SUBMISSION TO THE FTC HEARINGS ON COMPETITION AND CONSUMER PROTECTION IN THE 21ST CENTURY

"EMERGING COMPETITION, INNOVATION, AND MARKET STRUCTURE ISSUES AROUND ALGORITHMS, ARTIFICIAL INTELLIGENCE, AND PREDICTIVE ANALYTICS"

November 13-15, 2018

Nicolas Petit*

This submission comments on four competition policy questions linked to the deployment of algorithms and Artificial Intelligence ("AI") systems in markets. In light of the current state of technology development, my comments are necessarily speculative.

1. ***Are there antitrust or competition contexts in which differences between algorithms and AI are likely to matter?***

Yes. On the one hand, simple algorithms that follow structured "if-then" decision-making rules for pricing purposes represent few challenges for antitrust policy.[1] On the other hand, sophisticated AI systems like neural networks, deep learning, and reinforcement learning[2] that analyze sample data inputs and convert them into outputs generate bigger challenges, mostly linked to the so-called "*blackbox*" problem.[3] The blackbox problem means that neural networks, deep learning and reinforcement learning algorithms do not provide programmers and third parties – employees, managers, shareholders as well as regulatory agencies and courts – with information about the decision-making process leading to observed market results.[4] The blackbox problem has four prospective implications for antitrust policy. *First*, if the AI decisional outcome is anticompetitive, most hard questions relate to liability:[5] Is it

* Professor of law, University of Liege (Belgium), Research Professor, University of South Australia, Visiting Fellow, Stanford University Hoover Institution; petitn@stanford.edu and Nicolas.petit@uliege.be.

[1] For a definition, see OECD, *Algorithms and Collusion: Competition Policy in the Digital Age*, (2017), 8, available at http://www.oecd.org/daf/competition/Algorithms-and-colllusion-competition-policy-in-the-digital-age.pdf (hereafter OECD, *Algorithms and Collusion*) ("an algorithm is a sequence of rules that should be performed in an exact order to carry out a certain task"). By "Pricing purposes" I mean here pricing algorithms, but also customer scoring, data scraping and all other algorithms that assist firm pricing decisions.

[2] Deep learning is a "statistical technique" that links input to output, identifies patterns, and generates a classification given a more or less significant set of data.

[3] Gary Marcus, Deep Learning: A Critical Reappraisal, January 2, 2018, available at https://arxiv.org/ftp/arxiv/papers/1801/1801.00631.pdf (hereafter Marcus, Deep Learning) (noting that the resurgence in deep learning research dates back to 2012).

[4] Davide Castelvecchi, Can we open the blackbox of AI, Nature, October 5, 2016 (Deep learning and neural networks are said to be "*as opaque as the brain*").

[5] OECD, *Algorithms and Collusion*, 32 ("*by relying on deep learning, firms may be actually able to reach a collusive outcome without being aware of it, raising complex questions on whether any liability could ever be imposed on them should any infringement of the law be put in place by the deep learning algorithm*").

fair and efficient to impute liability to the firm, its managers and shareholders?;[6] What should be the impact, if any of AI outsourcing arrangements like leases, software as a service ("SaaS") or AI on the shelf?; Should we explore absolute strict liability regimes like in the field of product liability?[7] Or move towards regimes of joint liability with technology providers?

Second, if AI systems with high autonomy are ever deployed at scale across markets, effective antitrust enforcement may require to abandon decision rules that seek to infer whether conduct is (i) unilateral exclusion that is not competition on the merits (by reliance on no economic sense tests); or (ii) collusion that is not unilateral adaptation (by reliance on plus factors), and replace decisional rules of inference with pure level of harm doctrines or direct effects observations. At the end of the line, antitrust may have to rely less on rules that selectively proscribe bad conduct, than on no fault prohibitions of bad market outcomes (and perhaps regulation).[8] Recently, the Competition and Markets Authority ("CMA") in the UK referred to this idea noting that from an enforcement and regulatory perspective, it would be beneficial to understand whether and if a firm could know "*that its algorithm is implementing a collusive* <u>*outcome*</u> (emphasis added)".[9] This is distinct from the traditional approach which consists in asking if the firm knowingly partakes in a collusive *process* through communications and a meeting of the minds. Ultimately, we may embrace a pure economic approach to collusion and exclusion.[10]

Third, the blackbox problem creates a remedial difficulty. Computer scientist Gary Marcus talks of a "*debuggability*" problem.[11] Absent a procedure to diagnose the design or data defects that lead an AI system to generate anticompetitive outcomes, the remedial options for

---

[6] In most antitrust regimes, the firm, its managers and shareholders are in principle responsible for unlawful damage caused by their capital and labor, which arguable covers algorithms and AI systems.
[7] In EU law, contravening firms are held liable for penalty under a standard of fault or negligence.
[8] The situation is different for simple rule based pricing algorithms. Harrington and Parker explain that if the "rule determining price is written down in the algorithm", then effective antitrust enforcement could lead to adopt "a *per se* prohibition on certain pricing algorithms that support supracompetitive prices", based on an expert verification of the algorithm's code and operation. See Harrington Jr, Joseph E., Developing Competition Law for Collusion by Autonomous Price-Setting Agents (August 22, 2017). Available at SSRN: https://ssrn.com/abstract=3037818 or http://dx.doi.org/10.2139/ssrn.3037818
[9] CMA, *Pricing algorithms: Economic working paper on the use of algorithms to facilitate collusion and personalised pricing*, October 8, 2018, 52, available at https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/746353/Algorithms_econ_report.pdf (hereafter CMA, *Pricing Algorithms*).
[10] See, for this approach, Louis Kaplow, *Competition Policy and Price Fixing*, (1st edn, Princeton University Press 2013) (hereafter Kaplow, *Competition Policy and Price Fixing*).
[11] See Gary Marcus, Deep Learning.

both the firm and the antitrust structure are limited. Keeping decision making over prices to humans seems a first best.

Fourth, antitrust compliance by design is more difficult to ensure for autonomous AI agents than for rule-based algorithms. All readers of Isaac Asimov's three laws of robotics know this. Take a goal-based AI agent whose function is to "optimize prices" or to "maximize profits". An autonomous AI agent may read this code as an instruction to become the *price setter* in a market. In addition, techniques like "backpropagation" could lead the AI agent to autonomously adjust its goals, and consider that the best way to fulfill the set function is to eliminate rivals.[12] Now, assume that an additional compliance goal is written down in the AI agent's code: "in compliance with antitrust law" (an idea promoted by EU Commissioner Vestager and several antitrust scholars).[13] Compliance means conforming to a legal command. The AI agent may thus find it profit maximizing to design a strategy that exploits loopholes in antitrust law (for example, move from explicit collusion towards tacit collusion or vertically integrate to avoid application of rules on distribution agreements) or to change the legal command altogether by even corrupting lawmakers or government officials. Last, as in Asimov's novels, the AI system may recalibrate the relative value of each of the two goals in light of available data. Assume datasets suggest to the AI agent that the firm manager is not a risk averse individual but that shareholders are. The AI system may infer a managerial preference for short term gain versus a shareholder preference for compliance. Given that the AI agent may also have to "protect itself" by design – this is Asimov's third law and a requirement found in most AI ethics codes today – and that managers have life and death power over it, it may decide to maximize manager preferences, instead of shareholders one in breach of fiduciary duty.

Now, as much as science fiction makes pleasant readings, its influence on public policy should remain moderate. Today's AI research is still far from full autonomy. Even the most advanced AI systems lack autonomous goal setting. Most deep learning is supervised. And

[12] An alternative the AI instructing itself to always markup. This could be cured by introducing in the code an instruction to always markdown. But then, this instruction could lead to inefficiencies, because rising prices play a key role in adjusting supply to demand in the short term.

[13] See Margrethe Vestager, 'Algorithms and competition", Speech at the Bundeskartellamt 18th Conference on Competition, Berlin, 16 March 2017, available at https://ec.europa.eu/commission/commissioners/2014-2019/vestager/announcements/bundeskartellamt-18th-conference-competition-berlin-16-march-2017_en.
Vezzoso, Simonetta, Competition by Design (November 28, 2017). Prepared for Presentation at 12th ASCOLA Conference, Stockholm University, 15-17 June 2017. Available at
SSRN: https://ssrn.com/abstract=2986440 or http://dx.doi.org/10.2139/ssrn.2986440

back propagation – when AI systems recursively self-improve – is dependent on labeled datasets.[14] In the field, some experts even conjecture that AI may be hitting a wall.

Against this background, it is unsurprising at best that advanced AI systems have not yet been deployed in markets. Existing evidence suggests that sellers and buyers do not use deep learning at scale in markets to exploit consumers.[15] And even if deep learning was used, its potential to harm consumers is uncertain, and certainly not as existentially harmful as AI applications developed in other domains like lethal autonomous weapons.

The standard fact finding and evidence based approach of antitrust policy is the right one in so far as algorithms and AI systems are concerned. Antitrust fact finders should be vigilant to evidence of anticompetitive algorithmic design and/or data in discovery contexts.[16] But there is no rush for antitrust to address extravagant claims of widespread deep learning assisted tacit collusion or first degree price discrimination.

2. *What are the relevant competitive assets/potential barriers to entry in terms of algorithms and AI?*

In antitrust economics, a disputed empirical topic is whether there are increasing returns to scale in data. It is probably more right than wrong that scale in data feeds positive demand side network externalities, meaning increasing *economic* returns to scale. But we often fail to realize that scale in data displays diminishing *technological* returns, especially when data is used in support of algorithms and AI systems.[17] Firms need to incrementally invest in both fixed and variable assets to collect, curate, label and optimize data for use in algorithms and AI systems.

The rate of diminishing returns to data probably varies by application.[18] Some domains are more labor intensive than others, like content curation on social networks. Some classes of AI systems also exhibit more diminishing returns than others. Recently, AI legend Geoff Hinton, talked of a risk of "*exponential inefficiencies*" in relation to convolutional deep learning

---

[14] Even though in deep learning there is an absence of human supervised feature extraction. *See* Marcus, Deep Learning.

[15] CMA, *Pricing Algorithms*. See also European Commission, Consumer Market Study on Online Market Segmentation Through Personalised Pricing/Offers in the European Union, (June 2018), available at https://ec.europa.eu/info/sites/info/files/aid_development_cooperation_fundamental_rights/aid_and_development_by_topic/documents/synthesis_report_online_personalisation_study_final_0.pdf.

[16] CMA, *Pricing Algorithms*, 52.

[17] Sokol, D., & Comerford, R. (2016). Does antitrust have a role to play in regulating big data?; Lambrecht, Anja and Tucker, Catherine E., Can Big Data Protect a Firm from Competition? (December 18, 2015). Available at SSRN: https://ssrn.com/abstract=2705530 or http://dx.doi.org/10.2139/ssrn.2705530

[18] Susan Athey, Machine learning, government, and policy analysis, available at https://techpolicyinstitute.org/wp-content/uploads/2016/09/AtheyAI.pdf.

systems, noting that their reliance on large numbers of labelled examples may lead to their "*demise*".[19]

Regulatory initiatives could also erect barriers to entry in AI-intensive industries. Across the globe, countries are increasingly concerned by the ethical risks posed by the deployment of algorithms and AI systems in society. Regulatory compliance requirements and fragmentation in legal environments may raise the costs of doing business, limit entry and disproportionately increase the returns to compliance of large firms relative to smaller organisations.

3. *How will the upstream market to provide AI systems develop in the future, and is there anything that regulatory agencies should be on the lookout for from a competition perspective?*

Two technological properties of AI systems seem relevant when one thinks of market structure. First, AI systems constitute general purpose technologies ("GPT"). Like electricity or the steam engine, AI systems have applications in multiple sectors and display continuous innovational complementarities.[20] Second, AI systems are brittle, narrow and biased.[21] Scaling, transfers to other domains and product improvements encounter significant technological frontiers. These properties may help contextualize the ongoing merger wave observed in the computer software industry, and in particular Microsoft and IBM's respective acquisitions of GitHub and RedHat.

Let us start with GPTs. Economists usually distinguish two phases of GPT growth.[22] In the first phase, growth is driven by the adoption of the technology. In the second phase, growth is driven by investments in complementary technologies. Recall now that sophisticated AI systems are relatively recent (the latest breakthrough in deep learning date back to 2010-2012). Against this background, mergers between large tech firms and open source software providers may be looked at as a first phase strategy of technology adoption. Enlisting open

---

[19] Sabour, S., Frosst, N., & Hinton, G. E. (2017). Dynamic Routing Between Capsules. *arXiv*, cited by Marcus, Deep Learning.

[20] Timothy Bresnahan and Manuel Trajtenberg, *General Purpose Technologies: Engines of Growth*, 65 JOURNAL OF ECONOMETRICS, 83-108 (1995).

[21] Jason Ponti, Greedy, Brittle, Opaque, and shallow: the downsides to deep learning, (*Wired*, February 2, 2018), available at https://www.wired.com/story/greedy-brittle-opaque-and-shallow-the-downsides-to-deep-learning/?utm_campaign=News&utm_medium=Community&utm_source=DataCamp.com.

[22] Elhanan Helpman and Manuel Trajtenberg. A time to sow and a time to reap: Growth based on general purpose technologies. No. w4854. National Bureau of Economic Research, 1994; More generally, on the importance of time gaps and sequences in GPTs, see Susanto Basu and John Fernald, Information and communications technology as a general-purpose technology: evidence from US industry data, German Economic Review 8.2 (2007): 146-173.

source developers by M&A is a radically effective way to diffuse AI technologies within large open source ecosystems.[23]

Besides, the acqui-hiring of sizeable teams of open source developers makes sense for large computer software firms with significant sunk investments in AI systems. As scale increases and AI applications move from narrow to transversal, deep learning AI systems require exponential amount of manpower, engineering skills and troubleshooting talent.

4. What effect will the increasing reliance on algorithms and AI have on market structure across the U.S. economy?

Research on the impact of algorithms and AIs on market structure disproportionately focuses on the supply side, and much less on the demand side.[24] Buyers are often assumed to be price takers.

But algorithms and AI can increase buyer power. In 2017, the OECD noted that algorithms "*can recognise forms of co-ordination between suppliers (i.e. potentially identifying instances of collusive pricing) and diversify purchasing proportions to strengthen incentives for entry (i.e. help sponsoring new entrants). Another possible use could be the aggregation of demand by pooling consumers and creating buying platforms …*".[25] And for a long time, shopbots – also known as price comparison applications – have been deployed by strategic buyers in industries like air transport or hotels in order to profit from spatial and temporal price movements.[26]

Besides, without even trying to wield buyer power in the traditional economic sense, buyers on the demand side can deploy technological defenses and countermeasures to subvert sellers' algorithms and AI systems. Adversarial examples are a case in point. Deep learning algorithms are vulnerable to small perturbations of the input, that are imperceptible by

[23] Moreover, the addition of developers also serves the objective of early identification of novel applications use cases and sectors.

[24] Jochen Gönsch, Robert Klein, Michael Neugebauer, Claudius Steinhardt, Dynamic Pricing with Strategic Customers, (*SSRN*, September 21, 2012), available at http://ssrn.com/abstract=1979477 ("*most scientific publications on dynamic pricing, the modelling of customer behaviour traditionally only plays a minor role*".).

[25] OECD, *Algorithms and Collusion*.

[26] David Sarne, Competitive Shopbots-Mediated Markets [October 2012], ACM Transactions on Embedded Computing Systems.

humans but can cause misclassification.[27] Changing a pixel on a panda picture may cause an algorithm to identify a gibbon.[28]

While only little research and evidence exists on the possible deployment of adversarial examples in economic exchanges, a middleware market seems emerging that offers "*bot management*" and "*bot mitigation*" technology to hinder seller automated data scraping systems. Several small to mid-sized companies like Akamai, Luminati, Shieldsquare, Distil Networks, Infisecure Technologies and PerimeterX seem active in this space, alongside larger tech platforms like Amazon Web Services ("AWS").

Amongst others technology products, users can avail themselves of a variety of defense solutions like:

- CAPTCHA if actions are performed too fast on their website;
- Software that monitors where visitors click on buttons or the screen size;
- Software that seeks information on whether the query was issued from a mobile phone by a human by retrieving accelerometer and gyroscope information;
- Software that serves fake, randomly generated data, content or prices.

The tech press reports that there is intense competition in the middleware bot management market.[29] Antitrust agencies should make sure that this remains the case. And perhaps, they should raise awareness of customers to the opportunities generated by these cost-saving technologies.

---

[27] Margarita Osadchy, Julio Hernandez-Castro, Stuart Gibson, Orr Dunkelman, and Daniel Pérez-Cabo, No Bot Expects the DeepCAPTCHA! Introducing Immutable Adversarial Examples, With Applications to CAPTCHA Generation, IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY, VOL. 12, NO. 11, NOVEMBER 2017

[28] Ian Goodfellow, N. Papernot, S. Huang, Y. Duan, P. Abbeel, J. Clark, Attacking Machine Learning with Adversarial Examples, (*Open AI*, February 24, 2017), available at https://blog.openai.com/adversarial-example-research/.

[29] Klint Finley, 'Scraper' Bots and the secret internet arms race, (*Wired*, July 23, 2018), available at https://www.wired.com/story/scraper-bots-and-the-secret-internet-arms-race/. Wired reported in July that web security firm Akamai Technologies develops defenses for firms that wish to avoid scraping bots. Meanwhile, another firm called Luminati develops technology to mask bots (one of its application makes the bots appear to come from smartphones). Luminati can use smartphones without direct user consent, by indirect deals with app makers.