Editorial

Antitrust and Artificial Intelligence: A Research Agenda

Nicolas Petit[*]

Since last year, the hype around 'Artificial Intelligence' ('AI') has reached the antitrust community. A common thread to the emerging literature on Antitrust and AI ('AAI') is to describe the increasing use of algorithms on markets as a game changer.[1] In their page turner book *Virtual Competition*, Professors Ariel Ezrachi and Maurice Stucke prophesize the '*end of competition as we know it*', and advocate heightened government intervention.

The AAI literature makes three claims. *First*, algorithms will widen instances in which known forms of anticompetitive conduct occurs. The AAI scholarship conjectures that express and tacit collusion as well as almost perfect behavioural discrimination will be more common. *Second*, algorithmic markets will display new forms of anticompetitive conduct in non-price dimensions like data capture, extraction, and co-opetition (between '*super-platforms*' and applications developers) which challenge established antitrust doctrine. *Third*, deception is a design feature of algorithmic markets. Behind the '*façade*' of competition, consumers are nudged in exploitative transactions. In a telling metaphor, Ezrachi and Stucke compare us to the main character in the movie the *Truman show*.

AAI literature is the closest ever our field came to science-fiction. Like science-fiction, it is a lot of fun. And like science-fiction, its scenarios unearth fascinating research hypotheses for antitrust experts. Five areas deserve attention. *First*, the AAI literature essentially focuses on the *facilitating* role of algorithms on anticompetitive conduct. Those findings must now be complemented by a symmetrical investigation of the *destabilising* effect of algorithms on harm to competition. Take the tacit collusion scenarios. The AAI literature makes a convincing case that algorithms are a *plus factor* which renders tacit collusion more stable, durable, and versatile by facilitating detection and retaliation at lower levels of market concentration. Yet, the reciprocal hypothesis that oligopolists in high-frequency interaction may have stronger incentives to cheat is given shorter shrift. When transactions are customised—a feature of the digitalised economy—each bargain with a customer can be seen as a finite, one shot game which is incompatible with tacit collusion. Similarly, when personalised and dynamic pricing are combined, the range of price points over which oligopolists must coordinate is virtually infinite, because it is a function of the number of individual customers times the number of time units spent on digital markets. Last, when non-price competition on privacy and behavioural discrimination are introduced, there is more '*noise*' in the market, and detecting any punishing deviations may be significantly more costly.

Second, we still lack a proper understanding of countervailing strategies. To date, much of the AAI's literature focuses on B2C markets where sellers use algorithms to exploit boundedly rational consumers. But in B2B markets, sophisticated buyers may have the ability and incentives to make or buy countermeasures that undermine the operation of sellers' algorithms. Personally, I doubt that many car manufacturers will stay prey to input sellers' algorithms, even in the extreme scenario where the later are super-platforms like Google, Facebook, or Apple. The dieselgate is a bitter reminder of the automotive industry's technological capabilities. And the fast development of the cybersecurity industry suggests a non-trivial chance that we will witness the emergence of a market for countermeasure systems (data perturbation, masking, and randomisation sofware, for example).

Third, the AAI literature generates predictions on the basis of fairly strict assumptions, and more work is needed to understand if they are robust to circumstances. Tacit collusion is conceivably easier if one postulates that rival oligopolists use similar or homogeneous algorithms. Yet, as soon as the analysis is conducted under the assumption of algorithmic heterogeneity, a larger range of competitive outcomes becomes plausible. It is indeed uncontroversial that tacit collusion is less easy when oligopolists display asymmetries in costs, investments, structure, or market

---

[*] University of Liège, Liège Competition and Innovation Institute (LCII), Liège, Belgium; University of South Australia (UniSA), Adelaide, Australia

[1] A Ezrachi and ME Stucke. *Virtual Competition: The Promise and Perils of the Algorithm-Driven Economy* (2016). SK Mehra, 'Antitrust and the Robo-Seller: Competition in the Time of Algorithms' (2016) 100 Minn. L. Rev., 1323–75. G Surblyte, 'Data-Driven Economy and Artificial Intelligence: Emerging Competition Law Issues' (2017) 67(3) WuW, 120–7. R. Calo, 'Digital Market Manipulation' (2014), George Washington Law Review 82, 995–1034.

share, and we should attempt to understand the effects of algorithmic differentiation at preference specification (design) or construction (learning) stages. To put the point differently, because in the real world, algorithms are neither commodities—scientific progress in algorithm design is relentless—nor public goods—the ongoing EU search case against *Google* brings a powerful reminder—algorithmic asymmetry should be baseline hypothesis for antitrust policy. The same applies to the assumption that profit-maximising algorithms—unlike humans—do not fear detection and possible penalties. It is true that unlike humans, a computer cannot be incarcerated. But if we follow the assumption that an algorithm does not register losses, then there is no basis to consider that it can register the profits of anticompetitive activity. The point here is that a utilitarian algorithm is an agent that necessarily operates on behalf of someone else. And, therefore, a reasonable assumption is that a profit-maximising pricing algorithm will specify a fiduciary duty towards its vicarious governors, which will integrate constraints like antitrust compliance.

Fourth, an area where additional research is needed is *evidence*. Much of the early AAI papers report perverse instantiations of algorithmic exploitation. An often-heard example is about the book '*The Making of a Fly*', which sold once for $23 million on Amazon's platform. But one should not forget that customers happen to make bad deals on very competitive markets. And that circumstancial cases of consumer harm caused by pricing algorithms do not tell us much on whether market power is being exerted to an extent and intensity that deserves antitrust remediation. With this, it is unclear if the facts advanced in AAI literature denote '*a brief perturbation in competitive conditions*' as Judge Posner once wrote, or whether they constitute emerging proof of a market failure worthy of agency interest.

Last, but not least, the main hard question raised by AAI literature relates to the goals of antitrust. There remains a lot of ideological resistance in both the US and the EU to the idea that antitrust laws should address wealth transfers between sellers and buyers, and this could elevate an unsurmountable obstacle to the application of the competition rules to consumer exploitation through almost perfect behavioural discrimination and personal data extraction. Aware of that distributional controversy, Ezrachi and Stucke advance an additional—and profound—idea: virtual competition increases the '*deadweight loss by increasing distrust*'. Presenting the social costs of algorithmic exploitation in trust terms is appealing. Trust in strangers is a feature of modern economies. Third-party enforcement mechanisms like the courts system, regulation, and antitrust laws create trust and promote exchange amongst aliens. But is '*trust*' the core business of antitrust law? Taxation, war, or corruption all reduce trust, and inflict a deadweight loss on society. Yet, few would advance the proposition that antitrust laws should be used to address such harms. In my view, the social costs of algorithmic exploitation can be searched closer to established antitrust theory. When algorithms absorb most or all consumer surplus in a relevant market, they create an income constraint on consumers, which shifts the demand curve inward on an indeterminate number of other markets. This, in turn, reduces the sales opportunities of other producers, and shrinks a range of (*ir*-relevant) markets, which is a deadweight loss. From a policy perspective, this rationale could legitimise antitrust remediation against perfect behavioural discrimination (correcting for efficiencies), but would leave untouched personal data extraction, given the non-rival and imperfectly appropriable nature of data (no income constraint).

Most of the above-mentioned issues will need to be discussed, tested, and resolved before the scenarios of AAI literature can be integrated in policy environments. To date, the EU Commission is in observational mode. In March, EU Competition Commissioner Vestager said that we should '*keep a close eye on how algorithms are developing*', and learn from early experiences. She added: '*We certainly shouldn't panic about the way algorithms are affecting markets*'. Those words reflect a much welcome commitment to evidence-based antitrust policy, and certainly not that '*antitrust is dead*' as Judge Posner sarcastically put it a week ago.