February 2018

RBB Brief 55

RBB | Economics

Automatic Harm to Competition? Pricing algorithms and coordination

In a recent speech,[1] EU Commissioner Vestager expressed the following concerns regarding the use of pricing algorithms:[2]

"I think we need to make it very clear that companies can't escape responsibility for collusion by hiding behind a computer program."

Other competition agencies have echoed similar concerns for the risks that pricing algorithms pose for collusion. The topic has also featured in a recent OECD roundtable and in a number of academic papers.[3]

This Brief assesses the impact of different categories of pricing algorithms, identifies their links with coordination concerns, and evaluates some possible competition law enforcement responses. Pricing algorithms do raise some interesting issues, but the worst case scenarios for collusion have been overplayed, and some of the calls for increased intervention reveal a worrying gap in the understanding of the economics of oligopolistic markets.

Coordination and the oligopoly problem

Successful coordination happens when members of an oligopoly group find ways to restrict output and sustain prices above the competitive level for their mutual benefit.

Because any coordinated outcome is rendered inherently unstable by the ever-present individual incentive to undercut (i.e. "cheat"), in game theory terms this outcome is not an equilibrium in a one-shot game. But in repeated games it is possible that a mechanism can be found to supress this incentive to cheat. This generally requires a focal point (i.e. an obvious and compelling coordinating strategy to follow), monitoring and transparency (to permit members of the coordinating group to observe whether other members are adopting the strategy), and an effective punishment mechanism (to deter cheating).

Economics sees tacit and explicit coordination as the same phenomenon, though the legal consequences are often very different. Firms whose conduct breaches laws (such as Article 101) that prohibit anti-competitive agreements and concerted practices face huge fines and liability for damages, but such laws do not prohibit tacit collusion when coordination arises without any agreement between rivals.

This means that most existing competition laws are powerless to prevent some coordination from happening, possibly even some cases of coordination that suppress competition to such an extent that they result in the monopoly outcome. So the fact that coordination (and supra-competitive prices) might be observed in an industry cannot be used as a reliable indicator for whether an anti-competitive agreement has been reached.[4]

Regulators might well feel uncomfortable with this apparent "gap" in their enforcement powers, but there are sound reasons why most laws do not encroach on tacit collusion. First, it is generally very difficult for firms to sustain serious coordination without recourse to formal communication or agreements, given the inherent instability of such conduct. Second, in real world imperfect markets it is hugely complex to distinguish coordinated from non-coordinated outcomes, so attempts to intervene against such conduct carry a clear risk of false convictions and unintended chilling effects.

1. Vestager (2017), "Algorithms and Competition", Speech at the Bundeskartellamt 18th Conference on Competition, Berlin.

2. A pricing algorithm is a set of rules that convert inputs (e.g. market conditions, competitor prices, observable customer characteristics) into prices (or price recommendations for human review). Machine learning (or AI) refers to a situation where the rule that converts inputs to prices itself adapts (in an automated way) to a measure of performance.

3. OECD (2017), "Algorithms and Collusion: Competition Policy in the Digital Age". See also Ezrachi and Stucke (2016), "Virtual Competition: The Promise and Perils of the Algorithm Driven Economy", Harvard University Press.

4. Conversely, a number of cartel infringements involve conduct that, whilst unlawful, has failed in its attempts to elevate prices above the counterfactual level.

Pricing algorithms and coordination

Numerous types of pricing algorithms are in use in today's economy. Algorithms harness the processing power of computing and enable firms both to gather intelligence about the market in which they operate and to systematise their actions under different scenarios. It is evident that such tools confer substantial efficiency advantages and reduce the costs associated with market uncertainty. For example, algorithms have been utilised in the air transport and insurance sectors to solve problems such as clearing markets and ensuring that risk is evaluated and priced effectively. Nevertheless, there is a growing recognition that pricing algorithms can transform oligopoly conduct in a way that might contribute to coordinated outcomes.

A commonly observed pricing algorithm is a tool that monitors the pricing (or other conduct) of rival firms and incorporates an automated response, such as a commitment to match or undercut the prices of rivals when certain conditions are met. Such conduct may have clear consumer benefits. For example, if store A automatically matches discounts offered by store B, A's customers can benefit from B's lower prices without incurring the cost of switching to B. However, the flip side is that store B thereby stands to gain fewer new sales from customers who switch from A to B, and in some cases that may deter store B from offering the discount in the first place.[5]

Hence, the concern is that the use of such algorithms can tip the balance towards successful coordination by increasing transparency, reducing the gain from undercutting rivals, and making punishment more rapid. Further, since automation can readily deal with the computational complexity of scanning and adjusting multiple prices, there is a fear that algorithms bring coordinated outcomes into play in markets in which the variety and range of competitive parameters placed effective coordination beyond the reach of mere mortals.

The recent OECD staff working paper on the topic presents a simple game theory model that found a deterministic link between coordination and algorithms that monitor rival prices and match lower prices in real time. But that result applies only in a stylised model in which transparency is perfect and the price response immediate.[6] In some respects, the rarefied assumptions on which these theoretical results rely, when compared to the complexity of most real world markets, serves to emphasise that coordinated outcomes from price-matching algorithms may be the exception rather than the norm. For example, in the case of the rival stores using an algorithm to match prices, a price matching promise can apply only where both retailers sell identical products. To the extent that stores offer substitutable but different products (e.g. because they offer different competing brands, the same brand in different pack sizes and/or unbranded goods), price-matching clauses are more likely to affect the shape that price competition takes, rather than to eliminate rivalry to the detriment of consumers.

While these algorithms may just be simple rules for converting inputs (like competitor prices, market conditions, customer characteristics) into prices, the most sophisticated may employ machine learning and so embody an element of artificial intelligence. Algorithms can be tasked with solving complex problems and achieve this aim by adapting their underlying rules over time to learn from past outcomes (including from interaction with other participants such as rival firms). For example, in one case study a machine learning tool has been found to be capable of winning in a game of poker against expert human players, adjusting its conduct to mimic strategies such as bluffing.[7]

The true capabilities of such algorithms are not known with certainty, and they will no doubt change over time. The concern is either that they can be programmed by oligopoly members to solve the coordination problem after repeated interaction with rivals, or that they achieve a coordinated outcome even when given a more neutral objective of enhancing profits. But this does not render coordination inevitable. If these algorithms have power to constantly adapt, and since cheating from the coordinated outcome is always individually more profitable than adhering in the short term, how can participants in a game that is played out through algorithmic agents trust rivals not to cheat? Even if algorithmic pricing enables firms to take advantage of price transparency and increases the speed of retaliation to price cuts, firms would presumably have an incentive to seek ever more covert ways of cheating without being detected (for example via customer specific discrimination, secret discounts or the introduction of new products).

5. See The Economist, "Guaranteed Profits - Price-match guarantees prevent rather than provoke price wars." (2015)

6. See also Salcedo, B. (2015), "Pricing Algorithms and Tacit Collusion", Manuscript, Pennsylvania State University. The assumptions of the Salcedo model are even stronger – there is transparency not just over prices but over the algorithms themselves (algorithms can be decoded). There is also an assumption of short term commitment to an algorithm (algorithms cannot be revised quickly) and of long term flexibility (algorithms can be revised over time).

7. The *Libratus* programme, designed by Tuomas Sandholm and Naoh Brown at Carnegie Mellon University, cited in the OECD staff paper on Algorithms and Collusion (page 10).

8. See Online Pricing of Posters and Frames, CMA, 4 December 2015, https://www.gov.uk/government/news/online-seller-admits-breaking-competition-law

9. See "Commission accepts commitments by container liner shipping companies on price transparency", 7 July 2016, http://europa.eu/rapid/press-release_IP-16-2446_en.htm. This contrasts with the remedy extracted by the UK CMA, on 13 April 2016, in the cement market investigation case where similar conduct was addressed through a market study that did not allege any breach of existing competition law prohibitions, see https://www.gov.uk/government/news/cma-publishes-final-cement-market-data-order-and-undertakings

It might be argued that machines would learn that, in the long term, cheating is not worth it. But how would they learn that? If a punishment strategy were "grim" (i.e. punishment lasts for ever) then one mistake would be fatal to coordination. On the other hand, if punishment mechanisms were designed to permit a return to coordination (e.g. punishment is limited to the short or medium term) then numerous punishment strategies would potentially sustain coordination and so (as is often the case) the key issue is how to ensure that firms understand the exact terms of coordination, including the precise nature of the punishment strategy and how it evolves over time. Disagreements or misunderstandings over how to do so would undermine coordination.

This conclusion is further underlined when one considers the influence of the completely separate class of pricing algorithms that helps suppliers to implement demand-based pricing. By tailoring prices to meet the individual characteristics of each customer, these tools play an increasingly influential role in a wide range of industries, and have undoubtedly contributed to efficiency and innovation in the way that assets, goods and services are priced. Their ability to segment markets, implement price discrimination and remove the extent to which infra-marginal consumers can benefit from the actions of their marginal counterparts can also raise a number of public policy and consumer protection issues. There is, however, no suggestion that this class of algorithm enhances coordination concerns. On the contrary, they clearly act to increase the complexity of any coordination strategy, not least because they preclude or hugely complicate the task of identifying a meaningful focal point for coordination.

In short, the essential tension that exists within any coordinated outcome between collective and individual incentives is not in itself eliminated by the use of algorithms, and the indicators do not all point towards algorithms increasing the risks of successful coordination. They certainly do not suggest that coordinated outcomes are inevitable.

Policy responses

It is evident from statements such as Commissioner Vestager's that enforcers are keen to act against pricing algorithms, but the appropriate enforcement actions are not so clear.

There are some obvious misuses of pricing algorithms that seem to be contrary to existing competition laws against anti-competitive agreements and concerted practices. For example, an agreement between rivals to use a common algorithm in a specific attempt to coordinate their prices fits the criteria for a standard cartel. This situation arose in a recent investigation by the UK's CMA where firms selling posters and frames on Amazon agreed with each other what their algorithms should look like.[8]

Similarly, it is easy to see how a situation where rivals agree to outsource pricing decisions to a single entity which then finds itself in the position of controlling the price decisions of those rivals could fall into the category of joint selling arrangements that have been condemned in past cases.

However, other uses of pricing algorithms appear not to meet the current criteria for an unlawful anti-competitive agreement. For example, the parallel but uncoordinated adoption of price algorithms that incorporate price matching clauses, or situations where oligopoly members set an algorithm to solve for a profit-maximising objective, appear to mimic the kinds of parallel oligopoly conduct that is not currently sanctioned by these laws. Importantly, the fact that such conduct *might* lead to less competitive outcomes is not sufficient to show that the firms involved have breached existing laws against anti-competitive agreements.

The key question is whether, in this case, a change to the enforcement of existing laws, or perhaps the introduction of new laws, is justified to curb the potential evils of price algorithms. Given their stated determination to "do something", there is a real possibility that competition authorities will seek to increase their discretion to intervene.

One way to do so might be for them to expand their interpretation of what constitutes an anti-competitive "agreement". This process appears to be happening in any case, for example in cases such as the commitments extracted by the EU Commission concerning price announcements in the liner shipping industry.[9] However, the lessons from such cases do not suggest it is easy to provide a bright line test for what constitutes an anti-competitive agreement, or to provide the clarity on the remedies that would be required to justify the large fines and liability for damages that can arise from anti-competitive breaches of the law.

A second option might be for regulators to press for new powers to attack the use of algorithms that go beyond existing rules against anti-competitive agreements. But if regulators had licence to use an effects-based test to condemn all situations in which oligopolistic interaction resulted in successful tacit coordination, that would involve a substantial increase in the scope of competition policy. It would also create huge enforcement challenges if firms were condemned for doing no more than adapting intelligently to the oligopolistic markets in which they find themselves competing, or were required to set prices in a way that ignored the evident interdependence that characterises oligopoly.[10]

A third approach might be to explore the impact of algorithms on oligopolistic competition in the context of a less antagonistic setting, such as a market investigation. This approach might generate a deeper understanding of the impact of algorithms on competition without the sometimes oppressive influence exerted by the threat of heavy fines and liability to third party damages. But such enquiries can morph into unstructured fishing expeditions in which vague (if well-intentioned) attempts to intervene and remedy perceived problems are driven by an unpredictable array of policy objectives, and even political pressures.

Ultimately, the question is whether any mix of the above approaches to regulating pricing algorithms adequately addresses the objections that existed in the pre-algorithm world against attempts to regulate the evils associated with tacit coordination. The onus should be on regulators to explain how such powers might work before a leap in this direction can be justified. That would require two as-yet unmet conditions: first, a far more convincing story of the types of algorithms (and/or the market circumstances in which they are employed) that invariably lead to anti-competitive outcomes; and second, a clear policy rule that would successfully isolate such instances from the generality of pro-competitive applications of technology to pricing conduct, and identify a way to address those anti-competitive consequences without unintended chilling effects. A policy proposal that identified a clear competition problem, *and* specified an instrument that would be capable of fixing that problem without creating undue regulatory uncertainty, would justify serious attention. But the current spate of regulatory exhortations on the (possible) horrors of pricing algorithms fall well short of meeting either of these conditions.

Conclusions

Algorithms have a major influence on the way firms compete in today's economy, so it is appropriate for regulators and economists to study how they might change oligopoly conduct.

The adoption of pricing and other algorithms has undoubtedly led to huge increases in the information available to suppliers (and often to consumers too), and as such is likely to have exerted a powerful pro-competitive influence on many markets. Some aspects of algorithmic pricing nevertheless merit concern, where they facilitate coordinated market outcomes.

However, regulators' worst case fears of the impact of algorithmic pricing are not the same as a robust prediction that algorithms are inherently anti-competitive. If they were, it would be straightforward to define an administrable rule that dealt with the competition problem. But whilst the research in this area that has identified the *potential* for a link between algorithms and coordinated outcomes has been enthusiastically taken up by regulators in many jurisdictions, the uncertainty behind these adverse predictions, and the clear pro-competitive effects of these same tools, is often disregarded. The existing situation, in which some tacitly coordinated outcomes lie outside the reach of the main competition law prohibitions even if they can in principle result in the collective exercise of market power, is a reminder that there are real challenges to devising proportionate policy responses to address such conduct. It is therefore important that regulators do not over-reach either in the application of current competition laws to address the concerns that arise form algorithmic pricing, or in advocating new powers to address them.

10. For a broader discussion of whether it is realistic to expect firms to behave as if they competed in a one-shot game, when in the real world they typically face the same rivals on an ongoing basis, see the 2011 OFT Research Report on Conjectural Variations, at https://web.archive.org/web/20130818113513/http://oft.gov.uk/shared_oft/research/CV_Competition_Policy.pdf