



The University of Oxford Centre for Competition Law and Policy

www.competition-law.ox.ac.uk

Working Paper CCLP (L) 40

ARTIFICIAL INTELLIGENCE & COLLUSION: WHEN COMPUTERS INHIBIT COMPETITION

Ariel Ezrachi* & Maurice E. Stucke**

“We will not tolerate anticompetitive conduct, whether it occurs in a smoke-filled room or over the Internet using complex pricing algorithms. American consumers have the right to a free and fair marketplace online, as well as in brick and mortar businesses.”¹

INTRODUCTION

One may find it hard to imagine life without the power of computers. Indeed, all areas of our livelihood are affected and have benefited from technological development and an increasingly powerful computerised environment. In line with these developments, recent years have witnessed an ever increasing reliance on big data and big analytics and investment in the development of ‘smart’, ‘self-learning’ machines. These complex machines are set to assist

* Slaughter and May Professor of Competition Law, The University of Oxford. Director, Oxford University Centre for Competition Law and Policy.

** Associate Professor, University of Tennessee College of Law; Co-founder, Data Competition Institute.

¹ Assistant Attorney General Bill Baer of the Department of Justice’s Antitrust Division announcing the division’s first criminal prosecution against a conspiracy specifically targeting e-commerce. http://www.justice.gov/atr/public/press_releases/2015/313011.docx

in decision making, prediction, planning, trade, and logistics. They are also predicted to further enhance our more immediate living environment - the way we commute, shop and communicate.

Not surprisingly, the prospect of Artificial Intelligence (AI) has long fueled human imagination. The development of self-learning and independent computers raises challenging questions as to the future of the human race and the control, or lack of it, humans would exert over machines.²

Interestingly, these developments and the challenges raised by them are also relevant to the area of antitrust enforcement. Sophisticated computers are central to the competitiveness of present and future markets. With the accelerating development of AI, they are set to change the competitive landscape and the nature of competition restraints, which enforcement agencies will need to tackle.

This paper addresses these developments and considers the application of competition law to an advanced ‘computerised trade environment.’ Questions raised and discussed are neither futuristic nor speculative. The Department of Justice, for example, charged in 2015 a price-fixing scheme involving posters sold in the United States through Amazon Marketplace. To implement their agreements, the conspirators, according to the DOJ, “adopted specific pricing algorithms for the sale of certain posters with the goal of coordinating changes to their respective prices and wrote computer code that instructed algorithm-based software to set prices in conformity with this agreement.”³

With the present usage of computers and anticipated technological advancements, more prosecutions involving pricing algorithms are likely. Thus the questions raised in these cases are central to our current thinking on

² See J McCarthy, P Hayes, ‘Some Philosophical Problems from the Standpoint of Artificial Intelligence’ *Machine Intelligence* Vol 4 463-505 (1969), and Chapters 1 and 17 of G Luger, ‘Artificial Intelligence: Structures and Strategies for Complex Problem Solving’ (2005).

³ http://www.justice.gov/atr/public/press_releases/2015/313011.docx

antitrust enforcement and technological developments. Such questions concern, for example, the concept of agreement and intent in a computer dominated environment, the boundaries of legality and collusion, the antitrust liability of algorithms' creators and users, the ability to constrain AI, the relationship between humans and computers, and the possibility of creating ethical, law abiding, machines.

After discussing in Part I the way in which computerised technology is changing the competitive landscape, we explore in Part II possible ways in which computerised agents may be involved in anticompetitive collusion. We consider varying levels of technological development, which differ in the enforcement challenges they raise. Finally, Part III reviews the antitrust policy challenges raised by advanced computers and artificial intelligence.

I. The Changing Competition Landscape

The increased automation of computerised protocols and the rapid developments in technology have changed the way we interact, communicate and trade. Indeed, a look at the way in which we purchase goods and services reveals an increased reliance on the internet, computers and technology. These processes have accelerated the relative decline of the high street trade and the rise of digitalised markets. They have affected our competitive landscape as digitalised markets cover an ever increasing spectrum of commercial activities, from stock trading to the offer and purchase of online products and services.

These increasingly automated and digitalised transactions could create a more effective and transparent marketplace in which resources are allocated more efficiently and in which the best product or service, at the lowest price, triumphs. Alongside this pro-competitive promise, digitalised algorithm based markets are characterised by the ability of sellers to 'shadow' the

activities of users and harvest data on human behaviour. The new market environment provides sophisticated players with the capacity to monitor customers' activities, accumulate data and react to market changes with an ever-increasing speed. Computer algorithms may be used to optimise behavioural advertisements, individualised promotions and targeted, discriminatory pricing.

Indeed, with the rise of data-driven business models, companies are increasingly turning to computer algorithms that can learn from the data they process. Such algorithms operate by 'building a model from example inputs and using this to make predictions or decisions, rather than following strictly static program instructions'.⁴ The velocity at which data are generated, accessed, processed and analysed has increased⁵ and for some applications is now approaching real-time.⁶ Consequently, there is a 'growing potential for big data analytics to have an immediate effect on a person's surrounding environment or decisions being made about his or her life.'⁷ We see this with automated stock trading and other machine learning, where autonomous systems, through algorithms, can 'learn from data of previous situations and to autonomously make decisions based on the analysis of these data.'⁸

Online trade platforms, for example, have been using automatic sales-price determination algorithms for several years.⁹ These platforms enable

⁴ CM Bishop, *Pattern Recognition and Machine Learning* (2006).

⁵ McKinsey Report, *supra* note, at 98 ("More and more sensors are being embedded in physical devices—from assembly-line equipment to automobiles to mobile phones—that measure processes, the use of end products, and human behavior. Individual consumers, too, are creating and sharing a tremendous amount of data through blogging, status updates, and posting photos and videos. Much of these data can now be collected in real or near real time.").

⁶ White House Big Data Report, *supra* note, at 5.

⁷ White House Big Data Report, *supra* note, at 5 (giving as examples of high-velocity data "click-stream data that records users' online activities as they interact with web pages, GPS data from mobile devices that tracks location in real time, and social media that is shared broadly").

⁸ OECD Interim Synthesis Report, *supra* note, at 4.

⁹ Samuel B. Hwang, Sungho Kim 'Dynamic Pricing Algorithm for E-Commerce' in *Advances in Systems, Computing Sciences and Software Engineering* (2006) 149-155; N

sellers to segment the market by using dynamic pricing.¹⁰ This method of pricing is widely used in the travel industry, hotel booking, retail, sport and entertainment.¹¹ Pricing algorithms also dominate online sales of goods – optimising the price based on available stock and anticipated demand. Notably, such an algorithm, which was used by Amazon to optimise profitability, made headlines when it led to an unintended spiralling hike in price of Peter Lawrence’s book *The Making of a Fly*.¹² At its peak, the book was offered for sale at the price of \$23,698,655.93.¹³

Algorithms were also reported to be used in the insurance industry. For example, the so called ‘marketplace considerations’ algorithm used by the Allstate insurance company was set to optimise pricing by determining the likelihood that users would compare prices before purchasing insurance. The use of the algorithm was criticised as it facilitated non-risk-based selective pricing which ranged from up to 90% discount off the standard rate to an increasing of premiums by up to 800%.¹⁴

Interestingly, recent years have witnessed ground-breaking research and progress in the design and development of smart, self-learning machines. The field has attracted significant investment in deep-learning and AI by leading market players.¹⁵ In 2011, IBM’s Jeopardy!-winning Watson computer

Abe, T Kamba ‘A Web Marketing System With Automatic Pricing’ Computer Networks Vol 33 775-788 (2000); LM Minga, YQ Fend, YJ Li ‘Dynamic Pricing: E-commerce - Oriented Price Setting Algorithm’ International Conference on Machine Learning and Cybernetics Vol 2 (2003).

¹⁰ See Salil K. Mehra, Antitrust and the Robo-Seller: Competition in the Time of Algorithms (March 10, 2015). Minnesota Law Review, Vol. 100, Forthcoming. Available at SSRN: <http://ssrn.com/abstract=2576341> (discussing growth of pricing algorithms).

¹¹ See for instance providers of ‘dynamic pricing optimisers’ such as Boomerang Commerce, Prisync, Price Maker, RepricerExpress and others.

¹² <http://www.michaeleisen.org/blog/?p=358>

¹³ <http://www.digitaltrends.com/computing/why-did-amazon-charge-23698655-93-for-a-textbook/>

¹⁴ <http://www.consumerfed.org/news/840>

¹⁵ For example, on a year-over-year basis, funding to for AI start-ups jumped more than 300%. The most sizable deals included Sentient Technologies’ \$103.5M Series C financing from investors including Tata, Horizons Ventures and Access Industries and Vicarious

6

showcased the power of computers and used deep-learning techniques, which enable the computer to optimise its strategy following trials and feedback.¹⁶ Deep-learning techniques have also been implemented in day to day technology. For instance, the technology has been used by Microsoft in its Windows Phone and Bing voice search¹⁷ and by Audi in developing ‘driverless’ cars.¹⁸ More recently, the launch of Deep Q-network by Google showcased enhanced self-learning capacity. The computer was designed to play old fashioned Atari games. Importantly, it was not programmed to react to any possible move in the game. Rather, it relied on models which enabled it to ‘learn’ the game environment through trial and error and improve its performance over time. The technology mimics human learning by ‘changing the strength of simulated neural connections on the basis of experience. Google Brain, with about 1 million simulated neurons and 1 billion simulated connections, was ten times larger than any deep neural network before it.’¹⁹

These technological developments – the rise of computerised market environments, accumulation and harvesting of data, automation of protocols and machine learning – have far reaching consequences when considered in the context of the nature and characteristics of competition between firms and their interface with consumers.

Systems’ \$40M Series B led by Formation 8. ABB Technology Ventures later extended the round by another \$12M; <https://www.cbinsights.com/blog/artificial-intelligence-venture-capital-2014/>

¹⁶ <http://www.techrepublic.com/article/ibm-watson-the-inside-story-of-how-the-jeopardy-winning-supercomputer-was-born-and-what-it-wants-to-do-next/>

¹⁷ <http://www.technologyreview.com/featuredstory/513696/deep-learning/>.

¹⁸ <http://www.technologyreview.com/news/533936/ces-2015-nvidia-demos-a-car-computer-trained-with-deep-learning/>

¹⁹Antonio Regalado, Is Google Cornering the Market on Deep Learning?, MIT Technology Review, January 29, 2014, <http://www.technologyreview.com/news/524026/is-google-cornering-the-market-on-deep-learning/>; <http://www.nature.com/news/computer-science-the-learning-machines-1.14481>

II. THE SPECTRUM OF POSSIBLE ILLICIT CONDUCT

Competition enforcement typically focuses on possible illicit agreements among competitors, anticompetitive vertical restraints (such as resale price maintenance), the abuse of market power, and mergers that may substantially lessen competition. Our focus here is on collusion, which competition authorities across the world condemn.

While antitrust enforcement predominantly targets corporations, the law considers the nature of illicit conduct through a ‘human’ prism. Accordingly, the focal point for intervention is the presence of an agreement or understanding which reflects a concurrence of wills between the colluding companies’ agents. Illegality is triggered when companies, through their employees, directors, agents or shareholders, operate in concert to limit or distort competition.

Interestingly, when computer algorithms and machines take over the role of market players, the spectrum of possible infringements may go beyond traditional collusion. Computers may limit competition not only through agreement or concerted practice, but also through more subtle means. For example, this may be the case when similar computer algorithms promote a stable market environment in which they predict each other’s reaction and dominant strategy. Such a digitalised environment may be more predictable and controllable. Furthermore, it does not suffer from behavioural biases and is less susceptible to possible deterrent effects generated through antitrust enforcement.

In what follows we consider varying levels of technological development and use of computer algorithms, each raising different enforcement challenges. We identify four non-exclusive categories of collusion – the ‘*Messenger*’, ‘*Hub and Spoke*’, ‘*Predictable Agent*’ and ‘*Autonomous Machine*’. For each category, we consider the presence of two important legal concepts: [1] evidence of intent and a horizontal agreement, and [2]

potential liability:

The first category – *Messenger* – concerns the use of computers to execute the will of humans in their quest to collude and restrict competition. Under this basic scenario, humans agree to the cartel and use the computer to assist in implementing, monitoring, and policing the cartel. From an enforcement perspective, the legal concept of agreement can be applied straightforwardly, and prosecutors, with sufficient evidence, will have no difficulty in condemning the use of machines to facilitate coordination. Subsequently, intent evidence plays a limited role in this category.

The second category – *Hub and Spoke* – concerns the use of a single algorithm to determine the market price charged by numerous users. In this scenario, a single vertical agreement by itself may not necessarily generate anticompetitive effects and does not necessarily reflect an attempt to distort market prices. Yet, a cluster of similar vertical agreements with many of the industries' competitors may give rise to a classic hub-and-spoke conspiracy, whereby the developer (as the hub) helps orchestrate industry-wide collusion, leading to higher prices. Since evidence of the competitive effects of these vertical agreements may be mixed, intent evidence can help the competition officials to assess the agreement's purpose and likely competitive effects (i.e., did the companies agree to use the single algorithm to raise prices).

The third category – the '*Predictable Agent*' – presents a more complex scenario. Here, humans unilaterally design the machine to deliver predictable outcomes and react in a given way to changing market conditions. In this category, there is insufficient evidence of any agreement (either vertical or horizontal). Each operator is developing its machine unilaterally, with awareness of likely developments of other machines used by its competitors. An industry-wide adoption of similar algorithms may lead in this case to anticompetitive effect through the creation of interdependent action. Without agreement as such, the market may exhibit the conditions for tacit

collusion/conscious parallelism. As tacit collusion is not in itself illegal, proof of intent to change market dynamics is central in this scenario.

The fourth category – the *Autonomous Machine* – is the trickiest. Here the competitors unilaterally create and use computer algorithms to achieve a given target, such as profit maximisation. The machines, through self-learning and experiment, determine independently the means to optimise profit. Noticeably, under this category neither legal concept – intent nor agreement – apply. The computer executes whichever strategy it deems optimal, based on learning and ongoing feedback collected from the market. Issues of liability, as we will discuss, raise challenging legal and ethical issues.

Before elaborating on each of the categories, the following table summarises the key distinctions between them:

	Agreement	Intent	Liability
Category 1: <i>Messenger</i>	Strong evidence	Limited role	Per Se Illegal
Category 2: <i>Hub & Spoke</i>	Mixed evidence	Evidence used to clarify purpose and likely effect	Per Se / Rule of Reason
Category 3: <i>Predictable Agent</i>	No evidence	Evidence used to show motive and awareness in facilitating tacit collusion	Maybe under FTC Act § 5 or Article 102
Category 4: <i>Autonomous Machine</i>	No evidence	No evidence	Unclear

A. *First Category: The Computer as Messenger*

In this simple scenario, humans use computers to directly execute their instructions. Such use may be subjected to a traditional enforcement approach. An agreement or concerted practice may be established as humans collude through the medium of computers. In this category, humans are the masters, who map out the cartel, while the computer algorithms serve as the messenger, in that they are programmed to help effectuate the cartel, and monitor and punish any cheating. To illustrate: in a classic cartel agreement, executives from rival firms secretly agree to fix prices, allocate markets or bids, or reduce output.²⁰ Here, the executives, after colluding in secrecy, leave it to their computer algorithms to monitor and enforce the agreement.

Competition enforcers can rely on the case law involving an illicit *agreement* or concerted practice and use the concept of ‘object’²¹ or ‘per-se’ illegality.²² The implementation and monitoring of the agreement by the computer may reflect the scope of the agreement and its harm, but the computers’ failure to effectuate the agreement does not affect the agreement’s illegality.²³

²⁰ Scott D. Hammond, Dir. of Criminal Enforcement, U.S. Dep’t of Justice, Antitrust Div., ‘The Fly On The Wall Has Been Bugged—Catching An International Cartel In The Act’ (May 15, 2001), available at <http://www.justice.gov/atr/public/speeches/8280.htm> (ADM case).

²¹ Article 101(1) TFEU provides that agreements, decisions or concerted practices which have as their object or effect the direct or indirect fixing of selling price would be prohibited. The European courts and Commission have generally treated price fixing, market sharing and bid rigging arrangements as having the object of restricting competition.

²² Agreements among competitors that “tamper” with price structure are per se illegal. *United States v. Socony-Vacuum Oil Co.*, 310 U.S. 150, 221 (1940) (“Even though members of the price-fixing group were in no position to control the market, to the extent that they raised, lowered, or stabilized prices they would be directly interfering with the free play of market forces.”).

²³ *Power Conversion, Inc. v. Saft Am., Inc.*, 672 F. Supp. 224, 227 (D. Md. 1987) (“Price-fixing is per se illegal regardless of whether the objective is to raise or lower market prices, whether the agreement is successful or not, and whether the prices were reasonable or not.”). Thus, the Sherman Act reaches combinations formed for the purpose, and with the effect, of raising, depressing, fixing, pegging, or stabilizing prices. Antitrust plaintiffs need not prove that defendants fixed prices directly or controlled a substantial part of the

The stronger the evidence of an anticompetitive agreement in Category I, the lesser the need for intent evidence to establish the concurrence of wills. Still, the intent of the cartel members may play a significant role in establishing the violation and as such merits more detailed consideration.

The law has long considered a person's intent for specific actions.²⁴ The requisite evidence of *intent* for criminally prosecuted per-se illegal offences, such as price fixing, is relatively modest. Lower U.S. courts have held that when the challenged activity is per se illegal under the antitrust laws, the government in criminal cases need only prove the existence of an agreement and that the defendant knowingly entered into the alleged agreement or conspiracy.²⁵ Defendants' altruistic motives are legally irrelevant when the conduct is per se illegal.²⁶

commodity, no competition remained, or prices as a result were uniform, inflexible, or unreasonable. *Socony-Vacuum*, 310 U.S. at 222, 224 n.59.

²⁴ *Morissette v. United States*, 342 U.S. 246, 250-51 (1952); see also *United States v. U.S. Gypsum Co.*, 438 U.S. 422, 436 (1978) ("We start with the familiar proposition that '[t]he existence of a mens rea is the rule of, rather than the exception to, the principles of Anglo-American criminal jurisprudence.'" (quoting *Dennis v. United States*, 341 U.S. 494, 500 (1951))); LYNN STOUT, *CULTIVATING CONSCIENCE: HOW GOOD LAWS MAKE GOOD LAWS* 206 (2011) ("Intent is so central to criminal liability that a person with bad intent can be sent to jail even if she harms no one."). The U.S. Supreme Court has also recognised the relevance of the antitrust defendant's intent, which can be inferred from the defendant's anticompetitive conduct or lack of a valid non-pretexual justification.

²⁵ See, e.g., *United States v. Gillen*, 599 F.2d 541, 545 (3d Cir. 1979) (holding that "in price-fixing conspiracies, where the conduct is illegal per se, no inquiry has to be made on the issue of intent beyond proof that one joined or formed the conspiracy"). The government need not prove the "perpetrator's knowledge of the anticipated consequences" (*Gypsum*, 438 U.S. at 446) or intent to produce the anticompetitive effects. Instead, "a finding of intent to conspire to commit the offence is sufficient; a requirement that intent go further and envision actual anti-competitive results would reopen the very questions of reasonableness which the per se rule is designed to avoid." *United States v. Brown*, 936 F.2d 1042, 1046 (9th Cir. 1991) (quoting *United States v. Koppers Co.*, 652 F.2d 290, 296 n.6 (2d Cir. 1981)) (agreeing "with the express holdings of six other circuits, and the intimations of another, that Gypsum does not require proof of a defendant's intent to produce anti-competitive effects where the defendant is charged with a per se violation of the Sherman Act").

²⁶ *United States v. U.S. Gypsum Co.*, 340 U.S. 76, 87 (1950) ("Good intentions, proceeding under plans designed solely for the purpose of exploiting patents, are no defense against a charge of violation by admitted concerted action to fix prices for a producer's products, whether or not those products are validly patented devices."); *Giboney v. Empire Storage & Ice Co.*, 336 U.S. 490, 496 (1949) ("More than thirty years ago this Court said . . . 'It is too late in the day to assert against statutes which forbid combinations of competing

One example of the first category is the use of computers to facilitate collusion through ticket reservation systems. In the *Airline Tariff Publishing* case, the United States alleged that the defendant airlines used their computerised fare dissemination services to freely negotiate among themselves supra-competitive fares in multiple markets.²⁷ No one questioned that the defendants' computerised fare dissemination system had a pro-competitive purpose in supplying travel agents with basic information about the airline fares for specific routes. However, the antitrust risks arose when the defendant airlines also used this system as a forum to exchange information that was of limited or no use to consumers, but was important to the other airlines in communicating and agreeing upon supra-competitive fares.

The Antitrust Division asserted that the defendant airlines essentially signalled their concurrence or disagreement to entreaties to raise fares and/or eliminate discounted fares through the First and Last Ticket Dates. Essentially, the defendant airlines communicated among themselves relatively costless proposals to change fares through these footnote designators with First and Last Ticket Dates. They employed sophisticated computer programs to process all this fare information, which enabled them to monitor and analyse their competitors' responses to current and future fares on certain routes. These negotiations at times would link fare changes among different routes, and continue for several weeks until all the airlines had indicated their commitment to the fare increases by filing the same fares

companies that a particular combination was induced by good intentions.” (quoting *International Harvester Co. v. Missouri*, 234 U.S. 199, 209 (1914)); *United States v. Socony-Vacuum Oil Co.*, 310 U.S. 150, 221-22 (1940) (noting that the Sherman Act “has no more allowed genuine or fancied competitive abuses as a legal justification for such schemes than it has the good intentions of the members of the combination”); *Nash v. United States*, 229 U.S. 373, 377 (1913) (“The very meaning of the fiction of implied malice in such cases at common law was, that a man might have to answer with his life for consequences which he neither intended nor foresaw.”).

²⁷ *United States v. Airline Tariff Publ'g Co.*, 836 F. Supp. 9, 12 (D.D.C. 1993).

in the same markets with the same First Ticket Date. Likewise, the airlines used the Last Ticket Dates in connection with the footnote designators to communicate proposals to eliminate discounted fares currently being offered to consumers. Not only did this computerised fare dissemination system enable the defendants to negotiate supra-competitive fares, it importantly enabled them to verify that such fares would stick, and signal retaliatory measures against any airline that did not go along with specific fares for specific routes.²⁸

In an updated scenario of the above case, the airline executives would agree broadly not to compete along certain routes, and program their computers to ensure that each airline is allocated its set of customers, to monitor any deviations, and to react automatically to any defections. Importantly, the computers here are used to execute the task which they were set, using pre-loaded data and orders. While faster than their creators, the computer algorithms reflect and are limited by the amalgamation of human

²⁸ This information exchange greatly facilitated tacit collusion, and as noted by the Division, it was of little benefit to consumers. Some defendants disputed this claim, submitting numerous affidavits from travel agents praising the airlines' policy of advanced notice, and arguing that such signalling was employed in geographic markets where only one airline had market power. The travel agents did not, however, have access to some of this information (such as the footnote designators), and thus could not readily determine all of the airlines' contemplated changes to fares. Nor could the agents (unlike the airlines) readily determine the relationships between proposed fare increases for certain routes and the elimination of discounted fares on other routes. Moreover, the pricing information, asserted the Division, was unreliable and misleading, in particular because the airlines changed the ticket dates often. The Division's consent decrees attempted to shift the lever toward promoting information of use to the consumers. The decrees did not prohibit the posting of airfare pricing; rather, the defendants were prohibited from posting fare information of little significance to the consumer, namely Last Ticket Dates, with the exception of those used in advertised promotions, and First Ticket Dates. Thus, the airline's posted fares would have some significance for the consumer, as the travel agents could immediately purchase the ticket that day for that fare. Likewise, by restricting the airlines from using Last Ticket Dates except under advertised commitments, the decrees ended the "costless communication" among the defendants about which discounts should be removed. The decrees did not eliminate the possibility of tacit coordination. Rather, they made such negotiations costlier for the airlines by imposing some risk on the price leader. Moreover, when one airline recently violated this decree by signalling a price increase through a prohibited mechanism, it resulted in a \$3 million civil penalty.

instructions. The computers simply help execute the humans' anticompetitive agreement.

B. Second Category: Hub and Spoke

A second related category that may lead to horizontal collusion involves the use of a single algorithm to determine the market price or react to market changes. Such an algorithm – when used to service many traders – may generate anticompetitive effects. The common algorithm which traders use as a vertical input leads to horizontal alignment.

To illustrate the possible anticompetitive effect, consider the price algorithm Uber uses to determine the contract pricing for taxi services.²⁹ That algorithm has been referred to as 'algorithmic monopoly' as it is controlled by Uber and may mimic a perceived competitive price rather than the true market price.³⁰ As more drivers use Uber's algorithm, one may wonder what its effect on price may be. Reported instances in which the algorithm has pushed the price up raise challenging questions as to the possible manipulation via the algorithm of the perceived market price. With a growing number of users and providers, the alternative universe created by the algorithm may provide an opportunity for exploitation and coordinated price increases.

In this second category, the presence of a vertical agreement between the algorithm developer and user is not contested. The competitors – while agreeing to use the algorithm – did not necessarily agree to fix the prices for taxi services, etc. It is the parallel use of the same algorithm which may give rise to concerns.

²⁹<http://www.quora.com/How-does-Ubers-dispatch-algorithm-work;>
[http://www.technologyreview.com/review/529961/in-praise-of-efficient-price-gouging/;](http://www.technologyreview.com/review/529961/in-praise-of-efficient-price-gouging/)
http://www.slate.com/articles/news_and_politics/view_from_chicago/2015/01/uber_surge_pricing_federal_regulation_over_taxis_and_car_ride_services.html

³⁰ Id.

From an enforcement perspective this category is challenging, in particular, as it requires one to delve into the heart of the algorithm and establish whether it is designed in a way that would or may lead to exploitation. If the algorithm is designed to facilitate collusion among the users, then we would have the classic hub-and-spoke conspiracy; a similar review as the one discussed in our first category would apply.³¹

Absent such anticompetitive design, the competition authority could explore the possible adverse effect of these vertical agreements to use the algorithm under the more forgiving rule of reason standard.

Here, intent evidence may be used to assess the nature of the agreement (i.e., is it purely vertical or is it effectively a horizontal agreement among competitors), its likely competitive consequences, whether to categorise the conduct as a hard-core offence, and whether to prosecute civilly or criminally.³²

In evaluating collaboration among competitors, the agencies consider intent evidence, which “may aid in evaluating market power, the likelihood of anticompetitive harm, and claimed procompetitive justifications where an agreement’s effects are otherwise ambiguous.”³³ Thus, in determining

³¹ See for instance *Tesco v Office of Fair Trading* (Case 1188/1/1/11) Competition Appeal Tribunal, [2012] CAT 31. Indirect information exchange through a third party will be condemned where two phases are present: [I] Retailer A discloses to supplier B its future pricing, with the intention that B will pass that information to other retailers in order to influence market conditions. [II] C receives the information from B, knowing the circumstances in which it was disclosed by A to B; and makes use of that information in determining its own future pricing intentions.

³² U.S. Dep’t of Justice, Antitrust Division, Antitrust Division Manual ch. III-12 (5th ed. last updated March 2014) (noting how the Department of Justice would not prosecute the offence criminally if “here is clear evidence that the subjects of the investigation were not aware of, or did not appreciate, the consequences of their action”).

³³ Federal Trade Commission & U.S. Dep’t of Justice, Antitrust Guidelines for Collaborations Among Competitors 12 n.35 (Apr. 2000), <http://www.justice.gov/atr/public/guidelines/index.html>. Likewise, the European Commission assesses “whether or not an agreement has as its object the restriction of competition,” based on “a number of factors,” including evidence of the parties’ subjective intent. Communication from the Commission, Notice Guidelines on the Application of Article 81(3) of the Treaty (2004/C 101/08) Official Journal of the European Union C

antitrust liability, courts will consider the firms' intent in using the algorithms, i.e., whether they (i) intended a clearly illegal result, such as agreeing to fix prices or (ii) acted with knowledge that illegal results, which actually occurred, were 'probable.'³⁴

C. Third Category: Predictable Agent

The use of a Predictable Agent reflects a scenario where each firm unilaterally uses the computer as part of a more subtle strategy to enhance market transparency and predict behaviour. The industry-wide use of algorithms transforms the market reality to enable conscious parallelism and higher prices. In these new market conditions, the agents may more easily reach a tacit agreement, detect breaches and punish deviations from the common policy.

Unlike our first and second categories, the firms in Category III have not jointly agreed to anything. The firms – in unilaterally creating and implementing the algorithms – did not intend a clearly illegal result, such as agreeing to fix prices. Each firm had an independent economic self-interest to develop and rely on the algorithms; indeed, it may be contrary to the firm's self-interest to rely on human pricing or trading.

To illustrate this possibility, imagine an oligopolistic market in which transparency is limited and therefore conscious parallelism cannot be sustained. Absent tacit collusion, prices on the market will remain

101/97, ¶ 22 (Apr. 27, 2004). Also note: U.S. Dep't of Justice & Federal Trade Commission, Horizontal Merger Guidelines § 2.2.1 (August 19, 2010) ("Explicit or implicit evidence that the merging parties intend to raise prices, reduce output or capacity, reduce product quality or variety, withdraw products or delay their introduction, or curtail research and development efforts after the merger, or explicit or implicit evidence that the ability to engage in such conduct motivated the merger, can be highly informative in evaluating the likely effects of a merger.").

³⁴ United States v. U.S. Gypsum Co., 438 U.S. 422, 444-46 (1978) (concluding that "action undertaken with knowledge of its probable consequences and having the requisite anticompetitive effects can be a sufficient predicate for a finding of criminal liability under the antitrust laws").

competitive. Now, think of the basic conditions for tacit collusion/conscious parallelism. By utilising algorithms, competitors are able to stabilise the market and increase transparency.³⁵ In doing so, they bring the market reality closer to that necessary for conscious parallelism, leading to higher prices. Importantly, the price increase is not the result of express collusion (Category I), but rather the natural outcome of tacit collusion. While the latter is not itself illegal – as it concerns rational reaction to market characteristics – one may ask whether its creation should give rise to antitrust intervention.

This scenario raises several enforcement challenges. In essence, conscious parallelism takes place at two levels. First, when configuring the machines, each human, independently and without collusion, knows that when possible, a dominant strategy may be to follow the price increase of others. Furthermore, each person knows that if other firms settle for a similar program, an equilibrium may be established above competitive levels.

This conscious parallelism at the human level leads to the programming of machines which are aware of possible conscious parallelism at the market level. The computer is therefore set up to monitor the market and explore the likelihood of establishing interdependence of action, without venturing into concerted practice or illicit agreement. The computer is also programmed to punish deviations from a possible tacit agreement and to identify maverick firms which depart from the equilibrium.

In what follows we further elaborate on the market dynamics and enforcement challenges that Category III presents.

Market dynamics

Our scenario concerns similarly designed algorithms, which were developed independently, and are used to monitor activity on the market and

³⁵ See Salil (fn 8) on how pricing algorithms can promote tacit collusion under a Cournot model.

rationally follow price leadership. That activity may stabilise interdependence on a market, subsequently leading to higher prices. Several key features are noteworthy in our algorithm-led marketplace.

First, markets are typically more vulnerable to coordinated conduct when a firm's 'significant competitive initiatives can be promptly and confidently observed by that firm's rivals.'³⁶ This is more likely when 'the terms offered to customers are relatively transparent.'³⁷ In our scenario, for the computer programs to optimise pricing, key market data must be digitalised and accessible. Each firm programs its computer to maximise profit by reacting to other movements on the market. One may, for instance, imagine the use of historic data to calibrate the strategy of the computer and its dominant strategy. As such, when operating in a digitalised market environment they may be able to foster greater transparency and anticipate each other's moves. In such a scenario, computers can rapidly calculate the profit implications of myriad moves and counter-moves, the ability to police deviations and the strategies to punish deviations, and subsequently sustain parallel behaviour.

Furthermore, computer algorithms are quicker to observe price and demand changes, and respond (including tit-for-tat) in adjusting prices for relatively homogeneous products. Moreover, computers, to the extent that they are plugged into their customers' warehouses and amassing other data (such as shipment records etc) can identify if competitors are increasing sales (including expanding into serving new categories of customers, such as institutional buyers, or new territories). Thus computers, in quickly processing their market and customers' proprietary data, may be more effective in monitoring rivals' prices or customers, which not only increases transparency but the risk of coordination.

Second, markets are typically more vulnerable to coordinated conduct "if

³⁶ 2010 Merger Guidelines, *supra* note, at § 7.2.

³⁷ 2010 Merger Guidelines, *supra* note, at § 7.2.

a firm's prospective competitive reward from attracting customers away from its rivals will be significantly diminished by likely responses of those rivals," which "is more likely to be the case, the stronger and faster are the responses the firm anticipates from its rivals."³⁸ Here, computer algorithms can process the pricing-related data quickly to determine price. In markets where customers can switch easily between suppliers and where the goods are homogenous, computer algorithms can quickly detect price reductions by a rival and effectively deprive the rival of any significant sales. The greater the price transparency, the quicker the competitive response, the less likely the first-mover will benefit, and the less likely the price reduction. Industry-wide use of 'meeting-competition' clauses is likely to further increase the likelihood of assimilation through machine learning.

Thus markets are typically more vulnerable to coordinated conduct when each firm would be unlikely to profit from its competitive initiatives. For example, suppose Firm A's computer lowers the price. If Firm A's rivals immediately access the data and adjust prices downward, then Firm A would be unlikely to benefit with additional sales. Given the velocity with which the pricing algorithms can adjust, Firm A would unlikely develop amongst its customers the reputation of a price discounter. Accordingly, Firm A will have less incentive to lower price. Of course, companies can program the algorithms to behave as a maverick, such as rewarding market share growth over profitability within certain bounds, so as to enable them to quickly expand. But even here, the rival firms' computers may develop counter-strategies that ultimately thwart market share growth and foster instead coordinated behaviour.

Third, given the velocity of pricing decisions, firms would no longer have to rely on lengthy (eg 30-day) price announcements, where they wait and see

³⁸ 2010 Merger Guidelines, *supra* note, at § 7.2.

what the competitive response is to decide whether to raise price (and to what extent). Computers can have multiple rounds whereby one firm signals a price increase and the rival computers respond immediately and without the risk that that the firm that initiates the price increase will lose many customers to rivals. Essentially, companies may need now only seconds, rather than days, to signal price increases to foster collusion.

Fourth, the stability needed for tacit collusion is further enhanced by the fact that computer algorithms are not likely to exhibit human biases. Human biases can always be reflected in the programming code, but if some biases are minimised (such as loss aversion, sunk cost fallacy, framing effects), then the algorithm acts consistently on System 2 thinking, rather than System 1.³⁹ The computer does not fear detection and possible financial penalties or incarceration; nor does it respond in anger. Moreover, the universe may be closed with each algorithm sharing a common interest (profits) and inputs (same data), that may lead to stable, durable tacit collusion among a larger number of players as long as they can detect and appreciate the type of algorithm others are using. As computers assimilate, this becomes easier to predict.

Enforcement challenges

The main enforcement challenge in this category concerns the legality of conscious parallelism. A rational reaction by competitors to market dynamics, in itself, is not illegal. When such legal behaviour, absent communication or collusion, leads to an equilibrium being established above competitive levels – it does not trigger antitrust intervention. After all, one cannot condemn a firm for behaving rationally and independently on the market.⁴⁰

³⁹ Daniel Kahneman, *Thinking, Fast and Slow* (Farrar, Straus and Giroux 2011).

⁴⁰ See for example: Case C-199/92 P Hüls AG v Commission, EU Court of Justice,

The question is therefore whether one may condemn and challenge the creation of market conditions which led to sustaining tacit collusion – the creation of a transparent market in which monitoring and punishment mechanisms are present. Should the man-made formation of the conditions for tacit collusion through the use of advanced algorithms be condemned as illegal? And if so, under what conditions? Can the competition agency impute the presence of an illicit agreement or understanding among the competitors to use similar algorithms to dampen competition?⁴¹

Traditional competition provisions in most jurisdictions will require proof of agreement between the parties to change the market dynamics. Such proof may, however, be hard to obtain. This may be so particularly as the strategy to develop the algorithm, to begin with, has been a result of conscious parallelism. Evidence of exchange and sharing of information, or personnel movement from one company to another, may facilitate the finding of illicit concerted practice. One should acknowledge, however, that evolution dictates that the stronger, more powerful algorithms will likely prevail and dominate the technology market. This reality naturally fosters assimilation of systems between various computer developers and companies. A decision not to opt for the most advanced algorithm may be irrational. It would be as if a stock firm would want to rely on human floor traders, when most trading is automated. The use of similar algorithms may further facilitate the

[1999] ECR I-4287, [1999] 5 CMLR 1016; Joined Cases C-89, 104, 114, 116, 117, 125, 129/85 *Ahlström Osakeyhtiö and others v Commission (Wood Pulp II)* EU Court of Justice, [1993] ECR I-1307, [1993] 4 CMLR 407; Cases T-442/08 *CISAC v Commission*, EU General Court, [2013] 5 CMLR 15

⁴¹ The plaintiff can allege that the defendant firms collectively agreed to use these algorithms, namely their collective agreement to use a facilitating device that fosters tacit collusion. See *Todd v. Exxon Corp.*, 275 F.3d 191 (2d Cir. 2001). The benefit of this approach is that it may be easier to prove that the industry agreed to use algorithms (especially in order to ensure their interoperability) and knew that its rival firms' algorithms had similar reward structures than it is to prove an agreement to fix prices. The downsides of this approach are the cost, duration, and unpredictability of a rule of reason case, and the difficulty for the court in weighing the pro-competitive benefits of product developments with the anticompetitive effects.

stabilisation of the market environment: computers can more easily detect the market behaviours of competitors, anticipate the rivals algorithms' likely reactions to different competitive responses, and opt the path that given the competitive reactions will maximize profits, which may often be the path toward conscious parallelism.

Absent the presence of an agreement to change market dynamics, most competition agencies may lack enforcement tools, outside merger control, that could effectively deal with the change of market dynamics to facilitate tacit collusion through algorithms. In some instances one may consider the use of alternative provisions which do not require the presence of agreement to trigger their application. In the United States, for instance, the FTC can bring this claim under section 5 of FTC Act, which does not require an agreement, only a showing of an "unfair practice;" many states have a similar statute. But the FTC has been unsuccessful in bringing these types of claims as evident in *Boise Cascade* and *Ethyl*. If the court adopts the standard in *Ethyl*, the FTC would need to show either: 1. evidence that defendants tacitly or expressly agreed to facilitating device to avoid competition, or 2. oppressiveness, such as (a) evidence of defendants' anticompetitive intent or purpose or (b) the absence of an independent legitimate business reason for defendants' conduct. Accordingly, in Category III, the defendants may be liable if, when developing the algorithms or in seeing the effects, they were: (1) motivated to achieve an anticompetitive outcome, or (2) aware of their actions' natural and probable anticompetitive consequences.

D. Fourth Category: Autonomous Machine – Optimising Performance

The third category—in removing the legal concept of agreement—restricted the range of enforcement tools. The application of section 5 of FTC Act, for example, was contingent on anticompetitive motive or intent. In our last category, we completely remove the legal concept of intent. In doing so,

we exclude section 5 of FTC from the available enforcement tool-box. We consider the possibility that the computer developers foresee tacit collusion as one of many possible outcomes – but not necessarily the likeliest outcome. Smart machines may independently optimise profitability by reaching conscious parallelism—or they may not. Thus the algorithm developers are not necessarily motivated to achieve tacit collusion; nor could they predict when, how long, and how likely it is that the industry-wide use of algorithms would yield tacit collusion.

In this last category we assume that the computer is set a target such as the maximisation of profit, optimisation of performance etc. The algorithm then operates autonomously to achieve the target. The actions of the algorithm are governed by limiting principles which prohibit illegal activity such as price fixing or market sharing.⁴² They do, however, allow self-learning and experimentation.

In this category, we consider the possibility that a self-learning machine may find the optimal strategy is to enhance market transparency and thereby sustain conscious parallelism or foster price increases. Importantly, tacit coordination--when executed--is not the fruit of explicit human design but rather the outcome of evolution, self-learning and independent machine execution.

As noted earlier, conscious parallelism is legal. The question is whether such practices, when implemented by smart machines in a predictable digitalised environment, ought to be condemned. Note that this category differs from the third category in that there is no attempt by the user of the algorithm to facilitate conscious parallelism. The firm ‘merely’ relies on artificial intelligence. With machines rapidly adjusting to new data and competitive scenarios, the users and designers may know that increased

⁴² Absent such limiting principles, the scenario would be similar to the First Category of ‘messenger’.

transparency and supra-competitive prices may occur, but cannot predict ex ante when, for how long, and to what extent.

One should acknowledge the different levels of sophistication which characterise different machine learning algorithms, different agents and different market players. Faster and smarter operators may benefit from market transparency which is not available to others. Furthermore, their ability to react swiftly to change may leave others outside the market, thus increasing barriers to entry. Slower agents may be pushed outside the inner transparency circle and only gain access to it when leading agents opt not to react.⁴³

As with the third category, the ability to collude may have been facilitated by the presence of similar minded agents operating on the market. A self-learning machine may find it easier to tacitly collude with similar machines. It may be easier for such computers to anticipate and understand moves made by other machines which are designed along similar lines. Programmes and computers are easily duplicated – a reality in which a market is operated by similar minded agents should be anticipated.

Interestingly, in a market reality in which such future collusion is possible, the programmers designers may favour the use of similar algorithms. This seemingly benign decision may have significant implications once learning has taken place. The similar machines are more likely to ‘understand’ each other and stabilise a collusive outcome. Importantly, recall that on the ‘factory floor’ these computers have no specific commands which may trigger collusion. It is the self-learning in a transparent market occupied by similar minded agents with the same profit maximising goal which leads to collusion.

⁴³ Ultimately, as recognised by competition authorities, ‘[t]he ability of rival firms to engage in coordinated conduct depends on the strength and predictability of rivals’ responses to a price change or other competitive initiative.’ 2010 Merger Guidelines, *supra* note, at § 7.

Despite similar effects as in the third category, the lack of evidence of an anticompetitive agreement or intent in the fourth category may result in AI self learning escaping legal scrutiny.

III. REFLECTIONS AND POLICY CONSIDERATIONS

Coordinated, accommodating, or interdependent responses among computers raise challenging technical, enforcement and ethical questions. Evidently, these questions differ, depending on the category of technological implementation.

In the ‘*Messenger*’ category, the computer is used as the long arm of cartel members. Here it is merely providing an implementation platform and thus raises few challenges as to the presence of agreement (or intent).

The ‘*Hub and Spoke*’ category concerns the use of a vertical input to facilitate horizontal collusion. The first challenge concerns the technological decoding of the algorithm and related documentary evidence to determine whether it is designed to skew the market prices. In the affirmative, the scenario resembles the Messenger category. If not, the effects of such a network on price, usage and quality should be considered.

The ‘*Predictable Agent*’ category raises challenging questions as to the ability to condemn the creation or strengthening of conscious parallelism through a sophisticated algorithm. Could superior technology which enhances transparency be targeted and condemned, without the risk of chilling innovation and investment?

The ‘*Autonomous Machine*’ category raises similar difficulties with respect to conscious parallelism, but increases the complexity of identifying intent and distinguishing between the operation of the machine and that of its designer.

In what follows we review some of the legal and analytical challenges raised by the ‘*Predictable Agent*’ and ‘*Autonomous Machine*’ categories.

A. *Determining the Primary Purpose for Increasing Transparency*

Market transparency serves as a central variable that facilitates conscious parallelism in Categories III ‘*Predictable Agent*’ and IV ‘*Autonomous Machine*.’ Yet, market transparency is one of the central pillars of effective competitive markets. Greater transparency enables information to flow freely and thus enhances the competitive pressure to the benefit of consumers. After all, the model of perfect competition assumes that buyers will have full information on prices and product characteristics, and the model equilibrium predicts uniform and competitive prices for comparable goods. In a digitalised environment such as the Internet, greater price transparency may reduce the buyers’ search costs in finding the best deal, whether for airline tickets or chainsaws. It may reduce the sellers’ ability to price discriminate.

Thus if the algorithms increase market transparency, one challenge confronting the courts and competition authorities is that the defendants will often have an independent legitimate business reason for their conduct. Courts and the enforcement agencies may be reluctant to restrict this free flow of information in the marketplace. Its dissemination, observed the Supreme Court, “is normally an aid to commerce”⁴⁴ and “can in certain circumstances increase economic efficiency and render markets more, rather than less, competitive.”⁴⁵ Indeed, concerted action to reduce price transparency may itself be an antitrust violation.⁴⁶

⁴⁴ *Sugar Institute, Inc. v. United States*, 297 U.S. 553, 598 (1936).

⁴⁵ *United States v. United States Gypsum Co.*, 438 U.S. 422, 441 n.16 (1978); see also Richard A. Posner, *Antitrust Law* 160 (2d ed. 2001) (generally, the more information sellers have about their competitors’ prices and output, the more efficiently the market will operate).

⁴⁶ See, e.g., Press Release, Federal Trade Commission, Virginia Board of Funeral Directors & Embalmers, FTC 041-0014 (Aug. 16, 2004), available at <http://www.ftc.gov/opa/2004/08/vafuneral.htm> (board’s prohibition on licensed funeral directors advertising discounts deprived consumers of truthful information); Press Release, Federal Trade Commission, Arizona Automobile Dealers Association, FTC C-3497 (Feb. 25, 1994), available at 1994 WL 184107 (trade association illegally agreed with members to

Thus, one may find it difficult to fine-tune the enforcement policy to condemn ‘excessive’ market transparency. This may be particularly challenging when the information and data are otherwise available to consumers and traders and it is the intelligent use of that information which facilitates conscious parallelism.

So, if humans program the computer to optimise profit and know that by reacting to changing market conditions through self-learning, the computer will likely find collusion as the dominant strategy, are they liable? Perhaps – if there is very strong evidence of anticompetitive intent. If the executives, for example, call their algorithm *Gravy*, and tinker with it to better manipulate the market, and boast about this in their internal e-mails – as was the Securities and Exchange Commission’s (SEC) case against Athena Capital Research – then liability is likely.⁴⁷

In 2014, the SEC for the first time sanctioned a high frequency trading firm, Athena Capital Research, for manipulating the market by “placing a large number of aggressive, rapid-fire trades in the final two seconds of almost every trading day during a six-month period to manipulate the closing prices of thousands of NASDAQ-listed stocks.”⁴⁸ The SEC found that Athena used complex computer programs to manipulate the closing prices of thousands of NASDAQ-listed stocks over a six-month period.⁴⁹ The sophisticated algorithm, code-named *Gravy*, engaged in a practice known as “marking the close” in which stocks were bought or sold near the close of

restrict nondeceptive comparative and discount advertising and advertisements concerning the terms and availability of consumer credit); OECD DAFPE/CLP (2001) 22, *supra* note, at 183, 185–86 (citing examples of U.S. enforcement agencies seeking to increase price transparency); but see *InterVest, Inc. v. Bloomberg, L.P.*, 340 F.3d 144 (3d Cir. 2003) (lack of price transparency in bond market not illegal if consistent with unilateral conduct).

⁴⁷ Available online: <http://www.sec.gov/litigation/admin/2014/34-73369.pdf>

⁴⁸

<http://www.sec.gov/News/PressRelease/Detail/PressRelease/1370543184457#.VEOZlfdV8E>

⁴⁹ <http://www.sec.gov/litigation/admin/2014/34-73369.pdf>.

trading to affect the closing price: “The massive volumes of Athena’s last-second trades allowed Athena to overwhelm the market’s available liquidity and artificially push the market price – and therefore the closing price – in Athena’s favour.”⁵⁰

Athena’s employees, the SEC alleged, were “acutely aware of the price impact of its algorithmic trading, calling it ‘owning the game’ in internal e-mails.”⁵¹ Athena employees “knew and expected that Gravy impacted the price of shares it traded, and at times Athena monitored the extent to which it did. For example, in August 2008, Athena employees compiled a spreadsheet containing information on the price movements caused by an early version of Gravy.”⁵² Athena configured its algorithm Gravy “so that it

⁵⁰

<http://www.sec.gov/News/PressRelease/Detail/PressRelease/1370543184457#.VEOZlfdV8E>

⁵¹

<http://www.sec.gov/News/PressRelease/Detail/PressRelease/1370543184457#.VEOZlfdV8E>. As the SEC alleged Athena’s manipulative scheme focused on trading in order imbalances in securities at the close of the trading day:

Imbalances occur when there are more orders to buy shares than to sell shares (or vice versa) at the close for any given stock. Every day at the close of trading, NASDAQ runs a closing auction to fill all on-close orders at the best price, one that is not too distant from the price of the stock just before the close. Athena placed orders to fill imbalances in securities at the close of trading, and then traded or “accumulated” shares on the continuous market on the opposite side of its order.

According to the SEC’s order, Athena’s algorithmic strategies became increasingly focused on ensuring that the firm was the dominant firm – and sometimes the only one – trading desirable stock imbalances at the end of each trading day. The firm implemented additional algorithms known as “Collars” to ensure that Athena’s orders received priority over other orders when trading imbalances. These eventually resulted in Athena’s imbalance-on-close orders being at least partially filled more than 98 percent of the time. Athena’s ability to predict that it would get filled on almost every imbalance order allowed the firm to unleash its manipulative Gravy algorithm to trade tens of thousands of stocks right before the close of trading. As a result, these stocks traded at artificial prices that NASDAQ then used to set the closing prices for on-close orders as part of its closing auction. Athena’s high frequency trading scheme enabled its orders to be executed at more favorable prices.

⁵² SEC Order, *supra* note, at para 34.

would have a price impact.”⁵³ In calling its market-manipulation algorithm “Gravy,” and exchanging a string of incriminating e-mails, the company did not help its case. Without admitting guilt, Athena paid a \$1 million penalty.

This case is illustrative. Automated trading has the potential of increasing market transparency and efficiency, but can also lead to market manipulation.⁵⁴

Finding the predominant purpose for using an algorithm will not always be straightforward. Athena, for example, challenged the SEC’s allegations that it engaged in fraudulent activity: “While Athena does not deny the Commission’s charges, Athena believes that its trading activity helped satisfy market demand for liquidity during a period of unprecedented demand for such liquidity.”⁵⁵ A court might agree. Companies also can learn from *Athena* and be more circumspect in their e-mails.

More, reliance on intent evidence also does not help enforcers in Category IV, where consumers are still harmed by the conscious parallelism facilitated by industry-wide use of pricing algorithms.

B. Advanced Safeguards

A potential solution to increased transparency and cooperation may be to require firms to share the data used in their algorithms with the public. When the data exchange is asymmetrical, namely the data is not provided or valuable to the company’s and its competitors’ customers, the dissemination of such information among competitors, while not per se illegal, carries a high antitrust risk especially when its exchange is unlikely to promote overall

⁵³ SEC Order, *supra* note, at para 36.

⁵⁴ <http://dealbook.nytimes.com/2014/10/20/why-high-frequency-trading-is-so-hard-to-regulate/>

⁵⁵ <http://www.marketwatch.com/story/high-frequency-trading-firm-fined-for-wave-of-last-minute-trades-2014-10-16>

efficiency and is likely to (or in fact did) promote tacit collusion.⁵⁶

The problem is that information asymmetry, while relevant in the days when competitors exchanged printed price lists and e-mails, is less relevant with machine learning where the computers process a voluminous variety of commercially-available data. If the data is generally available, customers may be using it as well.

Another option is to program computers to ignore commercially sensitive information that, although publicly available, is of little or no value to customers but is very helpful in enabling the competitors to arrive at a *supra*-competitive price.⁵⁷ However, identifying such information is problematic. Part of the value of big data is data fusion, whereby computers link data sets from which new insights emerge.⁵⁸ Moreover, the data for some applications

⁵⁶ Why would competitors share a future price increase among themselves exclusively (or before announcing it publicly)? One possibility is to avoid the risk of losing customers as they negotiate through successive communications about how much to increase prices (or to decrease them, to threaten discounters). Moreover, by voluntarily sharing detailed transactional information with each other, the competitors can effectively police the price increase and detect any cheating. The customers, on the other hand, stand little, if anything, to gain by this increased price transparency among competitors. They are still largely in the dark about the future price increase (and thus cannot effectively adjust their purchases) or the prices that others have paid (and thus cannot leverage a better price with this information). It is questionable then how the marketplace is rendered more efficient and competitive by such asymmetric exchanges.

⁵⁷ As an example, in *Petroleum Products*, the defendant oil companies publicly announced, at times in advance of the effective date, the discounts (or decisions to withdraw discounts) to their franchisee gasoline stations. Coordinated Pretrial Proceedings in *Petroleum Products Antitrust Litig.*, 906 F.2d 432, 445 (9th Cir. 1990). The public dissemination of the discount information was of little value to the defendants' franchisees or the end consumer. The franchisees could not shop around for the best oil prices: they could only purchase from their franchisor. Nor did the consumers care what the gas station paid for the gasoline. They cared only about the retail price. The purpose and effect then of publicly announcing changes in discounts to the franchisees were, as several defendants' executives admitted, to quickly inform their competitors of the price change, in the express hope that these competitors would follow the move and restore their prices. Without such transparency, the other defendants might not have readily detected one defendant's withdrawal of its discount and followed accordingly, because the individual branded gas stations' retail prices varied considerably.

⁵⁸ Executive Office of the President, President's Council of Advisors on Science and Technology, Report to the President Big Data and Privacy: A Technological Perspective at x (May 2014) [hereinafter PCAST Report]; OECD Data-Driven Innovation, *supra* note, at 12 (observing that "In some cases, big data is defined by the capacity to analyse a variety of

– such as customers sharing their inventory data with suppliers – can promote efficiency while raising antitrust concerns.⁵⁹ Even if the customers seek to limit what information can be shared, the algorithms – by analysing a variety of data – could fill in the gaps. So it would likely be difficult to program the computers to ignore data sets without reducing efficiency.

C. Reconsidering the Relationship Between Humans and Machines

The consideration of the ‘autonomous agent’ raises ethical and policy questions on the relationship between humans and machines. In such instances, can the law attribute liability to companies for their computers’ actions? At what stage, if any, would the designer or operator relinquish responsibility over the acts of the machine?

Evidently, defining a benchmark for illegality in such cases is challenging. It requires close consideration of the relevant algorithm to establish whether any illegal action could have been anticipated or was predetermined. Such review requires consideration of the programming of the machine, available safeguards, its reward structure, and the scope of its activities. The ability to identify the strand which is of direct relevance is questionable. The complexity of the algorithms’ data-processing and self-learning increases the risk that enforcers, in undertaking this daunting undertaking, stray far afield of rule of ideals, such as transparency, objectivity, predictability, and accuracy.

Further, one must consider the extent to which humans may truly control self-learning machines. Humans design the initial algorithm, determine to

mostly unstructured data sets from sources as diverse as web logs, social media, mobile communications, sensors and financial transactions. This requires the capability to link data sets; this can be essential as information is highly context-dependent and may not be of value out of the right context. It also requires the capability to extract information from unstructured data, i.e. data that lack a predefined (explicit or implicit) model.”)

⁵⁹ <http://www.gsb.stanford.edu/insights/sharing-information-boost-bottom-line>.

launch it, and arguably could shut the computer down. But between its creation and death, the computer can undertake many strategies. Could a self-learning machine choose strategies that indirectly circumvent safeguards imposed to protect consumer welfare and operate independently to maximise profit? Could machines simply override the safeguards? Such questions draw attention to wider ethical and moral issues, which may affect the way in which a future society may evolve and the ability of humans to control and restrict such developments.⁶⁰

In the context of competition and markets, friction between profit maximisation, ethical trading and consumer welfare exists. As algorithms, through reinforcement learning, identify ingenious solutions, consumers and society can benefit. But as algorithms extend to everyday business decisions, such as fixing the prices for goods and services, there is also the possibility that computers – in order to maximise profits – engage in coordinated, accommodating, or interdependent behaviour. Importantly, they may do so through self-learning and rational decision making, in a deterrent-free-environment, bypassing safeguards which inhibit traditional price fixing or collusion.

To illustrate the multiple ethical decisions ‘smart’ computers may have to make, consider, for example, the design of smart “autonomous” cars. In designing these cars, car makers have to consider whether algorithms should replicate ethical human decision making. Such may be the case, for instance,

⁶⁰ Wendell Wallach and Colin Allen, *Moral Machines: Teaching Robots Right from Wrong* (OUP 2008) 6; James H Moor, ‘The Nature, Importance and Difficulty of Machine Ethics’ (2006) 21 *IEEE Intelligent Systems* 18; Colin Allen, Iva Smit and Wendell Wallach, ‘Artificial morality: Top-down, bottom-up, and hybrid approaches’ (2005) 7 *Ethics and Information Technology* 149; Samir Chopra and Lawrence White, ‘Artificial Agents – Personhood in Law and Philosophy’ (2004) 16 *European Conference on Artificial Intelligence* 635; Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (OUP 2014); Michael Anderson and Susan Leigh Anderson, ‘Machine Ethics: Creating an Ethical Intelligent Agent’ (2007) 28 *AI Magazine* 15; Isaac Asimov, *I, Robot* (Reprint, 2013 Harper Voyager)

when the computer identifies an imminent crash and needs to consider evasive action. Alternative crash outcomes may include, for example, severely injuring a child, killing an old person or damaging property. The ethical decision as to the least harmful action cannot be taken lightly.⁶¹

By analogy from the machine ethics debate, one may pose the question as to the way in which one could integrate ethics and legality into a computer programme. Outside the clear instance of collusion through communications, how could one constrain the computer's actions to avoid less competitive outcomes? Should such a move focus on the competitiveness of the market as a limiting benchmark or rather on illegality? As the two do not necessarily overlap, an explicit prohibition may not resolve the problem. While, there is no doubt that legality may be computed into any machine, our concern stems from the ability of the machine to change the competitive landscape and thus reduce competition.

In the area of ethics and morality, a rule-based approach to artificial intelligence has been criticised for its unsuitability and has proven to be insufficiently robust for most real-world tasks.⁶² When considered in the context of facilitating tacit collusion, one may wonder whether it may provide any workable rules to follow. This brings us back to the legal question of liability. To what extent should or could liability be imputed on the creator or operator of the machine? Should the human and machines be viewed as independent from each other or rather treated as one? In instances in which the machine does not act on instructions of the designer or operator, can liability be imputed? Can the use of a self-learning machine be condemned?

⁶¹ Chris Bryant, Driverless cars must learn to take ethical route, Financial Times, March 1, 2015 3:27 pm, <http://www.ft.com/cms/s/0/4ab2cc1e-b752-11e4-981d-00144feab7de.html#slide0>.

⁶² Colin Allen, Iva Smit and Wendell Wallach, 'Artificial morality: Top-down, bottom-up, and hybrid approaches' (2005) 7 Ethics and Information Technology 149, 149-150.

D. Deterrence and Liability

A possible solution may concern the imposition of liability once defendants become aware of the coordinated, accommodating behaviour among the rivals' computers. That is not uncommon in other areas of the law, whereby one is liable for continuing to knowingly benefit from an illegal source of income. For example, there is liability under the Proceeds of Crime Act 2002 for dealing with funds that one suspects of being criminal property, even if the defendant was not involved in the creation of those funds in the first place.

But would such an approach provide a workable intervention principle? Suppose the company recognises that its computerised pricing structure is optimising profits. It is plausible that profits are increasing as the computer programs are reducing costs (such as better utilising resources) or using price discrimination and strategic discounting. It would be unlikely for the company to ascertain precisely to what extent the increase in profits is attributable to coordinated, accommodating, or interdependent behaviour. It may be impossible to determine the extent to which the computer is reacting to market dynamics or shaping them. Would a hunch by executives that some of the profits are coming from such accommodating tacit agreement be enough to impose antitrust liability?

In addition, if the defendants become aware of tacit collusion, what can they do about it? Depending how far the pricing mechanism is integrated with other functions, they could not necessarily turn off the computer. Nor could they necessarily override the algorithm's price with their own price, as it may be logistically impossible to enter and update the pricing across markets. Moreover, the adjusted pricing may still be inflated.

Assuming that the computers are programmed to refrain from any violation of the competition laws, the company may have done all that it can to ensure compliance. The facilitation of conscious parallelism through

rational independent action may well fall outside the compliance with competition law. Furthermore, programming compliance is challenging, in particular when one attempts to capture the creation of market dynamic such as conscious parallelism. A command not to fix price may be simple to execute, but under reinforcement learning, the algorithm will seek ingenious solutions including, as the competition authorities recognise, the myriad possibilities of coordinated interaction, not all of which are illegal.⁶³

E. Active Intervention

Considering the above, it may be challenging for law makers to identify a clear, enforceable triggering event for intervention which would prevent the change of market dynamics. Furthermore, it is likely to be challenging for the competition agencies to enforce such a provision.

At the legislative level, one could consider the possibility of utilising an ex-ante approach by which under certain market conditions, companies will be required to report the use of certain algorithms. Such a regulatory mechanism is likely to trigger costs at agency and company levels. It may also prove difficult to implement, especially in cases involving the fourth scenario. Competition authorities would have a difficult time overseeing firms' attempts to design a machine to optimise performance while instructing it to ignore or to respond irrationally to market information and competitors' moves, or to pursue inefficient outcomes.

An ex post approach may trigger intervention when markets seem to operate in concert. A market review or inquiry may require companies to reveal the nature of algorithms used, in an attempt to ascertain whether these algorithms create excessive transparency or lead to interdependence. The more selective intervention under an ex-post regime may have more limited

⁶³ On this point see: 2010 Merger Guidelines, *supra* note, at § 7.

cost implications. It may also limit the adverse effects on innovation and investment, as it is only once tacit collusion is detected that the market is subjected to a monitoring exercise. An ex-post monitoring exercise would require the legislator to determine whether liability ought to be imputed on the companies involved. Taking our earlier discussion into account, one may favour no liability rule (Category IV) absent clear evidence of intent (Category III).

Be it an ex-post or ex-ante, regime, one still has to confront the challenge of identifying the adequate level of intervention, if such exists, when dealing with the creation of market conditions for conscious parallelism. A remedy which requires an algorithm to ignore market prices or not to react to market changes may well undermine competition. The enforcer's efforts to reduce price transparency may similarly harm consumers. Undoubtedly, intervention would require careful technological and policy fine tuning to avoid these pitfalls. Some may argue that these challenges should even tilt the balance in favour of nonintervention. Such an approach, however, risks creating a lacuna which may well be exploited by market players. Assuming that technology can provide benchmarks and tools for intervention, one should not ignore the possible rise of a new form of anticompetitive conduct.

CONCLUSION

Computer algorithms have transformed the way we trade and will continue to do so in an increasing pace. The creation of fast moving digitalised markets yields many benefits, yet it also changes the dynamics of competition and may limit it.

Our discussion explored four categories of technological use to inhibit competition. We identify as most challenging, from both legal and enforcement perspectives, instances in which algorithms may be used to

facilitate conscious parallelism and are not likely to be challenged under current laws.

The fourth category, which concerns the use of Autonomous Machines, further challenges our thinking as it raises questions as to the relationship between humans and machines. These questions may become increasingly prominent as technology advances and AI becomes an integral part of our surroundings.

The possible detachment between the actions of the algorithm and those of the human operators raise challenges regarding the ability to attribute liability to its operators, who may escape scrutiny due to the unforeseen nature of self-learning. Rule of law concerns include how to differentiate between express agreement and accommodating behaviour, and greater subjectivity over whether and when computers “agreed.” Ethical concerns include to what extent should humans be held accountable for low probability or hard to predict events? With no human intent and immoral conduct, there is a greater risk of jury nullification.

The detachment between the algorithm and its operators also reveals a potential failure to deter as algorithms are not susceptible to traditional deterrents, such as jail, monetary fines, and shaming.

In a digitalised universe in which the law’s moral fabric is inapplicable, any game theories are constantly modelled until a rational and predictable outcome has been identified. Given the transparent nature of these markets, algorithms may change the market dynamics and facilitate tacit collusion, higher prices, and greater wealth inequality. In such a reality, firms may have a distinct incentive to shift pricing decisions from humans to algorithms. Humans will more likely wash themselves of any moral concerns, in denying any relationship and responsibilities between them and the computer.

Thus, policymakers must recognize the dwindling relevance of

38

traditional antitrust concepts of “agreement” and “intent” in the age of Big Data and Big Analytics. Rather than redefining agreement or intent, perhaps policymakers need to introduce checks and balances into the original pricing algorithm and a monitoring function.

NO COMPUTERS WERE HARMED IN THE MAKING OF THIS PAPER