

1 The Power Chapter

Principle: Examine Power

Data feminism begins by analyzing how power operates in the world.

When tennis star Serena Williams disappeared from Instagram in early September 2017, her six million followers assumed they knew what had happened. Several months earlier, in March of that year, Williams had accidentally announced her pregnancy to the world via a bathing suit selfie and a caption that was hard to misinterpret: “20 weeks.” Now, they thought, her baby had finally arrived.

But then they waited, and waited some more. Two weeks later, Williams finally reappeared, announcing the birth of her daughter and inviting her followers to watch a video that welcomed Alexis Olympia Ohanian Jr. to the world.¹ The video was a montage of baby bump pics interspersed with clips of a pregnant Williams playing tennis and having cute conversations with her husband, Reddit cofounder Alexis Ohanian, and then, finally, the shot that her fans had been waiting for: the first clip of baby Olympia. Williams was narrating: “So we’re leaving the hospital,” she explains. “It’s been a long time. We had a lot of complications. But look who we got!” The scene fades to white, and the video ends with a set of stats: Olympia’s date of birth, birth weight, and number of grand slam titles: 1. (Williams, as it turned out, was already eight weeks pregnant when she won the Australian Open earlier that year.)

Williams’s Instagram followers were, for the most part, enchanted. But soon, the enthusiastic congratulations were superseded by a very different conversation. A number of her followers—many of them Black women like Williams herself—fixated on the comment she’d made as she was heading home from the hospital with her baby girl. Those “complications” that Williams experienced—other women had had them too. In Williams’s case, the complications had been life-threatening, and her self-advocacy in the hospital played a major role in her survival.

On Williams's Instagram feed, dozens of women began posting their own experiences of childbirth gone horribly wrong. A few months later, Williams returned to social media—Facebook, this time—to continue the conversation (figure 1.1). Citing a 2017 statement from the US Centers for Disease Control and Prevention (CDC), Williams wrote that “Black women are over 3 times more likely than white women to die from pregnancy- or childbirth-related causes.”²

These disparities were already well-known to Black-women-led reproductive justice groups like SisterSong, the Black Mamas Matter Alliance, and Raising Our Sisters Everywhere (ROSE), some of whom had been working on the maternal health crisis for decades. Williams helped to shine a national spotlight on them. The mainstream media also recently had begun to pay more attention to the crisis as well. A few months earlier, Nina Martin of the investigative journalism outfit ProPublica, working with Renee Montagne of NPR, had reported on the same phenomenon.³ “Nothing Protects Black Women from Dying in Pregnancy and Childbirth,” the headline read. In addition to the study cited by Williams, Martin and Montagne cited a second study from 2016, which showed that neither education nor income level—the

Serena Williams January 15, 2018 · Facebook Creator ·

I didn't expect that sharing our family's story of Olympia's birth and all of complications after giving birth would start such an outpouring of discussion from women — especially black women — who have faced similar complications and women whose problems go unaddressed.

These aren't just stories: according to the CDC, (Center for Disease Control) black women are over 3 times more likely than White women to die from pregnancy- or childbirth-related causes. We have a lot of work to do as a nation and I hope my story can inspire a conversation that gets us to close this gap.

Let me be clear: EVERY mother, regardless of race, or background deserves to have a healthy pregnancy and childbirth. I personally want all women of all colors to have the best experience they can have. My personal experience was not great but it was MY experience and I'm happy it happened to me. It made me stronger and it made me appreciate women -- both women with and without kids -- even more. We are powerful!!!

I want to thank all of you who have opened up through online comments and other platforms to tell your story. I encourage you to continue to tell those stories. This helps. We can help others. Our voices are our power.

27K 1.6K Comments 1.4K Shares 840K Views

Share

Figure 1.1

A Facebook post by Serena Williams responding to her Instagram followers who had shared their stories of pregnancy and childbirth-related complications with her. Image from Serena Williams, January 15, 2018.

factors usually invoked when attempting to account for healthcare outcomes that diverge along racial lines—impacted the fates of Black women giving birth.⁴ On the contrary, the data showed that Black women with college degrees suffered more severe complications of pregnancy and childbirth than white women without high school diplomas.

So what were these complications, more precisely? And how many women had actually died as a result? Nobody was counting. A 2014 United Nations report, coauthored by SisterSong, described the state of data collection on maternal mortality in the United States as “particularly weak.”⁵ The situation hadn’t improved in 2017, when ProPublica began its reporting. In 2018, *USA Today* investigated these racial disparities, and found what was an even more fundamental problem: there was still no national system for tracking complications sustained in pregnancy and childbirth, even though similar systems had long been in place for tracking any number of other health issues, such as teen pregnancy, hip replacements, or heart attacks.⁶ They also found that there was still no reporting mechanism for ensuring that hospitals follow national safety standards, as is required for both hip surgery and cardiac care. “Our maternal data is embarrassing,” stated Stacie Geller, a professor of obstetrics and gynecology at the University of Illinois, when asked for comment. The chief of the CDC’s Maternal and Infant Health branch, William Callaghan, makes the significance of this “embarrassing” data more clear: “What we choose to measure is a statement of what we value in health,” he explains.⁷ We might edit his statement to add that it’s a measure of *who* we value in health, too.⁸

Why did it take the near-death of an international sports superstar for the media to begin paying attention to an issue that less famous Black women had been experiencing and organizing around for decades? Why did it take reporting by the predominantly white mainstream press for US cities and states to begin collecting data on the issue?⁹ Why are those data still not viewed as big enough, statistically significant enough, or of high enough quality for those cities and states, and other public institutions, to justify taking action? And why didn’t those institutions just #believeblack-women in the first place?¹⁰

The answers to these questions are directly connected to larger issues of power and privilege. Williams recognized as much when asked by *Glamour* magazine about the fact that she had to demand that her medical team perform additional tests in order to diagnose her own postnatal complications—and because she was Serena Williams, twenty-three-time grand slam champion, they complied.¹¹ “If I wasn’t who I am, it could have been me,” she told *Glamour*, referring to the fact that the privilege she experienced as a tennis star intersected with the oppression she experienced as a Black

woman, enabling her to avoid becoming a statistic herself. As Williams asserted, “that’s not fair.”¹²

Needless to say, Williams is right. It’s absolutely not fair. So how do we mitigate this unfairness? We begin by examining systems of power and how they intersect—like how the influences of racism, sexism, and celebrity came together first to send Williams into a medical crisis and then, thankfully, to keep her alive. The complexity of these intersections is the reason that *examine power* is the first principle of data feminism, and the focus of this chapter. Examining power means naming and explaining the forces of oppression that are so baked into our daily lives—and into our datasets, our databases, and our algorithms—that we often don’t even see them. Seeing oppression is especially hard for those of us who occupy positions of privilege. But once we identify these forces and begin to understand how they exert their potent force, then many of the additional principles of data feminism—like challenging power (chapter 2), embracing emotion (chapter 3), and making labor visible (chapter 7)—become easier to undertake.

Power and the Matrix of Domination

But first, what do we mean by *power*? We use the term *power* to describe the current configuration of structural privilege and structural oppression, in which some groups experience unearned advantages—because various systems have been designed by people like them and work for people like them—and other groups experience systematic disadvantages—because those same systems were not designed by them or with people like them in mind. These mechanisms are complicated, and there are “few pure victims and oppressors,” notes influential sociologist Patricia Hill Collins. In her landmark text, *Black Feminist Thought*, first published in 1990, Collins proposes the concept of the *matrix of domination* to explain how systems of power are configured and experienced.¹³ It consists of four domains: the structural, the disciplinary, the hegemonic, and the interpersonal. Her emphasis is on the intersection of gender and race, but she makes clear that other dimensions of identity (sexuality, geography, ability, etc.) also result in unjust oppression, or unearned privilege, that become apparent across the same four domains.

The *structural domain* is the arena of laws and policies, along with schools and institutions that implement them. This domain organizes and codifies oppression. Take, for example, the history of voting rights in the United States. The US Constitution did not originally specify who was authorized to vote, so various states had different policies that reflected their local politics. Most had to do with owning property, which,

Table 1.1

The four domains of the matrix of domination

Structural domain Organizes oppression: laws and policies.	Disciplinary domain Administers and manages oppression. Implements and enforces laws and policies.
Hegemonic domain Circulates oppressive ideas: culture and media.	Interpersonal domain Individual experiences of oppression.

Chart based on concepts introduced by Patricia Hill Collins in *Black Feminist Thought: Knowledge, Consciousness, and the Politics of Empowerment*.

conveniently, only men could do. But with the passage of the Fourteenth Amendment in 1868, which granted the rights of US citizenship to those who had been enslaved, the nature of those rights—including voting—were required to be spelled out at the national level for the first time. More specifically, voting was defined as a right reserved for “male citizens.” This is a clear instance of codified oppression in the structural domain.

It would take until the passage of the Nineteenth Amendment in 1920 for most (but not all) women to be granted the right to vote.¹⁴ Even still, many state voting laws continued to include literacy tests, residency requirements, and other ways to indirectly exclude people who were not property-owning white men. These restrictions persist today, in the form of practices like dropping names from voter rolls, requiring photo IDs, and limits to early voting—the burdens of which are felt disproportionately by low-income people, people of color, and others who lack the time or resources to jump through these additional bureaucratic hoops.¹⁵ This is the *disciplinary domain* that Collins names: the domain that administers and manages oppression through bureaucracy and hierarchy, rather than through laws that explicitly encode inequality on the basis of someone’s identity.¹⁶

Neither of these domains would be possible without the *hegemonic domain*, which deals with the realm of culture, media, and ideas. Discriminatory policies and practices in voting can only be enacted in a world that already circulates oppressive ideas about, for example, who counts as a citizen in the first place. Consider an anti-suffragist pamphlet from the 1910s that proclaims, “You do not need a ballot to clean out your sink spout.”¹⁷ Pamphlets like these, designed to be literally passed from hand to hand, reinforced preexisting societal views about the place of women in society. Today, we have animated GIFs instead of paper pamphlets, but the hegemonic function is the same: to consolidate ideas about who is entitled to exercise power and who is not.

The final part of the matrix of domination is the *interpersonal domain*, which influences the everyday experience of individuals in the world. How would you feel if you were a woman who read that pamphlet, for example? Would it have more or less of an impact if a male family member gave it to you? Or, for a more recent example, how would you feel if you took time off from your hourly job to go cast your vote, only to discover when you got there that your name had been purged from the official voting roll or that there was a line so long that it would require that you miss half a day's pay, or stand for hours in the cold, or ... the list could go on. These are examples of how it *feels* to know that systems of power are not on your side and, at times, are actively seeking to take away the small amount of power that you do possess.¹⁸

The matrix of domination works to uphold the undue privilege of *dominant* groups while unfairly oppressing *minoritized* groups. What does this mean? Beginning in this chapter and continuing throughout the book, we use the term *minoritized* to describe groups of people who are positioned in opposition to a more powerful social group. While the term *minority* describes a social group that is comprised of fewer people, *minoritized* indicates that a social group is actively devalued and oppressed by a dominant group, one that holds more economic, social, and political power. With respect to gender, for example, men constitute the dominant group, while all other genders constitute minoritized groups. This remains true even as women actually constitute a majority of the world population. *Sexism* is the term that names this form of oppression. In relation to race, white people constitute the dominant group (racism); in relation to class, wealthy and educated people constitute the dominant group (classism); and so on.¹⁹

Using the concept of the matrix of domination and the distinction between dominant and minoritized groups, we can begin to examine how power unfolds in and around data. This often means asking uncomfortable questions: who is doing the work of data science (and who is not)? Whose goals are prioritized in data science (and whose are not)? And who benefits from data science (and who is either overlooked or actively harmed)?²⁰ These questions are uncomfortable because they unmask the inconvenient truth that there are groups of people who are disproportionately benefitting from data science, and there are groups of people who are disproportionately harmed. Asking these *who* questions allows us, as data scientists ourselves, to start to see how privilege is baked into our data practices and our data products.²¹

Data Science by Whom?

It is important to acknowledge the elephant in the server room: the demographics of data science (and related occupations like software engineering and artificial

intelligence research) do not represent the population as a whole. According to the most recent data from the US Bureau of Labor Statistics, released in 2018, only 26 percent of those in “computer and mathematical occupations” are women.²² And across all of those women, only 12 percent are Black or Latinx women, even though Black and Latinx women make up 22.5 percent of the US population.²³ A report by the research group AI Now about the diversity crisis in artificial intelligence notes that women comprise only 15 percent of AI research staff at Facebook and 10 percent at Google.²⁴ These numbers are probably not a surprise. The more surprising thing is that those numbers are getting worse, not better. According to a research report published by the American Association of University Women in 2015, women computer science graduates in the United States peaked in the mid-1980s at 37 percent, and we have seen a steady decline in the years since then to 26 percent today (figure 1.2).²⁵ As “data analysts” (low-status number crunchers) have become rebranded as “data scientists” (high status

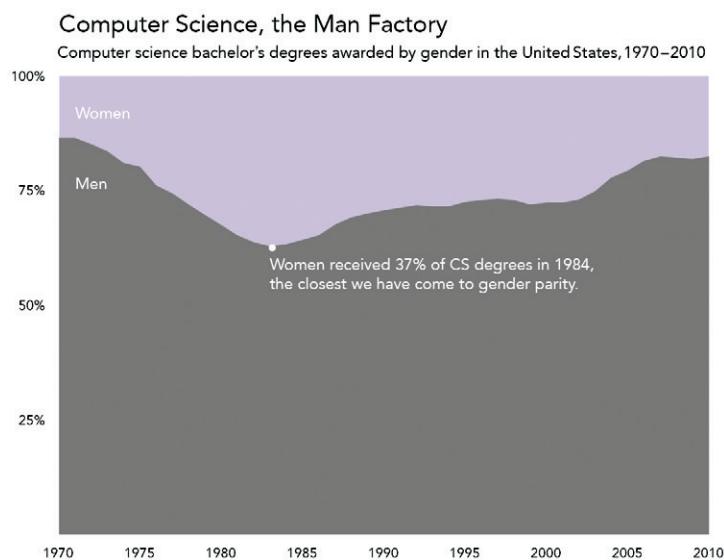


Figure 1.2

Computer science has always been dominated by men and the situation is worsening (even while many other scientific and technical fields have made significant strides toward gender parity). Women awarded bachelor's degrees in computer science in the United States peaked in the mid-1980s at 37 percent, and we have seen a steady increase in the ratio of men to women in the years since then. This particular report treated gender as a binary, so there was no data about nonbinary people. Graphic by Catherine D'Ignazio. Data from the National Center for Education Statistics.

researchers), women are being pushed out in order to make room for more highly valued and more highly compensated men.²⁶

There are not disparities only along gender lines in the higher education pipeline. The same report noted specific underrepresentation for Native American women, multiracial women, white women, and all Black and Latinx people. So is it really a surprise that each day brings a new example of data science being used to disempower and oppress minoritized groups? In 2018, it was revealed that Amazon had been developing an algorithm to screen its first-round job applicants. But because the model had been trained on the resumes of prior applicants, who were predominantly male, it developed an even stronger preference for male applicants. It downgraded resumes with the word *women* and graduates of women's colleges. Ultimately, Amazon had to cancel the project.²⁷ This example reinforces the work of Safiya Umoja Noble, whose book, *Algorithms of Oppression*, has shown how both gender and racial biases are encoded into some of the most pervasive data-driven systems—including Google search, which boasts over five billion unique web searches per day. Noble describes how, as recently as 2016, comparable searches for “three Black teenagers” and “three white teenagers” turned up wildly different representations of those teens. The former returned mugshots, while the latter returned wholesome stock photography.²⁸

The problems of gender and racial bias in our information systems are complex, but some of their key causes are plain as day: the data that shape them, and the models designed to put those data to use, are created by small groups of people and then scaled up to users around the globe. But those small groups are not at all representative of the globe as a whole, nor even of a single city in the United States. When data teams are primarily composed of people from dominant groups, those perspectives come to exert outsized influence on the decisions being made—to the exclusion of other identities and perspectives. This is not usually intentional; it comes from the ignorance of being on top. We describe this deficiency as a *privilege hazard*.

How does this come to pass? Let’s take a minute to imagine what life is like for someone who epitomizes the dominant group in data science: a straight, white, cisgender man with formal technical credentials who lives in the United States. When he looks for a home or applies for a credit card, people are eager for his business. People smile when he holds his girlfriend’s hand in public. His body doesn’t change due to child-birth or breastfeeding, so he does not need to think about workplace accommodations. He presents his social security number in jobs as a formality, but it never hinders his application from being processed or brings him unwanted attention. The ease with which he traverses the world is invisible to him because it has been designed for people

just like him. He does not think about how life might be different for everyone else. In fact, it is difficult for him to imagine that at all.

This is the *privilege hazard*: the phenomenon that makes those who occupy the most privileged positions among us—those with good educations, respected credentials, and professional accolades—so poorly equipped to recognize instances of oppression in the world.²⁹ They lack what Anita Gurumurthy, executive director of IT for Change, has called “the empiricism of lived experience.”³⁰ And this lack of lived experience—this evidence of how things truly *are*—profoundly limits their ability to foresee and prevent harm, to identify existing problems in the world, and to imagine possible solutions.

The privilege hazard occurs at the level of the individual—in the interpersonal domain of the matrix of domination—but it is much more harmful in aggregate because it reaches the hegemonic, disciplinary and structural domains as well. So it matters deeply that data science and artificial intelligence are dominated by elite white men because it means there is a collective privilege hazard so great that it would be a profound surprise if they could actually identify instances of bias prior to unleashing them onto the world. Social scientist Kate Crawford has advanced the idea that the biggest threat from artificial intelligence systems is not that they will become smarter than humans, but rather that they will hard-code sexism, racism, and other forms of discrimination into the digital infrastructure of our societies.³¹

What’s more, the same cis het white men responsible for designing those systems lack the ability to detect harms and biases in their systems once they’ve been released into the world.³² In the case of the “three teenagers” Google searches, for example, it was a young Black teenager that pointed out the problem and a Black scholar who wrote about the problem. The burden consistently falls upon those more intimately familiar with the privilege hazard—in data science as in life—to call out the creators of those systems for their limitations.

For example, Joy Buolamwini, a Ghanaian-American graduate student at MIT, was working on a class project using facial-analysis software.³³ But there was a problem—the software couldn’t “see” Buolamwini’s dark-skinned face (where “seeing” means that it detected a face in the image, like when a phone camera draws a square around a person’s face in the frame). It had no problem seeing her lighter-skinned collaborators. She tried drawing a face on her hand and putting it in front of the camera; it detected that. Finally, Buolamwini put on a white mask, essentially going in “whiteface” (figure 1.3).³⁴ The system detected the mask’s facial features perfectly.

Digging deeper into the code and benchmarking data behind these systems, Buolamwini discovered that the dataset on which many of facial-recognition algorithms are tested contains 78 percent male faces and 84 percent white faces. When she did



Figure 1.3

Joy Buolamwini found that she had to put on a white mask for the facial detection program to “see” her face. Buolamwini is now founder of the Algorithmic Justice League. Courtesy of Joy Buolamwini.

an intersectional breakdown of another test dataset—looking at gender and skin type together—only 4 percent of the faces in that dataset were women and dark-skinned. In their evaluation of three commercial systems, Buolamwini and computer scientist Timnit Gebru showed that darker-skinned women were up to forty-four times more likely to be misclassified than lighter-skinned males.³⁵ It’s no wonder that the software failed to detect Buolamwini’s face: both the training data and the benchmarking data relegate women of color to a tiny fraction of the overall dataset.³⁶

This is the privilege hazard in action—that no coder, tester, or user of the software had previously identified such a problem or even thought to look. Buolamwini’s work has been widely covered by the national media (by the *New York Times*, by CNN, by the *Economist*, by *Bloomberg BusinessWeek*, and others) in articles that typically contain a hint of shock.³⁷ This is a testament to the social, political, and technical importance of the work, as well as to how those in positions of power—not just in the field of data science, but in the mainstream media, in elected government, and at the heads of corporations—are so often surprised to learn that their “intelligent technologies” are not so intelligent after all. (They need to read data journalist Meredith Broussard’s book *Artificial Unintelligence*).³⁸ For another example, think back to the introduction of this

book, where we quoted Shetterly as reporting that Christine Darden’s white male manager was “shocked at the disparity” between the promotion rates of men and women. We can speculate that Darden herself wasn’t shocked, just as Buolamwini and Gebru likely were not entirely shocked at the outcome of their study either. When sexism, racism, and other forms of oppression are publicly unmasked, it is almost never surprising to those who experience them.

For people in positions of power and privilege, issues of race and gender and class and ability—to name only a few—are OPP: other people’s problems. Author and anti-racist educator Robin DiAngelo describes instances like the “shock” of Darden’s boss or the surprise in the media coverage of Buolamwini’s various projects as a symptom of the “racial innocence” of white people.³⁹ In other words, those who occupy positions of privilege in society are able to remain innocent of that privilege. Race becomes something that only people of color have. Gender becomes something that only women and nonbinary people have. Sexual orientation becomes something that all people *except* heterosexual people have. And so on. A personal anecdote might help illustrate this point. When we published the first draft of this book online, Catherine told a colleague about it. His earnestly enthusiastic response was, “Oh great! I’ll show it to my female graduate students!” To which Catherine rejoined, “You might want to show it to your other students, too.”

If things were different—if the 79 percent of engineers at Google who are male were specifically trained in structural oppression before building their data systems (as social workers are before they undertake social work)—then their overrepresentation might be very slightly less of a problem.⁴⁰ But in the meantime, the onus falls on the individuals who already feel the adverse effects of those systems of power to prove, over and over again, that racism and sexism exist—in datasets, in data systems, and in data science, as in everywhere else.

Buolamwini and Gebru identified how pale and male faces were overrepresented in facial detection training data. Could we just fix this problem by diversifying the data set? One solution to the problem would appear to be straightforward: create a more representative set of training and benchmarking data for facial detection models. In fact, tech companies are starting to do exactly this. In January 2019, IBM released a database of one million faces called Diversity in Faces (DiF).⁴¹ In another example, journalist Amy Hawkins details how CloudWalk, a startup in China in need of more images of faces of people of African descent, signed a deal with the Zimbabwean government for it to provide the images the company was lacking.⁴² In return for sharing its data, Zimbabwe will receive a national facial database and “smart” surveillance infrastructure that it can install in airports, railways, and bus stations.

It might sound like an even exchange, but Zimbabwe has a dismal record on human rights. Making things worse, CloudWalk provides facial recognition technologies to the Chinese police—a conflict of interest so great that the global nonprofit Human Rights Watch voiced its concern about the deal.⁴³ Face harvesting is happening in the US as well. Researchers Os Keyes, Nikki Stevens and Jacqueline Wernimont have shown how immigrants, abused children, and dead people are some of the groups whose faces have been used to train software—without their consent.⁴⁴ So is a diverse database of faces really a good idea? Voicing his concerns in response to the announcement of Buolamwini and Gebru's 2018 study on Twitter, an Indigenous Marine veteran shot back, "I hope facial recognition software has a problem identifying my face too. That'd come in handy when the police come rolling around with their facial recognition truck at peaceful demonstrations of dissent, cataloging all dissenters for 'safety and security.'"⁴⁵

Better detection of faces of color cannot be characterized as an unqualified good. More often than not, it is enlisted in the service of increased oppression, greater surveillance, and targeted violence. Buolamwini understands these potential harms and has developed an approach that works across all four domains of the matrix of domination to address the underlying issues of power that are playing out in facial analysis technology. Buolamwini and Gebru first quantified the disparities in the dataset—a technical audit, which falls in the disciplinary domain of the matrix of domination. Then, Buolamwini went on to launch the Algorithmic Justice League, an organization that works to highlight and intervene in instances of algorithmic bias. On behalf of the AJL, Buolamwini has produced viral poetry projects and given TED talks—taking action in the hegemonic domain, the realm of culture and ideas. She has advised on legislation and professional standards for the field of computer vision and called for a moratorium on facial analysis in policing on national media and in Congress.⁴⁶ These are actions operating in the structural domain of the matrix of domination—the realm of law and policy. Throughout these efforts, the AJL works with students and researchers to help guide and shape their own work—the interpersonal domain. Taken together, Buolamwini's various initiatives demonstrate how any “solution” to bias in algorithms and datasets must tackle more than technical limitations. In addition, they present a compelling model for the data scientist as public intellectual—who, yes, works on technical audits and fixes, but also works on cultural, legal, and political efforts too.

While equitable representation—in datasets and data science workforces—is important, it remains window dressing if we don't also transform the institutions that produce and reproduce those biased outcomes in the first place. As doctoral health student

Arrianna Planey, quoting Robert M. Young, states, “A racist society will give you a racist science.”⁴⁷ We cannot filter out the downstream effects of sexism and racism without also addressing their root cause.

Data Science for Whom?

One of the downstream effects of the privilege hazard—the risks incurred when people from dominant groups create most of our data products—is not only that datasets are biased or unrepresentative, but that they never get collected at all. Mimi Onuoha—an artist, designer, and educator—has long been asking *who questions* about data science. Her project, *The Library of Missing Datasets* (figure 1.4), is a list of datasets that one might expect to already exist in the world, because they help to address pressing social issues, but that in reality have never been created. The project exists as a website and as an art object. The latter consists of a file cabinet filled with folders labeled



Figure 1.4

The Library of Missing Datasets, by Mimi Onuoha (2016) is a list of datasets that are not collected because of bias, lack of social and political will, and structural disregard. Courtesy of Mimi Onuoha. Photo by Brandon Schulman.

with phrases like: “People excluded from public housing because of criminal records,” “Mobility for older adults with physical disabilities or cognitive impairments,” and “Total number of local and state police departments using stingray phone trackers (IMSI-catchers).” Visitors can tab through the folders and remove any particular folder of interest, only to reveal that it is empty. They all are. The datasets that should be there are “missing.”

By compiling a list of the datasets that are missing from our “otherwise data-saturated” world, Onuoha explains, “we find cultural and colloquial hints of what is deemed important” and what is not. “Spots that we’ve left blank reveal our hidden social biases and indifferences,” she continues. And by calling attention to these datasets as “missing,” she also calls attention to how the matrix of domination encodes these “social biases and indifferences” across all levels of society.⁴⁸ Along similar lines, foundations like Data2X and books like *Invisible Women* have advanced the idea of a systematic “gender data gap” due to the fact that the majority of research data in scientific studies is based around men’s bodies. The downstream effects of the gender data gap range from annoying—cell phones slightly too large for women’s hands, for example—to fatal. Until recently, crash test dummies were designed in the size and shape of men, an oversight that meant that women had a 47 percent higher chance of car injury than men.⁴⁹

The *who question* in this case is: Who benefits from data science and who is overlooked? Examining those gaps can sometimes mean calling out missing datasets, as Onuoha does; characterizing them, as *Invisible Women* does; and advocating for filling them, as Data2X does. At other times, it can mean collecting the missing data yourself. Lacking comprehensive data about women who die in childbirth, for example, ProPublica decided to resort to crowdsourcing to learn the names of the estimated seven hundred to nine hundred US women who died in 2016.⁵⁰ As of 2019, they’ve identified only 140. Or, for another example: in 1998, youth living in Roxbury—a neighborhood known as “the heart of Black culture in Boston”⁵¹—were sick and tired of inhaling polluted air. They led a march demanding clean air and better data collection, which led to the creation of the AirBeat community monitoring project.⁵²

Scholars have proposed various names for these instances of ground-up data collection, including *counterdata* or *agonistic data* collection, *data activism*, *statactivism*, and *citizen science* (when in the service of environmental justice).⁵³ Whatever it’s called, it’s been going on for a long time. In 1895, civil rights activist and pioneering data journalist Ida B. Wells assembled a set of statistics on the epidemic of lynching that was sweeping the United States.⁵⁴ She accompanied her data with a meticulous exposé of the fraudulent claims made by white people—typically, that a rape, theft, or assault

of some kind had occurred (which it hadn't in most cases) and that lynching was a justified response. Today, an organization named after Wells—the Ida B. Wells Society for Investigative Reporting—continues her mission by training up a new generation of journalists of color in the skills of data collection and analysis.⁵⁵

A counterdata initiative in the spirit of Wells is taking place just south of the US border, in Mexico, where a single woman is compiling a comprehensive dataset on femicides—gender-related killings of women and girls.⁵⁶ María Salguero, who also goes by the name Princesa, has logged more than five thousand cases of femicide since 2016.⁵⁷ Her work provides the most accessible information on the subject for journalists, activists, and victims' families seeking justice.

The issue of femicide in Mexico rose to global visibility in the mid-2000s with widespread media coverage about the deaths of poor and working-class women in Ciudad Juárez. A border town, Juárez is the site of more than three hundred *maquiladoras*: factories that employ women to assemble goods and electronics, often for low wages and in substandard working conditions. Between 1993 and 2005, nearly four hundred of these women were murdered, with around a third of those murders exhibiting signs of exceptional brutality or sexual violence. Convictions were made in only three of those deaths. In response, a number of activist groups like Ni Una Más (Not One More) and Nuestras Hijas de Regreso a Casa (Our Daughters Back Home) were formed, largely motivated by mothers demanding justice for their daughters, often at great personal risk to themselves.⁵⁸

These groups succeeded in gaining the attention of the Mexican government, which established a Special Commission on Femicide. But despite the commission and the fourteen volumes of information about femicide that it produced, and despite a 2009 ruling against the Mexican state by the Inter-American Human Rights Court, and despite a United Nations Symposium on Femicide in 2012, and despite the fact that sixteen Latin American countries have now passed laws defining femicide—despite all of this, deaths in Juárez have continued to rise.⁵⁹ In 2009 a report pointed out that one of the reasons that the issue had yet to be sufficiently addressed was the lack of data.⁶⁰ Needless to say, the problem remains.

How might we explain the missing data around femicides in relation to the four domains of power that constitute Collins's matrix of domination? As is true in so many cases of data collected (or not) about women and other minoritized groups, the collection environment is compromised by imbalances of power.

The most grave and urgent manifestation of the matrix of domination is within the interpersonal domain, in which cis and trans women become the victims of violence and murder at the hands of men. Although law and policy (the structural domain)

have recognized the crime of femicide, no specific policies have been implemented to ensure adequate information collection, either by federal agencies or local authorities. Thus the disciplinary domain, in which law and policy are enacted, is characterized by a deferral of responsibility, a failure to investigate, and victim blaming. This persists in a somewhat recursive fashion because there are no consequences imposed within the structural domain. For example, the Special Commission's definition of femicide as a "crime of the state" speaks volumes to how the government of Mexico is deeply complicit through inattention and indifference.⁶¹

Of course, this inaction would not have been tolerated without the assistance of the hegemonic domain—the realm of media and culture—which presents men as strong and women as subservient, men as public and women as private, trans people as deviating from "essential" norms, and nonbinary people as nonexistent altogether. Indeed, government agencies have used their public platforms to blame victims. Following the femicide of twenty-two-year-old Mexican student Lesvy Osorio in 2017, researcher María Rodríguez-Dominguez documented how the Public Prosecutor's Office of Mexico City shared on social media that the victim was an alcoholic and drug user who had been living out of wedlock with her boyfriend.⁶² This led to justified public backlash, and to the hashtag #SiMeMatan (If they kill me), which prompted sarcastic tweets such as "#SiMeMatan it's because I liked to go out at night and drink a lot of beer."⁶³

It is into this data collection environment, characterized by extremely asymmetrical power relations, that María Salguero has inserted her femicides map. Salguero manually plots a pin on the map for every femicide that she collects through media reports or through crowdsourced contributions (figure 1.5a). One of her goals is to "show that these victims [each] had a name and that they had a life," and so Salguero logs as many details as she can about each death. These include name, age, relationship with the perpetrator, mode and place of death, and whether the victim was transgender, as well as the full content of the news report that served as the source. Figure 1.5b shows a detailed view for a single report from an unidentified transfemicide, including the date, time, location, and media article about the killing. It can take Salguero three to four hours a day to do this unpaid work. She takes occasional breaks to preserve her mental health, and she typically has a backlog of a month's worth of femicides to add to the map.

Although media reportage and crowdsourcing are imperfect ways of collecting data, this particular map, created and maintained by a single person, fills a vacuum created by her national government. The map has been used to help find missing women, and Salguero herself has testified before Mexico's Congress about the scope of the problem.

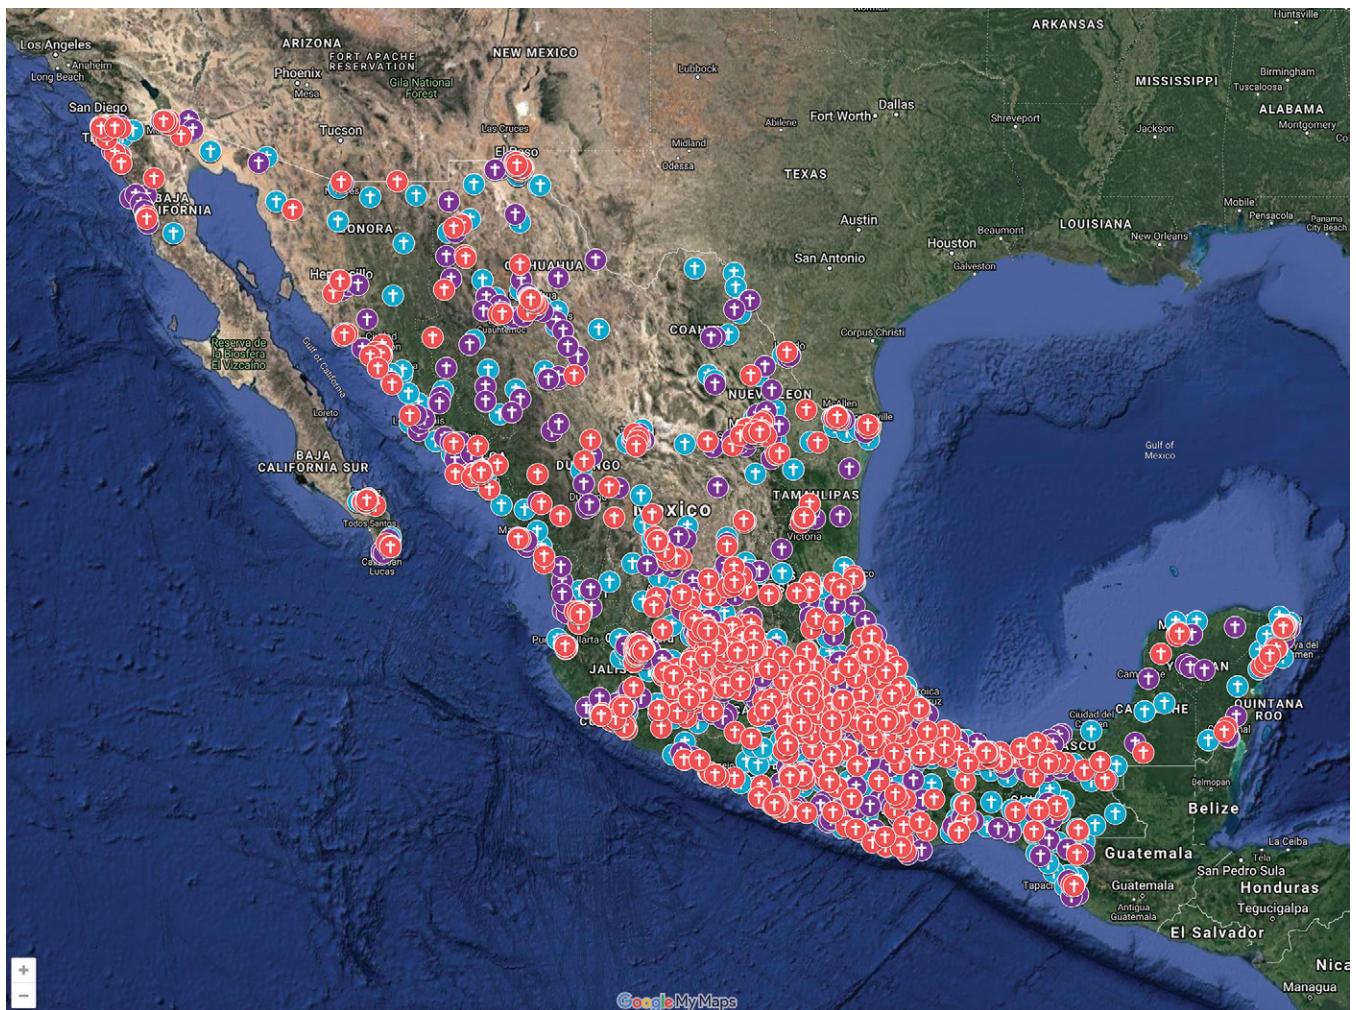


Figure 1.5

María Salguero's map of femicides in Mexico (2016–present) can be found at <https://feminicidiosmx.crowdmap.com/>. (a) Map extent showing the whole country. (b) A detailed view of Ciudad Juárez with a focus on a single report of an anonymous transfemicide. Salguero crowdsources points on the map based on reports in the press and reports from citizens to her. Courtesy of María Salguero.

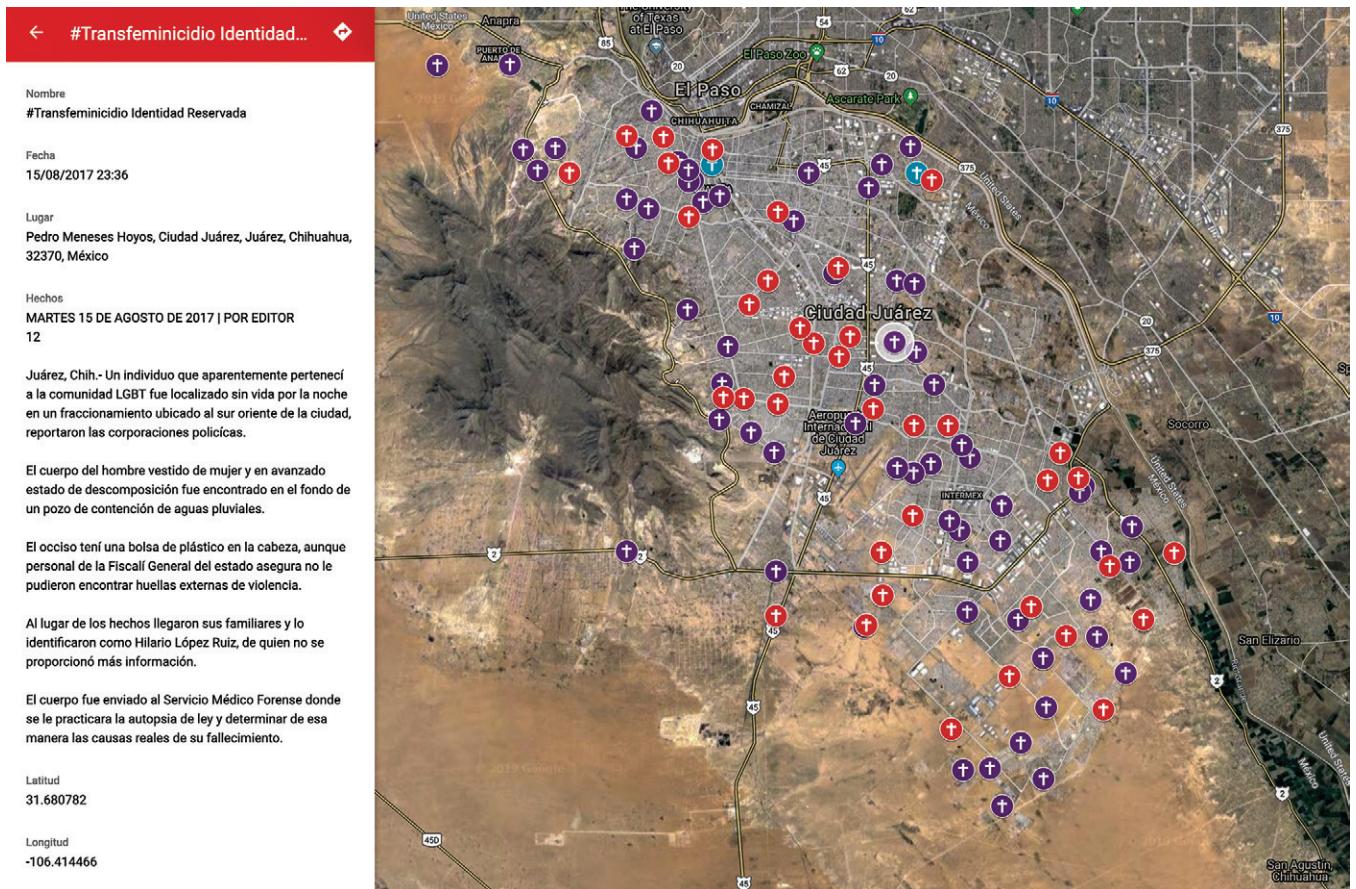


Figure 1.5 (continued)

Salguero is not affiliated with an activist group, but she makes her data available to activist groups for their efforts. Parents of victims have called her to give their thanks for making their daughters visible, and Salguero affirms this function as well: "This map seeks to make visible the sites where they are killing us, to find patterns, to bolster arguments about the problem, to georeference aid, to promote prevention and try to avoid femicides."

It is important to make clear that the example of missing data about femicides in Mexico is not an isolated case, either in terms of subject matter or geographic location. The phenomenon of missing data is a regular and expected outcome in all societies characterized by unequal power relations, in which a gendered, racialized order is

maintained through willful disregard, deferral of responsibility, and organized neglect for data and statistics about those minoritized bodies who do not hold power. So too are examples of individuals and communities using strategies like Salguero's to fill in the gaps left by these missing datasets—in the United States as around the world.⁶⁴ If “quantification is representation,” as data journalist Jonathan Stray asserts, then this offers one way to hold those in power accountable. Collecting counterdata demonstrates how data science can be enlisted on behalf of individuals and communities that need more power on their side.⁶⁵

Data Science with Whose Interests and Goals?

Far too often, the problem is not that data about minoritized groups are missing but the reverse: the databases and data systems of powerful institutions are built on the excessive surveillance of minoritized groups. This results in women, people of color, and poor people, among others, being overrepresented in the data that these systems are premised upon. In *Automating Inequality*, for example, Virginia Eubanks tells the story of the Allegheny County Office of Children, Youth, and Families in western Pennsylvania, which employs an algorithmic model to predict the risk of child abuse in any particular home.⁶⁶ The goal of the model is to remove children from potentially abusive households before it happens; this would appear to be a very worthy goal. As Eubanks shows, however, inequities result. For wealthier parents, who can more easily access private health care and mental health services, there is simply not that much data to pull into the model. For poor parents, who more often rely on public resources, the system scoops up records from child welfare services, drug and alcohol treatment programs, mental health services, Medicaid histories, and more. Because there are far more data about poor parents, they are oversampled in the model, and so their children are overtarged as being at risk for child abuse—a risk that results in children being removed from their families and homes. Eubanks argues that the model “confuse[s] parenting while poor with poor parenting.”

This model, like many, was designed under two flawed assumptions: (1) that more data is always better and (2) that the data are a neutral input. In practice, however, the reality is quite different. The higher proportion of poor parents in the database, with more complete data profiles, the more likely the model will be to find fault with poor parents. And data are never neutral; they are always the biased output of unequal social, historical, and economic conditions: this is the matrix of domination once again.⁶⁷ Governments can and do use biased data to marshal the power of the matrix of domination in ways that amplify its effects on the least powerful in society. In this

case, the model becomes a way to administer and manage classism in the disciplinary domain—with the consequence that poor parents' attempts to access resources and improve their lives, when compiled as data, become the same data that remove their children from their care.

So this raises our next *who question*: Whose goals are prioritized in data science (and whose are not)? In this case, the state of Pennsylvania prioritized its bureaucratic goal of efficiency, which is an oft-cited reason for coming up with a technical solution to a social and political dilemma. Viewed from the perspective of the state, there were simply not enough employees to handle all of the potential child abuse cases, so it needed a mechanism for efficiently deploying limited staff—or so the reasoning goes. This is what Eubanks has described as a *scarcity bias*: the idea that there are not enough resources for everyone so we should think small and allow technology to fill the gaps. Such thinking, and the technological “solutions” that result, often meet the goals of their creators—in this case, the Allegheny County Office of Children, Youth, and Families—but not the goals of the children and families that it purports to serve.

Corporations also place their own goals ahead of those of the people their products purport to serve, supported by their outsize wealth and the power that comes with it. For example, in 2012, the *New York Times* published an explosive article by Charles Duhigg, “How Companies Learn Your Secrets,”⁶⁸ which soon became the stuff of legend in data and privacy circles. Duhigg describes how Andrew Pole, a data scientist working at Target, was approached by men from the marketing department who asked, “If we wanted to figure out if a customer is pregnant, even if she didn’t want us to know, can you do that?”⁶⁹ He proceeded to synthesize customers’ purchasing histories with the timeline of those purchases to give each customer a so-called pregnancy prediction score (figure 1.6).⁷⁰ Evidently, pregnancy is the second major life event, after leaving for college, that determines whether a casual shopper will become a customer for life.

Target turned around and put Pole’s pregnancy detection model into action in an automated system that sent discount coupons to possibly pregnant customers. Win-win—or so the company thought, until a Minneapolis teenager’s dad saw the coupons for baby clothes that she was getting in the mail and marched into his local Target to read the manager the riot act. Why was his daughter getting coupons for pregnant women when she was only a teen?!

It turned out that the young woman was indeed pregnant. Pole’s model informed Target before the teenager informed her family. By analyzing the purchase dates of approximately twenty-five common products, such as unscented lotion and large bags of cotton balls, the model found a set of purchase patterns that were highly correlated



Figure 1.6

Screenshot from a video of statistician Andrew Pole's presentation at Predictive Analytics World about Target's pregnancy detection model in October 2010, titled "How Target Gets the Most out of Its Guest Data to Improve Marketing ROI." He discusses the model at 47:50. Image by Andrew Pole for Predictive Analytics World.

with pregnancy status and expected due date. But the win-win quickly became a lose-lose, as Target lost the trust of its customers in a PR disaster and the Minneapolis teenager lost far worse: her control over information related to her own body and her health.

This story has been told many times: first by Pole, the statistician; then by Duhigg, the *New York Times* journalist; then by many other commentators on personal privacy and corporate overreach. But it is not only a story about privacy: it is also a story about gender injustice—about how corporations approach data relating to women's bodies and lives, and about how corporations approach data relating to minoritized populations more generally. Whose goals are prioritized in this case? The corporation's, of course. For Target, the primary motivation was maximizing profit, and quarterly financial reports to the board are the measurement of success. Whose goals are *not* prioritized? The teenager's and those of every other pregnant woman out there.

How did we get to the point where data science is used almost exclusively in the service of profit (for a few), surveillance (of the minoritized), and efficiency (amidst scarcity)? It's worth stepping back to make an observation about the organization of the data economy: data are expensive and resource-intensive, so only already powerful

institutions—corporations, governments, and elite research universities—have the means to work with them at scale. These resource requirements result in data science that serves the primary goals of the institutions themselves. We can think of these goals as the *three Ss*: science (universities), surveillance (governments), and selling (corporations). This is not a normative judgment (e.g., “all science is bad”) but rather an observation about the organization of resources. If science, surveillance, and selling are the main goals that data are serving, because that’s who has the money, then what other goals and purposes are going underserved?

Let’s take “the cloud” as an example. As server farms have taken the place of paper archives, storing data has come to require large physical spaces. A project by the Center for Land Use Interpretation (CLUI) makes this last point plain (figure 1.7). In 2014, CLUI set out to map and photograph data centers around the United States, often in those seemingly empty in-between areas we now call *exurbs*. In so doing, it called attention to “a new kind of physical information architecture” sprawling across the United States: “windowless boxes, often with distinct design features such as an appliquéd surface graphics or a functional brutalism, surrounded by cooling systems.” The environmental impacts of the cloud—in the form of electricity and air conditioning—are enormous. A 2017 Greenpeace report estimated that the global IT sector, which is largely US-based, accounted for around 7 percent of the world’s energy use. This is more than some of largest countries in the world, including Russia, Brazil, and Japan.⁷¹ Unless that energy comes from renewable sources (which the Greenpeace report shows that it does not), the cloud has a significant accelerating impact on global climate change.

So the cloud is not light and it is not airy. And the cloud is not cheap. The cost of constructing Facebook’s newest data center in Los Lunas, New Mexico, is expected to reach \$1 billion.⁷² The electrical cost of that center alone is estimated at \$31 million per year.⁷³ These numbers return us to the question about financial resources: Who has the money to invest in centers like these? Only powerful corporations like Facebook and Target, along with wealthy governments and elite universities, have the resources to collect, store, maintain, analyze, and mobilize the largest amounts of data. Next, who is in charge of these well-resourced institutions? Disproportionately men, even more disproportionately white men, and even more than that, disproportionately rich white men. Want the data on that? Google’s Board of Directors is comprised of 82 percent white men. Facebook’s board is 78 percent male and 89 percent white. The 2018 US Congress was 79 percent male—actually a better percentage than in previous years—and with a median net worth of five times more than the average American household.⁷⁴ These are the people who experience the most privilege within the matrix



Figure 1.7

Photographs from *Networked Nation: The Landscape of the Internet in America*, an exhibition by the Center for Land Use Interpretation staged in 2013. The photos show four data centers located in North Bergen, NJ; Dalles, OR; Ashburn, VA; and Lockport, NY (counterclockwise from top right). They show how the “cloud” is housed in remote locations and office parks around the country. Images by the Center for Land Use Interpretation.



Figure 1.7 (continued)

of domination, and they are also the people who benefit the most from the current status quo.⁷⁵

In the past decade or so, many of these men at the top have described data as “the new oil.”⁷⁶ It’s a metaphor that resonates uncannily well—even more than they likely intended. The idea of data as some sort of untapped natural resource clearly points to the potential of data for power and profit once they are processed and refined, but it also helps highlight the exploitative dimensions of extracting data from their source—people—as well as their ecological cost. Just as the original oil barons were able to use their riches to wield outsized power in the world (think of John D. Rockefeller, J. Paul Getty, or, more recently, the Koch brothers), so too do the Targets of the world use their corporate gain to consolidate control over their customers. But unlike crude oil, which is extracted from the earth and then sold to people, data are both extracted from people and sold back to them—in the form of coupons like the one the Minneapolis teen received in the mail, or far worse.⁷⁷

This extractive system creates a profound asymmetry between who is collecting, storing, and analyzing data, and whose data are collected, stored, and analyzed.⁷⁸ The goals that drive this process are those of the corporations, governments, and well-resourced universities that are dominated by elite white men. And those goals are neither neutral nor democratic—in the sense of having undergone any kind of participatory, public process. On the contrary, focusing on those *three Ss*—science, surveillance, and selling—to the exclusion of other possible objectives results in significant oversights with life-altering consequences. Consider the Target example as the flip side of the missing data on maternal health outcomes. Put crudely, there is no profit to be made collecting data on the women who are dying in childbirth, but there is significant profit in knowing whether women are pregnant.

How might we prioritize different goals and different people in data science? How might data scientists undertake a feminist analysis of power in order to tackle bias at its source? Kimberly Seals Allers, a birth justice advocate and author, is on a mission to do exactly that in relation to maternal and infant care in the United States. She followed the Serena Williams story with great interest and watched as Congress passed the Preventing Maternal Deaths Act of 2018. This bill funded the creation of maternal health review committees in every state and, for the first time, uniform and comprehensive data collection at the federal level. But even as more data have begun to be collected about maternal mortality, Seals Allers has remained frustrated by the public conversation: “The statistics that are rightfully creating awareness around the Black maternal mortality crisis are also contributing to this gloom and doom deficit narrative. White



Figure 1.8

Irth is a mobile app and web platform focused on removing bias from birth (including prenatal, birth, and postpartum health care). Users post intersectional reviews of the care they received from individual nurses and doctors, as well as whole practices and hospitals. When parents to be are searching for providers, they can consult Irth to see what kind of care people like them received in the hands of specific caregivers. Wireframes from Irth's first prototype are shown here. Images by Kimberly Seals Allers and the Irth team, 2019.

people are like, ‘how can we save Black women?’ And that’s not the solution that we need the data to produce.”⁷⁹

Seals Allers—and her fifteen-year-old son, Michael—are working on their own data-driven contribution to the maternal and infant health conversation: a platform and app called Irth—from *birth*, but with the *b* for *bias* removed (figure 1.8). One of the major contributing factors to poor birth outcomes, as well as maternal and infant mortality, is biased care. Hospitals, clinics, and caregivers routinely disregard Black women’s expressions of pain and wishes for treatment.⁸⁰ As we saw, Serena Williams’s own story almost ended in this way, despite the fact that she is an international tennis star. To combat this, Irth operates like an intersectional Yelp for birth experiences. Users post ratings and reviews of their prenatal, postpartum, and birth experiences at specific hospitals and in the hands of specific caregivers. Their reviews include important

details like their race, religion, sexuality, and gender identity, as well as whether they felt that those identities were respected in the care that they received. The app also has a taxonomy of bias and asks users to tick boxes to indicate whether and how they may have experienced different types of bias. Irth allows parents who are seeking care to search for a review from someone like them—from a racial, ethnic, socioeconomic, and/or gender perspective—to see how they experienced a certain doctor or hospital.

Seals Allers's vision is that Irth will be both a public information platform, for individuals to find better care, and an accountability tool, to hold hospitals and providers responsible for systemic bias. Ultimately, she would like to present aggregated stories and data analyses from the platform to hospital networks to push for change grounded in women's and parents' lived experiences. "We keep telling the story of maternal mortality from the grave," she says. "We have to start preventing those deaths by sharing the stories of people who actually lived."⁸¹

Irth illustrates the fact that "doing good with data" requires being deeply attuned to the things that fall outside the dataset—and in particular to how datasets, and the data science they enable, too often reflect the structures of power of the world they draw from. In a world defined by unequal power relations, which shape both social norms and laws about how data are used and how data science is applied, it remains imperative to consider who gets to do the "good" and who, conversely, gets someone else's "good" done to them.

Examine Power

Data feminism begins by examining how power operates in the world today. This consists of asking *who questions* about data science: Who does the work (and who is pushed out)? Who benefits (and who is neglected or harmed)? Whose priorities get turned into products (and whose are overlooked)? These questions are relevant at the level of individuals and organizations, and are absolutely essential at the level of society. The current answer to most of these questions is "people from dominant groups," which has resulted in a *privilege hazard* so acute that it explains the near-daily revelations about another sexist or racist data product or algorithm. The *matrix of domination* helps us to understand how the privilege hazard—the result of unequal distributions of power—plays out in different domains. Ultimately, the goal of examining power is not only to understand it, but also to be able to challenge and change it. In the next chapter, we explore several approaches for challenging power with data science.