



O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. New York: Crown Publishing.

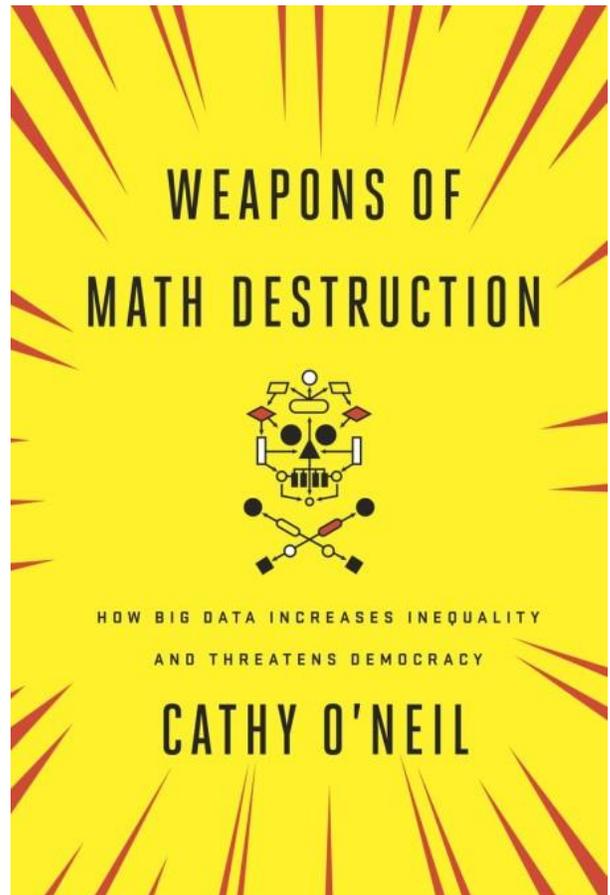
Pp. 259

ISBN: 978--0553041881-1

Reviewed by Francis Schrag
Emeritus, University of Wisconsin, Madison
United States

The effort to place education on a more scientific basis is at least a century old, involving a continuous effort to substitute “hard,” quantitative evidence for “soft,” qualitative judgment. Though schooling is only one of the domains Cathy O’Neil’s targets, her devastating critique of the increasing use of algorithms to evaluate teachers, rate universities, and identify college prospects makes *Weapons of Math Destruction* invaluable for educators and educational researchers. Moreover, they will find chapters dealing with other topics such as criminal justice, politics, employee selection, and the credit industry equally enlightening. Clearly written by a mathematician who has worked as a data scientist, *Weapons of Math Destruction* mounts a powerful critique of recent trends in educational accountability: the effort to improve schools and universities by deploying algorithms on large data sets, distorts the enterprise and subverts educational values. As O’Neil says in her conclusion:

Predictive models, are, increasingly, the tools we will be relying on to run our institutions, deploy our resources, and manage our lives...these models are constructed not just from data but from the choices we make about which data to pay



attention to—and which to leave out. Those choices are not just about logistics, profits, and efficiency. They are fundamentally moral. (p. 218)

One of O’Neil’s targets is value-added modeling, which measures teacher effectiveness on the basis of the rise (or fall) in student test scores at the end of the current year compared to scores obtained the previous year. These efforts are well intended, designed to reward the best and get rid of the worst teachers in a way that is free of bias or favoritism. But as O’Neil demonstrates, the system itself is based on questionable assumptions; it sometimes identifies the wrong teachers as failing (or succeeding), which can cause irreparable harm to individual teachers. So, you might say, no system is perfect. True, but what makes this kind of system so pernicious are four additional features which together constitute what O’Neil calls Weapons of Mass Destruction (WMDs):

1. There are no feedback loops by which the model can improve its effectiveness.
2. The assumptions of the models are masked by the appearance of value-free scientific neutrality.
3. The entire process is opaque making it almost impossible for an unjustly evaluated teacher to challenge her evaluation.
4. The victims are often the less advantaged in our society while the most advantaged go unscathed.

To illustrate the effects of WMDs, O’Neil discusses the case of a particular teacher who was almost certainly fired unjustly. She had received glowing reviews in previous years, and there is strong evidence that the high scores achieved by her entering class were doctored, making her failure to further raise them appear to be a result of her own inadequacy as a teacher. Since the teacher had no access to the algorithm on which she was judged wanting, she had no basis to challenge the evaluation. Now one might think a data scientist like O’Neil would urge that the value-added models be refined and perfected before use to eliminate whatever bugs they contain. Not so; instead, she tells us: Sometimes, it is all too clear from the get-go that certain WMDs are only primitive tools, which hammer complexity into simplicity, making it easier for managers to fire groups of people or to offer discounts to others.

Forget, at least for the next decade or two, about building tools to measure the effectiveness of teachers. It’s too complex to model, and the only available data are crude proxies. The model is simply not good enough yet to inform important decisions about the people we trust to teach our children. That’s a job that requires subtlety and context. Even in the age of Big Data, it remains a problem for humans to solve (pp. 208-209).

Members of the Society for Research on Educational Effectiveness and others who are attempting to develop so-called “evidence-based education” can easily ignore challenges to value-added teacher evaluation coming from those of us on the outside, especially those of us who are mathematically challenged; but they could not, or at least should not, ignore criticisms coming from someone who understands data science, computer science, and statistics from the inside.

O’Neil focuses her critical sights on two WMDs in post-secondary education, the annual *US News* ranking of colleges and universities, and the algorithm-driven efforts of for-profit colleges to recruit students from the most vulnerable populations. She demonstrates the assumptions behind the *US News* rating system, some of them plausible, some not, and she identifies the fundamental flaw behind the entire system: Since “educational excellence” is a concept that is “squishy,” contested, and resists quantification, quantifiable proxies must be selected to stand in for it. Some of these, such as acceptance rates, bear only an assumed relation to education quality, and one criterion that matters to most families who have limited assets, is entirely absent—cost.

But the problem is not so much the erroneous rankings that are produced by the *US News* algorithm as the deleterious side effects the rankings generate, notably an arms race among institutions to improve their position relative to other institutions, an arms race resulting from the fact that quantifiable proxies are much more amenable to manipulation. Here’s but one example O’Neil provides: colleges invest huge amounts of money in sports programs and amenities that have little to do with educational quality. Why, then do they do it? Because, as she shows, using Texas Christian University as an example, successful sports programs and amenities such as a \$100 million central mall and a new student union, are known to spur applications. And if more students apply, while TCU maintains the size of its entering class, that lowers its acceptance rate. Indeed, on the basis of expenditures such as these, TCU climbed 37 places in the rankings over a seven-year period. (p. 57)

Moreover, O’Neil shows how the enormous resources and energy going into this arms race, are accompanied by the growth of a huge market in consultants who help high school students from more affluent families gain acceptance to selective institutions. She cites one, Steven Ma, who has developed his own algorithms to predict the chances of a client being accepted at a specific college. If the student is successful, the family pays a fee that slides according to the odds of acceptance; if the student is rejected the fee is waived. The service is not cheap. O’Neil says a hypothetical student whose stats she provides might be charged over \$25,000 for admission to NYU.

Focusing on the less favored segments of society, O’Neil demonstrates the way for-profit colleges deploy algorithms to recruit students from the neediest populations. They use massive data sets and machine learning software programs to identify the specific vulnerabilities of individual prospective students in order to target their recruitment strategy to those vulnerabilities. They then lure these students with all kind of promises, even using fabricated job-placement data, for example. (p. 71) According to O’Neil, the Apollo Group, the parent company for the University of Phoenix spent \$2,225 per student on marketing, more than twice as much as the \$892 it spends on instruction. By comparison the Portland (OR) Community College’s marketing effort is only 1.2% of the \$5,993 spent on instruction on instruction per student. (p. 79)

I’ve characterized O’Neil’s critique of the deployment of algorithms and Big Data in the educational sector as devastating, but I must confess to some uneasiness about where this leaves us with respect to issues like teacher evaluation. Remember that the “evidence-based” movement was the result of legitimate criticisms of inferences based on classroom observations, even by veteran observers. As O’Neil herself acknowledges, “...the human brain runs internal models of its own, and they’re often tinged with prejudice or self-interest.” (p. 209) (Or educational ideology, I might add.) But even this acknowledgement understates the problem. It is now increasingly recognized that in making judgments and predictions under uncertainty, we humans are remarkably fallible. Thanks to the tradition of psychological investigation pioneered by Daniel Kahneman and Amos Tversky, ably recounted in Michael Lewis’ (2017) *The Undoing Project*, we have every reason to be suspicious of even the most expert judgment.

So, here’s the choice: On the one hand, algorithms which avoid some human proclivities leading to injustices but bake in others that do the same; and, on the other hand, the individual human judgments of “experts” that we have no reason to rely on and which historically have also led to injustices. On what basis should we decide between these two? The question cannot be answered by further research, because what we’re after in education cannot be reduced to quantitative proxies, or rather, any proxy selected by one side is likely to be contested by the other.

I think O’Neil has given us a couple of good reasons to be wary of the algorithmic solution despite its allure. First, it tends to operate behind the backs of those it judges, depriving its victims of the opportunity to confront those who judge them; and second, it reinforces a conception of education that conflates the well-educated person with the high-scoring test taker.

Let me conclude with a striking observation. The kind of society we live in, one in which we are manipulated by WMDs without our awareness, much less our consent, exhibits precisely the kind of “disciplinary” power limned 40 years ago by Michel Foucault. What makes this striking is that Cathy O’Neil has never read Foucault.¹ In *Discipline and Punish*, Foucault contrasted the disciplinary power emerging in the eighteenth and nineteenth centuries with that which existed previously in feudal societies.

In a disciplinary regime... individualization is ‘descending’: as power becomes more anonymous and more functional, those on whom it is exercised tend to be more strongly individualized; it is exercised by surveillance rather than ceremonies...by comparative measures that have the ‘norm’ as reference...by ‘gaps’ rather than by deeds (1977, p. 193).

Foucault taught us that the progressive political and moral narrative we like to tell ourselves, which highlights the liberation of individuals from abusive institutions of domination prevalent in the pre-modern period, is not the whole story. We must also recognize the more subtle forms of domination that have evolved to replace them. Cathy O’Neil has provided a superb demonstration of the dangers of our blind faith in digital technologies that only appear to benefit all but actually favor the already powerful.

References

- Foucault, M. (1977). *Discipline and punish: The birth of the prison*. Trans. A. Sheridan. New York: Pantheon.
- Lewis, M. (2017). *The undoing project: A friendship that changed our minds*. New York: W. W. Norton.
- O’Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. New York: Crown Publishing.

About the Reviewer

Francis Schrag is Professor Emeritus in the departments of Educational Policy Studies and Philosophy at the University of Wisconsin, Madison. He has published two books and numerous articles on a wide variety of topics in education and philosophy.

¹ Personal communication, October 22, 2016.



education review // reseñas educativas
a multi-lingual journal of book reviews

editors: **gustavo e. fischman /**
melissa cast-brede / gene v glass

Supported by the Mary Lou Fulton Teachers College, Arizona State University

 *Education Review/ Reseñas Educativas/ Resenhas Educativas* is supported by the edXchange initiative's Scholarly Communications Group at the Mary Lou Fulton Teachers College, Arizona State University. Copyright is retained by the first or sole author, who grants right of first publication to the *Education Review*. Readers are free to copy, display, and distribute this article, as long as the work is attributed to the author(s) and ***Education Review***, it is distributed for non-commercial purposes only, and no alteration or transformation is made in the work. More details of this Creative Commons license are available at <http://creativecommons.org/licenses/by-nc-sa/3.0/>. All other uses must be approved by the author(s) or ***Education Review***. ***Education Review*** is published by the Scholarly Communications Group of the Mary Lou Fulton Teachers College, Arizona State University.

Please contribute reviews at <http://www.edrev.info/contribute.html>.

Connect with *Education Review* on Facebook (<https://www.facebook.com/pages/Education-Review/178358222192644>) and on Twitter @EducReview