

CPI's North America Column Presents:

Summary of FTC Hearing #7: Algorithms, AI, and Predictive Analytics

*By Norman Armstrong, Jr. and Albert Kim
(King & Spalding)¹*



CPI COMPETITION POLICY
INTERNATIONAL

Copyright ©2018

Competition Policy International, Inc. for more information visit CompetitionPolicyInternational.com

December 2018

Introduction

The FTC continued its Hearings Initiative with a two-day hearing at Howard University School of Law in Washington, D.C. on November 13-14. In keeping with the FTC's dual enforcement priorities, the panels and presentations focused on both the consumer protection and antitrust implications of algorithms, artificial intelligence ("AI") and predictive analytics.

The hearing topics were complicated, broad and implicated many legal, regulatory and ethical issues. As Professor Andrew Gavil from Howard University School of Law noted in his introductory remarks, the hearing's agenda was designed to take a more interdisciplinary approach than prior hearings.

The panelists, drawn from universities, corporations, non-profits, trade associations and law firms, discussed the substantial progress made in the fields of algorithmic decision-making and machine learning, particularly in the last twenty-five years. This progress has created and is continuing to create significant benefits and opportunities, but at the same time raises profound challenges in terms of data privacy and security, as well as concerning questions of fairness, accountability and inclusiveness.

Introduction to Algorithms, Artificial Intelligence and Predictive Analytics

Given the complexity and technical nature of the subject matter, framing presentations by several academics established a useful foundation for the panel discussions. Professor John Dickerson from the University of Maryland provided a brief history of AI, describing how AI evolved from hard-coded expert knowledge databases with search routines (e.g., Deep Blue) that were characterized by "brittleness" and lack of an ability to truly learn probabilistic models that demonstrated a greater ability to analyze data and react appropriately through statistical machine "learning" (e.g., autonomous vehicles and neural networks used in AlphaGo). The current phase of AI focuses on issues such as identifying biased data and ways to debias, cooperative and adversarial strategies, and techniques in reinforcement learning.

Professor Michael Kearns from the University of Pennsylvania highlighted several problems that may emerge from the increased use of AI, machine learning, and algorithmic decision-making. He first noted that algorithmic decision-making used, for example, in credit scoring or online advertising, generally is not the result of programming but rather the product of machine learning on datasets that then train a decision-making model. Professor Kearns explained that although the algorithms themselves may often be relatively simple, complexity emerges as the algorithms interact with the data to create or tune a model. Furthermore, as a programmer optimizes the model for, say, accuracy, the model may generate unintended results in other areas by producing, for example, discriminatory or unfair outputs.

Professor Michael Jordan of the University of California, Berkeley, emphasized that the term "Artificial Intelligence" is sometimes an unhelpful term to use for the field of machine learning because it invokes ideas of robotic autonomy or human-imitative intelligence. From Professor Jordan's perspective, "deep learning" and "AI" are still statistical data analyses that have so far seen success in more mundane activities as fraud detection, supply-chain management, social media and online recommendation systems. Researchers now are refining data analysis techniques and incorporating decision-making algorithms for use across many disciplines, actors, and markets using enormous data sets.

Uses and Implications of Algorithmic Decision-Making

As computing power has grown and data storage costs have decreased, companies are turning to algorithmic decision-making both to increase the efficiency of existing business processes and to find solutions to new problems that formerly could not be solved by machine learning and algorithms. The panelists, including a number of corporate representatives, described several applications for algorithmic decision-making, such as image recommendations in photo databases (Adobe), consumer credit eligibility (Experian), fraud detection (Visa) and autonomous medical diagnosis (IDx Technologies).

As algorithmic decision-making increasingly makes its way into commercial, government and administrative uses, researchers and policy makers have raised corresponding notes of caution. The use of algorithms in areas such as employment, housing, credit scoring, crime prevention and criminal sentencing can have a considerable social impact: both positive, if the algorithms reduce human bias, and negative, if the algorithms or models themselves lead to discriminatory results. To be sure, however, data privacy and security concerns continue to be top-of-mind for the public and other stakeholders. In recent years, commentators have also raised the specter of algorithmic collusion, where algorithmic models act in ways that may generate anticompetitive results.

Ethical Considerations in Algorithmic Decision-Making

Panelists throughout the two-day hearing emphasized the need to incorporate societal norms at every stage of the design, testing and implementation of algorithmic decision-making programs and models. These principles embody values and considerations that include, but are broader than, traditional consumer protection concerns such as privacy, consent, fraud and security.

Referencing Microsoft's CEO Satya Nadella's principles for AI design, Jennifer Vaughn, a senior researcher at Microsoft, provided a set of goals for implementation of algorithms and machine learning, which included fairness, reliability, inclusiveness, safety, privacy and security, transparency and accountability—principles echoed by a number of panelists. As Ms. Vaughn noted (and as echoed by business representatives such as Irene Liu, General Counsel of Checkr), algorithms need to embody these values from the design stage. In Professor Kearns' words, the algorithms must “endogenize” the desired social principles as policing violations through government enforcement and regulations may be insufficient.

Consumer Protection and the Existing Regulatory Framework

In considering the consumer protection implications of algorithmic decision-making, panelists expressed some general consensus that there did not appear to be a need to draft a host of new laws. Professor Fred Cate of Indiana University's Maurer School of Law noted that Section 5 of the FTC Act is so broad (and at times amorphous) that it can be used to address almost any consumer protection scenario. Irene Liu of Checkr agreed on the flexibility of Section 5 and also noted that technology companies employing algorithmic decision-making already must comply with a range of different regulations, including obligations related to the Fair Credit Reporting Act and FDA and SEC regulations. Marianela Lopez-Galdos of the Computer & Communications Industry Association recommended that both the US and the EU maintain technology-agnostic approaches and not craft new laws aimed at specific technology.

At the same time, the existing regulatory framework may contain gaps, some of which are potentially very significant.² Professor Ryan Calo of the University of Washington School of Law noted both that the FTC likely has adequate enforcement tools at its disposal but also that it needs to be assertive and ask hard questions of companies that are using algorithmic decision-making in a potentially harmful manner. Irene Liu and others recommended that the FTC take a leadership role in issuing policies or drafting guidelines to assist companies in navigating the rapidly changing world of algorithms and AI.

As did speakers in other sessions, the panelists in the consumer protection session struggled with how traditional consumer protection values such as transparency, notice/consent, fairness and privacy might be safeguarded in algorithmic decision-making. Given the use of big data combined with opaque (either due to complexity or due to trade secrets) models, the panelists did not come to a consensus on implementation. Mr. Gillula recommended that enforcers should mandate content-specific transparency requirements and that the lack of transparency was not so much due to the complexity of “black box” algorithms but rather the fear of exposing trade secrets. Professor Cate noted that as notice and consent are routinely ignored by many consumers, a better route in some circumstances may be to require institutions to maintain records and produce them if there is consumer harm.

Algorithmic Collusion

In his introductory remarks, Bruce Hoffman, the Bureau of Competition Director at the FTC, noted that there is a particularly keen interest among non-US competition authorities on the possibility of algorithmic collusion or other anticompetitive harms from artificial intelligence. This is a hot-button issue, and a number of agencies are conducting research and/or increasing resources to address potential problems, such as the formation of the Data Analytics Unit in the Australian Competition and Consumer Commission.

Mr. Hoffman presented four possible theories of anticompetitive harm: (i) express collusion by and among AI, (ii) tacit collusion by and among AI, (iii) acceleration or enabling of unilateral strategies to limit competition and (iv) other potential harms not yet contemplated (e.g., price discrimination-based harm). Subsequent panelists discussing algorithmic collusion elaborated on the themes set out by Mr. Hoffman. Although there was some disagreement on the likelihood of algorithmic collusion, the panelists all agreed that further research and refinements of enforcement tools and techniques were required.

Professor Maurice Stucke described circumstances in which current tools may be insufficient. First, Professor Stucke outlined four possible collusion scenarios: (i) humans deliberately using algorithms as an instrument for express collusion, (ii) a “hub and spoke” arrangement in which competitors use the same common algorithm, technology platform or outsource competitive decisions to the same vendor, (iii) tacit collusion in which competitors unilaterally decide to use price-optimization algorithms knowing that such use may lead to tacit algorithmic collusion and (iv) no evidence of anticompetitive intent in utilizing pricing algorithms but a tacit collusion outcome. According to Professor Stucke, these last two tacit collusion scenarios present challenges for the existing enforcement framework. How can agencies identify where algorithmic tacit collusion occurs, particularly when pricing is dynamic? Should companies employing such software integrate principles of ethics and legality into the models? To what extent should software developers face liability if companies use their products to collude tacitly?

Dr. Ai Deng of Bates White and Professor Kai-Uwe Kuhn of the University of East Anglia both emphasized that designing algorithms to tacitly collude is a significant technical challenge. Both contended that a collusive equilibrium is generally not very stable without express communication, particularly when there are three or more market participants. Dr. Deng also noted that recent research indicates that greater information flow tends to destabilize cartels and that algorithmic decision-making generally speeds up the flow and use of market information.

In the same vein, Rosa Abrantes-Metz of Global Economics Group argued that pricing algorithms empirically are generally associated with increased competition rather than tacit collusion. While acknowledging that algorithms may theoretically facilitate signaling, price monitoring and competitive responses to deviations in oligopolistic markets, pricing algorithms are often used in less concentrated markets and lead to greater competition. Dr. Abrantes-Metz pointed to the example of financial instrument trading, comparing the shift from over-the-counter markets where trading is very opaque to trading on exchanges where markets are very transparent. Although trading on exchanges often employs pricing algorithms, competition has increased and collusion has decreased, as evidenced by narrower bid-ask spreads. Collusion still does occur, but according to Dr. Abrantes-Metz, these episodes of benchmark or auction rigging are primarily due to deficient structures, not the use of algorithmic decision-making.

Professor Joseph Harrington of the University of Pennsylvania contended that the existing enforcement framework may not be sufficient to regulate collusive algorithmic decision-making. In order to balance the need to regulate algorithms while not stifling innovation, Professor Harrington noted the imperative (i) to understand the risks of pricing algorithms, (ii) to improve techniques to detect algorithmic collusion and (iii) to refine tools for merger enforcement to better combat tacit collusion. Professor Harrington suggested testing algorithms to determine their pricing rules and competitive goals and recommended establishing a per se rule against using algorithms to limit competition, with the burden on companies to test and monitor their models. For this approach, the challenge would be to distinguish anticompetitive algorithms that are designed to reward and punish from efficiency-enhancing algorithms that, for example, simply adjust pricing in response to changes in demand data.

Conclusion

The FTC's two-day hearing continues the ongoing conversation between enforcers, researchers and the business community on the opportunities and challenges of algorithmic decision-making in everyday life. These are just initial steps, and the panelists agreed that there is a clear need for additional thinking and in particular, empirical research. As the use of algorithmic decision-making is already widespread and will continue to grow in breadth and sophistication, consumer protection and competition challenges correspondingly will increase.

¹ Norman Armstrong, Jr., Partner and Albert Kim, Counsel, King & Spalding LLP. All views are the authors' own. The authors would like to thank associates Meaghan Griffith and Mary Longenbaker for their valuable contributions.

² Although not a consumer protection problem, Jeremy Gillula of the Electronic Frontier Foundation noted that vendors are starting to market AI-based criminal justice risk-assessment tools and that EFF and other civil liberties groups do not believe that use of such tools is appropriate.