The Economics of Algorithmic Pricing: Is collusion really inevitable?*

Kai-Uwe Kühn
School of Economics
University of East Anglia
DICE and CEPR

Steven Tadelis
Haas School of Business
University of California, Berkeley
NBER, CEPR and CESIfo

December 2018

*** Work in Progress ***

Abstract

Concerns have recently been raised that rapid technological development in artificial intelligence and ecommerce would help facilitate collusive behavior and threaten competition. We build on several strands of research in economics to assess whether algorithmic pricing poses new threats to competition, and, more importantly, whether new enforcement tools are needed to respond to these potential threats. We argue that the concerns raised are overly aggressive because of a basic and fundamental difficulty of achieving coordination in real-world pricing games, which cannot be easily overcome by algorithms. We further argue that current policy and enforcement tools seem broadly adequate to address the threats of collusion that can be expected to be present when pricing algorithms are used, at least until more rigorous research shows otherwise.
JEL classifications: D43, K21, L13, L41.

* We are grateful to Marc Ivaldi, Carl Shapiro and Andy Skrzypacz for helpful discussions and comments.

1. Introduction

The staggering pace of innovation in the past two decades has been remarkable - from ecommerce, ridesharing, and streaming entertainment, to security, home control systems, and digital assistants, among others. Consumers have seen expanded access to information and lower prices for many goods and services, as well as more free time as countless tasks have become more convenient and efficient. At the same time, arguments have been made that consumers are being exploited by large technology companies, from concerns regarding the potential loss of privacy in a data driven economy, to identity theft and other forms of internet fraud. There also appears to exist a general suspicion that the sheer size of firms signals dominance to the detriment of consumers and fear that firms are using data to manipulate customers' purchases or to engage in unfair price discrimination.

In this context, legal scholars have raised the possibility that the rapid technological development and increasing data availability would also facilitate the greatest threat to competition: collusive behavior. The proposition put forward first by Ezrachi and Stucke (2015) and Mehra (2016) and developed further in Ezrachi and Stucke (2016), is that pricing algorithms such as those used by e-commerce companies or hotel websites would greatly facilitate—and even make inevitable—the emergence of collusion among them. The vision proposed is that companies could now reach and sustain collusive outcomes without any actual agreement or human interaction using lightnening fast pricing and machine learning algorithms. Competition regulators have reacted by exploring the need of new tools to address something that is perceived as a qualitatively new threat (see, e.g., McSweeny and O'dea, 2017).

In this paper we argue that algorithmic collusion is far from inevitable because of the fundamental difficulty of achieving coordination (i.e. "reaching an agreement") in real world pricing games. In particular, we turn to several bodies of research in economics that explore both collusion and algorithmic learning in order to assess the merit of the recent claims that algorithmic pricing will eventually lead to anticompetitive behavior. This literature offers a coherent framework through which we are able to assess whether the claim that algorithms in general, and machine learning in particular, significantly increase the danger of collusive outcomes in online markets in ways that cannot be addressed through traditional enforcement tools.

Importantly, the economics literature that explores the factors that facilitate collusion *assumes* that parties can coordinate on a specific equilibrium, and the analysis is concerned with whether they can *sustain* it. However, a more recent literature in economics demonstrates that coordination problems are in general hard to resolve and in most circumstances, do not appear to be resolvable without explicit communication.

Furthermore, we explain how the empirical evidence often cited to illustrate the danger from the quick responses of algorithms does *not* demonstrate likely collusion. Several papers on algorithmic collusion cite incidents of how algorithmic interaction went terribly wrong and has led to ridiculously high prices. For example, two retailers selling on Amazon's marketplace offered Peter Lawrence's textbook "The Making of a Fly" in the spring of 2011 for $18,651,718.08 and $23,698,655.93 respectively. This was the result of two pricing algorithms by two competitors that each conditioned their price on the price of their competitor. In particular, one seller named *profnath* set its price to be 0.9983 times the price offered by a competitor named *bordeebook*, while *bordeebook* was setting its price at 1.270589 times the price offered by *profnath*, leading to a divergent outcome of exponentially rising prices.[1] Such effects of price algorithms are clearly not in the interest of consumers. But notice that they are also not in the interests of sellers who see their sales shrink to zero.

This and similar examples illustrate the casual use of economic concepts when speculating about the competition effects of algorithms. Instead of illustrating the danger of collusion, the example does the exact opposite: it illustrates the abject failure of the algorithms used to come to a viable collusive outcome. After all, collusion is not about reaching the highest prices possible, but instead is meant to achieve a higher profit than under competition, i.e., firms need to achieve *a significant amount of sales*. Charging prices that lead to zero sales is not in the interest of sellers.

In fact, what failed in this example was the basic requirement of *successful* collusion, namely, an effective coordination of strategies that would lead to profitable outcomes for both firms. Failure to coordinate tends to lead to considerably lower profits for the firms. What is different between human actors and algorithms is the type of failure that occurs. With humans, the failure to coordinate almost invariably leads to competition and low prices. Humans will not continue to raise prices when profits start falling as a result, but simple algorithms might. Indeed, simple algorithms may often not converge to an equilibrium and instead diverge. The incidence in the example above is a typical outcome of divergent behavior of algorithms in games. However, such divergent algorithmic behavior is precisely *not collusive* but shows a clear failure of the algorithms to achieve a coordinated outcome.[2] It is precisely the issue of coordination, which the largely

[1] See Olivia Solon, "How a book about flies came to be priced $24 million on Amazon," *Wired Magazine*, 04/27/11 (https://www.wired.com/2011/04/amazon-flies-24-million/)

[2] Note also that the problem of divergent algorithms on a standard sales platform does not necessarily raise any regulatory concerns. Competing firms have a strong incentive to intervene and reset their algorithms in such a case if they want to increase their profits.

legal literature on algorithmic collusion has overlooked, that is central both for effective collusion and for enforcement efforts.

Based on the body of research we examine, we conclude that the conditions needed to effectively achieve collusive outcomes in complex environments in which thousands of goods are sold to billions of online shoppers are demanding beyond what algorithmic pricing and machine learning can achieve without engaging in direct or indirect communication between algorithms. In order to use algorithms in support of collusive outcomes, the amount of data and code sharing that firms must engage in would, in our assessment, constitute an agreement generating sufficient evidence that would suggest a direct violation of antitrust laws even under the current legal regimes.

Our conclusions are not in conflict with the simple intuition that has been used to spell out the concerns regarding algorithmic collusion because repeated interactions among competitors *can indeed* lead to collusive outcome. In the light of the repeated games literature it is impossible to argue that collusive outcomes can *never* arise, and we do not make any such claim. In fact, the repeated games literature upon which the intuition is based shows that almost *anythin*g can happen in equilibrium. Instead, we argue that collusive outcomes are *improbable* in the absence of explicit coordinating behavior, because tacit collusion based on focal points is much harder to achieve than often claimed.

Importantly, the simple intuition conveyed by Ezrachi and Stucke (2015, 2016), that rapid algorithmic pricing responses will quickly punish deviators and lead to collusion, does not distinguish the "incentive compatibility" problem from the "coordination problem." The nature of repeated interactions with high frequencies of actions and sophisticated detection mechanisms like those that can be achieved with algorithms do indeed help solve the incentive compatibility problem so that *in theory*, collusive agreements become easier to sustain and enforce. We argue, however, that *in practice* explicit coordination is necessary to achieve any robust form of collusion. Moreover, in practice the world is significantly noisier than the perfect information models used to show that collusion is more easily supported by frequent and rapid response strategies. Sannikov and Skrzypacz (20??) show that in environments with noisy outcomes, more frequent interactions can in fact hinder the ability of firms to collude. And he real world is a lot noisier and complex than the model they use to make this point.

In contrast to the claims of Ezrachi and Stucke, it does not follow that fully rational behavior—or in case of firms, common knowledge of profit maximization of all firms (or all algorithms)—leads to coordination. In practice, coordination problems need additional mechanisms that allows every player to come to a common understanding of how everyone else thinks about the commonly pursued price setting mechanism. Put differently, the incentive compatibility problem, which is at the heart of Ezrachi and Stucke's arguments, is all about *sustaining* an agreement that was somehow coordinated on already. The coordination problem is concerned with how firms *reach* an agreement.

Our goal is thus to build on several strands of the literature to identify the scope for collusion in markets, with and without algorithms. To do this, we turn to the economic literature on collusion in repeated games, both from a theoretical perspective as well as empirical evidence generated from a series of experiments that have explored the ability of players to collude with and without direct communication and agreements. The extreme assumptions needed to achieve collusion in theory, together with the results from attempted collusion in experiments, lead to the conclusion that the concerns raised about algorithmic pricing are neither well developed nor very plausible. Also, the algorithmic game theory literature points to the complexity of solving for equilibria even of simple games, suggesting a true uphill battle for algorithms to collude in complex real-world pricing games.

We believe that our analysis offers an important and potent counter argument to the growing alarmist concerns regarding the use of algorithms coming out of the recent legal literature. Automated tools and algorithms enable firms, both large and small across many industries, to increase efficiency and grow their business. Moreover, algorithms provide consumers with vauable tools to purchase goods at competitive prices. For example, as discussed in Gal and Elkin-Koren (2017), price comparisons websites and search algorithms can help consumers shop for the kind of goods and services they are seeking, at competitive prices, and with a variety of quality reassurances. Ths casues demand for any signle firm to be more elastic, which in turn makes collusion harder to sustain. Furthermore, with personalized discounts that can be offered to consumers when they log-in for a more personalized shopping experience, "secret" price cuts nd be offered that further hinder the ability to collude.

In the real world of commerce, prices have by and large fallen over the past few years, and no evidence of tacit collusion by online firms has been successfully presented. In contrast, in those cases where algorithm-based collusion has occurred, the difficulties that the involved firms had in achieving coordination is testimony to the fact that spontaneous algorithmic-driven collusion is highly unlikely. In our view, regulatory overreaction to speculative concerns about algorithmic collusion, and the noticeable fears and suspicions that such speculation seems to generate, would risk chilling the plethora of consumer benefits of automated pricing tools, and negatively impact future innovation in the rapidly growing ecommerce markets in particular, and in other technology applications more broadly.

To clarify that the concerns about insufficient enforcement instruments are not warranted, we discuss in Section 5 how the insights developed in this paper imply that current enforcement rules can successfully deal with those threats that are, in fact, credible. That said, we do not advocate against doing nothing. What is necessary is to study the development of pricing algorithms to beter understand how they work, what the true dangers may be using rigorous theries of harm, and if needed, offer guidance for regulatory solutions. For example, in financial markets one may worry about algorithms destabilizing the market and creating negative externalities. This is not collusion, and hence falls outside the boundaries of antitrust regulation, but still presents a concern that regulators may wish to address.

2. What Theoretical, Empirical and Experimental Research Implies About the Likelihood of Collusion

2.1 The conditions for cartel stability only address one part of the collusion problem

For the purpose of our analysis, we refer to collusive outcomes as those achieved through successfully coordinated behavior. For firms to successfully coordinate on a collusive outcome, they must first *agree* on the collusive actions and profit distribution, as well as on the consequences of "cheating," namely deviations from the agreed upon behavior. An agreement means that it is not enough that players think that it would be a good idea to raise the price above the competitive price. The game theoretic models used to demonstrate the sustainability of collusion assume that players believe that other players have the same idea of what prices should be set and what would happen in the case of deviations. Moreover, the models assume that players have common knowledge of the prescribed coordination. That is, they believe that the other players believe that they all share the same beliefs, and that they all believe that they believe that they share the same beliefs, and so on. As we discuss later, this is a complex coordination problem that cannot seem to be resolved even by the most extreme assumptions on rationality.

In fact, by and large, the formal game theoretic analysis of cartels and collusion does not concern itself with *how* firms come to a common understanding of what prices should be set and what happens when someone deviates from the pricing norm. Instead, the theory spells out the conditions needed to sustain collusion *given* an understanding of an agreed upon pricing and distribution norm. In this section, we give a brief overview of this part of the theory to later distinguish it better from the discussion of the coordination problem that is at the heart of our argument on algorithmic collusion.

The conditions under which an agreement (whether tacit or explicit) can actually be enforced involve on one hand the incentives to deviate from the collusive pricing norm and on the other hand the ability and the incentives to “punish” such deviations. Collusion is possible if the short run incentive to deviate from the agreement (which is the force that normally generates competition) is sufficiently outweighed by the long run loss from switching from a high profit collusive equilibrium to a low profit competitive equilibrium. Theoretical models of collusive agreement assume a common understanding of how to coordinate the switch from one possible equilibrium to another.[3]

Virtually all game theoretic treatments of the collusion problem deal with the question of how each one of these two aspects of the incentive balance are affected by market conditions. The theoretical analysis suggests two intuitive market conditions that make collusion easier: first, that the market be more transparent, and second, that price setting be more frequent. Transparency is important because without transparency it is not possible to determine whether one of the parties to the agreement deviated from the collusive price and thus, without sufficient transparency, punishment cannot be triggered. But if punishment cannot be triggered, then the incentive to compete in the short run cannot be deterred and any agreement that is reached cannot be sustained by the parties. Frequent price setting is important because the more frequent the price setting, the quicker the punishment response to a price cut by a deviating firm, and hence the smaller are the short-run gains that any firm can obtain from a deviation.

Despite the simple intuitive arguments that point to the role of transparency and frequent response times in facilitating collusion, once the details of collusion theory are examined more carefully, subtleties emerge that question the validity of even the simplest intuitions. For example, perfect transparency is hard to achieve for many reasons, such as natural demand fluctuations. In an imperfectly transparent market, Abreu et al. (1991) have shown that *longe*r rather than shorter time intervals between price revisions can facilitate collusion—in the sense discussed here—because slower price changes increase the ability to detect true deviations and distinguish them from random demand fluctuations. This allows competitors to trigger punishment with a higher level of confidence that an actual deviation occurred.

An interesting and related paper by Sanikov and Skrzypacz (2007) shows that it is impossible to

[3] Standard theory assumes common knowledge of equilibrium. This does in of itself prove that it is a necessary condition for collusion. For example, one player may believe that both players are playing a “tit-for-tat” pair of strategies, while the other believes that they are playing “grim trigger”. These inconsistent beliefs may keep both players on a path of collusive behavior, because both believe that there is some kind of threat. However, a realistic argument in which uncertainty and environmental changes will prove both players wrong, in turn leading to surprises and unplanned actions, which destroys any uncoordinated collusive behavior.

achieve collusion in a duopoly when (a) goods are homogenous and firms compete in quantities; (b) new, noisy information arrives continuously, without sudden events; and (c) firms are able to respond to new information quickly. What's important in their setup is that the flexibility to respond quickly to new information unravels any collusive scheme that the firms may try to implement. So, the assertion that rapid responses help facilitate cooperation is far from proven fact, and rigorous theoretical research shows that the opposite may be the case.

Besides the two mentioned conditions of transparency and frequent responses that supposedly help facilitate collusion, the check list of collusion-facilitating market conditions is long (see Ivaldi, Jullien, Rey, Seabright, and Tirole 2003, or Kühn 2008). It includes other conditions like the number of firms (more firms undermine collusion); demand growth (future growth facilitates collusion); business cycles and demand fluctuations (both hinder collusion); multi-market contact (helps facilitate collusion under restrictive conditions); innovative markets (make collusion harder); cross ownership (inconclusive results); and many more. Some folk wisdom has turned out to be incorrect (e.g. that collusion is easier with homogeneous goods, see Kühn and Rimler, 2009). But there is certainly one feature that consistently makes collusion harder: asymmetry between different firms in the market. The basic reason for this conclusion in the theoretical literature is that when firms are different, they face different incentives. The more different firms become, the harder it is to satisfy all the firms' incentives at the same time. For example, even if two firms sell the same product, they might have different inventory levels and hence, each may want to set different prices to clear different levels of inventory.

In contrast to the theoretical literature that highlights these incentive problems, the empirical literature has put the emphasis on the fact that asymmetries make it very hard to come to an agreement or common understanding in the first place. (See Levenstein and Suslow, 2006). Not only is there a question of coordination, but it is also unclear what prices should be set by each firm in a setting where there may be very different benefits of different price levels for the potentially colluding firms.

Regardless whether the difficulty is with reaching an agreement, or with providing the incentives necessary to sustain an agreement, from both aspects the conclusion is that asymmetry between firms severely limits the ability to collude, whatever the level of market transparency or the frequency of interaction. This result is of particular interest for algorithmic collusion because most competing digital firms differ from each other along many dimensions such as their cost structures, their product lines, their business models, their brand, etc. The literature therefore suggests that in this environment, collusion would be particularly hard to sustain, and unlikely to occur in the first place.[4]

It must also be stressed that the long list of factors that "facilitate" collusion only do so in the sense of relaxing the incentive conditions for maintaining a collusive pricing norm. They say nothing about the likelihood of collusion *actually* being achieved and cannot be taken as predictors for the likelihood of collusion *actually* occurring. The reason is that they indicate that collusion is more likely to be sustained (or make it sustainable at a higher price) once an agreement *is* reached, but they do not say much about the *likelihood* of an agreement being reached. This is stressed in Kühn (2008), who argues that the absence of market transparency and the observation of infrequent transaction could only be used as a negative test for collusion: markets with little transparency and with infrequent transactions make successful collusion unlikely even in the presence of agreements. However, the reverse does not follow. Even if markets are perfectly transparent and interaction takes place almost instantaneously, collusion does not necessarily arise.

In the next two subsections, we show that the economic literature in the last 15 years has accumulated consistent evidence for this observation, underlining the fact that coordination is hard to achieve and tacit collusion is therefore not a likely outcome – even in transparent markets with frequent interactions between firms. It is therefore central to the assessment of collusion to explore how coordination is actually achieved and in which sense algorithms have real advantages over humans in achieving it.

2.2 Perfect transparency by itself does not lead to collusion

Competition policy practitioners often regard price transparency of markets as a competition reducing feature. For example, in German antitrust proceedings there is often the claim that certain transparency measures undermine the strength of "secret competition" ("Geheimwettbewerb") and are therefore anticompetitive. This notion is not well grounded in economic theory. In models of information exchange with linear demand, expected prices are not systematically affected by an increase in information exchanged.[5] This suggests that there is not a first order effect of market transparency on the price level.

[4] Of course, asymmetries do not prevent firms from reaching agreements that increase joint profit, if they can make side payments and have enforceable contracts. But this of course requires exactly the kind of communication and coordination that current law classifies as illegal, which is not made easier by the use of algorithms.

[5] See Kühn, Kai-Uwe and Xavier Vives (1994), "Information exchanges between firms and their impact on competition", Report to DG IV, European Commission, Office for Publication of the European Communities, for an overview on this very large literature.

The only way in which market transparency can systematically increase prices is by facilitating the detection and punishment of deviations from a previously agreed upon collusive behavior. In the absence of cartel agreements, we should therefore only worry about market transparency if tacit collusion, i.e., the coordination on collusive behavior without communication and explicit coordination, were easy and would spontaneously arise.

In principle, the assumption that individuals can solve coordination problems and coordinate on one among a number of many equilibria is a very strong assumption. It does not follow from common knowledge of rationality, which itself is already a strong assumption. In fact, the common knowledge of rationality assumption does not resolve the problem of coordination, allowing a set of multiple prices to be "rationalizable" including the prices set in any Nash equilibrium.[6]

Nevertheless, some economists have suggested that the coordination problem could easily and spontaneously be solved. The view that coordination can be easily achieved and that tacit collusion is fairly pervasive has been part of the economics folklore for a long time. This general view goes back to a large literature that has attempted to refine the Nash equilibrium concept to come to a single prediction even when multiple equilibria exist. For example, Schelling (1960) in "Strategy of Conflict" suggested that coordination problems can often be resolved because obvious "focal points" will be recognized as such by individuals. It will be obvious to them that others will think in exactly the same way about the problem.

One rule for generating focal points that has immediate appeal is the Pareto criterion. In game theory, it has often been assumed that Pareto optimal outcomes should be focal and that one could simply refine sets of equilibrium outcomes to focus on Pareto optimal equilibria among the players.[7] Such a focus on Pareto optimality was criticized by Harsanyi and Selten (1988) who pointed out that some equilibria tend to be risky because the costs of non-coordination are particularly high should a player choose a strategy associated with the proposed Pareto optimal equilibrium. If the other players did not follow the reasoning that the Pareto optimal equilibrium

---

[6] Generally, in supermodular games (in which equilibria can be ranked by all the players' payoffs), all actions between those of the highest and lowest equilibrium are consistent with common knowledge of rationality (see Milgrom and Roberts, 1990). Note that this means that even in classes of games which always produce ranked equilibria in terms of profits, common knowledge of rationality does not lead to a resolution of the coordination problem.

[7] There are different versions of this notion, e.g., focusing on Pareto Undominated equilibria as in Bernheim and Whinston (1998) and the concept of "perfectly coalition-proof equilibria" of Bernheim et al. (1987). But even Pareto optimality will not resolve the conflict of interest inherent in collusion games, because there are many equilibria that are Pareto optimal among the colluding parties with very different distributions of the surplus. The problem we point out in the text is, however, even more fundamental.

should be played, the payoff consequences could be particularly dire to the player who does follow the "focal" strategy. Harsanyi and Selten (1988) called equilibria for which this was a problem "risk dominated."

While the analysis in Harsanyi and Selten (1988) is still in the tradition of attempting to distill an equilibrium refinement to select a more robust outcome, the existence of multiple relevant criteria for how to play a game with multiple equilibria seriously undermines the idea that players can spontaneously select among multiple equilibria without an explicit coordination mechanism. Spontaneous coordination would rely on simultaneously having common knowledge of not only rationality, but also of the assessment of strategic risk in such situations. This level of sophistication and conformity is obviously unrealistic. Furthermore, even moving from human actors to algorithms, there is no theoretical guidance on how algorithmic players select one of many possible equilibria. Just as humans cannot independently select among equilibria without communication, the same holds for algorithms.

Note that a profit-maximizing algorithm would not have information about other algorithms, in particular whether they are designed to be profit maximizing as well. This means that even an algorithm cannot resolve the multiple equilibrium problem without further help. Algorithms cannot "try out equilibria" as is often incorrectly stated in public discussions on the matter. As an independent algorithm, each would have to unilaterally choose decisions and strategies. Suppose that an algorithm was programmed to always try the strategy first, that corresponds to a Pareto optimal Nash equilibrium. What should the algorithm do if the other algorithmic (or non-algorithmic) player does not play the corresponding expected strategy? It depends on what this algorithm predicts about the behavior of the first algorithm. It could predict that the first algorithm tries the action again and as a result play its action that corresponds to the Pareto optimal equilibrium. But it could also believe that the first algorithm made a very low profit relative to having played a different strategy and will now adjust to increase profits. Both strategies would be perfectly rational. If the second algorithm predicts that the first algorithm will try something new, there is no point in matching the equilibrium price of the previous period. In fact, it may be very costly. This is what is referred to in game theory as *strategic risk*.

The problem with a lack of common knowledge of the strategies between the two algorithms is the same as that for two rational individuals. With multiple equilibria, it is impossible to forecast what the other player (whether human or algorithmic) will do, and if setting high prices is riskier than slightly lower ones it is difficult to coordinate on high prices spontaneously.[8] Obviously, what

[8] Note that we are not claiming that coordination without explicit communication can never happen. For example, if there has been a history of explicit collusion in an industry then it is perfectly reasonable to expect that market

is missing for effective collusions is for algorithms to know each other's price adjustment rules to allow them to reassure each other that they will stick to them. But since at least one algorithm has to adjust the decision rule, the mechanism by which the transition to coordination is achieved must lie outside of playing the pricing game.

The importance of strategic risk for the ability to coordinate on Pareto optimal equilibria has now been amply demonstrated in the experimental literature. Cooper et al. (1990) and Van Huyck et al. (1990) first demonstrated this in very similar games in which all equilibria are Pareto ranked, but the higher profit equilibria are risk dominated. In these games play converges to the worst of the Pareto ranked equilibria when there is no communication. In particularly interesting work, Brandts and Cooper (2007) studied how such coordination problems can be overcome in manager worker relationships. In their settings, experimental subjects experience coordination failure. They then explore potential mechanisms to lead players back to playing better equilibria. Brandts and Cooper (2007) show that even though incentive schemes that pay bonuses for coordinating on the good outcome were somewhat helpful in improving outcomes, the main solution to the coordination problems was communication. In particular, whenever the manager gave clear instructions to play the Pareto dominant equilibrium, workers were able to return to this outcome. The reassurance of explicit communication and clear agreement appears to be the main feature necessary to overcome even some of the simplest coordination problems in one shot games. The reason is that only such reassurance outside the game can resolve the indeterminacy of how rational profit maximizing individuals (or algorithms) can rationally determine their expectations about the behavior of competitors.

With respect to strategic risk, it is important to note that there is nothing that an algorithm can do that improves on the play of individual players. Algorithms that are designed to maximize expected profits must avoid costly strategic risk precisely because they are profit maximizing. It is not even enough to commit as an algorithm to the action that corresponds to the best equilibrium. It needs to be credibly announced that such an action will be played. Note that this is equivalent to what we learned is necessary in experiments to induce coordination. Clear communication from one party over what ought to be played generates coordination on the outcomes. However, this is precisely the type of coordinating communication that is generally illegal in antitrust law.

It has to be stressed, however, that coordination problems that involve selecting among collusive equilibria are far more complex. First, they involve games with repeated interaction. In order to make collusion incentive compatible, players need to coordinate on complex contingent

participants might see this as an obvious focal point. However, in that case the coordination problem has already been solved by prior communication.

strategies out of an infinite set of strategies. In fact, a collusion game generally has an infinite number of possible equilibrium outcomes, which can generally each be supported by an infinite number of strategies. Though it may be surprising that many of the older experimental studies on collusion have shown the emergence of tacit collusion, it has now become apparent that these results arise from the great simplicity of the constituent one period stage games that were studied.

In virtually all of the studies that find some significant collusive effects there are only two players who have only two actions available (see, e.g., Camera and Cassari, 2009). It is not too surprising that such settings encourage players to deviate from competitive behavior simply because there are always enough other players who try the "collusive" action. Individuals then get enough experience of jointly trying the collusive action to learn that it is jointly profitable. However, even in this setting one of the most careful studies by Dal Bo and Frechette (2011) has demonstrated that collusion is not easy to obtain. Unless players are extremely patient and collusion incentives are very large, players do not play the collusive outcome even where it is incentive compatible.

More recent research has debunked the notion that the collusive results from two by two games carry over to any setting that even remotely resembles reality. The first clear result on the effect of larger numbers of participants in oligopoly settings was established by Huck et al. (2004). They showed that in experimental duopoly Cournot (quantity-setting) games with perfect information output is reduced by about 10% below the one-shot Nash-Cournot prediction. But with 3 and 4 players no statistically significant deviation from the one-shot Nash-Cournot equilibrium can be detected, and for more players, repeated interaction led to even *higher* production levels. This is consistent with the theoretical result obtained by Vega Redondo (1996) regarding how individuals learn from how to play the Cournot game from observing the performance of others.

Fonseca and Normann (2011) run similar experiments for a simple Bertrand (price-setting) game where buyers have a constant valuation of 100 for the product. In the two-player game the experimental firms manage to set prices at 50 on average. But with 3 players the average prices drop to 6.3, and from that point onwards, prices correspond to no more than the typical noise in experimental behavior consistent with the competitive outcome.

But even such results have to be interpreted carefully since they assume that the non-collusive behavior would be well predicted by the Nash equilibrium of a one-shot game without repetition with the same players. It turns out that in Bertrand and Bertrand-Edgeworth games, players do not play as aggressively as the theory predicts. In ongoing work, Kühn and Normann (in progress) have conducted experiments aimed at studying the coordinated effects of mergers in a Bertrand-Edgeworth setting. While the behavior with repeated interactions in a symmetric triopoly is on average slightly above the prices one should observe in a one-shot game theoretically, the average price (and the distribution of prices) is indistinguishable from that of experiments for the same game played without repetition between the same individuals. No collusion effect and therefore no implicit coordination could be detected. This was even true for duopoly experiments.

It is important to note that these results, which show the difficulty in achieving tacit collusion in experimental settings, have all been shown in a context in which there is full transparency of market conditions. Every action that is taken by a competitor is seen by the other participants and can be reacted to in the next period. “Signaling” of pricing is just as much a possibility as “punishment” of a competitor who sets a low price. The games are often long. Even in games with more than 50 periods (followed by random ending of the game) tacit collusion is not reached. The spontaneous coordination on collusive prices and punishments for deviations therefore does not appear to arise in experimental settings outside some extreme over-simplified cases. In the real world, where the levels of complexity are high and algorithms have no direct insight into the ways in which their competitor algorithms are coded, getting to a collusive outcome would be significantly more challenging.

2.3 Short response times to price cuts of rivals do not lead to collusive outcomes

It is often argued that tacit collusion becomes easy when competitors can react to each other very quickly, which is emphasized in the arguments set forth by Ezrachi and Stucke (2015, 2016). The idea is that undercutting of a collusive price only generates a very short period of benefits because the reaction of the rival leads to punishment that will be relatively swift and, hence, the punishment outweighs the benefits. But as argued earlier, this intuition applies to the incentive condition for *sustaining* collusion. It is not clear that in the absence of a coordination mechanism the frequency of interaction plays any role at all in increasing the *likelihood* of collusion arising in the first place.

It is also worth noting that the experimental literature described in Section 2.2 above suggests that rapid interaction does not lead to collusive outcomes. The papers cited have extremely short reaction times of no more than 2 minutes because otherwise experiments with a sufficient number of repetitions could not be played within a reasonable amount of time. Furthermore, the experimental repeated games have minimal complexity in the market environment, which means that no time is needed to infer whether participants have deviated from collusive behavior. Even under these extreme conditions tacit collusion hardly arises spontaneously although the interaction is both rapid and the actions of players are perfectly observable.

But the lab settings are possibly not what those who believe that tacit collusion is easy in real markets have in mind. In these settings, participants in the experiments who represent firms set their prices simultaneously in every period. But many believe that “price leadership” would be extremely effective in resolving the coordination problem, so allowing players to play asynchronously may be important in helping sustain collusive outcomes. In fact, sequential moves are typical in markets in which there is rapid interaction of players. So the natural question is whether a setting with sequential move orders allows for tacit collusion that previous experimental studies overlook. Perhaps there is a simple indirect coordination mechanism through sequential price setting that supports collusive outcomes.

The gasoline retailing industry offers an excellent field setting in which to shed light on these questions. Over the years, there has been considerable concern that gasoline retailing might be highly collusive due to the ability to rapidly respond to price cuts and the high transparency of prices at the pump. Indeed, Ezrachi and Stucke (2017) have recently claimed that in gasoline markets with their rapid responses, increases in transparency have increased prices and margins in Chile, Germany, and Australia. Unfortunately, their discussion of gasoline pricing ignores much of the literature on the subject and relies exclusively on unpublished working papers that have yet to undergo peer review. The accumulated published literature from Canada, the US, and other countries overwhelmingly does not show increased prices and margins in gasoline markets that have rapid responses and are highly transparent.

There is now a large empirical literature on gasoline pricing. In this literature, it has become apparent that there is a fundamental difference between markets that exhibit asymmetric price cycles and those that either have fairly rigid prices or follow gasoline wholesale prices at fairly constant margins. Price cycles involve sharply increasing prices after which firms tend to undercut each other until prices get close to marginal cost, after which a new price jump is triggered. Some competition policy agencies and courts have sometimes interpreted price cycles as indications of tacitly collusive behavior.[9]

However, the view of price cycles being evidence of collusive behavior has come to be firmly rejected by the accumulated empirical evidence. Following Eckert (2002, 2003), Noel (2007a,b) was the first to show that price cycles occur in particularly fragmented and competitive markets and that average margins are particularly low in cycling markets. Further evidence that exploits natural experiments has shown that after short run supply shocks prices fall much faster in cycling markets due to the higher competition (Lewis 2009, Noel 2015), that cost pass through is higher

[9] See final report Bundeskartellamt, “Sektoruntersuchung Kraftstoffe – Abschlussbericht” of the German Federal Cartel Office into gasoline retailing in Germany. The claims in that report are not consistent with the economic literature.

in markets with cycles (Lewis and Noel 2010), and that customers mostly buy at the bottom of the cycle (Lewis and Wolak, 2017).

While many of these studies have been applied to U.S. and Canadian markets, there is similar evidence from many countries across the world. In Germany, for example, the evidence at the time of the sector inquiry of the Federal Cartel Office was that the pricing cycle length was no longer than a day and that cycling behavior had rapidly increased over time, which is expected if markets become more competitive. This implied that price revisions were faster than the changes in the wholesale prices faced by the gasoline stations. Being concerned about the price cycles as an indication of collusive price increases, the German Federal Cartel Office believed it could "improve" competition by making prices more transparent to consumers. Since the introduction of this increase in market transparency of unknown size (the market was highly transparent even before this), the frequency of cycles appears to have further increased.[10]

It is striking to note that the markets with the highest frequency of price cycling consistently experience the lowest margins and the highest responsiveness to supply shocks. The empirical studies therefore do not bear out the proposition that collusion is easier to achieve with a higher frequency of interaction.

The claims that the highly cyclical markets exhibit collusive behavior is also not plausible theoretically because it is known that collusion is much more difficult to sustain when the colluding firms are expecting declining prices (Haltiwanger and Harrington 1991). Hence, if the price increases were collusive, they would be particularly difficult to achieve because of the downward part of the cycle that immediately follows the price rise. Collusion would be much more easily sustainable if prices were kept constant and deviations punished by dramatic price decreases. Regular price cycles would be the least effective way to collude that we are aware of. It is therefore neither theoretically nor empirically plausible that gasoline retail markets are an example for rapid price setting and high transparency leading to tacitly collusive outcomes.

Note also that there are non-collusive theories that explain cyclical pricing in these markets from short term rigidities in price setting and unilateral market power. In particular, Maskin and Tirole (1988) have predicted in a duopoly setting the basic type of cycling that we see in gasoline

[10] Ezrachi and Stucke (2017) cite some papers selectively that suggest an increase in prices in gasoline markets in which transparency was increased. These papers seem to use the "difference in difference" technique to estimate the changes. However, the standard "common trends" assumption needed for identifying causal relationships using the difference-in-differences technique do not seem to match the setting in which these are applied given the very different trends in gasoline pricing that are due to local economic conditions in different European countries. In contrast, the literature we cite above from leading journals is extremely careful in showing what variation in the data identifies the claims of the authors.

markets.[11] While these equilibria sometimes generate very significant margins (unlike the experience in gasoline retailing), these equilibria are still the most competitive ones that can be generated in the setting with short time rigidities in price revision. What is important, though, is that neither increased transparency nor increased frequency in price setting increases average prices and margins in these settings. Neither does collusion emerge in these markets as a result of sequential price setting, nor do price setting rigidities generate higher prices under frequent price revision and high transparency settings. The claims in Ezrachi and Stucke (2017) basically contradict both theory and the empirical evidence from gasoline markets.[12]

2.4 Is strategic algorithmic price setting different?

We have shown based on both experimental and empirical literatures that overwhelming evidence points in the direction that tacit collusion is rare and difficult to achieve. There is a good theoretical reason for it: achieving a full common understanding of play is part and parcel of coordination, but coordination over infinite strategy spaces is not easily signaled through market behavior or obvious in terms of profit maximizing behavior. Indeed, the empirical evidence to date does not support the notion that quick responses lead to effective collusion, which undermines the basic premise of the arguments that tacit collusion is easy and therefore also easily achievable by algorithms. Rapid interaction and high market transparency do not resolve the *coordination problem*.

Since the evidence shows that rapid response is not what makes collusion easier, then the question becomes, what is it about artificially-intelligent algorithms that would make them more effective than humans as strategic players? (We discuss the challenges of non-strategic learning algorithms in section There is no reason to believe that algorithms are able to create focal points that are obvious to all players, including algorithmic ones. The basic issue is therefore that, acting through the lens of standard non-cooperative game theory, an algorithm has to change the beliefs of another algorithm (or many other algorithms) about its pricing strategy. The question is how an algorithm can simply through its pricing on the market signal a contingent strategy in an infinite dimensional space to another algorithm? Fundamentally, an algorithm cannot produce signals about the reaction to pricing events of a competitor if the competitor never uses those prices. However, the beauty of "cheap talk," i.e., conversations, is that hypothetical situations

[11] See Noel (2008) for extensions to triopoly, varying marginal costs, and other robustness analyses.

[12] The one theoretical paper we are aware of that claims collusive price outcomes due to quick responses is Bhaskar (1989). The paper attempts to formalize the intuition that quick responses lead to no competition by modelling the price setting game as a possibility of changing prices rapidly *before* any sales can occur. The crucial piece in this paper is that price adjustments stop as soon as one of the players does not change the price anymore. In that setting with no sales until price adjustments stop, the monopoly price is the only equilibrium outcome. This essentially imposes a commitment power to monopoly pricing that is not present in real markets. Essentially, Bhaskar (1989) models negotiation over the price charged and imposes that the agreement is implemented without the availability of a punishment mechanism. But these assumptions effectively introduce both a coordination mechanism and an enforcement mechanism as part of the model.

that matter for behavior but (almost) never arise can be taken account of in an agreement. This fundamental issue makes direct communication beyond price setting an extremely important instrument to achieve coordinated outcomes, which is discussed in the next section.

3. The Role of Communication for Coordination

As we demonstrated above, a series of experimental research papers provides robust evidence that tacit collusion appears almost impossible to generate in laboratory settings outside of classes of very simple stage games that do not capture even the most mundane complexities that characterizes real markets. We also described the evidence that argues against collusive arrangements in gasoline markets which have been thought by some to be subject to tacit collusion due to market transparency and rapid price changes. Coordination on collusive outcomes is therefore hard to achieve spontaneously in the absence of explicit agreements.

The recent research based on laboratory experiments has, however, also shown that the barriers to parties colluding on coordinated outcomes appear to collapse once communication is introduced. In Huck and Normann (2011), for example, once the parties are allowed to communicate then price setting with two players leads on average to above 90% of the monopoly price, almost the perfect collusive outcome. In the same settings, prices on average reach almost 80% of the monopoly price with four players, and exceed 50% with 8 players – a dramatic difference from the behavior without communication. Collusion theory appears to perform well when coordination through communication can take place – in sharp contrast to situations in which coordinating communication is not allowed.

The role of communication has been studied in more detail by Cooper and Kühn (2014), who have carefully evaluated the type of communication that is effective for collusion to be successful. They show that agreements on the collusive price are not decisive. Most of the time, experimental subjects who can talk, quite easily coordinate on the best collusive price when they can talk. However, what is central to collusion is that subjects have a clear view of how other parties to the agreement will react to possible cheating.

There are two mechanisms by which communication is effective in Cooper and Kühn (2014). First, the most reliable factor to achieve collusive outcomes is communication in which clear punishments are threatened when the explicit agreement on price is violated. It turns out that such agreements are only concluded after longer conversations in which it is clarified how the other party thinks about the problem and whether they actually really mean to punish deviating behavior. This fact shows how difficult it is even in regular language (which is in itself developed to resolve coordination problems) to ensure a common understanding of how players are expected to coordinate their behavior. But this observation is critical -- simple agreements just on the price to be set tend to be cheated on with high frequency.

The second mechanism that very strongly supports collusion is repeated conversation and feedback about past behavior. In particular, there is frequent and intense verbal punishment by players who complain about the cheating by their counterparts. The fact that such verbal expression is associated with much more collusion suggests a high degree of norm driven behavior and the effectiveness—in a large part of the experimental population—to react very strongly to social disapproval and approval. This responsiveness and social reaction to deviating behavior is much more powerful because it occurs without too much explanation about the actual punishment mechanism. Social punishments and rewards appear to be understood easily and are easily transferred to the collusion setting. These are clearly behavioral responses that make coordination through communication extremely successful in Kühn and Cooper (2014).

It should be noted that this result contradicts the conjecture in Ezrachi and Stucke (2017) that algorithms collude more effectively because they are not subject to human biases. The results reported above imply that the opposite is the case. It is precisely human biases that facilitate collusion in the experiments that have studied collusive behavior in the lab.

If it is so difficult for humans in real market and as experimental subjects to coordinate on collusive play spontaneously and without communication, what makes algorithms so different? Do they need to communicate about a common collusive norm to achieve coordinated outcomes or do they have other means to identify focal points (in the language of Schelling) that are not available to humans? How do algorithms reproduce the affect against cheating that appears to stabilize collusion in experimental settings? These questions seem even more salient because experimental subjects are often modeled to behave as if they were some form of adaptive learners in the lab. The reason is that they clearly need time to learn how to play the game, and their learning behavior appears to be adaptive. Of course, by their nature, algorithms are adaptive agents, and human players do have some form of a cognitive algorithm that they follow. Why, then, with behavior in the lab that is far from collusive, would automated algorithms be expected to automatically converge to some collusive outcome?

We show in the next section that the assertions by Ezrachi and Stucke (2015) and other authors in the law literature appear to rely on a misunderstanding or misinterpretation of how algorithms function and learn.

4. Challenges to Collusive Algorithms

As explained in the previous section, the conditions needed for tacit collusion to successfully be executed require not only sophisticated players that are protected from entry, but market conditions that are somewhat pristine in that there is minimal complexity in a stable and unchanging environment. But even in stable environments, rigorous theoretical research shows that adaptive behavior that is aimed at profit maximization will not easily – if at all – lead to collusive outcomes. In this section, we describe some of the most relevant theoretical models that bring to bear on the question of tacit collusion in pricing games.

4.1 Profit maximizing learning algorithms do not easily lead to collusion

Pricing algorithms are by their nature adaptive mechanisms that take as input the history of past outcomes and then devise best responses to these histories, either through hard-coded contingency plans, or using a variety of machine learning models. In order to predict whether adaptive mechanisms will result in collusive or competitive outcomes, it is paramount to understand how the nature of the game might cause adaptive mechanisms to perform.

In an extremely influential paper, Milgrom and Roberts (1990) study how adaptive behavior will lead to outcomes in an important class of games called "supermodular games." These are games that share an interesting and economically meaningful feature, namely, that each player's best response is increasing in the actions taken by its rivals. To see how this general idea applies to pricing games, if firms A and B are competitors, and firm B's raises its price (i.e., increases its action) then the best response of firm A is also to raise its price (its best response increases). Hence, a canonical price-competition game, often referred to as a "Bertrand competition game," falls neatly into the class of games that Milgrom and Roberts (1990) study.

Their study focuses on how outcomes will evolve from a wide class of adaptive learning mechanisms, which range from more naïve adaptive procedures to more sophisticated ones.[13] What their study reveals is that in supermodular games, these adaptive processes must lead to outcomes that, loosely speaking, are contained within the set of outcomes that can be justified by players playing a Nash equilibrium of the static game, as if it were not repeated. In particular, for a wide class of common demand functions (including linear, logit, or constant-elasticity of substitution specifications), they show that only one pure strategy Nash equilibrium in the

[13] These include (1) "best response dynamics" where each player in each round expects his competitors to do the same thing they did at the last round, (2) "fictitious play," where players assume their competitors play mixed strategies that coincide with the historical empirical distribution of the past play, and (3) "Bayesian learning," considered to be in the paradigm of rational learning.

(possibly asymmetric) Bertrand pricing game can be the outcome of an adaptive process. Furthermore, they show that any adaptive dynamic process leads to behavior that converges from *any starting point* to the unique equilibrium. These results provide a very strong foundation for adaptive processes leading to a non-cooperative equilibrium.

The approach taken by Milgrom and Roberts (1990) was to explore whether short run rational behavior (profit maximization while believing that everybody can be assumed to be a profit maximizer) as postulated in game theory could be learned by repeated play of the game by boundedly rational individuals.

One possibility of how such learning could be thrown off would be by one of the most well-known results in the literature on repeated games and collusion: the celebrated "Folk Theorem." It states that many different outcomes, with competitive pricing at one extreme to monopolistic pricing at the other, can be supported as equilibrium outcomes if players are patient enough and value future profits. This plethora of possibilities was a challenge to theory, because it implied that it had little predictive power. Instead, it was a "possibility" theory, which basically said that many things can happen when players are engaged in repeated play. In fact, the theory allows such a wide range of equilibria partly because players are assumed to be able to coordinate their expectations over an infinite future horizon and take infinite past histories of play into account.

Milgrom and Roberts (1990) restrict attention to more realistic modes of adaptive behavior than those studied in the classic collusion theory. They take seriously the limitations of adaptive learning that algorithmic pricing would have to abide by. Namely, algorithms result in prices adapting over time in a way that responds to the evolution of history, and they cannot create strategic replies to infinitely long histories (a limitation explored in Complexity Theory). For these reasons, the Milgrom and Roberts (1990) model provides an excellent framework through which to explore the workings of algorithms.

In summary, the important results derived by Milgrom and Roberts (1990) strongly suggest that an algorithm, which is basically playing adaptive dynamics of one sort or another, and is unable to create strategic reliance on infinitely long histories, would actually *not* converge to a collusive outcome and instead *will* converge to a non-collusive equilibrium outcome. Hence, much more restrictive environments, or strategic play that goes beyond adaptive procedures, are needed for algorithms to be able to converge to an outcome that is collusive.

4.2 Algorithmic collusion requires significant communication

A recent attempt to put rigorous structure to the conjecture that pricing algorithms will lead to collusive agreements was made by Salcedo (2015), who develops a model of competition between two competitors in continuous time, in which the two firms use algorithms (modeled as finite automata) to set prices over time. The paper follows a long tradition in economic theory that is aimed at uncovering the necessary assumptions needed to support a result of interest. Although the paper is written as an attempt to show that conditions exist under which a collusive agreement is inevitable, a careful consideration of the conditions needed to guarantee the collusive outcome shows how implausible these conditions are, and hence, how unlikely it is that tacit collusion would be the predicted outcome of algorithmic pricing.

In the model developed by Salcedo (2015), pricing algorithms respond to demand conditions and to a rival's prices. Importantly, at some exogenous and unpredictable times, the current pricing algorithm of each firm becomes completely apparent to the other firm. Shy of literally exchanging code, this assumption alone makes any serious consideration of the practicality of the result a futile exercise. Even the most sophisticated machine learning algorithms lack the ability to explore all possible contingencies.

Furthermore, there is a particular sequence of timing and events needed in the model for the result to be established. Namely, it is critical that firms take longer to learn about their rival's algorithms than the arrival of stochastic shocks to the demand from consumers. However, to learn about demand may itself be challenging, especially in an environment with demand uncertainty. For example, if demand follows a random process that is influenced by consumer tastes, income shocks, and the introduction of other products, the time to learn demand may or may not be shorter than the time it would take to learn about a rival's algorithm.[14]

One of the key insights from Salcedo (2015) is that it is in each firm's best interest to make its algorithm transparent to its rival so that it can be decoded easily and as quickly as possible, which helps the firms to coordinate on collusive outcomes. Each firm also has an incentive to create some commitment to their own algorithm and not be able to revise it too quickly, as this creates a form of market power in the game between the rivals. This implies that the best way for firms to create collusive agreements is for them to be completely transparent and share their code, while at the same time create some frictions in their abilities to revise their codes. The latter is a consequence of the needed assumption noted above that it is critical for firms to take longer to learn about their rival's algorithms than the arrival of stochastic shocks to the demand from consumers.

---

[14] It is important to note that in environments with stochastic demand, which is the realistic assumption to make, tacit collusion becomes significantly more difficult and price wars are inevitable. See the seminal work by Green and Porter (1984).

The need for complete and transparent code sharing suggests that current competition policy has significant bite. There is currently consensus that the private exchange of information about future pricing is a violation of antitrust law both in the U.S. and in Europe. If not treated as a cartel, it would generally be treated as either per se illegal or as an infringement by object (in the European context). The communication of an algorithm to a competitor would arguably fall under such a prohibition. Sharing algorithms does not just share plans about future pricing, but even shares the *rules* by which every possible future price change can be inferred. Sharing a pricing algorithm with a competitor therefore would logically fall a fortiori under such a rule.

In cases where algorithms are programmed to explicitly collude, or collusion is facilitated through a third-party algorithm provider (i.e., hub-and-spoke), collusion can be expected to be even easier to detect than before. Creating explicit programming in an algorithm will be known by a larger number of individuals than the key account managers that are privately involved in many cartels. The likelihood for information leakages is therefore high and the evidence for programmed cartel behavior relatively easy to verify.

To be successful, algorithmic collusion would still involve illicit communication, a "meeting of the minds" on final strategy, and the algorithms would need to be designed to send signals and to enforce these agreements in automated ways; invariably involving more people, greater resources, and a longer "paper trail" than non-algorithmic collusion. In fact, competition enforcers have already successfully detected and prosecuted price fixing agreements that were enforced through the use of algorithms.[15]

In fact, the difficulties in achieving coordination in the presence of algorithms is nicely illustrated by a recent case at the Competition and Markets Authority (CMA) in the UK.[16] The case concerns an agreement between Trod Ltd and GB eye Ltd not to undercut each on licensed sports and entertainment posters and frames on the UK version of Amazon Marketplace. The decision describes the difficulties in implementing this agreement for the parties. In fact, the parties found it hard to figure out whether the other party was cheating on the price matching agreement or whether problems with the software and their implementation of the agreement were to blame for deviations from the rule of not undercutting each other. The technicians literally had to communicate extensively about how their algorithms worked in order to generate a workable

---

[15] This is what happened in the case brought in 2015 by the U.S. DOJ against Topkins for colluding on prices of posters sold online. Topkins and its co-conspirators designed and shared their pricing algorithms, which were programmed to collude on prices according to their agreement. https://perma.cc/QMT6-ZQMN

[16] The decision on "Online sale of posters and frames, Case 50223" can be found at https://assets.publishing.service.gov.uk/media/57ee7c2740f0b606dc000018/case-50223-final-non-confidential-infringement-decision.pdf

collusive outcome. Exchanging information about the pricing rule and the algorithm was required to come to an understanding. The problem was exactly what one would expect based on theory. The pricing patterns observed cannot be inverted to find the underlying price setting strategies or infer the workings of the algorithm. A meta-language beyond price as a signaling device is necessary for coordination, i.e., a language that allows communication about the algorithms (i.e., the strategy itself) in order for coordination to work.

Some of the legal literature has speculated that in the future, AI could simply develop such a meta language without external human intervention because these algorithms can "learn." While this may be a theoretical possibility, it requires a multitude of steps. First, a language has to be developed, shared and coordinated between AI agents in order to fulfill a coordination function. Second, any firm would have to activate communication capabilities in its price setting tools, which is as such hard to justify as a function of a pricing tool. Alternatively, the use of an AI agent would have to be justified by, essentially, delegating large proportions of management functions that may require legitimate communications with competitors.

But note that the use of AI agents in such functions would not necessarily generate a problem for antitrust enforcement in principle. Communication can still be detected through a communications record between AI agents and there is no reason why a firm would not be liable and accountable for the actions and communications of AI agents just as much as they are for the collusive activities of an employee. The fundamentals of collusion enforcement are not changed by the fact that transparency is high and communication rapid. Current antitrust regimes have the adequacy to deal with AI agents that communicate in order to facilitate collusion.

4.3 Simple price matching algorithms

It turns out that collusive dangers that may arise from algorithms are not necessarily tied to the speed and sophistication of algorithmic prices. On the contrary: some concerns arise from the ability of hard-coded algorithms to commit to very simple rules that specifically do not learn from the past and do not adapt to the changing environment. The economic literature has shown that such commitments to simple response algorithms can indeed lead to collusive outcomes. We explain in this subsection, that the assumed ability to commit to simple rules is *not* a credible assumption in real world settings in general, and in ecommerce in particular.

The seminal work of Axelrod (1984) showed how simple imitation algorithms can lead to cooperation in the Prisoner's dilemma. In particular, if each of the two players adopts a "tit-for-tat" strategy in which each just copies what the former did in the last period, then these strategies are "stable" from an evolutionary perspective and they support the coordinated outcome that is best for the two players. Applying the logic from Axelrod's arguments to pricing games, one may wonder whether a simple tit-for-tat variation on prices may lead to the collusive outcome in which the firms split monopoly profits. And more importantly, one should wonder whether algorithms would learn to converge on the collusive outcome even if the simple tit-for-tat strategy is not hard wired into the algorithm.

We first consider a natural variant of tit-for-tat: each firm charges the monopoly price and if it detects any firm charging less, then it will revert to undercut the rival firm's last price until it reaches its marginal cost. In this case, any deviation from a pre-specified price, e.g., the monopoly price, would lead to a price war that ends in marginal cost pricing and stays there forever. If the firms happen to agree on a pre-specified price, and if the environment is stable, then they can indeed collude on the pre-specified price. However, any slight deviation due to some error, or to some changing circumstances, would immediately cause a downward spiral to the competitive price. This is precisely the type of problem that the colluding firms in the Trod/CB eye case discussed in the previous section encountered. If they did not carefully intervene in setting the starting point for the algorithm then price undercutting continued to occur.

Three things are worth noting. First, as in any collusive outcome firms would still have to agree on the price (and the penalty scheme) to achieve a collusive outcome. This would require communication that would violate current antitrust laws. Second, any entry by another firm would disrupt the equilibrium. Last, with more than two competitors, a "tremble" that causes a downward spiral of prices becomes ever more likely.

A second variant of tit-for-tat is one in which competing firms adopt a price-matching strategy as follows: each firm monitors its competitors, and matches the lowest price. Like the undercutting tit-for-tat strategies described above, the market price supported by such strategies is indeterminate, and firms would need to communicate in order to agree on the pre-specified price. Now imagine that somehow the firms managed to coordinate on a price ex ante. If in some period one particular firm charges a lower price, all other firms will immediately match it, resulting in a new market price that is lower. But if one firm raises its price then others will not follow, unless there are only 2 firms. That is, if there are only two firms then if one raises its price the other follows, and the new market price is higher. But with more than 2 firms there is no natural price leader, and there is no mechanism to raise prices. Hence, with more than two firms any profitable price can be supported by these strategies, and if there are some trembles, then like before, the market will likely converge to the competitive price.

In a game with at least three players, the only way to coordinate monopoly pricing is if the price-matching "graph" is fully connected: Imagine 6-players. If they are split into disjoint price-matching groups of 3 where A and B say “we’ll match C’s price” and D and E say “we’ll match F’s price” then the Nash-Bertrand (competitive price) is an equilibrium because C and F still need to coordinate. But if, for example, everyone but F says “we’ll match F’s price” then F has monopoly pricing incentives. This type of coordinated outcome is highly unlikely to occur by chance, but would rather require relatively sophisticated (and illegal) agreements. In this case, the only equilibrium is for all players to match the lowest price.

Note that the arguments laid out above make some very strong assumptions. First, they assume undifferentiated goods. When this assumption is dropped and there is some local monopoly power, this is less straightforward. Second, with different cost structures and randomly changing costs, the adoption of such schemes will surely violate the changing incentive constraints that firms face when their costs change, and profit maximization itself will be hampered by sticking to these simple algorithms.

Things may seem more manageable when there are only two firms. Imagine a variant of tit-for-tat where each firm just says that it will match the price of its competitor and, more importantly, each firm knows this about its competitor. In this case each firm knows that if it raises its price, then its competitor will match it. If the firms are symmetric in every way then this would lead to the monopoly price. No learning or AI is necessary, just a simple matching algorithm. However, several real-world issues undermine to naïve logic behind this simple example. First, demand and costs fluctuate randomly adding noise that undermines such coordination. Second, entry by a third firm would steal market share and unless there is some way to add the entrant to the collusive cartel, the algorithms would not be able to sustain collusion.

Last but not least, important issues of brand perception kick in. For example, take the well-known tag line “Walmart's Every Day Low Prices” which is the company’s commitment to be perceived as a low-cost retailer. If the company would engage in such a price creeping algorithm with its competitors, and entry would occur, then years of building a low-cost brand would be destroyed, which in turn will cause customer trust and loyalty to be lost.

It is important to note that for a company that wishes to build and maintain an every-day-low-price brand, simple monitoring price algorithms can actually help achieve this goal and help sustain competition. Imagine a retailer like Walmart that wants to adopt a low profit margin pricing strategy in order to sustain a low-price image and build brand loyalty around it. A simple algorithm can be used to do this as follows: (i) choose a mark-up rule for each product j, which given the product’s cost will result in a price $p_j^{\max}$; (ii) Monitor the prices of competitors and pick the lowest competing price, p_{comp}. (iii) if $p_j^{\max} > p_{\text{comp}}$ then charge p_{comp} and otherwise charge $p_j^{\max}$. This simple algorithm allows the retailer to commit never to charge more than its low-cost price $p_j^{\max}$ as long as others are charging no less than this level, and if any competitor charges less, the retailer will match the lowest competing price.[17]

On a final note, simple algorithms can also be used to engage in price signaling. Namely, a price leader can adopt an algorithm that makes public announcements about future pricing strategies in an attempt to signal price hikes to rivals who may then match the higher prices. However, algorithms offer no advantage over human actions along these lines, and are subject to the same regulatory controls that have been used to clamp down on this behavior.

4.4 Machines and Complexity

The focus of our analysis so far has drawn from the economics and game theory literatures, and draws from the basic theoretical frameworks and the empirical evidence from lab and field studies. But because we are concerned with computer algorithms, it is important to point to a large body of work in computer science that explores the limitations of algorithms in terms of the complexity they can actually handle and solve for within a practical time frame.

A large literature in theoretical computer science investigates whether certain problems can be solved relatively "fast" in polynomial time, or whether they are harder problems that cannot. This is the focus of complexity theory. About two decades ago, several computer scientists showed an interest in exploring the level of complexity needed in order to solve for equilibria in games. In an important paper, Daskalakis et al. (2009) showed that algorithms have real challenges in solving Nash equilibria even of simple static games. The upshot is that if solving for equilibria in simple finite one shot games is difficult, then solving for equilibria in complex environments like dynamic play in real world markets is obviously not computable in any reasonable amount of time. In a recent paper, Gal (2017) discusses the role that algorithms can play to facilitate joint profit maximization among competitors. The computer science literature strongly suggests that such endeavors, which require algorithms to compute equilibria of very complex games, is not likely to be solvable in the time that would make them beneficial for the companies that are trying to use them.

It is also important to note that another strand of literature in theoretical computer science sheds lots of doubt on the ability of computer algorithms and AI agents to solve problems that are complex and that change constantly with changes in the environment. Wolpert and Macready (2005) discuss the celebrated "no free lunch theorems" in computer science that, broadly speaking, demonstrate that if an algorithm performs well on a certain class of problems then it

[17] Of course, the retailer can add a price-floor in order to minimize losses in case a competing price is too low to justify matching.

necessarily will have degraded performance on the set of all remaining problems. That is, algorithms may solve one problem fine, but as the problem morphs with changes in the environment, the algorithm will not perform well and a new algorithm will be needed to adapt to the new environment. Hence, AI algorithms are a lot more limited than the lay person might infer from the public discussion.

5. Policy Implications

We have already discussed in section 2 that the typical check list criteria for the ease of sustaining collusive outcomes do not indicate that internet and e-commerce companies to be high on the list of industries in which collusion is likely. First, retailers are very asymmetric both in firm sizes, technologies used, and in variety of business models and cost structures. The multi-sided business models with very different degrees of integration of different activities makes it particularly hard to align incentives. As a result, the most robust criteria for the effectiveness of collusion (when an explicit agreement can be obtained) is not satisfied for making collusion stable.

Even if one argues that entry into online retail and internet services is generally difficult, the arguments above suggest that these industries are not particularly likely to generate collusion in the first place. That said, entry has become exceedingly easy given the advancements in information technology and web-based services, which have dramatically reduced the cost of setting up and operating a new business. With cloud computing, for example, new firms no longer have to invest in physical data storage or processing infrastructure. What used to be expensive upfront capital expenditure is now ongoing operating expenditure, and new entrants can flexibly pay to “ramp up” computing and data storage capacity as and when it is needed. Customer relations management, customer service, recruiting, and a variety of other functions that in the past required significant investment can now be scaled almost linearly given the plethora of web-based and outsourced services that are offered.

However, even ignoring the increased ease of entry and high asymmetry between players in electronic retail, we have demonstrated in this paper that algorithms do not magically overcome the most fundamental barrier to collusion: solving the coordination problem. Market transparency and rapid interactions do not solve the coordination problem. Algorithms also do not get to a collusive outcome simply if their objective is to maximize profits (if that is even possible given the complexity of the maximization problem). This is not really surprising. The Game Theory literature has shown that rationality, and even common knowledge of rationality, does not solve the coordination problem. There is more that algorithms would need to do to converge on collusive outcomes.

First, even the development of focal points requires a fundamental property that is not obvious AIs will have in the near future: a theory of mind regarding the solutions that stick out as obvious to other algorithms. In Schelling's discussion of focal points, social conventions used in other contexts and knowledge about the thinking and adherence to conventions of others in other contexts are critical to develop focal points. It is questionable whether this type of reasoning by analogy from other contexts can be expected from AI algorithms in the near future or whether the algorithmic tools used in pricing would have that type of experience of other algorithms.

Second, the experimental work by Cooper and Kühn (2014) has made clear that social context and social reaction are extremely important in supporting collusive outcomes. They appear to be more effective in generating collusion than the basic theoretical incentive mechanism. Such social context is missing for algorithms. Hence, the lack of human irrationality will hinder rather than help the establishment of collusive outcomes for AI agents

Third, even in situations where social context helps in coordination and where social rewards play an important role to incentivize cooperation, coordination without explicit communication about the required actions has been proven over and over again to be extremely difficult if not impossible. This means that to really be able to collude effectively, AI agents would need to develop a common language to come to an agreement in the sense of having a common understanding of expected behavior and reactions to deviations from expected behavior by the other algorithms. They would need to coordinate on "carrot and stick" strategies that can become complex in ways that algorithms cannot solve within a manageable time period.

The implementation of algorithms that eventually lead to collusion requires that the communications protocol would have to be programmed *into* the algorithm, which requires coordination on language, or coordination on developing communication capability and thus the ability to develop a language for an activity (pricing) that does not need communication to be performed. Humans, in contrast, have already coordinated on language and can therefore use a knowledge of terms and protocols on how to assure each other of reaching an agreement from other contexts. A pricing algorithm would need to develop a common language with other algorithms. This requires some facility, as well as the expectation that this happens simultaneously with at least some of the other hundreds of algorithms that are used by diverse competitors. Even to just learn about the feasibility of collusion, algorithms would need a device to correlate their strategies to learn about joint changes in prices in order to able to even evaluate such scenarios. For that, algorithms would need to understand when their counterparts are performing a joint experiment and when they are instead responding to changes in market conditions. As actual cases of algorithmic collusion in real life show, this is an enormous hurdle to take.

Nothing in this paper should be misread as claiming that collusion cannot be achieved. Instead, the claim is that the measures that need to be taken to prevent algorithmic collusion are virtually identical to the instruments currently used in competition policy practice to detect collusion and enforce anti cartel policy.

Cartel enforcement relies critically on the detection and punishment of behavior that could be part of coordinating activities. It appears that coordination between algorithms and even AI agents will require the disclosure, and even joint design, of algorithms between competitors. Coordinated programming efforts would be required to ensure the ability of algorithms to perform a handshake or include some fixed built in coordination agreement in order to achieve coordination. These are all activities that are no more easily concealable than a traditional cartel that works with meetings at freeway rest stops, phone calls, and e-mails, if not harder. There would be many computer engineers involved and these engineers would have to obtain the necessary knowledge about the algorithms of competing firms. Such evidence is just as detectable as that of meetings on price discussion. In fact, the programming of the algorithm would leave a digital trail and a track record of such activities in itself.

Anything that would constitute coordination in the writing of pricing algorithms or in the efforts to make pricing algorithms “compatible” would amount to coordination of prices. Whether a concrete price for a given period is agreed or a programmed pricing strategy is agreed makes no difference and appears to be well covered by the laws. The definitions in the law do not even seem to need stretching for these behaviors to be covered under standard cartel enforcement.

The same is true for conversations about the specific pricing algorithms used by competitors. Since firms are not allowed to communicate about prices they are planning to set under current antitrust rules, it should follow immediately that communication of the algorithm or about its design and workings with competitors violates current laws. We believe that this is basically covered by the prohibition of exchanges on future planned prices a fortiori. The algorithm, after all, is a rule that explicitly maps market conditions into planned prices. Hence, the exchange of the algorithm precisely specifies the planned pricing under all future circumstances that the algorithm can distinguish and therefore must be illegal under current antitrust rules.

We have also discussed that even where AI agents develop language, communicate, coordinate, and collude, this situation would not be at all different from current antitrust rules regarding antitrust compliance by firms. A firm is liable for any collusive behavior of its managers or decision makers. An AI agent would just as much be a decision maker that the firm is responsible for. If it decided to collude with another AI agent from another firm, there seems to be no question that a firm can be fined just as surely as when lower level managers colluded without the knowledge of top management. The evidence of a communications trail would still be there. This would be the case also because AI agents would not be exempt from requirements of documenting their communications and firms' compliance activities would monitor for contacts with competitors in exactly the same way as with humans.

We also see economically no fundamentally new phenomenon when coordination is achieved by a firm that writes algorithms for collusion and then sells them to competitors. The value of the algorithm is then generated specifically by the collusion programmed into it and therefore the sale of such an algorithm to competitors would be done with the intent to achieve coordinated pricing. It is another question whether this is a realistic scenario because there are many vendors of algorithms and coordination will only be achieved if there is common usage.

However, we believe that making explicit that all these types of coordination activities between algorithms fall squarely into the current prohibitions of agreements and potentially anticompetitive information exchange about future pricing would provide much needed legal certainty in the ecommerce sector and clarify what steps firms might want to take to assure antitrust compliance. Probably no firm would want to endow its pricing algorithms with the ability to talk to the competition in any case and thus preventive actions will be easy. Most companies would also not want to share their algorithms. It is precisely what many firms view as their competitive business advantage. But once it is clear that exchanges of and about implemented algorithms would trigger a cartel case, it is also clear that coordination between algorithms will be a very unlikely event.

Our policies therefore appear to be quite a robust for a long time to come because even self-aware algorithms will need to solve the coordination problem. We should however get concerned when they can use time travel to hide their communications.

References

Abreu, Dilip, Paul Milgrom, and David Pearce. 1991. “Information and Timing in Repeated Partnerships” Econometrica, vol. 59(6), pp. 1713-33

Axelrod, Robert. 1984 *The evolution of cooperation*. Basic Books

Bernheim, B. Douglas, and Michael D. Whinston. 1987 "Coalition-Proof Nash Equilibria ii. Applications." *Journal of Economic Theory* **42(1)**: 13-29.

Bernheim, B. Douglas and Michael D. Whinston, 1998 "Exclusive Dealing," *Journal of Political Economy* **106(1)**: 64-103.

Bhaskar, V. 1989. “Quick Response in duopoly ensures the monopoly outcome,” *Economics Letters*, **29(2)**:103-107.

Bigoni, Maria, Marco Casari, Andrzej Skrzypacz and Giancarlo Spagnolo. 2015. “Time Horizon and Cooperation in Continuous Time,” *Econometrica* **83(2)**:587-616.

Brandts, Jordi and Cooper, David J. 2007. “It's What You Say, Not What You Pay: An Experimental Study of Manager-Employee Relationships in Overcoming Coordination Failure,” *Journal of the European Economic Association*, **5(6)**:1223–1268

Camera, Gabriele and Marco Cassari. 2009. “Cooperation among strangers under the shadow of the future.” *The American Economic Review* **99(3)**:979-1005.

Cooper, Russell W., Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross. 1990. “Selection Criteria in Coordination Games: Some Experimental Results.” *The American Economic Review*, **80(1)**: 218-33.

Cooper, David J. and Kai-Uwe Kühn. 2014. “Communication, Renegotiation, and the Scope for Collusion,” *American Economic Journal: Microeconomics*, **6(2)**: 247-78

Dal Bo, Pedro and Guillaume Frechette (2011). “The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence”, American Economic Review, Vol. 101, pp. 411-429.

Daskalakis, Constantinos, Paul W. Goldberg, and Christos H. Papadimitriou. 2009. “The complexity of computing a Nash equilibrium,” *SIAM Journal on Computing* 39.1 (2009): 195-259.

Eckert, Andrew. 2002. “Retail price cycles and response asymmetry,” *Canadian Journal of Economics/Revue canadienne d'économique* **35(1)**: 52-77.

Eckert, Andrew. 2003. “Retail price cycles and the presence of small firms,” *International Journal of Industrial Organization* **21(2)**: 151-170.

Fonseca, Miguel A., and Hans-Theo Normann. 2012. “Explicit vs. tacit collusion—The impact of communication in oligopoly experiments,” *European Economic Review* **56(8)**: 1759-1772.

Gal, Michal S. 2017. “Algorithmic-Facilitated Coordination,” Prepared for tor the OECD forum on Algorithms and Collusion, June 2017.

Gal, Michal S. and Niva Elkin-Koren. 2017. “Algorithmic Consumers,” *Harvard Journal of Law & Technology*, **30**:1-45.

Green, Edward J., and Robert H. Porter. 1984. “Noncooperative Collusion under Imperfect Price Information.” *Econometrica*, **52(1)**:87–100.

Haltiwanger, John, and Joseph E. Harrington Jr. 1991. “The impact of cyclical demand movements on collusive behavior,” *The RAND Journal of Economics,* **22(1)**: 89-106.

Harsanyi, John C., and Reinhard Selten. 1988 *A general theory of equilibrium selection in games*. MIT Press.

Huck, Steffen, Hans-Theo Normann and Jorg Oechssler. 2004. “Two are Few and Four are Many,” *Journal of Economic Behavior and Organization,* **53(2)**:435-446.

Ivaldi, Marc, Bruno Jullien, Patrick Rey, Paul Seabright, and Jean Tirole (2003), “The Economics of Tacit Collusion”, Final Report for DG Competition, European Commission.

Kühn Kai-Uwe and Hans-Theo Normann. 2016. “The coordinated effects of mergers: an experimental analysis”, unpublished manuscript, DICE.

Kühn, Kai-Uwe and Michael S. Rimler. 2009. “The comparative statics of collusion models,” University of Michigan working paper.

Levin, Lawrence, Matthew S. Lewis and Frank Wolak. 2017. “High Frequency Evidence on the Demand for Gasoline,” *American Economic Journal: Policy*, (forthcoming).

Levenstein, Margaret C., and Valerie Y. Suslow. 2006. "What Determines Cartel Success?" *Journal of Economic Literature* **44(1)**:43-95.

McSweeny, Terrell and Brian O’dea. 2017. “The Implications of Algorithmic Pricing for Coordinated Effects, Analysis and Price Discrimination Markets in Antitrust Enforcement,” *Antitrust*, **32(1)**:75-81.

Mehra, Salil K. 2016. “Antitrust and the Robo-Seller: Competition in the Time of Algorithms,” *Minnesota Law Review* **100**:1323-1352.

Milgrom, Paul and John Roberts. 1990. “Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities,” *Econometrica* **58(6)**:1255-1277

Noel, Michael D. 2007a. “Edgeworth price cycles, cost-based pricing, and sticky pricing in retail gasoline markets,” *The Review of Economics and Statistics* **89(2)**: 324-334.

Noel, Michael D. 2007b. “Edgeworth price cycles: Evidence from the Toronto retail gasoline market,” *The Journal of Industrial Economics* **55(1)**: 69-92.

Noel, Michael D. 2008. “Edgeworth price cycles and focal prices: Computational dynamic Markov equilibria,” *Journal of Economics & Management Strategy,* **17(2)**: 345-377.

Salcedo, Bruno. 2015. “Pricing Algorithms and Tacit Collusion,” mimeo, Pennsylvania State University.

Sannikov, Yuliy and Andrzej Skrzypacz. 2007. “Impossibility of Collusion under Imperfect Monitoring with Flexible Production,” *American Economic Review* **97(5)**:1794-1823.

Stigler, George J. 1964. “Theory of Oligopoly,” *Journal of Political Economy* **72(1)**:44–61.

Stucke, Maurice E. and Ariel Ezrachi. 2015. “Artificial Intelligence & Collusion: When Computers Inhibit Competition,” Oxford Legal Studies Research Paper No. 18/2015.

Stucke, Maurice E. and Ariel Ezrachi. 2016. *Virtual Competition: The Promise and Perils of the Algorithm-Driven Economy*. Harvard University Press.

Stucke, Maurice E. and Ariel Ezrachi. 2017. “Two Artificial Neural Networks Meet in an Online Hub and Change the Future (of Competition, Market Dynamics and Society),” working paper, Oxford University.

Van Huyck, John B., Raymond C. Battalio, and Richard O. Beil. 1990. "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure." *The American Economic Review* **80(1)**: 234-48

Wolpert, David H., and Macready, William G. (2005) "Coevolutionary free lunches", IEEE Transactions on Evolutionary Computation, 9(6): 721-735

Nate:

1. Won’t make public statements when you have thousands of products. If you have a handful of products then you don’t need algorithms.
2. Colleen: what if everyone says we’re all matching Amazon. Nate: need thousands of companies, and if they all make th same statement then they should be sued.
3. Summary: even in the unlikely if you look at the conditions where it could happen it seems that algos are not likely to be in play and algos are used when there are thousands of products