



August 20, 2018

Donald Clark  
Federal Trade Commission  
Office of the Secretary  
600 Pennsylvania Avenue NW  
Suite CC-5610 (Annex C)  
Washington, DC 20580

RE: Competition and Consumer Protection in the 21st Century Hearings, Project Number P181201

Dear Secretary Clark,

The Center for Data Innovation is pleased to submit these comments in response to the request for comment from the Federal Trade Commission (FTC) on whether broad-based changes in the economy, evolving business practices, new technologies, or international developments might require adjustments to competition and consumer protection enforcement law, enforcement priorities, and policy.<sup>1</sup>

The Center for Data Innovation is the leading think tank studying the intersection of data, technology, and public policy. With staff in Washington, D.C., 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 predictive analytics, open data, cloud computing, and the Internet of Things. The Center is a non-profit, non-partisan research institute affiliated with the Information Technology and Innovation Foundation.

There is an inherent tension between regulators and innovators. Regulators seek to place limits on businesses to avoid risky scenarios that might result in harmful outcomes. In contrast, innovators take risks in their pursuit of new ideas. Unfortunately, as the economy has become more innovation-based, there has been considerably less focus on how regulatory agencies can both protect consumers from harm and avoid undermining incentives for innovation. We offer these comments on how the FTC can better find this balance.

---

<sup>1</sup> "Hearings on Competition and Consumer Protection in the 21<sup>st</sup> Century," Federal Trade Commission, n.d. <https://www.ftc.gov/policy/hearings-competition-consumer-protection>.



## THE CONSUMER WELFARE IMPLICATIONS ASSOCIATED WITH THE USE OF ALGORITHMIC DECISION TOOLS, ARTIFICIAL INTELLIGENCE, AND PREDICTIVE ANALYTICS

Some people are concerned that algorithmic decision-making will result in racial bias, such as financial institutions denying loans on the basis of race. However, in many cases, because flawed algorithms hurt the company using them, businesses have strong incentives to not use biased algorithms and regulators are unlikely to need to intervene. For example, banks making loans would be motivated to ensure their algorithms are not biased because, by definition, errors such as granting a loan to someone who should not receive one, or not granting a loan to someone who is qualified, costs banks money. In addition, even if some companies do not have a financial incentive to avoid biased algorithms, existing laws that prohibit such discrimination, such as the Fair Credit Reporting Act and the Equal Credit Opportunity Act, still apply.

Another argument regarding the inadequacy of privacy laws to protect consumer welfare is that the collection of large amounts of data allows companies to discriminate against consumers, including practicing price discrimination, charging different consumers different prices depending upon the likelihood that they will buy a product.<sup>2</sup>

Indeed, there is some evidence that companies are getting quite good at doing this.<sup>3</sup> This is often combined with the worry that disadvantaged groups will end up paying higher prices. But there are two reasons why price discrimination might not be a bad thing. First, to the extent that a platform has market power and can only set one price, its incentive is to raise prices on everyone and decrease supply. This allows the company to capture more value from the product and lowers the total benefit to society. If the company can charge different prices to different users, this social loss is reduced. Some consumers might still pay higher prices, but buyers will not purchase a product unless it makes them better off. Second, the ability to charge different prices is not limited to raising prices. Companies also have an incentive to lower prices for consumers who are reluctant to purchase the good.<sup>4</sup> This effect might actually be progressive. The company will charge a higher price to those

---

<sup>2</sup> Nathan Newman, "The Costs of Lost Privacy: Consumer Harm and Rising Economic Inequality in the Age of Google," *William Mitchell Law Review* 40, no. 2 (2014), <http://open.mitchellhamline.edu/cgi/viewcontent.cgi?article=1568&context=wmlr>.

<sup>3</sup> Burton G. Malkiel, "The Invisible Digital Hand," *The Wall Street Journal*, updated November 28, 2016, <http://www.wsj.com/articles/the-invisible-digital-hand-1479168252>

<sup>4</sup> Manne and Sperry, "The Problems and Perils of Bootstrapping Privacy and Data Into an Antitrust Framework," 7. "It is inconsistent with basic economic logic to suggest that a business relying on metrics would want to serve only those who can pay more by charging them a lower price, while charging those who cannot afford it a larger one."



users whose demand is inelastic. To the extent that lower-income consumers are more price responsive, they will benefit from price discrimination.<sup>5</sup>

More broadly, the FTC should recognize that consumers as a whole are going to benefit from greater use of algorithms, particularly artificial intelligence (AI). Though there are concerns about the potential harms that could arise from the use of AI, such as AI exacerbating unconscious human bias, the proposals that have gained popularity among consumer advocates to address these harms would be at best largely ineffective and at worst cause more harm than good. The two most popular ideas—requiring companies to disclose the source code to their algorithms and explain how they make decisions—would cause more harm than good by regulating the business models and the inner workings of the algorithms of companies using AI, rather than holding these companies accountable for outcomes.

The first idea—“algorithmic transparency”—would require companies to disclose the source code and data used in their AI systems. Beyond its simplicity, this idea lacks any real merits as a wide-scale solution. Many AI systems are too complex to fully understand by looking at source code alone. Some AI systems rely on millions of data points and thousands of lines of code, and decision models can change over time as they encounter new data. It is unrealistic to expect even the most motivated, resource-flush regulators or concerned citizens to be able to spot all potential malfeasance when that system’s developers may be unable to do so either.<sup>6</sup>

Additionally, not all companies have an open-source business model. Requiring them to disclose their source code reduces their incentive to invest in developing new algorithms, because it invites competitors to copy them. Bad actors in China, which is fiercely competing with the United States for AI dominance but routinely flouts intellectual property rights, would likely use transparency requirements to steal source code.<sup>7</sup>

The other idea—“algorithmic explainability”—would require companies to explain to consumers how their algorithms make decisions. The problem with this proposal is that there is often an inescapable trade-off between explainability and accuracy in AI systems. An algorithm’s accuracy typically scales with its complexity, so the more complex an algorithm is, the more difficult it is to explain. While this

---

<sup>5</sup> The White House, *Big Data and Differential Pricing* (Washington, DC: The White House, February 2015), 17, [https://obamawhitehouse.archives.gov/sites/default/files/whitehouse\\_files/docs/Big\\_Data\\_Report\\_None\\_mbargo\\_v2.pdf](https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/docs/Big_Data_Report_None_mbargo_v2.pdf).

<sup>6</sup> Will Knight, “The Dark Secret at the Heart of AI,” *MIT Technology Review*, April 11, 2017, <https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/>.

<sup>7</sup> Joe Uchill, “China Broke Hacking Pact Before New Tariff Fight,” *Axios*, April 10, 2018, <https://www.axios.com/china-broke-hacking-pact-before-new-tariff-tiff-d19f5604-f9ce-458a-a50a-2f906c8f12ab.html>.



could change in the future as research into explainable AI matures—DARPA devoted \$75 million in 2017 to this problem—for now, requirements for explainability would come at the cost of accuracy.<sup>8</sup> This is enormously dangerous. With autonomous vehicles, for example, is it more important to be able to explain an accident or avoid one? The cases where explanations are more important than accuracy are rare.

Fortunately, regulators have an alternative to these flawed approaches. Instead of pursuing heavy-handed regulations or ignoring these risks, they should adopt the tried-and-true approach of emphasizing light-touch regulation, with tailored rules for certain regulated sectors that fosters the growth of the algorithmic economy while minimizing potential harms. The challenge for regulators stems from the fact that innovation, by its very nature, involves risks and mistakes—the very things regulators inherently want to avoid. Yet, from a societal perspective, there is a significant difference between mistakes that harm consumers due to maleficence, negligence, willful neglect, or ineptitude on the part of the company, and those that harm consumers as a result of a company striving to innovate and benefit society. Likewise, there should be a distinction between a company’s actions that violate regulations and cause significant harm to consumers or competitors, and those that cause little or no harm. If regulators apply the same kind of blanket penalties regardless of intent or harm, the result will be less innovation.<sup>9</sup>

To achieve a balance, regulators should take a harms-based approach to protecting individuals, using a sliding scale of enforcement actions against companies that cause harm through their use of algorithms, with unintentional and harmless actions eliciting little or no penalty while intentional and harmful actions are punished more severely. Regulators should focus their oversight on operators, the parties responsible for deploying algorithms, rather than developers, because operators make the most important decisions about how their algorithms impact society.

This oversight should be built around algorithmic accountability—the principle that an algorithmic system should employ a variety of controls to ensure the operator can verify algorithms work in accordance with its intentions and identify and rectify harmful outcomes. When an algorithm causes harm, regulators should use the principle of algorithmic accountability to evaluate whether the operator can demonstrate that, in deploying the algorithm, the operator was not acting with intent to harm or with negligence, and to determine if an operator acted responsibly in its efforts to minimize harms from the use of its algorithm. This assessment should guide their determination of whether,

---

<sup>8</sup> Cliff Kuang, “Can A.I. Be Taught to Explain Itself?,” *New York Times Magazine*, November 21, 2018, <https://www.nytimes.com/2017/11/21/magazine/can-ai-be-taught-to-explain-itself.html>.

<sup>9</sup> Daniel Castro and Alan McQuinn, “How and When Regulators Should Intervene,” (Information Technology and Innovation Foundation, February 2016), <http://www2.itif.org/2015-how-when-regulators-intervene.pdf>.



and to what degree, the algorithm's operator should be sanctioned. Defining algorithmic accountability in this way also gives operators an incentive to protect consumers from harm and the flexibility to manage their regulatory risk exposure without hampering their ability to innovate.

This approach would effectively guard against algorithms producing harmful outcomes, without subjecting the public- and private-sector organizations that use the algorithms to overly burdensome regulations that limit the benefits algorithms can offer.



Sincerely,

Daniel Castro  
Director  
Center for Data Innovation  
[dcastro@datainnovation.org](mailto:dcastro@datainnovation.org)

Joshua New  
Senior Policy Analyst  
Center for Data Innovation  
[jnew@datainnovation.org](mailto:jnew@datainnovation.org)