

Competition Policy and the Tech Industry – What’s at stake?

Remarks at Computer & Communications Industry Association
April 12, 2018

D. Bruce Hoffman
Acting Director, Bureau of Competition, U.S. Federal Trade Commission

Thank you to the Computer and Communications Industry Association for asking me to speak today, to Marianela Lopez-Galdos for organizing this event, and to my staff—in particular, Kelly Signs—for helping me prepare these remarks.

Before I begin, I need to provide the disclaimer that my remarks today do not necessarily represent the views of the Federal Trade Commission, any Commissioner, or any prospective Commissioner.

“Competition policy and the technology industry” is today’s hot headline. It is a hot-button issue at conferences, in the media, and in the enforcement world. It is a broad topic. In fact, it’s often not clear whether commentators discussing the “technology industry” mean to include in that “industry” firms in the “Online” space (which could be further divided into search, ratings/reviews, messaging/communication, social media, retail, and many other categories); hardware; software; intellectual property; and/or data (which could also be further divided into many categories, including collection, sale, and use).

For the purposes of today’s discussion I will largely skip past this definitional issue, though it’s important to note that it can be crucial in antitrust analysis. I will direct my remarks to technology broadly, and will indicate where I have in mind some more specific subset of the broader technology industry.

The antitrust, or competition, policy discussion about the technology industry has tended to revolve around certain broad themes. One of these themes is a generalized perception that the “technology industry”—often, though not always, focused on the online space, or even social media in particular—is too big or too powerful in some broad sense. There are also concerns about the collection or use of data, either as a distinct source of harm, or as contributing to market power. Other concerns have been raised about artificial or machine intelligence, or the use of algorithms that could affect firm behavior. Then there are concerns, which are not particularly unique to the technology industry, that firms might take actions to reduce competition, such as buying potential or nascent competitors, or seeking to harm or disadvantage their rivals.

Much of the debate that has arisen in these contexts involves two threads of discussion. The first thread is the contention that antitrust law is dated and inadequate and can’t handle the new technology industry. The second and closely related thread is that antitrust enforcers

need to rethink their current focus, which is on protecting the competitive process and promoting consumer welfare, as opposed to, for example, protecting competitors.

Before I specifically address these issues, and speaking broadly, I think it's important and helpful to take seriously these critiques of antitrust enforcement. At the FTC, we welcome debate and discussion. I think that it's healthy to reexamine your fundamental assumptions for any regulatory process or intellectual pursuit on a regular basis. That way, you are forced to think through rigorously what you're doing, why you're doing it, and whether what you're doing is correct. I didn't necessarily expect, when Acting Chairman Ohlhausen asked me to return to the FTC, to be in midst of a discussion over whether a 40-year bipartisan consensus on the appropriate role and standards for antitrust policy is off-base or misguided. But I think that discussion is worth having, and is something from which we can learn.

Having said that, though, as our former Chairman Bill Kovacic noted recently, the consumer welfare standard that is currently the touchstone for antitrust enforcement in the United States—and also in Europe—is the result of decades of experience in the United States with failed standards, such as protecting competitors at the expense of consumers. This prolonged experiment is somewhat unique to United States, largely because we've had antitrust laws on the books and actively enforced them for so long that we've had lots of opportunities to get our approach wrong—opportunities that we have often taken. But having made those mistakes, we want to make sure we don't repeat them. And, I think that our experience with failed approaches to antitrust standards is probably worth close study by those considering alternative approaches going forward. After all, while it's important and useful to question the status quo, it is also good to learn from experience.

Let me turn now to the first “thread” I mentioned before—the idea that antitrust isn't up to the challenge of technology, that something in our antitrust laws should be radically changed in order properly to address the tech industry. Interestingly, this is hardly a new issue. To illustrate that point, I would like to rewind 20 years, to a speech by Joel Klein, who was serving at that time as the Assistant Attorney General in charge of the Antitrust Division of the United States Department of Justice. Assistant Attorney General Klein, in 1998, gave a speech that addressed the application of antitrust law to the technology industry, and which included the following quote:

[Some] question whether the existing antitrust laws can possibly be relevant to today's economy. The Sherman Act was passed in 1890 in response to the nationwide industrial trusts that the railroads had made possible, and the Clayton Act was passed in 1914 and was aimed largely at retailing and wholesaling practices in localized markets. How, then, can these ancient statutes be relevant to a 21st Century, information-based, economy? I get asked that question, especially by non-antitrust-lawyers, probably more than any other. And I answer, unhesitatingly, that the laws are just fine, precisely because, unlike most contemporary

statutes, they are common-law provisions and, therefore, they are not locked in text or time.¹

Last fall, Johannes Laitenberger, Director-General for Competition of the European Commission said something quite similar—that there is no need to reinvent antitrust to address the modern economy, because competition laws in Europe have been remarkably adaptable. Those laws, he pointed out, focus on broad policy prescriptions that leave room for analyses to be refined, tests to be modified, and guidance to be changed based on developments in markets and technology and in learning.²

I agree. The antitrust enterprise in the United States, and now elsewhere, has been a study in learning and thinking and evolution. For the most part, that effort has highly beneficial, including as we take our principles and apply them to new developments, such as today's technology industry.

With those larger points in mind, let me now turn to current views of the role of antitrust in the technology industry. Obviously (and as I noted earlier) I can't speak for everybody, but it's worthwhile starting with what the United States government has actually said on this issue. In June of 2015, the United States government made a submission to the OECD competition committee on this topic, which stated, in part:

Innovation is the hallmark of a dynamic and competitive economy, but can pose challenges for legislative and regulatory bodies trying to keep pace with rapidly evolving businesses. Disruptive innovation (including new products, services, and business models), in particular, often results in new, better, and/or lower-priced products and services to consumers, but may not fit within existing regulatory frameworks, and thus, can raise challenges for regulators. Competition authorities can play an important role shaping the inevitable transitions caused by disruptive innovation, by advocating for regulatory responses that do not unduly restrain competition, enforcing competition rules to ensure that incumbents do not foreclose new rivals from the market, and using studies and other research methods to foster greater understanding of new technologies and business models.³

¹ “The Importance of Antitrust Enforcement in the New Economy,” remarks of Joel I. Klein, Assistant Attorney General of the Antitrust Division, Department of Justice before the New York State Bar Association, New York, Jan. 29, 1998, <https://www.justice.gov/atr/speech/importance-antitrust-enforcement-new-economy>.

² “EU competition law in innovation and digital markets: fairness and the consumer welfare perspective,” remarks of Johannes Laitenberger, Director-General for Competition, European Commission, Brussels, Oct. 10, 2017, http://ec.europa.eu/competition/speeches/text/sp2017_15_en.pdf.

³ Submission of the United States, Hearing on Disruptive Innovation (DAF/COMP/WD (2015)54, 2 https://www.ftc.gov/system/files/attachments/us-submissions-oecd-other-international-competition-fora/1507disruptive_innovation_us.pdf.

That 2015 statement is a highly pro-technology, pro-innovation statement of antitrust philosophy. It basically boils down to this: don't let incumbent industry participants and existing regulatory structures block innovation and deter the advancement of new technology. But does that statement mean that technology companies get a free pass because innovation is generally pro-competitive? Absolutely not.

With that in mind, I want to turn to how antitrust enforcers incorporate technology and on-going innovation in our competition work. First, I'm going to talk about mergers and how we think about mergers insofar as they impact technology industry companies. I will then address issues relating to the conduct of technology firms.

Mergers

It is no surprise to anybody here that the Federal Trade Commission and the Department of Justice believe in vigorous enforcement of merger law. We use the best available legal and economic tools for horizontal mergers. We continue to apply the Horizontal Merger Guidelines, which fully apply to technology company mergers.

One criticism of using our standard approach in technology mergers is that it is overly focused on price effects. I would agree that we always want to assess whether a particular merger is likely to have an effect on price, and the importance of preventing merger-induced price effects is reflected both in U.S. law as well as the Horizontal Merger Guidelines. So yes, price is important. Price also has the virtue of being measurable, which makes it something that we can actually grapple with and economists can rely on to run sophisticated econometric models.

But we also consider other likely effects of mergers. The Horizontal Merger Guidelines spell out in some detail not only price effects but also how to address potential quality, output, and innovation effects made possible by a merger. In fact, we have a long history in merger review of considering those issues. Sometimes we shorthand other effects into our measurement of price effects because you could certainly think of price, output, quality, even innovation as being aspects of the bundle of features that make up a product or service. Ultimately, however, we are trying to assess all the ways in which consumers benefit from or prefer a product or service, and then think about how the merger might enhance or diminish those benefits for consumers. This includes vertical as well as horizontal mergers.

When we look at mergers in the technology industry, I think it's important to note two things. First, we use the best available analytical tools to evaluate technology industry mergers, just as do mergers in any other sector of the economy. And, as you would expect, when we apply those tools in the extremely fact-intensive context of a merger review, the results we obtain vary with the facts. Sometimes we determine that there are no aspects of competition that are likely to diminish after the merger, despite spending extensive time and resources in our investigations. For example, the Commission investigated Zillow's acquisition of Trulia, two

consumer-facing real estate technology platforms. After a significant investigation, the Commission determined there was no coherent theory of anticompetitive effect.⁴

We came to a similar conclusion when we looked at Amazon's acquisition of Whole Foods last summer.⁵ In some respects, the Amazon/Whole Foods merger was arguably vertical; in others, arguably horizontal; in others, arguably conglomerate. The deal has often been characterized as a technology industry merger because Amazon is viewed as a quintessential technology company. But in other ways, the merger involved the antithesis of what most people would think of as the technology industry—the purchase of a grocery store chain. In any event, we did not find any evidence that supported a theory that the merger would violate Section 7 of the Clayton Act or Section 5 of the FTC Act.

But the fact that those two mergers—based on the facts available us at the time—did not present facts that would support a merger challenge under the statutes we enforce does not mean that other technology mergers will also pass muster. Mergers between technology companies can and do raise competitive concerns, as a couple of examples demonstrate.

Recently, the Commission filed litigation to block the merger of CDK and Auto/Mate.⁶ The transaction involved a merger between auto dealer software platforms. These are systems that auto dealers use to track their services, their prices, and other crucial functionalities. They also offer platforms that allow third party apps to integrate with the system, which give dealerships even more customized functionality. Despite Auto/Mate's limited penetration in the market to date—with less than 6 percent share—it was having an outsized impact on competition with other platforms, especially CDK. Moreover, the evidence indicated that Auto/Mate's market share not only understated its current competitive significance, but, importantly, its future competitive significance (a point I will return to in a few minutes). After the Commission announced that it would seek to block the deal, the parties abandoned their plans. Earlier last year, the FTC also sued to stop the proposed merger of DraftKings and FanDuel, two online platforms for sports contests. The Commission alleged that the two firms were locked in a battle to out-do each other, with innovative contests and large prizes, and that the merger would eliminate the incentives to continue that type of direct price and non-price competition.⁷ That merger, too, was abandoned during litigation process.

In reality, many of our merger cases involve concerns about non-price aspects of competition that would be lost, including on-going innovation and product development efforts. At this moment, we have four merger cases in litigation, and at least one involves technology issues.

⁴ See Statement of Commissioners Ohlhausen, Wright and McSweeney Concerning Zillow, Inc./Trulia, Inc., FTC File No. 141-0214 (Feb. 19, 2015), https://www.ftc.gov/system/files/documents/public_statements/625671/150219zillowmko-jdw-tmstmt.pdf.

⁵ See Statement of Acting Director of the Bureau of Competition on the Agency's Review of Amazon.com, Inc.'s Acquisition of Whole Foods Market Inc., Aug. 23, 2017.

⁶ In re CDK Global, Dkt. 9382 (complaint filed Mar. 20, 2018).

⁷ In re DraftKings, Inc., Dkt. 9375 (complaint filed Jun. 19, 2017).

What to do about data?

Let me turn to some specific issues that have come up in the popular press and in the antitrust conversation relating to technology industry mergers. First, and this came up in connection with Amazon/Whole Foods, some argue that we should pay particular attention when a merger results in one firm gaining more data or information about its customers. The argument is that amassing data is the equivalent of a price increase or quality decrease. In other words, in exchange for offering something to a consumer—for example, a free platform for shopping and ordering merchandise—the technology firm obtains data from that consumer experience, so you could think of that as the equivalent of a price, and if the firm acquires more data per transaction or interaction, as a price increase (or a quality decrease).

This is an interesting theory. However, our understanding is that at present there's neither a theoretical nor an empirical basis for assuming in every case that a firm acquiring more data about customers is imposing the equivalent of a price increase or quality decrease. The actual implications of an exchange of data between a consumer and a platform or other type of technology is much more complicated, and the antitrust implications of such transactions are highly fact dependent—and are an area about which we still have a great deal to learn.

Let me elaborate on a few ways in which data is likely different from a simple payment of money. As an initial matter, it's not necessarily clear that many consumers place much economic value on data about them, or that different consumers value such data in the same ways.⁸ Obviously, different consumers may place different values on money, but the valuation that consumers place on money is bounded in a way that the valuation they place on data may not be. The issues involving firms' access to and use of consumer data are more analogous to a differentiated product, where different consumers may have different values associated with each aspect of product available in the market, and each consumer makes a purchase choice based on his or her own utility. This is reflected in the price the consumer is willing to pay. Empirically, consumers at large have not seemed to place greater value on firms that collect less data, and we are not aware of literature that clearly establishes any level of consumer valuation of data. Of course, this could change—and, given current events, it could change rapidly—but right now, there's no good reason to think that consumers value data about themselves in the same way that they value money in their bank accounts.

Moreover, unlike money, sometimes data can be transferred repeatedly to multiple firms, i.e., it is non-rivalrous. That means that even if firm A acquires a piece of data about me, I may still be able to provide the same data to firm B, C, D, or ad infinitum. That's not always true, but it's often true, and complicates any simple analogy between data and price.

⁸ Of course, many consumers care about how certain kinds of sensitive data (such as financial account information) may be used or misused by others. That is different than valuing the data itself. In contrast, when a consumer pays for a product, she probably cares little about how the company uses the money provided she receives the product.

In other cases, older data may be more valuable than recent data, or vice versa; data is not necessarily fungible, and may change in value in idiosyncratic ways depending on the data at issue and the interests of the firm seeking to acquire the data. In contrast, money is always money—a dollar is worth a dollar, no matter how old it is—and the same dollar cannot be shared with more than one person at a time.

Further, there's an immediate positive return to consumers for data transfer. This is not directly analogous to the indirect potential benefit that you get from paying money to somebody where that money may or may not be plowed back into improving the quality of the product you're offering. When you give data to a tech firm, often that data is immediately used to improve the service that the firm is providing to you. In fact, this exact issue—that increased data has direct positive effects on the customers providing the data—has sometimes been cited as a competition concern, on the grounds that this effect could make the services data-rich firms provide too good for competitors to match (whether this is a valid competition concern is an entirely different issue). This is simply not true for an increase in price, and it also is clearly distinct from a reduction in quality. In fact, in some ways, this effect could be viewed as an increase in quality or a reduction in price.

Now, it is fair to point out that there are significant information asymmetries between consumers and technology platforms concerning data collection and use. Carnegie Mellon researchers found that it would take an individual consumer 76 days to read all the privacy policies of the websites the average consumer visits in a single year.⁹ Obviously, this is impossible—or rather, no consumer is going to spend the time it takes to do this. But the presence of an information asymmetry doesn't lead inexorably to the conclusion that the market cannot address data issues, nor to the conclusion that antitrust enforcers should determine what value consumers should place on data. In fact, information asymmetries are fairly common in competitive economies and they're often addressed by competition. Firms compete by identifying the information asymmetries and offering solutions to them. We see this occurring now in the technology industry, where some firms are differentiating themselves by offering better privacy protections, or better data security than their rivals. If consumers value data, there's no reason to assume that competition on this point cannot result in data being valued appropriately.¹⁰

None of this means that data can't ever be thought of like price or quality. In the right case, we could very well have antitrust concerns about a transaction that requires consumers to yield

⁹ “Reading the Privacy Policies You Encounter in a Year Would Take 76 Work Days,” *The Atlantic*, Mar. 1, 2012.

¹⁰ If competition does not reward firms that offer more privacy or data protection, it is possible that some market failure is preventing competition from addressing data issues—though we would want to see evidence identifying such a market failure. It is important, though, to note that it is also possible that not enough consumers value data protection sufficiently for offering additional protection to become a valuable product attribute; or, it could also be that consumers find that the value they obtain from providing data outweighs the benefits of withholding it. If the latter is true—the market is not rewarding additional data protection because consumers do not place sufficient value on such protection—it is not clear that antitrust enforcers could or should substitute their judgment for consumers’.

more data.¹¹ But as the foregoing discussion indicates, this raises very fact-specific issues in an area where theory is unclear and the state of our knowledge is very limited. We cannot presently determine that consumers should value data in certain ways, or should not want to provide that data to firms in the same way that we can generally assume that consumers would not want to pay higher prices or accept lower quality. We are thus not in a position to adopt a general view that increases in data collection are the precise equivalent of a price increase or quality reduction. The specific facts will matter quite a lot, and also, hopefully, the state of research and analysis on this issue will continue to develop and improve.¹²

In this context, I would like to briefly return to a point that I touched upon earlier. In the general discussion about data, occasionally we hear an argument that an acquisition of data may make the acquiring firm too effective a competitor. But it is important to remember that acquisitions that improve competitiveness are not anticompetitive. When a firm acquires an asset that enables it to provide better service and by doing so, becomes more attractive to customers—rendering its rivals’ jobs that much more difficult—that’s not normally an antitrust problem. In fact, that outcome would typically be viewed as a procompetitive aspect of a merger. For example, if a merger allowed two firms to combine, say, production or research and development capabilities in a way that rendered the new firm more competitive—with lower costs, or with a superior product—that would normally be considered a benefit of the merger. So too with data; absent other facts, if a merger allowed the merged firm to combine two data sets to achieve a synergistic outcome so its customer service or products become substantially better for its customers, antitrust policy would generally view that result as procompetitive. Again, this does not mean that acquisitions of data can never be anticompetitive—but it is vital to avoid confusing harm to competitors with harm to competition.¹³

Acquisitions of Innovator Firms

I want to turn to another recurring fact pattern in the discussion about antitrust policy with regards to the technology industry—the acquisition of nascent competitors. The idea here is that large technology firms have developed a tendency to buy start-ups, and by so doing, are foreclosing the development of emerging rivals that might ultimately unseat them. In my view, this is a completely legitimate and real theory of competitive harm (and, of course, it is not

¹¹ And, of course, our Bureau of Consumer Protection addresses a wide range of data privacy and data security issues, and has substantial enforcement authority in that arena.

¹² For a more detailed description of many of these issues and survey of the economic literature surrounding the value of privacy, see Alessandro Acquisti, Curtis Taylor, and Liad Wagman, “The Economics of Privacy,” *Journal of Economic Literature* 54, vol. 2, pp. 442-92 (June 2016). For a description of a methodology to value the loss consumers suffer when their data is misappropriated as a result of a firm’s failure to provide a promised level of data protection, see Dan Hanner, Ginger Zhe Jin, and Marc Luppino, “Economics at the FTC: Horizontal Mergers and Data Security,” *Review of Industrial Organization* 49, vol. 4, pp. 613-631 (Dec. 2016) (the relevant discussion is at pp. 627-31).

¹³ I do not address here other, more straightforward applications of antitrust law to mergers in which data may be relevant. Some of those scenarios are catalogued in Terrell McSweeney & Brian O’Dea, “Data, Innovation and Potential Competition in Digital Markets—Looking Beyond Short-Term Price Effects In Merger Analysis,” *Antitrust Chronicle*, Vol. 2, Winter 2018, pp. 7-13.

unique to the technology industry). It's something that we think about a lot and we pay attention to in transactions where it appears to present possible issues. In fact, as I mentioned earlier, this issue was a driving concern behind the Commission's decision to challenge the CDK/AutoMate merger. Auto/Mate, the firm being acquired, had a fairly small share. However, the evidence showed that looking solely at current market shares would miss a major issue—that Auto/Mate appeared to be on the cusp of becoming a much more important and vibrant competitor. The merger would have snuffed out that developing competition, and that played a major role in the Commission's decision to file suit to stop that merger.

However, while we take these cases very seriously, our inquiry into firms' ability to become future competitors is bounded by our ability to predict what is likely to happen with and without the merger. We have to have an evidentiary and economic basis to determine that the firm being acquired is really likely, in some reasonable time horizon, to be a significant competitor. Otherwise, we don't have a case to bring. Without an evidentiary and economic basis for making this sort of prediction, doing so would simply be the arbitrary judgment of antitrust enforcers (like me). That's not good enforcement; it's not the rule of law.

It's also worth noting that there could be consequences to taking an aggressive (and less evidence-based) approach to acquisitions of startups or nascent competitors. Two potential consequences are fairly obvious (and, of course, there could be consequences we haven't even thought of).

First, there's a risk of immediate consumer harm, because a larger incumbent firm acquiring a startup offers the startup a quick path to market by leveraging the acquirer's capital, technology, distribution, or other assets. Without the acquisition, the startup would have to develop those assets itself, which it might or might not be able to do, and which may take longer than would happen with the acquisition. If that development is delayed or never occurs, consumers could suffer a loss.

Second, there could conceivably be a negative effect on the capital market for startups. To the extent exit strategies for startups involve acquisitions, if such acquisition opportunities are constrained the capital available for startups may fall. That, in turn, could result in fewer startups. In other words, a strategy designed to preserve the independence of startups could result in a decline in startup activity. This, of course, is a very difficult area to quantify, let alone to make any kind of confident predictions, and I won't venture to do so. But, as we consider policy in this area, these are some of the issues we need to think about thoughtfully and carefully.

Conduct Concerns related to Technology Companies

Finally, let me turn to potential anticompetitive conduct by technology companies. I will start by briefly talking about collusion, including agreements not to compete. Outside of hard-core cartel cases, which are prosecuted by the Department of Justice, the FTC is concerned about agreements that don't quite rise to the level where jail is an appropriate outcome. There is

nothing particularly unique about technology firms when it comes to collusion. The incentives to profit by agreeing with your competitors are the same as in other industries, and the FTC, over the years, has pursued numerous collusion cases in technology and related areas, such as intellectual property (including the many cases the FTC has brought involving agreements between pharmaceutical companies not to compete, and reverse payment patent settlements between branded and generic firms). Of course, the FTC benefits from our research abilities, and in the past few years, the FTC has issued reports on big data,¹⁴ the sharing economy,¹⁵ and the Internet of Things.¹⁶

One hot topic in the collusion area is the potential effect of the increasing use of machine algorithms for pricing or other competitive decisions, or even artificial intelligence or AI. At the FTC, we have given a great deal of thought to these issues, including analyses reflected in speeches and papers by our Acting Chairman Ohlhausen¹⁷ and Commissioner McSweeney.¹⁸ I would also recommend that you read our recent FTC/DOJ paper to the OECD on the use of algorithms, which provides a good overview of thinking on these questions.

Fundamentally, at this early point in the development of these kinds of systems, we do not know if, in general, algorithms and AI will facilitate collusion, or if instead they will sharpen competition, or both, or neither. Current empirical results do not point in any particular direction, and theory doesn't supply a definitive answer.

Algorithms are simply problem-solving systems. Machines can perform algorithms more quickly than humans, or can handle algorithms that are computationally intractable for humans, but that does not necessarily mean that the algorithm will be better or worse at reaching an anticompetitive outcome than humans. Artificial intelligence and machine learning do more than simple algorithms—including adapting behavior to changing environments, and potentially doing so in ways that humans can neither predict nor, potentially, deconstruct. But there are obvious potential obstacles to algorithms, or AI, “colluding” in any way that would be more likely to succeed than humans would on their own. In the interests of time and space, I will not elaborate on this in depth, but here are few points.

¹⁴ See FED. TRADE COMM'N, BIG DATA: A TOOL FOR INCLUSION OR EXCLUSION (Jan. 2016), <https://www.ftc.gov/system/files/documents/reports/big-data-tool-inclusion-or-exclusion-understanding-issues/160106big-data-rpt.pdf>.

¹⁵ FED. TRADE COMM'N, THE “SHARING” ECONOMY: ISSUES FACING PLATFORMS, PARTICIPANTS & REGULATORS: A FEDERAL TRADE COMMISSION STAFF REPORT 3 (2016), <https://www.ftc.gov/reports/sharing-economy-issues-facing-platforms-participants-regulators-federal-trade-commission>.

¹⁶ FED. TRADE COMM'N, INTERNET OF THINGS: PRIVACY AND SECURITY IN A CONNECTED WORLD (2015), <https://www.ftc.gov/reports/federal-trade-commission-staff-report-november-2013-workshop-entitled-internetthings> (Staff Report).

¹⁷ Acting Chairman Maureen Ohlhausen, Remarks to the Concurrences Conference on Antitrust in the Financial Sector: “Should We Fear The Things That Go Beep In the Night? Some Initial Thoughts on the Intersection of Antitrust Law and Algorithmic Pricing” (May 23, 2017), <https://www.ftc.gov/public-statements/2017/05/should-we-fear-things-go-beep-night-some-initial-thoughts-intersection>.

¹⁸ Terrell McSweeney & Brian O’Dea, “The Implications of Algorithmic Pricing for Coordinated Effects Analysis and Price Discrimination Markets in Antitrust Enforcement,” Antitrust, Fall 2017.

It's reasonable to think of firms using AI or sophisticated algorithms to make pricing or other competitive decisions as, in effect, playing complex games. But complex multistage games have large, if not infinite, solution sets—there can be multiple Nash equilibria in such games.¹⁹ As a result, it's not clear that machines would be able to determine the collusive equilibrium or price outcome. Additionally, in the real world, there is a serious noise problem—for a machine attempting to find a collusive outcome, how will it know if a rival has signaled it by cutting production, or if instead the rival's factory has flooded, temporarily restricting capacity but not because the rival has any plan to keep capacity down? Further, in some mathematical systems, as the number of participants increases solutions become unmanageable.²⁰ This is an area in which academic and empirical work is going on, and which could benefit from more work so that we can have better guidance.

One piece of academic work that is worth reviewing is some experimental work by Kai-Uwe Kühn and Steve Tadelis on machine collusion. In their experiments, machines were not able to collude without human intervention, for some of the reasons I've mentioned above. As they put it (in a somewhat tongue-in-cheek way) "Two Artificial Neural Networks meet a multidimensional continuum of subgame-perfect Nash equilibria in an online hub. With unbelievable speed the two Artificial Neural Networks react and say "huh?"²¹ Overall, while humans clearly can use algorithms as tools to facilitate collusion, we currently do not know whether algorithms—or AI—may become able to collude themselves, or to more (or less) perfectly reach oligopoly outcomes.

As for unilateral conduct by technology firms, this is an area we look at closely, and are very interested in pursuing. The last time I served at the Commission (with, among others, Joe Simons and Susan Creighton as Bureau Directors, and Tim Muris and Debbie Majoras as Chairmen) we aggressively pursued unilateral cases, such as the Bristol Myers/Squibb Orange Book case,²² which was an abuse of government process case; Rambus,²³ which alleged patent hold-up involving a standard setting organization; and Unocal, which involved both types of conduct.²⁴ And, of course, the FTC is a bipartisan agency, and has under the prior

¹⁹ Technically, in an infinitely repeated game, under the Folk theorem virtually any outcome could be achievable—and in a finite game, by backward induction the logical outcome may be to play the one-shot game Nash equilibrium. In a simple prisoner's dilemma, that outcome happens to be to defect—or to not collude. Thus, it is at least theoretically possible that machines will be very effective at not colluding.

²⁰ This is most well-known in the n-body problem in the physics of celestial mechanics, but similar issues can arise in other areas (including, apparently, some aspects of machine learning). Of course, turning back to game theory, in some repeat-game systems as the number of players increase the outcomes can asymptotically converge to certain equilibria, but I am not aware of any literature that would allow us to make a general conclusion about the likelihood of this outcome across all sectors of the economy.

²¹ Kühn and Tadelis, *Algorithmic Collusion* (2017), available at http://www.cresse.info/uploadfiles/2017_sps5_pr2.pdf.

²² *In re Bristol-Myers Squibb Co.*, C-4076 (complaint filed Mar. 7, 2003) <https://www.ftc.gov/enforcement/cases-proceedings/0110046/bristol-myers-squibb-company-matter>.

²³ *In re Rambus, Inc.*, Dkt. 9302 (Commission order Aug. 2, 2006); *rev'd Rambus, Inc. v. FTC*, 522 F.3d 456 (D.C.Cir. 2008).

²⁴ See Statement of the Commission, *In re Union Oil Co. of California*, Dkt. 9305 (complaint filed Mar. 4, 2003) https://www.ftc.gov/system/files/documents/public_statements/568351/050610statement9305_0.pdf

administration launched technology industry cases involving vertical contracts and unilateral conduct—notably, the pending case involving Qualcomm.²⁵ If I were in the technology industry, I would not assume that there will be any reluctance about identifying and pursuing unilateral conduct cases. But it's important to remember that we have a legal framework that requires that we be able to show that competition—not just competitors—has been harmed by the challenged conduct, and, in most cases, that the firm engaging in the conduct is a monopolist or is likely to become one.²⁶ That is a legal constraint on what we do, and it wisely keeps us focused on avoiding making an unwarranted inference of competitive harm from mere tough competition.

Conclusion

In conclusion, the antitrust laws are robust, forward-looking, and demonstrably capable of evolving with the times. Those laws are fully applicable to the technology industry. At the FTC, we believe in aggressive but thoughtful and fact-based enforcement. We believe in the continued evolution of antitrust law along with developments in economic learning. We engage in antitrust R&D, and retrospectives, and we're not afraid to bring hard cases.

²⁵ FTC v. Qualcomm (N.D. Cal. filed Jan. 17, 2017).

²⁶ Verizon Commc'ns Inc. v. Law Offices of Curtis V. Trinko, LLP, 540 U.S. 398, 408 (2004).