

ANNALS OF TECHNOLOGY

NEW YORK CITY'S BOLD, FLAWED ATTEMPT TO MAKE ALGORITHMS ACCOUNTABLE

By Julia Powles

December 20, 2017



Automated systems guide the allocation of everything from firehouses to food stamps. So why don't we know more about them? Photograph by Mario Tama / Getty

The end of a politician's time in office often inspires a turn toward the existential, but few causes are as quixotic as the one chosen by James Vacca, who this month hits his three-term limit as a New York City Council member, representing the East Bronx. Vacca's nearly four decades in local government could well be defined by a bill that he introduced in August, and that passed last Monday by a unanimous vote. Once signed into law by Mayor Bill de Blasio, the legislation will establish a task force to examine the city's "automated decision systems"—the computerized algorithms that guide the allocation of everything from police officers and firehouses to public housing and food stamps—with an eye toward making them fairer and more open to scrutiny. In mid-October, I and some of my colleagues from a group at Cornell Tech that works on algorithmic accountability attended a hearing of the Council's technology committee to offer testimony on the bill. As Vacca, who chairs the committee, declared at the time, "If we're going to be governed by machines and algorithms and data, well, they better be transparent." Many of his constituents, he said, felt that "some inhuman computer is spitting them out and telling them where to go, and, if you don't like it, lump it."

Algorithms intersect with the daily lives of New Yorkers in countless ways, matching students with schools, assessing teacher performance, rooting out Medicaid fraud, and helping building inspectors manage their workloads. Vacca first became interested in the issue in the context of policing; he felt that his local precinct in the Bronx couldn't adequately explain the "criteria and formula" behind its staffing decisions. "That always annoyed me, and I felt that I was not being given a lot of the answers I wanted," Vacca told me, earlier this month. Starting in May, he and a couple of enterprising young staffers, Zachary Hecht and Malaika Jabali, took a stab at addressing the problem. They wrote a draft bill of about a hundred words, which became the focus of the October hearing. A tiny, intriguing, ambitious thing, it proposed that whenever a city agency wished to use an

automated system to apportion policing, penalties, or services, the agency would be required to make the source code—the system’s inner workings—available to the public. It would also be required to simulate the algorithm’s real-world performance using data submitted by New Yorkers.

Very quickly, this version of the legislation proved to be a long shot. “Many stakeholders communicated that, since we are going into unknown terrain, they wanted to go a little slower than my original bill anticipated,” Vacca said. The final law ditches the original draft’s disclosure requirements and sets up a fact-finding task force in their place. Convened by de Blasio, the task force will develop recommendations on a range of issues, including which types of algorithms should be regulated, how private citizens can “meaningfully assess” the algorithms’ functions and gain an explanation of decisions that affect them personally, and how the government can address “instances in which a person is harmed” by algorithmic bias. The only relic of the original draft’s requirements is an oblique reference to “making technical information . . . publicly available where appropriate.”

The task force will be the first city-led initiative of its kind in the country, and it is likely to have a significant impact, nationally and internationally, when it reports its findings, in late 2019. There is no doubt, however, that the final law represents a scaling back of Vacca’s early ambitions. One of the main stumbling blocks in the first draft, according to testimony at the October hearing and a number of sources involved in the negotiations, was the requirement to make source code fully public. This invited strong resistance from some policy experts, who warned that such openness might create security risks and give bad actors an easy way to game the public-benefits system, and from tech companies, which argued that it would force them to disclose proprietary information, supposedly harming their competitive advantage.

In our testimony at the October hearing, Helen Nissenbaum, Thomas Ristenpart, and I warned the technology committee that the proprietary-information

argument might well thwart any attempt at algorithmic transparency, giving companies too much leeway to advance “broad and baseless” claims to corporate secrecy. We proposed a qualified solution—less than total disclosure of the source code, more than nothing at all—with a particular emphasis on the data that drives the city’s systems. But the administration wasn’t persuaded. As Freddi Goldstein, a spokeswoman for de Blasio’s office, told me, “Publishing the proprietary information of a company with whom we contract would not only violate our agreement, it would also prohibit other companies from ever doing business with us, which would prevent us from trying innovative solutions to solve everyday problems through technology.”

In seeking to address these concerns, the final law introduces a couple of problems of its own. Currently, Vacca said, the Council is “impeded in doing our oversight function” by a lack of access to basic knowledge. There is no readily accessible public information on how much the city spends on algorithmic services, for instance, or how much of New Yorkers’ data it shares with outside contractors. Given the Council’s own struggle to find answers, the question now is whether the task force will do any better. Can it develop good recommendations, and fulfill its mandate, without the close coöperation of agencies and contractors? An intermediate draft of Vacca’s bill included extensive reporting requirements, which would have compelled agencies to provide the task force with relevant information. But that draft, like the August version, was rejected by the city administration, and now the task force will have to rely on voluntary disclosures as it studies how automated systems are designed, procured, and audited. For a government body without real legal powers, this will be a Herculean, or perhaps Sisyphean, undertaking.

The law’s second apparent failing is that it doesn’t address how the city government, and those who advise it, can exercise some muscle in their dealings with the companies that create automated-decision systems. As my discussion with Goldstein made clear, the administration is committed to protecting the

contractual and proprietary interests of technology vendors. But, when I spoke with Ellen Goodman, a professor at Rutgers Law School who has been researching city-vender contracts for predictive-algorithm programs, she argued that “we expect our representatives to push back on the universe of what is claimed as proprietary.” This is especially the case in New York, whose size, wealth, and high-quality demographic data make it a more desirable client than most cities. “For many of these vendors, it’s the biggest customer they’ll get,” Goodman said. “If New York doesn’t use that power to make systems accountable, who will?”

Frank Pasquale, a law professor at the University of Maryland who has advocated for qualified transparency as a means of balancing commercial and public interests, told me much the same. “While the terms of past contracts are hard to revisit, New York City should commit to demanding openness in all future contracts with vendors of these algorithmic services,” he said. “They have the leverage here, not the firms. Secrecy may incentivize tiny gains in efficiency, but those are not worth the erosion of legitimacy and public confidence in government. It’s a dereliction of duty to allow vital decisions to be made by a black box.”

Whatever the new law’s inadequacies, many of the people I spoke with saw it as an opportunity for greater engagement on important questions. “Think of this bill as an experiment in the world of algorithmic accountability, sent out much like Captain Picard, from ‘Star Trek,’ would send out a probe to explore a wormhole,” Cathy O’Neil, the author of “Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy,” told me. “What we’re finding is that the world of algorithms is one ugly wormhole.” In insulating algorithms and their creators from public scrutiny, rather than responding to civic concerns about bias and discrimination, the existing system “propagates the myth that those algorithms are objective and fair,” O’Neil said. “There’s no reason to believe either.”

Julia Powles researches technology law and policy at Cornell Tech and New York University.

More: [Algorithms](#) [Algorithmic bias](#)

THE DAILY

The best of *The New Yorker*, every day, in your in-box, plus occasional alerts
when we publish major stories.

E-mail address

Your e-mail address

Sign up

By signing up, you agree to our [User Agreement](#) and [Privacy Policy & Cookie Statement](#).

Read More



THE DOT POLITIC

By Jacqueline Feldman



CURRENCY

HOW AN OLD HACKING LAW HAMPER'S THE FIGHT AGAINST ONLINE DISCRIMINATION

By E. Tammy Kim

By Gideon Lewis-Kraus



INFINITE SCROLL

IS A.I. ART STEALING FROM ARTISTS?

According to the lawyer behind a new class-action suit, every image that a generative tool produces “is an infringing, derivative work.”

By Kyle Chayka

Your Privacy Choices