CENTER FOR DATA INNOVATION

November 21, 2022

Federal Trade Commission
Ms. April Tabor
Office of the Secretary
400 7th St. SW, 5th Floor, Suite 5610
Washington, DC 20024

Re: Trade Regulation Rule on Commercial Surveillance and Data Security

Dear Ms. Tabor,

On behalf of the Center for Data Innovation (datainnovation.org), I am pleased to submit this response to the Federal Trade Commission's (FTC) request for comments on its advance notice of proposed rulemaking (ANPR).[1]

The Center for Data Innovation studies the intersection of data, technology, and public policy. With staff in Washington, London, and Brussels, the Center formulates and promotes pragmatic public policies designed to maximize the benefits of data-driven innovation in the public and private sectors. It educates policymakers and the public about the opportunities and challenges associated with data, as well as technology trends such as open data, artificial intelligence, and the Internet of Things. The Center is part of the Information Technology and Innovation Foundation (ITIF), a nonprofit, nonpartisan think tank.

OVERVIEW

As explained in the ANPR, the Commission is considering whether it should implement new regulatory measures concerning the ways in which companies collect, aggregate, protect, use, analyze, and retain consumer data, as well as transfer, share, sell, or otherwise monetize that data in ways that are unfair or deceptive.

The Commission is particularly concerned about "companies' growing reliance on automated systems is creating new forms and mechanisms for discrimination based on statutorily protected categories, including in critical areas such as housing, employment, and healthcare."[2] To address these concerns, the Commission appears to be focused on regulatory measures that would minimize

[1] "A Proposed Rule by the Federal Trade Commission" Trade Regulation Rule on Commercial Surveillance and Data Security," Federal Register, August 22, 2022, https://www.federalregister.gov/documents/2022/08/22/2022-17752/trade-regulation-rule-on-commercial-surveillance-and-data-security.
[2] Ibid.

algorithmic error. The Commission also appears to be under the impression that companies have the ability and responsibility to minimize these errors through their own individual actions.
The comments below explain why algorithmic fairness is not simply a function of minimizing error rates and why the FTC should consider the effect of any particular decision system (whether algorithmic or human) on inequality as a whole rather than focusing exclusively on automated systems.

Please find our responses to the following questions in the document below.

A. How prevalent is algorithmic error? To what extent is algorithmic error inevitable? If it is inevitable, what are the benefits and costs of allowing companies to employ automated decision-making systems in critical areas, such as housing, credit, and employment? To what extent can companies mitigate algorithmic error in the absence of new trade regulation rules?
B. To what extent, if at all, should new rules require companies to take specific steps to prevent algorithmic errors? If so, which steps? To what extent, if at all, should the Commission require firms to evaluate and certify that their reliance on automated decision-making meets clear standards concerning accuracy, validity, reliability, or error? If so, how? Who should set those standards, the FTC or a third-party entity? Or should new rules require businesses to evaluate and certify that the accuracy, validity, or reliability of their commercial surveillance practices are in accordance with their own published business policies?
C. To what extent, if at all, do consumers benefit from automated decision-making systems? Who is most likely to benefit? Who is most likely to be harmed or disadvantaged? To what extent do such practices violate Section 5 of the FTC Act?
D. If new rules restrict certain automated decision-making practices, which alternatives, if any, would take their place? Would these alternative techniques be less prone to error than the automated decision-making they replace?

Sincerely,

Hodan Omaar
Senior Policy Analyst
Center for Data Innovation
homaar@datainnovation.org

A. How prevalent is algorithmic error? To what extent is algorithmic error inevitable? If it is inevitable, what are the benefits and costs of allowing companies to employ automated decision-making systems in critical areas, such as housing, credit, and employment? To what extent can companies mitigate algorithmic error in the absence of new trade regulation rules?

Every system—human or computer—can have errors, but what counts as an error is a matter of perspective and whether an error harms or helps consumers depends on context. Suppose a bank has certain conditions for offering customers a loan. If the bank offers a loan to someone who does not meet those conditions, and that customer pays the loan back on time, did the bank make an error in offering the loan or in setting the conditions for the loan? Conversely, if a bank offers a loan to someone who eventually defaults on their payments, did the bank make an error in issuing the loan? In all of these cases, whether the bank used a human loan officer or an algorithm to make these decisions is mostly irrelevant. And in general, companies operating in the free market have an incentive to reduce errors (whether human or computer-generated) because errors hurt their businesses (e.g., by losing customers, overlooking better job applicants, or making mistakes they must later remedy).

The FTC has not defined what it means by algorithmic error, but if the FTC's goal is to ensure outcomes are fair when businesses use automated systems in decision-making, striving for systems to be error-free is not the chief means to that end. To see why, consider two housing authorities that are using AI systems to achieve the same policy goal: allocating financial support through housing assistance programs. One housing authority uses a system whose objective function is to minimize the total number of families that experience eviction. The other housing authority uses a system whose objective function is to first provide the family who is most likely to be evicted with as much assistance as possible, then move on to the next, until the budget is exhausted. Assume both systems are completely error-free. As researchers from Harvard, Cornell, and Princeton University show, the objective function chosen in this sort of scenario can target very different groups of people.[3] Even if the systems are error-free and working completely as intended, they would have disparate outcomes because they formalize the problem in different ways.

The FTC should recognize that algorithmic fairness is not simply a function of an error rate and it cannot effectively address unfair outcomes by merely encouraging companies to mitigate algorithmic error. Rather than focusing on how to get companies to optimize their algorithms to individual fixed notions of fairness, the FTC should be considering the effect of any particular decision system (whether algorithmic or human) on inequality as a whole.

[3] Rediet Abebe, Jon Kleinberg, & S. Matthew Weinberg, "Subsidy Allocations in the Presence of Income Shocks," *Proceedings of the AAAI Conference on Artificial Intelligence* (2020): 34(05), 7032-7039, https://doi.org/10.1609/aaai.v34i05.6188.

B. To what extent, if at all, should new rules require companies to take specific steps to prevent algorithmic errors? If so, which steps? To what extent, if at all, should the Commission require firms to evaluate and certify that their reliance on automated decision-making meets clear standards concerning accuracy, validity, reliability, or error? If so, how? Who should set those standards, the FTC or a third-party entity? Or should new rules require businesses to evaluate and certify that the accuracy, validity, or reliability of their commercial surveillance practices are in accordance with their own published business policies?

The underlying assumption in this question is that individual companies have the ability and responsibility to correct for algorithmic errors, and that if every company ensured their own actions minimized and prevented algorithmic errors, overall welfare would be maximized. However, recent research suggests that in some contexts, there are factors affecting error that are outside any individual company's control.

This outcome is a potential consequence of algorithmic monoculture, which is when multiple decision-makers (or firms) deploy the same systems, or systems that share components such as datasets and models.[4] For instance, imagine multiple firms using the same algorithmic model to screen resumes of job candidates. This scenario is close to the real-world context, more than 700 companies including over 30 percent of Fortune 100 companies rely on a single vendor's tools for resume screening. What recent research suggests is that even if the algorithmic screening tool is more accurate than human evaluators and less error-prone than other tools on the market, accuracy may become worse when multiple firms use the same ones.[5]

This counterintuitive result is somewhat like the Braess paradox, an observation German mathematician Dietrich Braess made that illustrates how individual entities choosing their most rational option can lead to lower overall welfare when collective interaction is involved. The paradox states that when one adds a new road to a road network it can slow down overall traffic flow rather than speeding it up because individual drivers act selfishly. Drivers want to get from point A to point B in the fastest time, so if the new route is the most efficient way to get to their destination, all drivers will choose to take it. Choosing the new route would be optimal if only one driver did it, but if they all do it, the route becomes suboptimal. Similarly, companies seeking to fill job vacancies want to choose the best performing hiring tools. But it could be the case that if every company chooses to use the same system, there are more errors overall even though the system is the most accurate one on the market.

---

[4] Jon Kleinberg and Manish Raghavan, "Algorithmic monoculture and social welfare," *PNAS* (2021), https://doi.org/10.1073/pnas.2018340118.
[5] Ibid.

The reason this can happen in the hiring context is based on the probabilistic properties of rankings. Rather than diving into the math of it, instead consider that for each firm there is some "true ordering" of best to worst candidates for a job role and when using a system (or humans), a firm is using a ranking that best approximates that true ordering. When two firms use a single system to screen candidates, they rely on a common ranking that is a single approximation. Research suggests that even if a single approximation is more accurate in isolation, it can create more errors overall if multiple entities use it.[6] It is better, they say, for multiple entities to use different approximations, even if those approximations are less accurate. The key takeaway is that in the hiring context, independence can be more important than accuracy for reducing errors.

Importantly, this may not be the case for all settings. Algorithmic monoculture could be desirable in some settings as the authors themselves posit. It may be the case that in other high-risk areas, multiple decision-makers using a single centralized algorithmic system may reduce errors. In education, for instance, economists have found outcomes have improved as algorithms for school assignment have become more centralized.[7] Perhaps in healthcare, the allocation of scarce resource by different hospitals would be best done if they all used the same algorithmic systems. Perhaps not. We do not know because it has not been studied yet.

Again, the FTC should recognize that minimizing error rates is not the only factor—or even the most important factor—for encouraging fair outcomes in every context. Before rushing to regulate and potentially causing negative unintended consequences, the FTC should investigate how different factors affect fair outcomes in different contexts.

C. To what extent, if at all, do consumers benefit from automated decision-making systems? Who is most likely to benefit? Who is most likely to be harmed or disadvantaged? To what extent do such practices violate Section 5 of the FTC Act?

Algorithms offer several benefits to consumers. While much of the reporting on algorithmic decision-making systems has focused on these systems' potential for bias, there are many ways these systems can actually help reduce inequality and improve individuals' access to opportunity, especially for individuals for whom opportunities have been historically limited. For instance, using algorithms can improve equity in decisions about allocating scarce resources in healthcare settings and expanding access to educational opportunities. Algorithms can also replace biased human

[6] Ibid.

[7] Atila Abdulkadiroglu Nikhil Agarwal & Parag A. Pathak, "The Welfare Effects of Coordinated Assignment: Evidence from the NYC HS Match," *National Bureau of Economic Research* (2015), https://www.nber.org/papers/w21046.

decision-making for tasks like home appraisals.[8] However, even when algorithms can do good by making existing processes more efficient and equitable for consumers, public backlash and opaque implementations can erode the trust needed for them to achieve impact.

Take the example of the Boston public school system, which in 2018 proposed using an algorithmic system to improve school busing in ways that would cut costs by millions of dollars a year, help the environment, and better serve students, teachers, and parents.[9] The district had two aims, the first of which was to cut transportation costs. More than 10 percent of the public school system's budget goes toward busing children to and from school—the district's annual cost per student is the second highest in the United States.[10] The district's second goal was to to reconfigure school start times so that high school students could get more sleep, as early school starts for teenagers has been linked to serious health issues such as decreased cognitive ability, increased obesity, depression, and increased traffic accidents. Indeed, the American Academy of Pediatrics recommends that teenagers not start their school day before 8:30 AM, but only 17 percent of U.S. high schools comply according to a 2015 report from the Center of Disease Control and Prevention (CDC).[11]

Boston public school officials engaged researchers from the Massachusetts Institute of Technology (MIT) to build an algorithm to achieve its twin goals, which they did. The *Boston Globe* called their solution a "marvel."[12] The algorithm helped the district optimize bus routes, cutting 50 of the 650 school buses used, $5 million off the budget, and 20,000 pounds of carbon emissions each day while also optimizing bell times. Importantly, the algorithm's solution for bell times redressed equity. In the past, the district manually staggered start and end times, but its approach predominantly provided wealthier and whiter schools with later start times while schools with poorer and minority students disproportionately shouldered earlier times. In contrast, the algorithm's solution distributed advantageous start times equally across major racial groups, while significantly improving them for students in all of those groups.[13]

---

[8] Debra Kamin, "Home Appraised With a Black Owner: $472,000. With a White Owner: $750,000," *The New York Times,* August 25, 2022, https://www.nytimes.com/2022/08/18/realestate/housing-discrimination-maryland.html.
[9] David Scharfenberg, "Computers Can Solve Your Problem. You May Not Like The Answer," *The Boston Globe*, September 21, 2018, https://apps.bostonglobe.com/ideas/graphics/2018/09/equity-machine/.
[10] Ellen P. Goodman, "The Challenge of Equitable Algorithmic Change," *The Regulatory Review* (2019), https://www.theregreview.org/wp-content/uploads/2019/02/Goodman-The-Challenge-of-Equitable-Algorithmic-Change.pdf.
[11] Centers for Disease Control and Prevention, "Most US middle and high schools start the school day too early," news release, August 6, 2015, https://www.cdc.gov/media/releases/2015/p0806-school-sleep.html
[12] David Scharfenberg, "Computers Can Solve Your Problem. You May Not Like The Answer."
[13] Ellen P. Goodman, "The Challenge of Equitable Algorithmic Change."

Despite everything the algorithm offered, the district had to scrap the algorithm due to the swift and strong public pushback. Rutgers law professor Ellen Goodman describes how disgruntled parents carried signs at a school committee meeting that read “families over algorithms,” and “students are not widgets” in her 2019 paper *The Challenge of Equitable Algorithmic Change*.[14] But it is not clear the algorithm was really the problem. Indeed, Goodman describes the pushback as a case of “algorithmic scapegoating,” which Cornell researchers explain is where the algorithm “stood in for substantive issues around equity and disruptive change that were really at stake (though potentially more contentious to discuss) and might well have been at stake even without an algorithm in the picture. The tragedy of the case is that the algorithm could have provided the flexibility to involve the public in choosing among multiple trade-offs. If implemented, it might have created a more equitable system than what existed originally.”[15]

The takeaway for the Commission from this episode is twofold: One, algorithmic systems can reduce inequality from human decision-making when they are designed well. Two, these beneficial solutions “gain legitimacy not through their mathematical exactitude but through community engagement throughout the process,” as Goodman puts it.[16]

D. If new rules restrict certain automated decision-making practices, which alternatives, if any, would take their place? Would these alternative techniques be less prone to error than the automated decision-making they replace?

Focusing solely on restricting automated decision-making systems as a solution to complex, messy, and long-standing social problems is what American philosopher Charles West Churchman called “taming the growl,” in 1967.[17] These sorts of solutions he notes, consist “of ‘carving off’ a piece of a problem and finding a rational and feasible solution to this piece….the taming of the growl may deceive the innocent into believing that the wicked problem is completely tamed.”[18]

Reframing systemic problems into algorithmic ones can divert attention away from more insidious causes of these problems. Indeed, many use AI as a convenient target for moral indignation about systemic social problems but in a way that does not advance momentum toward change. For instance, there are widespread calls to ban the use of AI-enabled risk assessment tools that decide if an accused person should be allowed bail by predicting the likelihood they will miss a future court

---

[14] Ibid.

[15] Karen Levy, Kyla E. Chasalow, & Sarah Riley, “Algorithms and Decision-Making in the Public Sector,” A*nnual Review of Law and Social Science* (2021): Volume 17, https://www.researchgate.net/publication/352208766_Algorithms_and_Decision-Making_in_the_Public_Sector.

[16] Ellen P. Goodman, “The Challenge of Equitable Algorithmic Change.”

[17] C. West Churchman, “Wicked Problems,” *Management Science* (1967): Vol. 14, No. 4, https://punkrockor.files.wordpress.com/2014/10/wicked-problems-churchman-1967.pdf.

[18] Ibid.

appointment, but the underlying social problem, which is that many people cannot afford bail so they must remain in jail for weeks or months while awaiting trial, is not one that AI created nor is it one that transparency into algorithms and data can solve. As Cornell researchers put it, "the part of a problem that can be addressed algorithmically may steal political oxygen from nonalgorithmic reforms or palliate concerns that "something must be done" but without addressing root causes."[19] Not only does restricting automated systems in cases like the one described in the criminal justice context fail to address the root causes behind unfairness, it narrows the scope of what is on the table for politicians to change and changes the values used to evaluate those options.[20]

[19] Karen Levy et al., "Algorithms and Decision-Making in the Public Sector."
[20] Ibid.