entirety will benefit from the numerous examples that Keohane provides to show how the market mechanism can result in more just or more benevolent outcomes with only small interventions to overcome market failures. Here Keohane shows that the right sort of capitalism can be deployed to ethical ends.

Filled with clear examples and case studies, *Capital and the Common Good* is a superior introduction to the world of innovative finance. Keohane's attention to technology and financial theory balanced with a summary of challenges that each strategy faces results in a well-rounded treatment of the subject. It offers much for theologians, ethicists, and economists to consider.

—D. Glenn Butner, Jr.
Sterling College, Kansas

Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy

Cathy O'Neil

New York: Penguin Random House, 2016 (272 pages)

Weapons of Math Destruction exposes the tremendous power that data, the analytics of data, and the use of analytics yield over many aspects of our lives. Building on the shocking title, the author creatively uses the imagery of bombs, weapons, and war to organize the book and emphasize the real dangers of data analytics in our world.

The book begins with a chapter titled, "Bomb Parts," in which O'Neil describes the three measures she uses throughout the book to evaluate whether a data analytical system is a Weapon of Math Destruction (thematically abbreviated WMD) or is instead "benign." First, the more *opacity* a model has within it, the more dangerous it is. O'Neil considers *opacity* to be the level of transparency within the algorithm to those using it and affected by it as well as how accurately the statistics used in the analysis actually represent the desired outcome. Second, the larger the *scale* of the utilization and thus impact of the model, the more dangerous it is, simply because it affects more people. In addition, scale also includes the expanded applications of the model beyond its initial purpose. Third, the more *damage* a model has caused or has the potential to cause increases the likelihood of its diagnosis as a WMD. Within the analysis of the *damage* caused, O'Neil considers who exactly is hurt by the model and especially weighs the impact on vulnerable populations such as the poor or minorities. These three measures are best understood when operationalized in the dozens of examples presented in the subsequent chapters, of which we will explore several to get the gist of the dangers.

The sport of baseball is filled with statistics, statisticians, and sophisticated analyses that drive decisions in nearly every aspect of the sport. Baseball teams today rely heavily on these analyses to determine which players to draft, to retain, and to trade; where to position their defenses against various batters; which pitcher to start or bring in depending on the game, inning, upcoming batters, and so forth; and countless other decisions.

In evaluating the extensive use of models and analysis in baseball against the WMD criteria, we can see that these models are benign. First, the models are quite transparent with clear, specific, highly relevant data collected and used in the analysis. For example, anyone with enough time, interest, and statistical proclivity could collect and analyze data on a certain batter to predict his tendencies to hit the ball to certain areas of the field in a given inning, against specific teams or pitchers, or even in certain weather. The statistics used in this analysis are specific and clearly correlated to the desired outcome, such as the number of times a certain batter hit the ball to left field in a given situation in the past year. Second, the baseball models are not *scaled* beyond their intention and have no real risk of being applied to situations or people beyond their scope. Third, while some may argue that data have taken the fun out of the sport of baseball or that a player was mistakenly passed over in the draft due to the modeling, any actual *damage* would be narrowly focused and usually only affect the team using the analysis. Furthermore, a critical point is that when mistakes are made, the models will be corrected due to the constant feedback of data from all teams and players being continually fed into the system. For example, if a team passes over a draft pick and that player goes on to become an all-star for another team, the team that passed him over now has new data to add to their models so as not to make the same mistake in the future. Given this assessment of the use of data analytics in baseball, these models do not qualify as WMD.

A similar use of analytics that is widely used in the business and academic worlds is automated screening systems for human resource departments and college admissions offices. While explored separately in the book, these systems have similar issues that cause alarm and move them into the category of WMD. Unlike the nonopaque baseball models, these systems use proxy statistics to predict future success and are typically not transparent to the applicant and sometimes not transparent to the company using them either. The algorithms used in these applicant-screening systems use proxy statistics such as number of years at a current position, level of education, score on an exam, and strength of references or essays to determine whether a candidate gets a job or gets into a school. These are considered proxy statistics because they do not necessarily correlate to a successful future employee or student. Furthermore, these systems, unlike the baseball systems, do not have large-*scale*, interorganizational feedback informing their systems and improving the models. For example, if a company's screening system rejects an applicant who goes on to become a star salesperson and later a vice president for another company, the rejecting company's models will never be updated to make sure they accept a candidate like that in the future. Additionally, these systems fail in the measure of *opacity* due to the inability of applicants to know the algorithm being used to evaluate them. For example, if candidates are rejected in one college application system, they will likely be rejected by other systems, and they may have no way to discover why they are being rejected. Whether these systems have the *scale* and *damage* levels to further qualify them as WMD may be debatable, but they provide a clear example of failing the *opacity* measure, are risky for organizations to use, and represent threats to applicants being screened.

Other examples of WMD presented by O'Neil are systems predicting criminal recidivism and crime locations, evaluating risk of mortgage-backed bonds, and ranking the top colleges and universities in *U.S. News and World Report*. These systems all fail in one or more of the three WMD measures and are classified as dangerous threats. The recidivism and crime prediction systems provide the clearest illustration of the WMD measures.

The recidivism and crime prediction systems, while attempting to overcome biased human predictions, end up basing their analysis on statistics that do, in fact, introduce the very bias they are trying to avoid into the model. Among the data considered as part of the recidivism models are criminal activity of acquaintances and information about a person's living situation and hometown. These types of questions guarantee that people from inner cities will score tougher sentences than those from suburban or rural areas. These systems also have pernicious feedback loops, in which the data fed back into the models only affirm and reinforce the model, even if the statistics are not predictive and it is the model itself that is resulting in the "accuracy" of the system. For example, if the crime prediction system predicts that crimes will be more likely in certain neighborhoods and thus the police presence increases in those areas, it is guaranteed that there will be more prosecuted crimes in those areas—even if there are just as many crimes happening in other areas that are now going undetected due to lack of police presence in those places. In addition to these issues resulting in concerning *opacity* and *damage* measures, the *scale* of these types of systems is large and continuing to grow due to the underfunded and understaffed penal systems and the perceived effectiveness and efficiency of these systems.

O'Neil's theoretical and practical experience and understanding of data modeling and analysis as well as her concern about the impact of these models on individuals, organizations, and the world make for a compelling and eye-opening case for using great caution when building, using, and trusting data analytical systems. While some of O'Neil's specific conclusions may seem agenda-driven, her core argument is well made and supported. Of the three WMD measures, *opacity* and *scale* are more objective and clear and could be used alone to determine the risk of a given analytical model. However, bringing a measure of *damage* into the model rightly adds an assessment of the systems' impact on individuals and communities, even while it adds potential subjectivity into the model.

It is clear that no analytical model is perfect and no model should be blindly trusted without humans understanding the model and having a key role in interpreting the findings. As with most things, building and using these tools without a properly formed moral conscience leads to injustice and disorder.

—Cory Maloney (e-mail: cmaloney@franciscan.edu)
Franciscan University of Steubenville, Ohio