the Mount St. Mary's University case, and to the extent that non-cognitive evaluations are increasingly common in retention and completion efforts they are likely to be present across "Drown the Bunnies" models. For the sake of generosity and parsimony, in analyzing the Mount St. Mary's University survey, we can assume that the CES-D Scale (Radloff 1977), Ten Item Personality Measure (Gosling et al. 2003), and Interpersonal Reactivity Index (Davis 1983) are included for the stated purposes of self-discovery and support, as these have the most tenuous connections to educational success. The new courses and programs section, with its response option in which students would pay a premium for the program, does raise some possibility that they could be used to identify students with potential economic value to the university, but merits benefits of the doubt in favor of its apparent purpose of general market research because of the difficulty of capturing that value for students who are already enrolled. This leaves a number of sections, however, that test academic behavior and personality traits (most notably "grit" or "resilience") that are purported to be associated with academic success. The use of this or similar surveys in predictive retention models is thus as likely to test whether students see themselves in the role of a student, especially where they also use early first-term academic data that is more reflective of students' preconceptions of what higher education is like than of their capacity to learn and adapt to the reality of the institution.

To be sure, this approach is prejudicial to those who see their primary role as something other than a student. But this is not a dispositive consideration. Most institutions will have unique elements to their institutional cultures into which some students will fit better than others, and fit into the institutional culture is likely to be a factor in retention and completion. At selective universities, most students who fail to return for their second year are unlikely to have simply dropped out of school;

they are likely to have transferred to an institution where they are more comfortable—where the institutional culture is a better fit for them. At some institutions, this might play to the favor of those who do not see themselves as students first, where a culture that supports working students with existing social networks proves alienating to those who take on a more traditional student role. It is not unreasonable to think that an institution might want to help incoming students make this decision quickly, and such an approach would in fact respect the individuality of students. Mount St. Mary's University might be criticized for an inept approach to this—the purpose would be better served by administering the survey as part of the admissions process or before students arrive on campus so that students could choose more appropriate institutions for the fall. But since respect for the principle of individuality is dependent on the specifics of the approach, individuality does not resolve concerns about "Drown the Bunnies" models generally.

6.2.3 Autonomy

In principle, students are understood to be adults, capable of making their own decisions: they are autonomous individuals who can critically think and act to further their own good. Student analytics can be deeply challenging to this autonomy. Rarely, such systems are outright coercive, but one could imagine developing such systems by, for instance, linking student activity data from a learning management system to financial aid awards. Rather than relying on end of semester grades, an institution might condition aid on keeping up on work performed over the course of the semester: reading materials accessed, assignments completed, etc. In conjunction with an "Aid Like a Paycheck" disbursement schedule (MDRC 2016), this would essentially condition students' immediate ability to meet their financial needs on compliance with course requirements.

More common than fully coercive systems are those that operate more paternalistically. Austin Peay State University developed Degree Compass based on the belief that students made poor choices regarding which classes to take. This system was designed to counter perceived lack of information and laziness among students: "They choose on the basis of course descriptions or to avoid having to wake up for an 8 a.m. class on Monday," Provost Tristan Denley says enrollment prior to implementing the system. Such considerations led to unwise (in the view of the administration) choices about what courses to take in order to graduate as quickly as possible (Parry 2012). The solution is for the institution to make choices for the students, albeit by nudging rather than coercing them, steering them to the wise decision that the administration would like them to make. A more complex version of this uses the control of minutiae and the possibility of constant surveillance—core features of the typical marriage of LMS and analytics suite—to create a disciplinary environment (Foucault 1995, pp. 146–149) in which the structure of the situation leads students comply with the institution's preferences on their own. Systems that classify students on their course activity and communicate that classification to the student and the instructor (Parry 2011) operate within a framework of disciplinary power rather than a classically coercive one but nonetheless severely constrain student autonomy and agency.

The application of autonomy principles to education, however, is eternally problematic. The authoritarian structure of the just *polis* in Plato's Republic emerges from the need to educate the guardian class so that they are gentle to friends and ferocious to enemies; the titular character of Rousseau's treatise on education, *Emile*, concludes his education by telling his educator, "I have decided to become as you have made me," as ambiguous an endorsement of autonomy as one can imagine. Education is traditionally premised on the idea that the educators know better, in some fashion, than those they educate, and that the students must submit themselves, in some fashion, to their educators' guidance. Neither paternalism nor discipline is inherently unethical, though most notions of autonomy would at least forbid deception of the student. It is certainly unreasonable for an institution that maintains a formal commitment to the good of its students—especially where that commitment has legally binding effect and the institution received benefits from doing so—to seek its own benefit in the form of more favorable federally reported data under the guise of paternalism. But this is merely to say that acts of paternalism must be genuinely paternalistic. It is not a *prima facie* unreasonable argument to contend that bunnies should be drowned where the institution knows that the bunnies would not survive on their own and where they are drowned only in the sense that they are directed to more suitable alternatives. This is essentially Scherer's and Anson's (2014) argument against open access community colleges (an argument that is, to be sure, deeply problematic in many respects). One could hardly argue that these conditions are met in the Mount St. Mary's University case, but this is, again, only to argue that the university did a poor job of drowning its bunnies.

6.2.4 Discrimination

One of the emerging issues in big data more broadly is its potential for "algorithmic discrimination" and "digital redlining" (Gilliard 2016). Biases in data and in the social practices underlying it are increasingly seen as incorporating bias into analytics processes. As Cegłowski puts it, "We have turned to machine learning, an ingenious way of disclaiming responsibility for anything. Machine learning is like money laundering for bias" (Cegłowski 2016). A series of ProPublica investigations have found that the same algorithm, trained on data from a variety of different news sources, produced wildly different results. When the algorithm used by Google to identify synonyms for searches is trained on data from web sites of the left-leaning *Huffington Post* and *The Nation*, the most common search synonym for "BlackLivesMatter" is "hashtag" (i.e., one receives results for "BlackLivesMatter," "#BlackLivesMatter," and "BlackLivesMatter hashtag"). But when trained on data from the right-leaning *Daily Caller* and *Breitbart*, the most common synonym is "AllLivesMatter," an antonym rather than a synonym for the original search (Larson et al. 2016).

Thus when a search is executed, the differently trained algorithms will generate different search results. Both algorithms will show results presenting claims of institutionalized racism in the criminal justice system made by the Black Lives Matter movement. But the conservatively trained algorithm automatically replies that "All Lives Matter," a reply originally crafted to discredit the arguments of the movement, while the liberally trained one will lead to social media that present a widespread, grassroots social movement built around "#BlackLivesMatter" that may or may not be representative of activism offline. In either case, the training data has given the algorithm a form of political agency. This can be deliberate but it need not be; the investigations also found that Asians were twice as likely to be offered a higher than average price for Princeton Review SAT prep courses, which the company argues is an incidental effect of its geographic pricing algorithm (Angwin et al. 2015). In this case, geography likely became an effective proxy for race, producing a disparate effect presumably without intentional discrimination. Of course, when that disparate effect is protected by unsound claims of the absolute soundness of the method, the line between disparate effect and intent becomes severely blurred.

These effects, Scholes (2016) argues, are readily seen in predictive student analytics as well, especially where group membership is used to predict risk scores. She contends that effective instructional design can mitigate these risks by designing analytics around effort measurements and dynamic factors. While this is laudable as an effort to overcome demographic discrimination, the recent debate over the nature of "grit" as a character trait promoting success suggests effort measures are no better. As reviewed recently by Ris (2015), grit is alternately seen as precisely the kind of dynamic, effort-based measure that Scholes advocates or as a product of a safe, White, privileged upbringing (e.g., Calarco 2014) in which opportunities to recover from failure abound. This makes grit a proxy for traditional race and class in victim-blaming explanations for failure. To the extent that this critique of grit is true, using an instrument that measures grit in an algorithm writes racial discrimination into educational practices by laundering it through "a clean, mathematical apparatus that gives the status quo the aura of logical inevitability" (Cegłowski 2016).

Grit is a central feature of the Mount St. Mary's University survey. Four sections with a total of 28 items come from instruments measuring grit (the Short Grit Scale), resilience (the Connor-Davidson Resilience Scale), or attitudes often discussed in the grit literature (the Internal-External Locus of Control Scale and the Delaying Gratification Inventory). These items make up half of the psychometric aspects of the survey. To the extent that these traits are artifacts of diverging racial experiences in which Whites exist as autonomous selves to non-White, and especially Black, others who are acted on by raced structures, it is entirely nonsensical to expect non-Whites and Whites to hold, at the same rate, attitudes reflecting an internal rather than external locus of control, to see failure as a growth opportunity rather than an occasion to expect systemic punishment (especially in the context of a race-loaded school-to-prison pipeline), or to believe that working hard and delaying gratification will lead to better outcomes in the future rather than subject their rewards to arbitrary seizure, as through discriminatory civil forfeiture process. Moreover, there is at least some evidence to suggest that grit can be taught (Pratt

2014), in which case it makes little sense to dismiss students based on grit. The Mount St. Mary's University survey is, from this perspective, very much a case of digital redlining. Given that grit is a key concept in non-cognitive student assessment generally, it is hard to see how such assessments can avoid algorithmic discrimination. The ethical perspective on predictive student analytics thus challenges, in strong terms, using grit as part of a "Drown the Bunnies" model, and calls attention more generally to the need to ensure the data and models that are used in predictive student analytics are not sleight-of-math ways of repeating social inequalities.

None of this is to say that predictive student analytics are inherently just, unjust, or lack justice considerations at all. Rather, the Mount St. Mary's University case shows that student analytics are not simply value-neutral math based on objective data; they are systems of classification and management, and moments of ethical agency. Justice in student analytics is not just a matter of good faith and good intentions. Each time an analytics system is implemented, those implementing it face unavoidable ethical questions. For example, autonomy is a paramount value in American society, but there are some circumstances in which it can be violated. One must therefore answer the question of whether a particular implementation of student analytics violates autonomy and, if so, whether the particular violation is permissible. One must then consider whether the tradeoff between transforming a human student into an inforg and the gains to that student from the analytics process are consistent with the institution's responsibility to its students. One must evaluate whether the data chosen to underlie the algorithm on which the intervention is based are more than just the biases of the collectors and analysts. And one must question whether other values might take precedence over autonomy, as an absolute claim for the primacy of autonomy (or privacy, or efficiency, etc.) is by no means uncontroversial. To raise these questions is not to criticize learning analytics and argue for a blanket injunction against the field. Questions like these are inherent in every analytics process, ones that must be answered uniquely in each case.

6.3 Structural Justice

The preceding analysis, unfortunately, does little to suggest that ethical frameworks alone are adequate to the challenge of building predictive student analytics processes that do not raise deep concerns. But even if it did, there would remain a deeper challenge. At Mount St. Mary's University, knowledge of what is ethical—clearly present in abundance among everyone involved other than Newman—was insufficient by itself to produce a just outcome. This raises a consideration that ethics, with its focus on the justification of individual action, is poorly suited to engage. The decisive factor at Mount St. Mary's University was the balance of political power: the authority of the university president and the inherent gaps in the principal-agent relationship between him and his employees. This reflects fundamentally a changed conception of the terrain of student analytics: a movement from pedagogy to politics. We think of educational institutions as rather far removed from such questions, especially when considering questions of pedagogy. But it is easy to overlook that the actors driving the adoption of student analytics are rarely instructors; most commonly a coalition of administrators and vendors introduce analytics process to the institution, and then address use by instructors as seeking "buy-in" (implicitly, to a *fait accompli*) from faculty. This should encourage the view that predictive student analytics is a management process as much as an instructional or student support one. It is thus necessary to examine not only the myriad ethical considerations presented by "Drown the Bunnies" models but also these political relationships from the perspective of a coherent conception of justice itself.

Rawls holds that justice is "the first virtue of *social institutions*" (Rawls 2005, p. 3, emphasis added), most commonly of political institutions associated with the nation-state. Political philosophers have taken two approaches to justice. The most common, distributive justice, considers the social institutions and practices of a community (which would include, for many philosophers, voluntary associational relationships like higher education institutions) to be just if they reflect a just distribution of material and moral goods (for example, rights, liberties, or authority). There are myriad theories of distributive justice varying both by focus on processes versus outcomes and by standards for determining a just distribution of goods.

Young (1990), however, argues that distributive approaches are inadequate when the question is one of relationships among people or groups rather than material goods. Such approaches fail to understand how social institutions shape "action, decisions about action, and provision of the means to develop and exercise capacities" (1990, p. 16); these are as much a consequence of structural factors as they are distributions of material and moral goods. Instead, Young presents a structural theory of justice: a society is just to the extent that social structures and relationships facilitate both the capacity to develop and exercise one's human capacities (that is, self-development) and supports one's participation in determining their actions (self-determination). Likewise, injustice has two corresponding forms. The denial of self-development is oppression; that of self-determination, domination. These conditions are not matters of either distribution or of individual ethics; they are part of the social structure:

> Oppression in this sense is structural, rather than the result of a few people's choices or policies. Its causes are embedded in unquestioned norms, habits, and symbols, in the assumptions underlying institutional rules and the collective consequences of following those rules. . . . In this extended structural sense oppression refers to the vast and deep injustices some groups suffer as a consequence of often unconscious assumptions and reactions of well-meaning people in ordinary interactions, media and cultural stereotypes, and structural features of bureaucratic hierarchies and market mechanisms—in short, the normal processes of everyday life. We cannot eliminate this structural oppression by getting rid of the rulers or making some new laws, because oppressions are systematically reproduced in major economic, political, and cultural institutions. (Young 1990, p. 41)

Young uses this framework to understand and articulate the claims of injustice posed by various emancipatory social movements, suggesting that it is a particularly valuable approach for information justice.

Simultaneously protecting both self-development and self-determination is difficult. There are always some constraints on self-development and self-determination; especially in education, structural justice is constrained by the sometimes inherent conflict between the two. Most pedagogical approaches assume that learning requires systematic guidance and thus limit students' self-determination in order to further their capacities for self-development. Moreover, the two considerations are mutually constructive. Self-development begets self-determination, and vice versa, and both dimensions of structural justice create the self as they reflect it. Perrotta and Williamson show that "Methods used for the classification and measurement of online education are partially involved in the creation of the realities they claim to measure" (Perrotta and Williamson 2016, p. 2). Students exist as students in part because we choose to measure them as students; measuring them as students creates the role of student discussed above as well as reflects a role that students bring themselves. At many institutions, for instance, continuing education students are not included in reported enrollment data. All institutions exclude non-degree-seeking students when reporting graduation rates to the federal IPEDS data system, but they report such students in the IPEDS enrollment survey. These decisions shape who institutions think of—and plan for—when they refer to their students as well as which attendees at the institution think of themselves as students. The decision of whom to count as students plays a substantial role in creating the identity of "student." As such, predictive student analytics, like all educational practices, do not just constrain or facilitate self-development or self-determination; they in fact contribute to the creation of the selves that seek development and determination.

Structural justice is thus a far more complex form of justice than distributive justice, matching the complexity of the underlying relationships it seeks to evaluate. This is why most structural approaches to justice argue not for a maximal distribution of self-development and self-determination but rather against social structures that limit these, and see justice as a political rather than rational process, as the outcome of a negotiation among conflicting groups that itself respects self-development and self-determination rather than the result of an algorithmic determination (as, for instance, Rawls' two principles of justice would have one perform). Just as being governed by the winner of a free and fair election who respects the right of opposition and the rule of law does not deprive one of one's self-determination, neither does a student who willing submits to the tutelage of an expert who is committed to that student's self-development do so. It is the capacity of the student to self-determine that recognizing the expertise of the teacher is the best path toward the student's own self-development, and their participation in institutional decisions that will exercise students' capacities for self-determination collectively (e.g., by making "all students shalt" policies), that enables a just relationship between teacher and student. Likewise, that exercise of self-determination requires a prerequisite level of self-development. This interplay is best captured by interactions through social institutions that respect the self-determination and self-development of all, that is, by political processes.

The concept of justice being a structural condition (or at least of having a structural dimension as a major component) that must be addressed through political processes shows why ethical approaches alone are unlikely to resolve concerns about the "Drown the Bunnies" model: the model—indeed all models—consists not just of individual practices that should be proscribed or mandated by ethical "law" (a law that lacks sanction, one notes), but also of capacities and concepts embedded in the institutional conditions of the university. If those processes do not respect self-development and self-determination, no amount of ethical reasoning will be effective, as we saw in the Mount St. Mary's University case. The key to sound predictive student analytics, then, is attention to the structural conditions in them that determine the extent to which student analytics supports the students' self-development and self-determination. The structural determinants of the Mount St. Mary's University case, directions characteristic of most instances of predictive student analytics, take many different forms.

6.3.1 Organizational Structures

The organizational authority structure of the university is the most readily apparent structure that influences predictive student analytics. The decisions about how to implement student analytics take place in organizations. The most obvious aspect of this is the organizational hierarchies of the university. When discussing justice in student analytics, there is a tendency to focus immediately on the expected intervention, which draws attention to the intent and competence of actors and the organizational roles that define those actors. Simon Newman's beliefs about the best way to manage student retention were irrelevant to the fortunes of Mount St. Mary's University students until he became President Simon Newman and could translate those beliefs into intentions and institutional actions. And they remained relevant only as long as he had the backing of other organizational structures (the Board of Trustees), who in turn could provide that support only as long as they were supported by empowered stakeholders such as donors who could vote with their pocketbooks and students who could vote with their feet.

Newman's demand that the university dismiss students unlikely to be retained reflected his (perceived) position of authority in the institution, a perception informed by a view of the university as a business organization. As president he acted as if he had full authority over the institution and was responsible for maximizing return on investment, as he would in the private equity firms from which he came to Mount St. Mary's university. We see this not only in the demand that the survey be used for this purpose without informing students but also in his disregard for institutional policies regarding dismissal procedures and in his firing of faculty members for disloyalty. This left students with neither the opportunity to participate in the decision (either directly or through the faculty members who were voicing the students' interests) nor the ability to make an informed decision about participation. This domination would have led to oppression of at least the 20–25 students Newman intended to dismiss, not just in that they were denied the educational opportunity at Mount St. Mary's University but that they were denied so in a way that prevented them from pursuing other educational opportunities. Had the students had meaningful input, the university might well have seen the virtue of using student analytics in the admissions process rather than after students arrived, a solution many institutions and commercial analytics vendors are now touting.

And yet, even with Newman's assertion of absolute authority backed by the Board of Trustees, the president did not get his way. Organizational hierarchies are more complex than might appear. Newman made the assumption of a bureaucratic hierarchy working toward the accomplishment of a monistic, self-interested aim characteristic of the business world. Failing to question that norm may well have prevented him from understanding the other norms at work in the university, especially those of shared governance and responsibility to serve the students rather than the university. These norms empowered the faculty and other senior administrators to oppose Newman through the organizational authorities, protections, and relationships created by shared governance principles. Tenured faculty members received a notice of termination but were ultimately not terminated in part because of the legal protections the university's tenure process gave them. Administrators lost their administrative positons, which at many institutions are held at the pleasure of a top-level institutional leader, but tenure protected their faculty positions. Those protections enabled resistance as much as the authorities of their positions. And that resistance allowed opponents of the process to take advantage of the inherent gaps that exist in any system of bureaucratic authority, creating slippage between nominal principles and nominal agents (see, e.g., Kassim and Menon 2003) that thwarted Newman's aims.

Ultimately, the "Drown the Bunnies" model is an exercise of managerial authority. Student analytics is a management process, one that affirms the authority of the institution—a social structure in which the student participates—over the student. Student analytics first makes students "legible" (See Scott 1998) to the institution so that the administration can describe and understand the environment within which the institution acts: The Mount St. Mary's University survey identifies different types of students who have varying needs and will likely follow different paths in their academic careers. Analytics then makes the behavior of the student (at least appear to be) predictable to the institution: Degree Compass generates a predicted likelihood of a student passing a course in order to generate a recommendation. And finally analytics is used to control the environment itself by providing a basis for reliable action: EAB targets student support resources on the "Murky Middle" so that institutions can maximize their retention and graduation rates. All of these efforts ultimately make the institution's actions more reliable and more likely to achieve their ends, enhancing its capacity to act on—thus to exercise authority over—its students. This is a significant shift in self-determination, and to the extent that it is driven by institutional rather than student interests, a significant limitation on students' self-development.

6.3.2 Political and Economic Structures

Bureaucratic hierarchy is by no means the only organizational structure that shapes the justice of "Drown the Bunnies" models. The institution, the state, and the political economy all structure the content and use of student analytics. Analytics processes may use hundreds of variables, but institutions and the state have chosen what variables will be available to the system by choosing what data to collect, how to store it, and what to make public. This puts emphasis on certain highly visible data points and may prevent analysis of others. Within institutions, these data structures often leave students entirely uninvolved or only include often unrepresentative student organizations. This is unquestionably a major factor in the Mount St. Mary's University case. Newman appears to have been driven solely by the desire to improve the federally reported first-year retention rate, and pursued a method that did so only because of the federal formula for identifying the cohort. If the cohort were defined differently—for instance, if it included all newly admitted students who enrolled before the first day of classes—Newman's approach would have done nothing to change the retention rate; every student who took the survey would have been included in the cohort. Had Newman been president of a public university with an early census date he would not have had time to dismiss students or the ability to push the institutional reporting date back as he did. And had he been obligated to consider input from students, parents, or other stakeholders who ultimately opposed his process he might have had to consider whether the better response would be to develop an institutional narrative on the weaknesses of the retention rate for the university, as many institutions do for the highly unrepresentative IPEDS graduation rate. The federal policy context, especially to the extent that it is influenced by non-democratic factors such as industry lobbying and policy biases toward traditional students from traditional families attending traditional institutions (Institutional Effectiveness and Planning 2015) rather than the contemporary non-traditional student majority, incentivizes institutions to undermine student self-determination.

Newman's efforts to shape retention rates also followed major pushes by the Obama administration to enact a ratings process for institutions (first the Postsecondary Institution Ratings System and then a revamped College Scorecard) that gave high priority to retention rates and thus increased pressure on institutions to report higher rates. In this respect, Mount St. Mary's University is typical of the commercial analytics vendors being used as part of "Drown the Bunnies" models, whose products are usually designed to support students included in the IPEDS graduation rate cohort and not retention of new students generally. Students, in the face of pressures on institutions that are to be met with commercial analytics products, serve as a means to an end that is in the institution's interest, a resource to be used to shape the retention rate that the institution reports, rather than as the end that the retention rate measures: There is all the difference in the world between improving student retention and improving IPEDS retention rates, and the focus on the latter at the expense of students' interests—or of defining the latter as exclusively the former—is a form of structural oppression.

Beyond the institutions themselves the political economy of predictive analytics both situates systems within intellectual property law that makes them "black boxes" opaque to examination and, as development is often a commercialization of one institution's system, makes generalizability an assumption rather than demonstrating it: Political and economic power thus reinforce scientism. The failure to examine these structures makes an ethics-only approach likely to fail. For example, although Boon (2016) recommends sharing data with students to empower them in informed, data-driven decision-making and involves students directly in decisions about student analytics, much of the student involvement in her model takes place in an environment that has already been strongly constrained by institutional decisions and systems. Students have a modicum of autonomy within a much deeper system of constraint that makes "informed decision-making" much more a matter of institutionally driven disciplinarity.

This was not an issue in the Mount St. Mary's University case, as they did not reach the point of using commercial student analytics packages. But it must be a serious concern for the "Drown the Bunnies" model more broadly. Austin Peay State University commercialized Degree Compass, first expanding its use to other institutions in Tennessee—demonstrating it at "schools with a broad cross-section of curricular structures and student populations" where "their data offer us an opportunity to further refine the predictive modeling"—and then selling it to Desire2Learn where it could be integrated into its learning management system and connect to institutional data. While developers tout Degree Compass as using "modeling techniques could calibrate themselves to differing institutional settings and student populations" (Denley 2013). But it does not provide access to the model itself. A competing product in use at another institution is structured so that retention predictions can only be accessed individually, preventing the kind of bulk data extraction necessary for the institution to test the accuracy of its predictions for the campus. The institution's representatives were told this was to prevent reverse engineering the algorithm, which was protected intellectual property. Under such conditions, institutions—and the students on whom their policies act—are unable to interrogate these systems' algorithms or results and understand the basis of its recommendations not because of technical limitations, but because of how a profit-making entity uses law to protect its economic interests, undermining the students' informed participation in a process that could easily lead to their dismissal from the university.

6.3.3 Knowledge Structures

Organizational structures, whether within the university, the state and society, or the economy, are certainly significant to achieving justice in predictive student analytics. But this focus on organizational structures obscures knowledge structures involved in predictive student analytics. These structures, found at intersections of policy, science, and technology, are equally important in protecting students from oppression and domination. The broader structural context of analytics algorithms is built on the assumption that predictive learning analytics is, in fact, predictive. This belief is upheld in many cases by scientism, the ideology that science is the only path to true knowledge and that scientific knowledge is inherently and unquestionably objective (Hyslop-Margison and Naseem 2007; Peterson 2003). In predictive analytics, scientism especially reflects an extreme version of traditional positivist science. Observation and law-like generalization are foundational to information science in spite of decades of challenges to this approach in the social sciences and education. It is, for example, common for analytics reports to quote naïve error rates rather than proportional reduction in error measures and to attribute causation even when using models that do not support causal interpretation (Baradwaj and Pal 2011; Delavari et al. 2008; Llorente and Morant 2011; Parry 2012; Thomas and Galambos 2004; Vialardi et al. 2009). Model choice depends on assumptions about reality and intent, but these are rarely interrogated because of hyperpositivist beliefs about the efficacy of predictive analytics.

Scientism then hides other knowledge structures in which justice interests are found. An analytics process is part of a nexus in which problems, data and models, and interventions mutually support and inform each other. The underlying data is not an objective representation of reality but rather the end result of a translation process that is as much technical as it is social. These regimes of knowledge are structures laden with questions of justice, such as the recognition of particular racial or ethnic groups which in turn allows such groups to be represented in—or excluded from—decision-making by being included in or excluded from the data (Johnson 2016b). These structures confer intellectual authority on developers while shutting down critical inquiry with flippant injunctions against arguing with facts and dismissive contrasts between sound data and unfounded instinct. This is not a consequence of bad data or bad models; these challenges are inherent in predictive analytics and are protected—along with those who use them to further their ends—by the structure of scientific knowledge claims itself.

This may be the most oppressive aspect of the "Drown the Bunnies" model: We believe our methods accurately tell us which bunnies to drown, for we have science. Those who suggest otherwise—for example, faculty members at Mount St. Mary's University who claimed that the data were inadequate to the intended use or students who can speak to considerations that are not quantifiable or even simply not collected because it didn't occur to anyone to do so—are denied the legitimacy of their claims by structures of knowledge that exclude information external to the process as non-knowledge. Students' self-development is hampered because data science says, with high confidence in precise quantitative scores and an acceptably low error (or, in Newman's words, "collateral damage") rate, that they are unsuccessful students. They believe that they are not in control of their lives. They lack the resilience to keep going after failure. They are introverted. They are depressed. Essentialized into who they were at the moment they took a survey—an assumption unsupportable at the most basic level in survey research (Zaller 1992)—the problem-model-intervention nexus cuts off, seemingly without human action (though of course not in fact so, since a human wrote the algorithm and decision rules that the model applies), an avenue for self-development.

6.3.4 A Structurally (More) Just Alternative

Alternatives are the heart of structural justice. The recurring theme of this book is Young's "It does not have to be this way. It could be otherwise." The "Drown the Bunnies" model is not the only way that universities can use predictive analytics, and these alternatives can be far more just.[3] UVU is in the process of adopting a student analytics system that has significantly more promise than those rooted in the "Drown the Bunnies" approach. While details are still being determined as the system is implemented, enough is now known to provide a useful contrast. UVU uses a commercial vendor providing an analytics platform but not a fully-developed algorithm; the algorithm is based on data only for UVU's students. It is a long-term tracking system that provides advisors with information to prioritize those students most in need of advising support rather than using a triage model, and will provide data on most UVU students rather than just the first-time, full-time cohort. It has far more robust privacy controls to ensure that data provided by students to their advisors is not available to others on campus or to the algorithm driving the analytics engine. It will also allow UVU to evaluate the success of student support programs and target those programs to the students who will benefit most from specific programs (Student Success and Retention 2017).

This system is more just than the "Drown the Bunnies" model on two levels. Immediately, it is apparent that this system is designed to support self-determination and self-development. There is a different relationship between the university and the student in the UVU system, a mentoring and development one rather than one in which students are resources with which to pursue the university's interests. Students are given guidance from professional advisors whose aim is to help students succeed and are put in positions of making informed decisions rather than informed of institutional determinations designed to encourage students to fail in ways that help the university. And all students are valued in the UVU system, not just those that can improve GRS-based retention and graduation rates. It can be seen philosophically as an implementation of the Kantian categorical imperative—very much a structural concept of justice in that it establishes a relationship among people—to always treat persons as ends in themselves and never as means only.

On a deeper level, the structural justice in UVU's student analytics process reflects broader structural justices at the university. UVU's mission and core themes are strongly embedded in the university's culture (Utah Valley University 2014). One of these core themes is "Inclusive," which puts great emphasis on UVU's open admission practices. There is thus an institution to counter both the GRS-based metrics and the political economy of software development, introducing ambiguity and complexity into the logics of appropriateness. A system that includes all students rather than those who help retention rates is a foreseeable—though by no means inevitable—consequence of the institutionalized mission. UVU could not

[3] To some extent I must question whether they can be fully just; the simple act of acting on a statistical prediction seems inherently unjust in that it assumes that current conditions are permanent traits and therefore that individuals have little capacity to change themselves. That is, to my mind at least, antithetical to the entire purpose of education.

conceivably adopt a triage strategy without undermining its mission. That in turn leads to a recognition that UVU's students are also quite different from a typical institution, driving UVU's efforts to find systems that develop algorithms specific to the university's students that go beyond an "algorithm-in-a-box," and reshaping the intervention strategy to providing information that would support a strong advisor-student relationship. UVU's organizational culture supports these structural advantages. The university has a more participatory decision-making culture than many institutions. This drove a procurement and implementation process that included—indeed, was led by—advising and student affairs rather than just information technology and institutional research, and thus shaped a very different set of values than is seen in "Drown the Bunnies" approaches.

6.4 Conclusion

This structural analysis suggests that the "Drown the Bunnies" model fails because it is structurally unjust; it oppresses and dominates students. Students' self-determination is undermined by organizational forms that establish paternalistic—literally, *in loco parentis*—authority over them. Using this authority, students' self-development is subordinated to the needs of institutions, governments, and vendors. The well-noted ethical concerns are most often a consequence of the organizational, political, and knowledge structures of student analytics: Privacy, individuality, autonomy, and discrimination are likely to be addressed most effectively where analytics processes aim at self-development and support self-determination.

Unfortunately, the analysis above suggests that the scope for individual practitioners to influence analytics in order to secure information justice is limited. While they can make some choices as best they can—particularly if there is sensitivity to structural concerns and a desire to support students' self-development and self-determination rather than just institutional performance funding formulae and key performance indicators—justice in predictive student analytics is a fundamentally structural challenge.

References

Angwin, J., Mattu, S., & Larson, J. (2015, September 1). The tiger mom tax: Asians are nearly twice as likely to get a higher price from Princeton review. *ProPublica*. https://www.propublica.org/article/asians-nearly-twice-as-likely-to-get-higher-price-from-princeton-review. Accessed 11 Nov 2016.

Astin, A. W., Astin, H. S., & Lindholm, J. A. (2010). *Cultivating the spirit: How college can enhance students' inner lives*. San Francisco: Jossey-Bass.

Baradwaj, B. K., & Pal, S. (2011). Mining educational data to analyze students' performance. *International Journal of Advanced Computer Science and Applications, 2*(6), 63–69.

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.

Calarco, J. M. (2014). Coached for the classroom: Parents' cultural transmission and children's reproduction of educational inequalities. *American Sociological Review, 79*(5), 1015–1037. https://doi.org/10.1177/0003122414546931.
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.
Christin, A., Rosenblat, A., & Boyd, Danah. (2015). *Courts and predictive algorithm*. Presented at the data & civil rights: A new era of policing and justice, Washington, DC. http://www.datacivilrights.org/pubs/2015-1027/Courts_and_Predictive_Algorithms.pdf
Connor, K. M., & Davidson, J. R. T. (2003). Development of a new resilience scale: The Connor-Davidson Resilience Scale (CD-RISC). *Depression and Anxiety, 18*(2), 76–82. https://doi.org/10.1002/da.10113.
Davis, M. H. (1983). Measuring individual differences in empathy: Evidence for a multidimensional approach. *Journal of Personality and Social Psychology, 44*(1), 113–126. https://doi.org/10.1037/0022-3514.44.1.113.
Delavari, N., Phon-Amnuaisuk, S., & Beizadeh, M. R. (2008). Data mining application in higher learning institutions. *Informatics in Education, 7*(1), 31–54.
Denley, T. (2013, September 4). Degree compass: A course recommendation system. *Educause Review Online*. Accessed 21 Feb 2017.
Duckworth, A. L., & Quinn, P. D. (2009). Development and validation of the short grit scale (Grit–S). *Journal of Personality Assessment, 91*(2), 166–174. https://doi.org/10.1080/00223890802634290.
EAB Student Success Collaborative. (2015, January 21). *The Murky middle project*. https://www.eab.com/technology/student-success-collaborative/members/white-papers/the-murky-middle-project. Accessed 18 Feb 2017.
Floridi, L. (2010). *Information: A very short introduction*. Oxford: Oxford University Press.
Foucault, M. (1995). *Discipline and punish: The birth of the prison* (2nd ed.). New York: Vintage Books.
Gilliard, C. (2016). *From redlining to digital redlining*. Presented at the OLC Innovate, New Orleans, Louisiana.
Gosling, S. D., Rentfrow, P. J., & Swann, W. B. (2003). A very brief measure of the Big-Five personality domains. *Journal of Research in Personality, 37*(6), 504–528. https://doi.org/10.1016/S0092-6566(03)00046-1.
Hoerger, M., Quirk, S. W., & Weed, N. C. (2011). Development and validation of the delaying gratification inventory. *Psychological Assessment, 23*(3), 725–738. https://doi.org/10.1037/a0023286.
Hyslop-Margison, E. J., & Naseem, M. A. (2007). *Scientism and education empirical research as neo-liberal ideology*. Dordrecht: Springer. http://public.eblib.com/EBLPublic/PublicView.do?ptiID=337528. Accessed 21 February 2014.
Institutional Effectiveness and Planning. (2015). *Postsecondary Institution Ratings System (PIRS) evaluation*. Orem: Utah Valley University.
Jaschnik, S. (2016a, January 20). Are at-risk students bunnies to be drowned? *Inside Higher Ed*. https://www.insidehighered.com/news/2016/01/20/furor-mount-st-marys-over-presidents-alleged-plan-cull-students. Accessed 14 Feb 2017.
Jaschnik, S. (2016b, February 26). Tough questions for Mount St. Mary's. *Inside Higher Ed*. https://www.insidehighered.com/news/2016/02/26/accreditor-demands-answers-mount-st-marys-numerous-standards. Accessed 14 Feb 2017.
Johnson, J. A. (2014). The ethics of big data in higher education. *International Review of Information Ethics, 21*, 3–10.
Johnson, J. A. (2016a). *The value—and limits—of distributive justice in information privacy*. Presented at the Eastern Sociological Society 2016 Annual Meeting Digital Sociology Mini-Conference, Boston, Massachusetts.
Johnson, J. A. (2016b). Representing "Inforgs" in data-driven decisions. In J. Daniels, K. Gregory, & T. McMillan Cottom (Eds.), *Digital sociologies*. Bristol: Policy Press.

Joseph, Y., & McPhate, M. (2016, February 29). Mount St. Mary's president quits after firings seen as retaliatory. *The New York Times*. https://www.nytimes.com/2016/03/02/us/simon-newman-resigns-as-president-of-mount-st-marys.html?_r=0. Accessed 14 Feb 2017.

Kassim, H., & Menon, A. (2003). The principal-agent approach and the study of the European Union: Promise unfulfilled? *Journal of European Public Policy, 10*(1), 121–139. https://doi.org/10.1080/1350176032000046976.

Larson, J., Angwin, J., & Parris Jr., T. (2016, October 19). Breaking the black box: How machines learn to be racist. *ProPublica*. https://www.propublica.org/article/breaking-the-black-box-how-machines-learn-to-be-racist?word=blacklivesmatter. Accessed 11 Nov 2016.

Lee, P. Y. (2016, February 10). "Drown the bunnies": Mount St. Mary's president fires faculty for backlash against his "put a Glock to their heads" freshman retention plan. *Salon*. http://www.salon.com/2016/02/10/drown_the_bunnies_mount_st_marys_president_fires_faculty_for_backlash_against_his_put_a_glock_to_their_heads_freshman_retention_plan/. Accessed 14 Feb 2017.

Llorente, R., & Morant, M. (2011). Data mining in higher education. In K. Funatsu (Ed.), *New fundamental technologies in data mining* (pp. 201–220). New York: InTech. http://www.intechopen.com/books/new-fundamental-technologies-in-data-mining/data-mining-in-higher-education.

MDRC. (2016). Aid like a paycheck. *MDRC*. http://www.mdrc.org/project/aid-paycheck. Accessed 11 Nov 2016.

Mount St. Mary's University. (2016, August). *Mount St. Mary's University Class of 2019 survey*. https://www.scribd.com/doc/299020003/Class-of-2019-survey

Nissenbaum, H. (2010). *Privacy in context: technology, policy, and the integrity of social life*. Stanford: Stanford Law Books.

Parry, M. (2011, December 11). Colleges mine data to tailor students' experience. *The Chronicle of Higher Education*. https://chronicle.com/article/A-Moneyball-Approach-to/130062/

Parry, M. (2012, July 18). College degrees, designed by the numbers. *The Chronicle of Higher Education*. https://chronicle.com/article/College-Degrees-Designed-by/132945/.

Perrotta, C., & Williamson, B. (2016). The social life of learning analytics: Cluster analysis and the "performance" of algorithmic education. *Learning, Media and Technology*, 1–14. https://doi.org/10.1080/17439884.2016.1182927.

Peterson, G. R. (2003). Demarcation and the scientistic fallacy. *Zygon, 38*(4), 751–761. https://doi.org/10.1111/j.1467-9744.2003.00536.x.

Pratt, A. A. (2014, May 9). *University retention and resilience: Resilience and hardiness in undergraduate students* (Undergraduate Thesis). Orem: Utah Valley University. Retrieved from http://contentdm.uvu.edu/cdm/ref/collection/UVUTheses/id/648

Radloff, L. S. (1977). The CES-D scale: A self-report depression scale for research in the general population. *Applied Psychological Measurement, 1*(3), 385–401. https://doi.org/10.1177/014662167700100306.

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

Ris, E. W. (2015). Grit: A short history of a useful concept. *Journal of Educational Controversy, 10*(1), 3.

Rotter, J. B. (1966). Generalized expectancies for internal versus external control of reinforcement. *Psychological Monographs: General and Applied, 80*(1), 1–28. https://doi.org/10.1037/h0092976.

Scherer, J. L., & Anson, M. L. (2014). *Community colleges and the access effect: Why open admissions suppresses achievement*. New York: Palgrave Macmillan.

Schisler, R., & Golden, R. (2016, January 19). Mount president's attempt to improve retention rate included seeking dismissal of 20–25 first-year students. *The Mountain Echo*. Mount St. Mary's University, Emmitsburg. http://msmecho.com/2016/01/19/mount-presidents-attempt-to-improve-retention-rate-included-seeking-dismissal-of-20-25-first-year-students/. Accessed 14 Feb 2017.

Scholes, V. (2016). The ethics of using learning analytics to categorize students on risk. *Educational Technology Research and Development, 64*(5), 939–955. https://doi.org/10.1007/s11423-016-9458-1.

Scott, J. C. (1998). *Seeing like a state: How certain schemes to improve the human condition have failed*. New Haven: Yale University Press.
Singer, N. (2014, April 22). InBloom student data repository to close. *The New York Times*, New York, p. B2.
Solove, D. J. (2008). *Understanding privacy*. Cambridge, MA: Harvard University Press.
Student Success and Retention. (2017, April 26). *Student analytics demonstration*. Presentation at Utah Valley University, Orem, Utah.
Svrluga, S. (2016a, January 19). University president allegedly says struggling freshmen are bunnies that should be drowned. *The Washington Post: Grade Point*. https://www.washingtonpost.com/news/grade-point/wp/2016/01/19/university-president-allegedly-says-struggling-freshmen-are-bunnies-that-should-be-drowned-that-a-glock-should-be-put-to-their-heads/. Accessed 17 Nov 2016.
Svrluga, S. (2016b, February 12). The freshman survey that rang alarm bells for some at Mount St. Mary's. *The Washington Post Grade Point*. https://www.washingtonpost.com/news/grade-point/wp/2016/02/12/the-freshman-survey-that-rang-alarm-bells-for-some-at-mount-st-marys/?utm_term=.84f0075063f9. Accessed 22 Feb 2017.
Thomas, E., & Galambos, N. (2004). What satisfies students? Mining student-opinion data with regression and decision tree analysis. *Research in Higher Education, 45*(3), 251–269.
Two Crows Corporation. (2005). *Introduction to data mining and knowledge discovery* (3rd ed.). Potomac: Two Crows Corporation. http://www.twocrows.com/intro-dm.pdf.
Utah Valley University. (2014, September). *Mid-cycle self-evaluation report*. https://www.uvu.edu/accreditation/docs/uvu_mid_cycle_evaluation.pdf. Accessed 16 May 2017.
Vialardi, C., Bravo, J., Shafti, L., & Ortigosa, A. (2009). Recommendation in higher education using data mining techniques. In T. Barnes, M. Desmarais, C. Romero, & S. Ventura (Eds.), *Educational data mining 2009: 2nd international conference on educational data mining, proceedings* (pp. 190–199). Cordoba: International Working Group on Educational Data Mining. http://www.educationaldatamining.org/EDM2009/uploads/proceedings/vialardi.pdf.
Young, I. M. (1990). *Justice and the politics of difference*. Princeton: Princeton University Press.
Zaller, J. (1992). *The nature and origins of mass opinion*. Cambridge: Cambridge University Press.

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.