


Let Them Eat Efficiency

STEPHANIE WYKSTRA

The past few years have seen an upsurge in warnings about biases embedded in technological tools. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*, by Virginia Eubanks, a political scientist at the State University of New York at Albany, joins this body of work. The book calls attention to ways that society has given short shrift to people who are in need, while raising important questions about whether new tools make people better or worse off than they would be otherwise.

The three main chapters of the book comprise three case studies, each illustrating aspects of what Eubanks calls “the digital poorhouse.” She begins by laying out the history of the (physical) poorhouse in the United States. These were facilities built as “homes of last resort” for indigent people, including abandoned children, older adults, and handicapped people. The first was built in Boston in 1662, and poorhouses became common in the 1800s. Eubanks describes many as having “horrid conditions,” in which inhabitants suffered from neglect and abuse.

As welfare benefits evolved in the twentieth century, new rules relegated large swaths of people—including people of color and never-married mothers—to stingy and punitive forms of financial relief. The mid-1970s saw the inception of new electronic databases at the same time that policy-makers started to cut welfare benefits. Eubanks writes that in 1973 nearly half of people under the poverty line received federal aid; a decade later, after digital tools were introduced, that rate had dropped to 30%. Thus, she writes, “the revolt against welfare rights birthed



Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor

by Virginia Eubanks. New York, NY: St. Martin's Press, 2017, 272 pp.

the digital poorhouse.”

The first case that Eubanks examines is the Indiana Family and Social Services Administration's (FSSA) turbulent effort to “modernize” its processes. FSSA is charged with doling out benefits, including Medicaid and the Supplemental Nutrition Assistance Program (SNAP, or what used to be called food stamps), to those who are eligible. In 2006, the state's Republican governor, Mitch Daniels, called for reforms to “clean up welfare waste,” calling FSSA a “monstrous bureaucracy.” The state hired IBM to “modernize” the system, shifting to call center workers rather than face-to-face case workers, adopting an online application, and requiring anyone currently receiving

benefits to resubmit paperwork.

After a few months, it was clear that the new system was a disaster. Eubanks describes accounts of call center employees with little experience bursting into tears because they couldn't answer caller questions, and documents laboriously faxed by applicants and then lost forever. The number of people receiving benefits plunged, including many who were eligible. Among those who lost benefits was a woman who ended up unable to obtain insulin for her diabetes for seven months. In 2010, Indiana sued IBM for breach of contract, arguing that the company had misled the state about its capacity to develop the new system. A 2017 judgment against the computer company that awarded Indiana \$78 million in damages has been appealed by IBM.

Eubanks's second case looks at homelessness in Los Angeles, with a focus on a new coordinated-entry system adopted by the city in 2013. The system involves an assessment tool that enables administrators to match homeless people with available housing. Survey questions include questions about mental health crises, drug activities, and potential for self-harm, as well as personally identifying information such as the applicant's full name, social security number, and birth date. The higher the score, the greater the likelihood that the person needs emergency mental health or medical care. A second algorithm then matches the person to housing for which he or she may be eligible. Whereas the old system was haphazard, sometimes rewarding only “the most functional people,” the new system matches people with housing that is a better fit for their circumstances.

The new system may enable Los Angeles to do a better job of allocating resources to those who are most urgently in need than it would in the absence of the tool. However, as Eubanks points out, 75% of the homeless in South Los Angeles remain unsheltered. Eubanks's main argument is that the new technology, without a big increase in spending, won't solve the problem. As she writes, "While Los Angeles residents have agreed to pay a little more to address the problem ... they don't want to spend the kind of money it would take to really solve the housing crisis." She also critiques the decision to store for years personally identifying and sensitive information about people, even if they remain unmatched to housing, arguing that there's a risk that this information will be used by the police to criminalize the homeless.

Although the Los Angeles and Indiana situations differ in many ways, Eubanks's critique of both has strong similarities. In both cases, many people in need were denied help. In Indiana this was due to incompetence; in Los Angeles there are simply not enough resources to adequately address the homeless problem. Eubanks writes that the technological aspects of both systems served to conceal or paint the decisions as reasonable ones, made under the guise of cutting "waste" or efficiently allocating resources. Instead, Eubanks suggests, welfare modernization was "a high-tech Trojan horse concealing a very old-fashioned effort to cut public benefits."

The third case focuses on an algorithm used in Allegheny County, Pennsylvania, by the Office of Children, Youth, and Families. The office is responsible for investigating potential cases of child neglect and abuse, which could lead to the removal of children from their parents and placement into foster care. The county recently adopted a risk assessment instrument called the Allegheny Family Screening Tool (AFST). The tool aids the process of deciding whether to conduct an

investigation following a call about potential mistreatment. AFST uses data from past cases to analyze which factors are correlated with the removal of a child from a home and with subsequent referral calls within two years.

AFST itself does not make decisions about whether to remove the child. Instead, it produces a "risk score" when someone calls in a potential case, drawing on data such as demographic information about the family, past welfare interaction, criminal history, and other information from government institutions and services. The risk score is one piece of information that call screeners use to decide whether to "screen in" the case, which means that it

To really help people, society needs to dedicate adequate resources to helping them, not just divvy out inadequate help more efficiently.

will advance to the next stage, in which a social worker will visit the home. (When the score is above a certain level, there's a mandatory screen-in, though this process is not automatic and is often overridden by supervisors).

Debate about risk assessments, such as those done using AFST, is taking place in a number of fields, including areas beyond family services, such as criminal justice decision-making. There is a lot of controversy centered around questions such as whether it is acceptable to use particular inputs (often correlated with race and socioeconomic status) or what constitutes "fairness" in algorithmic design. Advocates for risk assessment, including the designers in Allegheny County who developed AFST, argue that the tools can help combat bias and

increase fairness by providing more standardized and accurate predictions than human decision-makers on their own.

In Allegheny County in particular, the process of tool development was remarkably transparent: it involved the wider community in the design and was attentive to the potential for racial and other kinds of bias. Eubanks acknowledges this, writing that AFST is the "best-case scenario for predictive risk modeling in child welfare." She also said in an interview with *Jacobin* magazine that "in Allegheny County, they have done every single thing that progressive folks who talk about these algorithms have asked people to do." In spite of that, she said in the interview, "I still believe the system is one of the most dangerous things I've ever seen."

Why is this? Eubanks critiques the algorithm in several ways. First, she questions what the algorithm is designed to predict, which are cases of child removal and subsequent referral calls within two years. These are proxy outcomes and not cases of verified abuse, which are too rare for the algorithm to predict with any reasonable degree of accuracy. But this, Eubanks argues, means that the algorithm isn't necessarily predicting outcomes that society really cares about, but is instead predicting the behavior of the agency and of the community that makes referral calls. Eubanks is also concerned that there is unfairness in the inputs. The data on which the algorithm relies are, for the most part, more available for low-income families than for middle-class families, given that the latter interact with public services far less often. Finally, Eubanks thinks there's a false sense of transparency in the use of predictive algorithms. Even though Allegheny County does release quite a bit of information, including the inputs and outputs of the tool, Eubanks writes, "I find the philosophy that sees human beings as unknowable black boxes and machines as transparent deeply troubling." A computer-based prediction

system may carry a “cloak of evidence-based objectivity and fallibility,” she argues, and may also be difficult to change once in place.

These concerns make sense. And yet there remains a question about what she would suggest as the way forward. Is there a way for this tool to meet the appropriate burden of proof that it is making things fairer and more accurate (in the relevant senses) compared with human judgment alone, according to Eubanks? Are there particular changes that she’d want to see? Or does Eubanks think it is simply a bad idea for Allegheny County to use any algorithm in its call-referral process, regardless of how it’s done (and if so, why)? The chapter doesn’t provide clear answers to these questions.

This brings us to a larger point about the book. On first read, Eubanks’s overarching theme might seem to be that the technology has made things worse than they would have been otherwise. But she doesn’t spend a lot of time parsing the role played by human decision-making from the role of technology, or analyzing their relative impacts. This is because she is very likely making a different point: that technology as currently deployed in these cases isn’t making things *much better* for those who are suffering and in need. To really help people, society needs to dedicate adequate resources to helping them, not just divvy out inadequate help more efficiently.

This argument is clear from her discussion of the Indiana and Los Angeles cases, in which algorithms are used to distribute insufficient resources with greater efficiency. And even in the Allegheny County case, which may at first seem to be mostly about the algorithm, Eubanks often returns to the role of poverty in child welfare investigations. Are parents really neglectful, she asks, or are they suffering from poverty and at risk of having their children removed because it is so hard for them to get the resources they need? Her argument is that society needs a

much more dramatic change in how the poor are treated: new technologies, when accompanied by the same (stingy and punitive) policies, can’t be the way forward.

In the end, the strength of Eubanks’s work lies in her compassionate and close attention to the lives of the many people whom society has failed to help. In her conclusion, she quotes a 1968 speech by Dr. Martin Luther King Jr., who imagines telling “the God of history” about all the things we’ve accomplished through scientific and technological progress, and receiving the answer, “That was not enough!”

Stephanie Wykstra (@swykstr) is a freelance writer and researcher with a focus on criminal justice reform and research transparency. Her work has appeared in Slate, Vox, the Stanford Social Innovation Review, Aeon, and other places.


Money Well Spent?

IRWIN FELLER

Measuring the Economic Value of Research: The Case of Food Safety
edited by Kaye Husbands Fealing, Julia I. Lane, John L. King, and Stanley R. Johnson. New York, NY: Cambridge University Press, 2018, 204 pp.

Measuring the Economic Value of Research is an important book, edited by a distinguished group of researchers who focus on the science of science policy, an emerging interdisciplinary approach to evaluating the scientific enterprise. The book offers a case study of government

support of academic research on food safety, mainly by the US Department of Agriculture, the National Institutes of Health, and the National Science Foundation. The case study showcases the analytical and empirical prowess of a new data platform called UMETRICS (Universities: Measuring the Impacts of Research on Innovation, Competitiveness, and Science). The UMETRICS platform is a conceptually and technically innovative approach for obtaining, refining, and integrating heterogeneous data sets that can be used to better inform decisions relating



to the level, apportionment, human capital needs, and physical facility requirements of public investments in research and development (R&D) and the impacts or returns that these investments have on policy objectives.

The food safety case study is a