ALGORISTS DREAM OF OBJECTIVITY

Peter Galison

Peter Galison *is a science historian; Joseph Pellegrino University Professor and co-founder of the Black Hole Initiative, Harvard University; and the author of* Einstein's Clocks and Poincaré's Maps: Empires of Time.

In his second-best book, the great medieval mathematician al-Khwarizmi described the new place-based Indian form of arithmetic. His name, soon sonically linked to "algorismus" (in late medieval Latin) came to designate procedures acting upon numbers—eventually wending its way through "algorithm," (on the model of "logarithm"), into French and on into English. But I like the idea of a modern algorist, even if my spellcheck does not. I mean by it someone profoundly suspicious of the intervention of human judgment, someone who takes that judgment to violate the fundamental norms of what it is to be objective (and therefore scientific).

Near the end of the 20th century, a paper by two University of Minnesota psychologists summarized a vast literature that had long roiled the waters of prediction. One side, they judged, had for all too long held resolutely—and ultimately unethically—to the "clinical method" of prediction, which prized all that was *subjective*: "informal," "in-the-head," and "impressionistic." These clinicians were people (so said the psychologists) who thought they could study their subjects with meticulous care, gather in committees, and make judgment-based predictions about criminal recidivism, college success, medical outcomes, and the like. The other side, the psychologists continued, embodied everything the clinicians did not, embracing the *objective*: "formal," "mechanical," "algorithmic." This the authors took to stand at the root of the whole triumph of post-Galilean science. Not only did science benefit from the actuarial; to a great extent, science *was* the mechanical-actuarial. Breezing through 136 studies of predictions, across domains from sentencing to psychiatry, the authors showed that in 128 of them, predictions using actuarial tables, a multiple-regression equation, or an algorithmic judgment equalled or exceeded in accuracy those using the subjective approach. They went on to catalog seventeen fallacious justifications for clinging to the clinical. There were the self-interested foot-draggers who feared losing their jobs to machines. Others lacked the education to follow statistical arguments. One group mistrusted the formalization of mathematics; another excoriated what they took to be the actuarial "dehumanizing;" yet others said that the aim was to understand, not to predict. But whatever the motivations, the review concluded that it was downright immoral to withhold the power of the objective over the subjective, the algorithmic over expert judgment.[1]

---

[1] William M. Grove & Paul E. Meehl, "Comparative efficiency of informal (subjective, impressionistic) and formal (mechanical, algorithmic) prediction procedures: The Clinical-Statistical Controversy," *Psychology, Public Policy, and Law,* 2:2. 293-323 (1996).

The algorist view has gained strength. Anne Milgram served as Attorney General of the State of New Jersey from 2007 to 2010. When she took office, she wanted to know who the state was arresting, charging, and jailing, and for what crimes. At the time, she reports in a recent TED Talk, she could find almost no data or analytics. By imposing statistical prediction, she continues, law enforcement in Camden during her tenure was able to reduce murders by 41 percent, saving thirty-seven lives, while dropping the total crime rate by 26 percent. After joining the Arnold Foundation as its vice president for criminal justice, she established a team of data scientists and statisticians to create a risk-assessment tool; fundamentally, she construed the team's mission as deciding how to put "dangerous people" in jail while releasing the non-dangerous. "The reason for this," Milgram contended, "is the way we make decisions. Judges have the best intentions when they make these decisions about risk, but they're making them subjectively. They're like the baseball scouts twenty years ago who were using their instinct and their experience to try to decide what risk someone poses. They're being subjective, and we know what happens with subjective decision making, which is that we are often wrong." Her team established nine hundred-plus risk factors, of which nine were most predictive. The questions, the most urgent questions, for the team were: Will a person commit a new crime? Will that person commit a violent act? Will someone come back to court? We need, concluded Milgram, an "objective measure of risk" that should be inflected by judges' judgment. We know the algorithmic statistical process works. That, she says, is "why Google is Google" and why moneyball wins games.[2]

Algorists have triumphed. We have grown accustomed to the idea that protocols and data can and should guide us in everyday action, from reminders about where we probably want to go next to the likely occurrence of crime. By now, according to the literature, the legal, ethical, formal, and economic dimensions of algorithms are all quasi-infinite. I'd like to focus on one particular siren song of the algorithm: its promise of objectivity.

Scientific objectivity has a history. That might seem surprising. Isn't the notion—expressed above by the Minnesota psychologists—right? Isn't objectivity coextensive with science itself? Here it's worth stepping back to reflect on all the epistemic virtues we might value in scientific work. Quantification seems like a good thing to have; so, too, do prediction, explanation, unification, precision, accuracy, certainty, and pedagogical utility. In the best of all possible worlds these epistemic virtues would all pull in the same direction. But they do not—not any more than our ethical virtues necessarily coincide. Rewarding people according to their need may very well conflict with rewarding people according to their ability. Equality, fairness, meritocracy—ethics, in a sense, is all about the adjudication of conflicting goods. Too often we forget that this conflict exists in science, too. Design an instrument to be as sensitive as possible and it often fluctuates wildly, making repetition of a measurement impossible.

"Scientific objectivity" entered both the practice and the nomenclature of science after the first third of the 19th century. One sees this clearly in the scientific atlases that provided scientists with the basic objects of their specialty: There were (and are) atlases

---

[2] TED Talk, January 2014, https://www.ted.com/speakers/anne_milgram.

of the hand, atlases of the skull, atlases of clouds, crystals, flowers, bubble-chamber pictures, nuclear emulsions, and diseases of the eye. In the 18th century, it was obvious that you would not depict this particular, sun-scorched, caterpillar-chewed clover found outside your house in an atlas. No, you aimed—if you were a genius natural philosopher like Goethe, Albinus, or Cheselden—to observe nature but then to perfect the object in question, to abstract it visually to the ideal. Take a skeleton, view it through a camera lucida, draw it with care. Then correct the "imperfections." The advantage of this parting of the curtains of mere experience was clear: It provided a universal guide, one not attached to the vagaries of individual variation.

As the sciences grew in scope, and scientists grew in number, the downside of idealization became clearer. It was one thing to have Goethe depict the "ur-plant" or "ur-insect." It was quite another to have a myriad of different scientists each fixing their images in different and sometimes contradictory ways. Gradually, from around the 1830s forward, one begins to see something new: a claim that the image making was done with a minimum of human intervention, that protocols were followed. This could mean tracing a leaf with a pencil or pressing it into ink that was transferred to the page. It meant, too, that one suddenly was proud of depicting the view through a microscope of a natural object even with its imperfections. This was a radical idea: snowflakes shown without perfect hexagonal symmetry, color distortion near the edge of a microscope lens, tissue torn around the edges in the process of its preparation.

Scientific objectivity came to mean that our representations of things were executed by holding back from intervention—even if it meant reproducing the yellow color near the edge of the image under the microscope, despite the fact that the scientist *knew* that the discoloration was from the lens, not a feature of the object of inquiry. The advantage of objectivity was clear: It superseded the desire to see a theory realized or a generally accepted view confirmed. But objectivity came at a cost. You lost that precise, easily teachable, colored, full depth-of-field, artist's rendition of a dissected corpse. You got a blurry, bad depth-of-field, black-and-white photograph that no medical student (nor even many medical colleagues) could use to learn and compare cases. Still, for a long stretch of the 19th century, the virtue of hands-off, self-restraining objectivity was on the rise.

Starting in the 1930s, the hardline scientific objectivity in scientific representation began running into trouble. In cataloging stellar spectra, for example, no algorithm could compete with highly trained observers who could sort them with far greater accuracy and replicability than any purely rule-following procedure. By the late 1940s, doctors had begun learning how to read electroencephalograms. Expert judgment was needed to sort out different kinds of seizure readings, while none of the early attempts to use frequency analysis could match that judgment. Solar magnetograms—mapping the magnetic fields across the Sun—required the trained expert to pry the real signal from artifacts that emerged from the measuring instruments. Even particle physicists recognized that they could not program a computer to sort certain kinds of tracks into the right bins; judgment, trained judgment, was needed.

There should be no confusion here: This was not a return to the invoked genius of an 18th-century idealizer. No one thought you could train to be a Goethe who alone among scientists could pick out the universal, ideal form of a plant, insect, or cloud. Expertise could be learned—you could take a course to learn to make expert judgments about electroencephalograms, stellar spectra, or bubble-chamber tracks; alas, no one has ever thought you could take a course that would lead to the mastery of exceptional insight. There can be no royal road to becoming Goethe. In scientific atlas after scientific atlas, one sees explicit argument that "subjective" factors had to be part of the scientific work needed to create, classify, and interpret scientific images.

What we see in so many of the algorists' claims is a tremendous desire to find scientific objectivity precisely by abandoning judgment and relying on mechanical procedures—in the name of scientific objectivity. Many American states have legislated the use of sentencing and parole algorithms. Better a machine, it is argued, than the vagaries of a judge's judgment.

So here is a warning from the sciences. Hands-off algorithmic proceduralism did indeed have its heyday in the 19th century, and of course still plays a role in many of the most successful technical and scientific endeavors. But the idea that mechanical objectivity, construed as binding self-restraint, follows a simple, monotonic curve increasing from the bad impressionistic clinician to the good externalized actuary simply does not answer to the more interesting and nuanced history of the sciences.

There is a more important lesson from the sciences. Mechanical objectivity is a scientific virtue among others, and the hard sciences learned that lesson often. We must do the same in the legal and social scientific domains. What happens, for example, when the secret, proprietary algorithm sends one person to prison for ten years and another for five years, for the same crime? Rebecca Wexler, Visiting Fellow at the Yale Law School Information Society Project, has explored that question, and the tremendous cost that trade-secret algorithms impose on the possibility of a fair legal defense.[3] Indeed, for a variety of reasons, law enforcement may not want to share the algorithms used to make DNA, chemical, or fingerprint identifications, which puts the defense in a much weakened position to make its case. In the courtroom, objectivity, trade secrets, and judicial transparency may pull in opposite directions. It reminds me of a moment in the history of physics. Just after World War II, the film giants Kodak and Ilford perfected a film that could be used to reveal the interactions and decays of elementary particles. The physicists were thrilled, of course—until the film companies told them that the composition of the film was a trade secret, so the scientists would never gain complete confidence that they understood the processes they were studying. Proving things with unopenable black boxes can be a dangerous game for scientists, and doubly so for criminal justice.

Other critics have underscored how perilous it is to rely on an accused (or convicted) person's address or other variables that can easily become, inside the black box of algorithmic sentencing, a proxy for race. By dint of everyday experience, we have grown used to the fact that airport security is different for children under the age of twelve and adults over the age of seventy-five. What factors do we want the algorists to have in their often hidden procedures? Education? Income? Employment history? What one has read, watched, visited, or bought? Prior contact with law enforcement? How do we want algorists to weight those factors? Predictive analytics predicated on mechanical objectivity comes at a price. Sometimes it may be a price worth paying; sometimes that price would be devastating for the just society we want to have.

---

[3] Rebecca Wexler, "Life, Liberty, and Trade Secrets: Intellectual Property in the Criminal Justice System," 70 *Stanford Law Review,* XXX (2018).

More generally, as the convergence of algorithms and Big Data governs a greater and greater part of our lives, it would be well worth keeping in mind these two lessons from the history of the sciences: Judgment is not the discarded husk of a now pure objectivity of self-restraint. And mechanical objectivity is a virtue competing among others, not the defining essence of the scientific enterprise. They are lessons to bear in mind, even if algorists dream of objectivity.