The Difference Engine: Perpetuating Poverty Through Algorithms

Author : Rebecca Tushnet

Date : July 18, 2018

Virginia Eubanks, **Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor** (2018).

We have a problem with poverty, which we have converted into a problem with poor people. Policymakers tout technology as a way to make social programs more efficient, but they end up encoding the social problems they were designed to solve, thus entrenching poverty and over-policing of the poor. *In Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*, Virginia Eubanks uses three core examples—welfare reform software in Indiana, homelessness service unification in Los Angeles, and child abuse prediction in Pennsylvania—and shows that while they vary in how screwed up they are (Indiana terribly, Los Angeles a bit, and Pennsylvania very hard to tell), they all rely on assumptions that leave poor people more exposed to coercive state control. That state control both results from and contributes to the assumption that poor people's problems are their own fault. The book is a compelling read and a distressing work, mainly because I have little faith that the problems Eubanks so persuasively identifies can be corrected.

Eubanks writes:

> Across the country, poor and working-class people are targeted by new tools of digital poverty management and face life-threatening consequences as a result. Automated eligibility systems discourage them from claiming public resources that they need to survive and thrive. Complex integrated databases collect their most personal information, with few safeguards for privacy or data security, while offering almost nothing in return. Predictive models and algorithms tag them as risky investments and problematic parents. Vast complexes of social service, law enforcement, and neighborhood surveillance make their every move visible and offer up their behavior for government, commercial, and public scrutiny.

As Eubanks points out, the poor are test subjects because they offer "'low rights environments' where there are few expectations of political accountability and transparency." Even those who do not care about poverty should be paying attention, however, because "systems first designed for the poor will eventually be used on everyone."

Eubanks' recommendation, even as more punitive measures are being enacted, is for more resources and fewer requirements. Homelessness isn't a data problem, it's a carpentry problem, and a universal basic income or universal health insurance would allocate care far better than a gauntlet of automated forms. Eubanks points out that automation, despite its promised efficiencies, has coincided with kicking people off of assistance programs. In 1973, nearly half of people under the poverty line received AFDC (Aid to Families with Dependent Children), but a decade later that was 30 percent (coinciding with the introduction of the computerized Welfare Management System) and now it's less than 10 percent. Automated management is a tool of plausible deniability, allowing elites to believe that the most worthy of the poor are being taken care of and that the unworthy don't deserve care, as evidenced by the fact that they failed to behave as they were asked to do in complying with various requirements to submit information and be subjected to surveillance.

Eubanks begins with the most obvious disaster: Indiana's expensive contract with IBM to get rid of most caseworkers and automate medical coverage. Thousands of people were wrongly denied coverage, creating trauma for medically vulnerable people even when the denials were ultimately reversed. Indiana's failure to create a working centralized system led to some backlash. Eubanks quotes people who suggest that the result from the backlash was a hybrid human-computer system, which restored almost enough caseworkers to deal with the people who make noise, but not enough for those who can't. Of course, human caseworkers have their own problems—accounts of implicit and even explicit racial bias abound—but discrimination is easily ported to statistical models, such that states with higher African-American populations have "tougher rules, more stringent work requirements, and higher sanction rates." And Indiana's automated experiment disproportionately drove African Americans off the TANF (Temporary Assistance for Needy Families) rolls, perhaps in part because the system treated any error (including those made by the system itself) as deliberate noncompliance, and many people simply gave up.

The Los Angeles homelessness story is different, but not different enough. It provides a useful contrast of a "progressive" use of data and computerization. The idea was to create "coordinated entry," so that homeless people who contacted any service provider would be connected with the right resources, sorting between the short-term and long-term homeless, who need different services, some of which can be less than helpful if given to the wrong groups. There's a lot of good there, including the idea of "housing first": rather than limiting housing only to those who are sober, employed, etc., the aim is to get people housed because of how hard all those other things are without housing. Eubanks profiles a woman for whom coordinated entry was a godsend.

But Eubanks also identifies two core problems: (1) The system itself is under-resourced; all the coordination in the world won't help when there are only 10 beds for every 100 people in need of them. (2) The information collected is invasive and contributes to the criminalization and pathologization of poor people. The data are kept with minimal security and no protection against police scrutiny, which is particularly significant because, as Eubanks rephrases Anatole France, "so many of the basic conditions of being homeless—having nowhere to sleep, nowhere to put your stuff, and nowhere to go to the bathroom—are also officially crimes." Homeless people can rarely pay tickets, and so the unpaid fines turn into warrants (turning into days in jail when they can't afford bail, even though these kinds of nuisance charges are usually dismissed once in front of a judge). People in the database turn into fugitives.

These two problems reinforce each other. Given the low chance of getting help, people are less willing to explain their circumstances, often stories of escalating misfortune and humiliation, to the representative of the state's computer. The resource crunch also contributes to workers' felt imperative to find the most deserving and thus to scrutinize every applicant for appropriate levels of dysfunctionality. Too little trauma, and services might be deemed unnecessary. But too much dysfunctionality can also be disqualifying—the housing authority might determine that a client is incapable of living independently. One group of caseworkers Eubanks discusses "counsel their clients to treat the interview at the housing authority like a court proceeding." They also see vulnerable clients rejected by landlords; Section 8 vouchers to pay for housing are nice, but still require a willing landlord, and the vouchers expire after six months, meaning that a lot of clients just give up. Meanwhile, "[s]ince 1950, more than 13,000 units of low-income housing have been removed from Skid Row, enough for them all." It's also worth noting how much discretion remains with humans, despite the appearance of Olympian objectivity in a housing need score: clients are assessed based on self-reports, and they won't always tell people they haven't grown to trust about circumstances bearing on their needs, including trauma.

What really mattered to getting resources devoted to addressing homelessness in Los Angeles, Eubanks argued, was rights, not data. Court rulings found that routine police practices—barring sleeping in public and confiscating and destroying the property of homeless people found in areas where they were considered undesirable—were unconstitutional. Once that happened, tent cities sprung up in places visible to people with money and power. Better data helped in identifying what resources were needed where, but tent cities were the driver of reform.

Finally, the experience of child welfare prediction software in Allegheny County, Pennsylvania, has continuities with and divergences from the other two stories. The software is at the moment used just to back up individual caseworkers' determinations of whether to further investigate child abuse based on a call to the child welfare hotline, though Eubanks already saw caseworkers tweaking their own estimates of risk to match the model's, an instance of automation bias that ought to alarm us. Some of the problems were statistical: the number of child deaths and near-deaths in the county is thankfully very low, and you can't do a good model with a handful of cases a year for a population of 1.23 million.

Setting the base-rate problem aside, you can't actually measure levels of child abuse. You can measure proxies, such as how many calls to CPS (Child Protective Services) are made and how many children CPS removes from a home. As a result, the automated system ends up predicting "decisions made by the community (which families will be reported to the hotline) and by the agency and the family courts (which children will be removed from their families), not which children will be harmed." Unfortunately, those proxies are precisely the ones we know are infected with persistent racial and class bias, so that bias is baked into the predictions. This is the same problem explained so well in Cathy O'Neil's [*Weapons of Math Destruction*](), a [good book to read]() along with this one.

In Allegheny County itself, "the great majority of [racial] disproportionality in the county's child welfare services arises from referral bias, not screening bias." Sometimes this arises from perceptions of neighborhoods being bad, so the threshold for reporting someone from those neighborhoods is lower—which in the US means minority neighborhoods. But the prediction system "focuses all its predictive power and computational might on call screening, the step it can experimentally control, rather than concentrating on referral, the step where racial disproportionality is actually entering the system." And it gets worse: the model is evaluated for whether it predicts future referrals. "[T]he activity that introduces the most racial bias into the system is the very way the model defines maltreatment."

In rural or suburban areas, where witnesses are rarer, no one may call the hotline. Families with enough resources use private services for mental health or addiction treatment and thus don't create a record available to the state (if they don't directly talk about child abuse in a way that triggers mandatory reporting). Either way, those disproportionately whiter and wealthier families stay out of the system for conduct that would, if they were visible to the system, increase their risk score. The system can provide very useful services, but those services then become part of the public record, helping define a family as at-risk. A child whose parents were investigated by CPS now has a record of interaction with the system that, when she becomes a mother, will increase her risk score if someone reports her. Likewise, use of public services is coded as a risk factor. A quarter of the predictive variables in the model are "direct measures of poverty"—TANF, SSI (Supplemental Security Income), SNAP (Supplemental Nutrition Assistance Program), and county medical assistance. Another quarter of the predictive variables measure "interaction with juvenile probation" and the child welfare agency itself, when "professional middle-class families have more privacy, interact with fewer mandated reporters, and enjoy more cultural approval of their parenting" than poorer families. Nuisance calls by people with grudges are also a real problem.

Even if that didn't bother you, consider this: of 15,000 abuse reports in 2016, at its current rate of (proxy-defined) accuracy, the system would produce 3,600 incorrect predictions. And the planned model is supposed to be "run on a daily or weekly basis on all babies born in Allegheny County." This is a big step forward not just in extending the tech to everyone, but also in commitment to prediction. Prediction is about guessing how poor people might behave in the future based on data from their networks, not just about judging their past individual behavior, and thus it can infect entire communities and generations. At the same time, "digital poorhouses," as Eubanks calls the networks into which data about poor people are fed, are hard to see and hard to understand, making them harder to organize against.

Eubanks also points out that parents can naturally resent outside scrutiny and often feel that once the child welfare system is involved the standards keep getting raised on them, no matter what they try to do. And caseworkers interpret resistance and resentment as danger signs. While these reactions aren't directly dependent on the technology, they are human behaviors that change what the technology does in the world.

In theory, big data could increase transparency and decrease discrimination where that comes from the humans in the system. Unfortunately, that doesn't seem to be what's happening. Among other things, the purported "transparency" of algorithms, even putting [trade secrets]() aside, is very much a transparency for the elite who can figure the code out, not for ordinary participants in democratic governance, who basically have to take experts' explanations on faith.

In addition, Eubanks finds:

> the philosophy that sees human beings as unknowable black boxes and machines as transparent…deeply

> troubling. It seems to me a worldview that surrenders any attempt at empathy and forecloses the possibility of ethical development. The presumption that human decision-making is opaque and inaccessible is an admission that we have abandoned a social commitment to try to understand each other. Poor and working-class people in Allegheny County want and deserve more: a recognition of their humanity, an understanding of their context, and the potential for connection and community.

This sounds great, but I wonder if it is fully convincing, in the fallen world in which we live. On the other hand, given that there are other interventions that wouldn't sort the "worthy" from the "unworthy" in the ways that current underfunded services are forced to do, it is certainly persuasive to argue that we shouldn't try to move from biased caseworkers to biased algorithms.

Along with non-technical solutions, Eubanks offers some ethics for designers, focusing on whether the tools they make increase the self-determination and agency capabilities of the poor, and whether they'd be tolerated if targeted at the non-poor. I think she's overly optimistic about the latter criterion, at least as applied to private corporate targeting, which we barely resist. The example of TSA airport screening is also depressing. Perhaps I'd suggest the modification that, if we expect wealthier people to buy their way out of the system, as they can with TSA Pre-check and CLEAR Global Entry (at least if they're not Muslim), then there is a problem with the system. Informed consent and designing with histories of oppression in mind, rather than assuming that equity and good intentions are the default baselines, are central to her vision of good technological design.

Like the far more caustic Evgeny Morozov, Eubanks contends that we have turned to technology to solve human problems in ways that are both corrupting and self-defeating. And Eubanks doesn't focus the blame on Silicon Valley. The call for automation is coming from inside the polity. In fact, while IBM comes in for substantial criticism for overpromising in the Indiana example, the real drivers in Eubanks' story are the policy wonks who are either trying to shrink the system until it can be drowned in the bathtub (Indiana), or sincerely trying to build something helpful while resources are continually being drained from the system (Los Angeles and Pennsylvania).

Ultimately, Eubanks argues, the problem is that we're in denial about poverty, an experience that will happen to the majority of Americans for at least a year between the ages of 20 and 65, while two-thirds of us will use a means-tested public benefit such as TANF, SNAP, Medicaid, or SSI. But we persist in pretending that poverty is "a puzzling aberration that happens only to a tiny minority of pathological people." We pass a suffering man on the street and fail to ask him if he needs help. We don't keep our tormented child in an isolated place, as they do in Omelas. Instead of walking away, we walk by—but we don't meet each other's eyes as we do so. This denial is expensive in so many ways—morally, monetarily, and even physically, as we build entire highways, suburbs, private schools, and prisons so that richer people don't have to share in the lives of poorer people. It rots politics: "people who cannot meet eachothers' eyes will find it very difficult to collectively govern." Eubanks asks us to admit that, as Dan Kahan and his colleagues have repeatedly demonstrated in work on cultural cognition, our ideological problems won't be solved with data, no matter how well formed the algorithm.

Cite as: Rebecca Tushnet, *The Difference Engine: Perpetuating Poverty Through Algorithms*, JOTWELL (July 18, 2018) (reviewing Virginia Eubanks, **Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor** (2018)), https://cyber.jotwell.com/the-difference-engine-perpetuating-poverty-through-algorithms/.