# Fighting the "Privilege Hazard": A Feminist Data Intervention

**Article by Catherine D'Ignazio, Lauren Klein**
July 14, 2021

In a world where power is wielded unequally, data science is no exception. When it is largely the most privileged who are gathering the data and writing the codes within a corporate culture of profit maximisation, it is unsurprising that the results reproduce social and economic inequalities. But it is not inevitable, argue the authors of _Data Feminism_, Catherine D'Ignazio and Lauren Klein, who are taking a stand for data justice. Rooted in intersectional feminism, power analysis and a commitment to act are central to their approach.

**Green European Journal: What is data feminism?**

**Catherine D'Ignazio**: Data feminism is an approach to thinking about all aspects of data science – everything from data collection to cleaning, analysis, visualisation and communication, and the deployment of data-driven systems – that is informed by intersectional feminist thought. Broadly speaking, the question that _Data Feminism_ tries to answer is: "What would feminist data science look like?"

**Lauren Klein**: By looking to intersectional feminism as our model, we aim to bring together ways of engaging with issues that are both theoretically informed and backed up by action. Feminism has many different meanings for many different people, not all of them unequivocally good. Certain feminisms have been very exclusionary in their approach to equality and how they define their goals. We deliberately anchor our approach in intersectional feminism because it comes with a conceptual understanding of how power works in the world. This is coupled with a commitment to act, to try to rebalance the unequal power relations that we encounter both individually and collectively.

**Why is it so important to bring a feminist perspective to data science?**

**Lauren Klein**: We need an intervention into data science and systems because today they're everywhere. They're making decisions for us of personal, national, and international importance, and they're doing so in ways which are unequal along the lines of gender, race, class and more. This is because these systems are created within our unequal society, and as a result they reflect that inequality.

The starting point is to recognise that these systems are out there, that they are hugely influential, and that they are perpetuating rather than ameliorating unequal power relations. But what can we do about them? Intersectional feminism has been working over decades, if not centuries, to challenge unequal systems of power, and we can apply the lessons of intersectional feminism to data. Put another way, we try to learn from how activists and thinkers have opposed unequal systems of power in the past, and we try to

apply those ways of challenging power to present-day data systems.

> *Data systems are created within our unequal society, and as a result they reflect that inequality.*

**Can you give some examples of how data is marginalising minorities?**

**Catherine D'Ignazio:** My MIT colleague Joy Buolamwini has shown that <u>facial recognition systems are failing</u>, with especially high error rates for women of colour. When she dug into the data sets used to train the systems, she found that they overwhelmingly relied on the faces of white men.

There are other cases as well. Amazon, for example, was developing an internal resume screening tool that would allow HR to automatically sift through large amounts of resumes to "filter up" the good candidates. But because Amazon was using the resumes of their current employees to train their system, and their employees were predominantly men, the system had a built-in gender bias. It was systematically demoting applicants who had attended all-women schools, like Wellesley College [a prestigious women's liberal arts college in the US], or applicants who had done women's sports or other activities that included "women's" as an adjective. In the end, they scrapped the system. They couldn't figure out how to purge the gender bias from it.

**Lauren Klein**: Another example is the <u>algorithm that was used to grade students in the United Kingdom when A-level exams were cancelled</u> due to the pandemic. The UK government introduced an algorithm that was intended to predict the scores that each student would have received had they sat the exam. But the model didn't just consider students' individual records; it also factored in their school's historical performance. So, if a student attended an under-resourced school that historically hadn't performed as well as other schools, their individual score was automatically lowered. If a student attended a well-resourced school that had historically done very well, their score would automatically be bumped up. Students took to the streets protesting, holding signs saying, "Fuck the algorithm". In the end, the government had to rescind the use of the algorithm.

**How could your book's insights be applied to deal with the current problems?**

**Catherine D'Ignazio**: Our book looks at how data science can be used to challenge power, and we're seeing a good amount of that happening already. Some of the most exciting work at the intersection of data and justice is being done not in computer science departments, but in data journalism, community-based organisations, activism and social movements. It's coming from artists, librarians, public information professionals, and so on. One thing we argue for in *Data Feminism* is a broader definition of data science. If we define the field more broadly and include these civil society data practices, then we already have some great models for how we can use data to challenge and change power structures.

We also look at strategies like collecting counter-data. Some groups are collecting data as a way of counting and quantifying information about issues that institutions aren't

adequately measuring. A good example is the <u>Mexican activist María Salguero</u> who is counting feminicides, gender-based killings of women. Or the <u>Anti-Eviction Mapping Project</u> in the United States, which is using a broader definition to count evictions. As a result, it's recording more than the government has on record.

**Lauren Klein**: There's also auditing algorithms. A major issue is that the organisations that have the capacity to mobilise large amounts of data are usually run by the people who experience the most privilege in the world. Tech in general tends to be predominantly white, male, and centred in the Global North. All this leads to what we call the "privilege hazard": the idea that those in the positions to make decisions – about what is worth counting, what problems are worth addressing, or what the potential pitfalls might be – are unaware of the potential harms. They just don't know how to look for them.

> *Tech in general tends to be predominantly white, male, and centred in the Global North. All this leads to what we call the "privilege hazard".*

One example of the privilege hazard at work, and how an audit can help call attention to it, is an <u>algorithm designed for pre-trial risk assessment in the US justice system</u>. In the US, a judge usually decides whether those who stand accused of a crime and are awaiting trail can be released on bail. An algorithm was designed to replace the judge in these situations, and its designers thought that they were removing bias from the decision-making process. But, because they didn't think about the biased data they were feeding into the system, they were unaware of the clear and significant harms that their system would bring for minoritised communities. One of the roles of data journalists, academic researchers, and others is to test these algorithms to check that they're performing their claimed function and whether they're doing so in an unbiased way.

**The share of women in data science or IT jobs and related university programmes is still very low (and shrinking in the US, as your book shows). What are the reasons for this?**

**Lauren Klein:** There are many reasons. One is how the expertise of women – and for that matter, anyone who isn't coming from the majority perspective – has been historically devalued. And even if people come to the table with the same amount of expertise, those who occupy these minoritised positions find themselves questioned more and their ideas adopted less. Eventually this becomes very disheartening. So, people leave.

The other aspect is structural. There's a long history that shows that knowledge often originates from those who experience something most closely or directly. But as that knowledge gets more systematised and professionalised, those with experiential knowledge find themselves being pushed out of the fields by people who have studied the subject at university or received professional credential. The women computers are a classic example of this. The earliest computer programmers were women because they were essentially repurposed secretaries. Many of them had to learn how to use computers on the job, although some did have advanced degrees but could only find work as secretaries. <u>Grace Hopper</u> essentially came up with the idea to abstract instructions into algorithms to avoid

having to do the same task repeatedly. That is how the abstraction that underlies so much of computation came to be.

But the problem of women being pushed out of fields as they become professionalised or gain more prestige predates computation. A large body of work looks at how medicine became professionalised and, after the creation of medical schools, certain knowledge was taken from women, like midwives, who learned through experience and was given over to obstetricians, who were men with medical degrees. It's the same thing with home cooks and professional chefs. There are so many patterns where lived experience is devalued when the possibility of professional credential is presented. The same thing is happening in the computer science fields.

**In December 2020, the co-lead of Google's ethical AI team  <u>Timnit Gebru</u> was fired for jointly authoring a <u>critical paper</u> on the potential pitfalls of large language models. This wasn't the first time Google was in the spotlight for its unwillingness to change the status quo. Does this signal that there's been no substantial change in the work cultures of big tech companies?**

**Catherine D'Ignazio:** The firing of Timnit Gebru really sent shock waves through academia and industry. She was a high-profile hire and one of the very few Black women working in AI. She was in a position of co-leadership with a white woman, Margaret Mitchell, who then was also subsequently fired. Gebru's case exposes the limits of what's possible from inside a corporation. It pushes us to ask when to take further steps such as regulations to constrain corporate actions. We simply cannot leave it up to corporations to ensure that their products are ethical. Ultimately, a company is accountable to its shareholders, not the public. They will always make decisions to maximise profit, and they will only incorporate ethical principles when convenient.

Gebru's paper was not even that critical! It just pointed out that large language models have human bias, and that using those models runs the risk of re-inscribing those biases. Or that generating the models has environmental consequences. These are important things to say. The paper's authors were not necessarily going after Google. They were just pointing out the risks that apply to anyone working in this space.

> *We simply cannot leave it up to corporations to ensure that their products are ethical. Ultimately, a company is accountable to its shareholders, not the public.*

I have very little faith in the CEOs of corporations to check their own power, but I have optimism in workers. In tech companies, the workers are increasingly getting organised. Google's workers, for example, recently unionised. That's a really encouraging sign. Policymakers are also starting to get on board and think in a more future-facing direction, not only reacting to the technologies that corporations create but starting to think of innovative ways to regulate so that these technologies can work for everybody.

**Lauren Klein:** We're also seeing a desire to imagine what an ideal set of guidelines or

regulations to keep these large and powerful corporations in check might look like. The US is usually reactive in its governance; it tends to first be presented with a problem and only then does it try to figure out how to regulate against it. But in the EU – first with the General Data Protection Regulation and now with the Ethics Guidelines for Trustworthy AI – there's an effort to outline core principles that might  find implementation in the years ahead. It's really promising to have a recognition that these are questions of power and inequality, and thus they require a large-scale governmental response.

**Projects like the free online book database Gigapedia or Sci-Hub (which makes paywalled research papers accessible) have helped bring knowledge to people who otherwise couldn't afford it. There have been many attempts to close them down and both stand accused of copyright infringement. To what extent is improvement possible in the currently dominant political and economic systems?**

**Catherine D'Ignazio:** This recalls the work of my former advisor, Ethan Zuckerman, on what he calls the "digital public infrastructure". The current internet has centralised power in the hands of big corporations and platforms which have been able to mobilise large resources and attract large user bases. They essentially provide public services and a kind of public commons – we could even say a democratic public space – except, of course, it is not public. It's a privatised space subject to rules, regulations, policies, and the norms of corporations. I love hearing examples like those in the question, and I think the history of the internet is to an extent a history of these examples, all the way back to Napster. A lot of important questions have been raised by new models for knowledge sharing or production that challenge the existing business model for monetising information.

I see the concept of the digital public infrastructure as a call to action for government and civil society, and to the public more broadly, to claim a stake in what these systems can be. To actively think about them and not leave it up to corporations to decide which set of values their technologies and platforms should operate by. The American abolitionist Frederick Douglass once said, "Power concedes nothing without a demand." We need to make demands, and those demands need to come from different sectors of society. Maybe that results in the breaking up of the corporations. We need to fundamentally reimagine the public organisation of knowledge and social sharing. These are the public spaces of this century and I'm not OK with leaving public space up to Facebook.

**Do the insights of your work on data justice have broader political relevance?**

**Lauren Klein:** I'm inspired by how progressive political movements are usually composed of coalitions who represent different perspectives but who are aligned in their larger goals. This model fits very well with one of the central tenets of feminist thinking: the best and most complete form of knowledge doesn't come from any one individual source, but rather it's the aggregate picture that emerges when individuals or groups bring their perspectives to the table. Brought together, those perspectives assemble a more complete picture of the problem at hand.

Data, for the most part, captures a very small part of the richness of human experience. Because of that, data should never be used on its own. To make decisions, data needs to be coupled with personal experience, attention to history, culture, and seen in the context of the structural inequalities we've been talking about. Progressive political movements know

this instinctively, and they should remember it when encountering datasets and evaluating data systems. They should remind themselves that data systems work in the same ways as the rest of the world: they're influenced by the fact that power is wielded unequally in the world. So the same advice applies to data justice as to other forms of social justice. These problems are real, they're ongoing, and they need to be taken on. And in taking them on, political groups – even progressive ones – must remember to look to those who are most impacted by a particular issue for knowledge and guidance about how it should be addressed.

Catherine D'Ignazio is an assistant professor of urban science and planning at MIT. She is also director of the Data + Feminism Lab which uses data and computational methods to work towards gender and racial equity. Her book *Data Feminism* (MIT Press, 2020), co-authored with Lauren Klein, charts a course for more ethical and empowering data science practices.

Lauren Klein is an associate professor in the departments of English and quantitative theory and methods at Emory University, where she also directs the Digital Humanities Lab. She is the author of two books: *Data Feminism* (MIT Press, 2020), co-authored with Catherine D'Ignazio, and *An Archive of Taste: Race and Eating in the Early United States* (University of Minnesota Press, 2020).