FEDERAL TRADE COMMISSION

COMPETITION AND CONSUMER PROTECTION

IN THE 21ST CENTURY

Wednesday, November 14, 2018

9:00 a.m.

Howard University School of Law

2900 Van Ness Street, NW

Washington, D.C. 20008

FEDERAL TRADE COMMISSION

I N D E X

WELCOME AND INTRODUCTORY REMARKS

MR. HOFFMAN: Well, good morning, everybody, and welcome to the seventh FTC hearing on Competition and Consumer Protection in the 21st Century. I have been told I have about an hour and a half for these introductory remarks -- no, I'm just kidding. Don't worry, don't worry. I won't take nearly that long.

But let me welcome you. I think these are an incredibly important series of events. We have fantastic panelists who have really important and interesting things to say, and I think it's going to help us create a record that will be very useful for a long time to come.

Let me start by giving a couple of quick disclaimers. First, everything I say today in this brief introductory speech will be only my personal remarks, not necessarily the views of the Federal Trade Commission or any Commissioner. And let me also, by the way, thank Howard for hosting this event. It's a real pleasure to be here.

And, parenthetically, if there are any students who come into the audience or are watching or listen to any of this, you're thinking about careers in antitrust, I encourage that. Think about it hard. It is a great career, and call me.

The other disclaimer I wanted to give is for those of you who were not sure what those giant apparatus in the back were, they are cameras. This event is being photographed and webcast. It will be posted to the FTC website. And by participating in the event, you consent to these terms. So just to be clear, if anybody does not want to be on camera, now is the time to make your quick exit.

I thought I would start by just briefly talking about the purpose of the hearings, why are we doing hearings on competition and consumer protection in the 21st Century and why are we doing hearings on artificial intelligence? I know that Professor Gavil spoke about this, and I wanted to echo the educational purpose, the importance of the educational purpose of these hearings.

At the Federal Trade Commission, we are very much in study and learning mode on the issues of antitrust and its application to modern and developing technologies. We think debate and discussion is critical -- central to the advancement of knowledge and understanding of the development of good competition policy in these areas.

We recognize that we and probably everybody in the world have a lot to learn about these topics, a lot to think about. And it’s, we think, incredibly important to bring together thought leaders and experts on these issues so that we can have the kind of debate that will inform our decision-making. Facts are critical; understanding is critical. When you’re developing regulatory or enforcement philosophies, it’s vital that you have a robust foundation in fact and a robust foundation in theory.

And so as we began the process of putting hearings together, as we started looking around the landscape of the antitrust world these days, one of the things that was immediately apparent was there was an awful lot of discussion, but there was not a collection of thinking, a collection of fact, a collection of theory that would enable the development of policy on the kind of foundation that I talked about.

So recognizing that, that gap, I guess, in the underpinnings of enforcement, Chairman Simons thought one way to address it is, and Bilal obviously played a huge role in putting this together, was to convene hearings of this sort, hearings similar to those that Chairman Pitofsky put together.

Now, let me turn from that to algorithms, artificial intelligence, machine learning more specifically. To say that there's a robust debate about the role that these rapidly advancing technologies play in society at large in our everyday lives and in antitrust enforcement would greatly understate the issue. I actually spend a lot of time reading about this. I will confess to understanding almost nothing about it because the technologies are so sophisticated, but I read a lot about it.

A few days ago, The New York Times quoted Facebook's founder as stating that in the next five to ten years Facebook will develop artificial intelligence that outperforms humans in all human senses, including cognition. Data scientists at Google have made similar projections. And if you read Sapiens, a book that came out recently, you'll find at the end of it a discussion about whether or not humanity is on a path to replacing itself with some form of artificial intelligence, which has, of course, long been speculated about in science fiction, notably in Terminator, which we don't think is a huge issue right at this moment, but maybe the next set of hearings down the road, you know, 20, 30 years from now.

There's, of course, a lot of skepticism about this, and one of the things I found about artificial intelligence, I spoke at a conference in Brussels about a year ago, maybe 13 months ago, and there was a great deal of discussion among lawyers about the implications of artificial intelligence and algorithms. And I discovered from talking about them that I think there was literally no one in the room who understood anything about how those technologies worked or what their actual capabilities were.

And in the course of that, one of the panelists referenced a paper that had been written actually by Kai-Uwe Kuhn and his coauthor Professor Tadelis, that talked about empirical work on artificial intelligence and what algorithms and artificial intelligence were actually capable of doing at the time, which was considerably intentioned with the views of the lawyers about what it can do, which frankly I think we're largely informed by Terminator.

So that, to me, reemphasized the importance of actually developing a foundation and understanding of what these technologies can do, and with that I'm going to turn a little bit to some discussion of the technologies and their implications. Now, when I talk about these technologies, I'm going to use the term "technologies" broadly, or I might use "algorithms," but I mean by it to group algorithms, artificial intelligence, and machine learning together.

I recognize that doing that is inaccurate. These are not the same things. They arguably represent points on a continuum of machine learning or machine approaches to solving problems, but there's actually very considerable differences between machine learning and simple algorithms, between artificial intelligence and different kinds of artificial intelligence, and they may have different implications for policy.

But for purposes of today's brief remarks I'm not going to try to delve into those differences. I'm going to treat them sort of monolithically. We heard yesterday at the hearings about companies and experts involved in the technological side of this about how some of these technologies are used in the marketplace, what some of them do, what some consumer protection implications of these issues are.

Today, we're going to talk more about competition policy. The first panel today is going to talk about whether algorithms can collude or might be able to do so in the future. We're going to have another panel that's going to talk about competition, innovation, and market structure questions that revolve around the use of these technologies. And then we're going to have a panel that wraps up that talks about legal and regulatory issues going forward.

Now, these are hot issues around the world. I think I obviously get a lot of literature or bulletins on upcoming conferences. And I think it would be fair to say that 95 percent of the upcoming competition law conferences involve, at least in part, panels on algorithms, artificial intelligence, machine learning, and technological implications for antitrust policy.

We, being the United States antitrust agency, submitted a paper to the OECD Competition Committee last year that provides an overview and discussion of some of our thinking on these topics and in particular on algorithms and collusion. But we also noted in that paper that consumers have benefitted a lot from these advances in technology, not just because they drive economic growth, but because they provide low-cost services, they provide higher quality goods and services, more choices, and innovative new products.

So is this a one-way street? Are these technologies merely beneficial? Is there really any basis for any particular competition policy concern? Clearly, there is. Despite the benefits these technologies can bring to consumers, it's easy to see at least possibilities in which competitive dynamics could be put in play by the technologies.

Let me talk about a couple specific examples. Number one, is it possible that machine intelligence, artificial intelligence, could actually collude by itself? So imagine that you have -- and algorithms, I think, won't suffice for this -- but imagine that I have artificial intelligence where I have machines that are engaging in cognition in some sense, I mean, leaving aside the almost metaphysical question of what cognition actually means, but is it possible that machines could collude in the sense of explicitly agreeing on price, output, customer allocation, market allocation? And, if so, what does that mean for antitrust policy? Can you put a machine in jail for example?

Second, and I think arguably you have much more shorter terms, much more short-term significance, is it possible for machines to reach the oligopoly outcomes more quickly or more sustainably than humans can? And let me just digress for one second on that. One of the foundational principles of merger policy is that we want to prevent mergers that result in firms acquiring the ability to achieve an oligopoly outcome and pricer output.

And what I mean by that is in a noncooperative oligopoly, you could nonetheless have a situation arise where output is reduced or prices increase towards the cartel outcome or towards the monopoly outcome because relatively small numbers of firms can reach the conclusion that it is in all their interests to restrict output or raise price and that the cumulative effect of doing so is beneficial to all. So the payoff is good, in essence, if you collude without colluding.

And this does not involve direct communication; it doesn't involve meeting in the back rooms of restaurants in New York like the book publishers did, for example, in the e-books case. It doesn't involve the kind of thing that you could be put in jail for. So this is a big concern of merger policy because once a merger occurs that creates this kind of condition there's not much we can do about it. Section 1 of the Sherman Act doesn't reach it anymore. So we spend a lot of time thinking about mergers that would enable that outcome to occur so we could prevent it.

So a question is, well, can algorithms collude in this sense, in the sense of independently and without communicating with each other reaching a price-raising or output-reducing outcome better than humans can?

A third possibility is could machine intelligence, algorithms, technology achieve or cement market power by enabling unilateral strategies to acquire, for example, or to destroy competitors before they become a threat? Is it possible that the use of sophisticated technology to survey the landscape and to monitor activity will enable dominant firms to identify threats and extinguish them before they become real threats in some way that is superior to what humans currently could do, and, if so, what do we do about it? And I'll come back to that last point in a second.

And, then, of course, there's other, right? There's a broad category here of things that could happen that we don't really know about. Could, for example, algorithms improve price discrimination? Price discrimination is not necessarily a bad ting. In a lot of contexts, it's welfare-enhancing, but also it has some other implication.

So I think also when you think about all these issues you then have to say to yourself, and if so, let's assume any of these things is possible, what would we do about it? And let me just tackle the noncooperative oligopoly outcome point briefly in this. Let's assume that it was, in fact, possible for algorithms to independently determine that the best outcome for each of their independent firms was a pricing or output strategy that caused prices to rise or output to fall towards a monopoly-type outcome or a cartel type outcome. But each algorithm is simply implementing the most rational economic choice for the company that's using it at any given time.

Is our solution for that to require companies to program their algorithms to behave irrationally, to make bad decisions? Is that really a logical consequence of antitrust policy? Is it a necessary consequence? I raise that not because I think that's actually the right outcome or the right set of choices that we would have but simply to suggest that it's not enough to identify potential problems but you also have to think about what are possible solutions and what are the implications of those solutions, assuming the problem even exists.

Now, fundamentally, at this stage, this is an early, early stage in the development of these technologies. I have in my pocket here two iPhones because I've got the government-issued phone and my personal phone. This technology is basically about ten years old. It’s ubiquitous -- a smartphone, that is. It makes use of a series of other technologies which are, in many cases, less than ten years old. It’s really difficult to see where all this is going to go in the next 10, 20 years. We don’t even fully understand it today. And that, in fact, is the purpose of this panel -- this series of panels and the hearings that we’re doing in this to determine, as best we can, are these technologies likely to sharpen competition, reduce competition, or do both or neither, and, if so, how do we address these issues?

I think also one last point on this. There is some real grounds for caution here. We want to be very careful not to regulate or enforce without the kind of empirical, factual, and theoretical framework that I mentioned earlier. Ignorance is not a path to wise policy. I’ve heard suggestions occasionally that we don’t really understand technology, we don’t understand artificial intelligence, we don’t know what it’s going to do and, therefore, we should regulate it. That may be so in the sector or regulatory context, but I think it’s terrible competition policy.

For competition policy, what we need and what we have historically emphasized, and this is a point that Bill Kovacic, a former Chair at the FTC made, and I'll circle back to this in a second, is we have tried to do the R&D first to figure out the issues first and then develop policy on that kind of foundation, and that parenthetically is an incremental process. We're always learning and always trying to improve what we do, but we don't act before we have some understanding. Bill called it the R&D of competition policy as part of the NDA of what we do in antitrust. I think it's critically important. That is what these hearings are all about.

And on that, let me thank all of our -- on that note, let me thank our panelists in advance. Let me say that I think the -- as I said at the beginning, the record that this is going to generate will provide the foundation for the policies that we need to consider in the future, and I'm very grateful to everybody for making the time to be here today. Thank you.

(Applause.)

ALGORITHMIC COLLUSION

MR. RHILINGER: Great, Bruce. Thank you very much for that introduction. Much appreciated way to get us started.

Now we're going to start our panel discussion on algorithmic collusion. Good morning, everyone, and thanks again for being here. My name is James Rhilinger. I'm a Deputy Assistant Director in the Mergers II Division at the FTC's Bureau of Competition. My comoderator is Ellen Connelly, an Attorney Advisor in the Office of Policy Planning at the FTC. We want to welcome you to our panel. We have a very accomplished group of panelists today. Bruce referenced the robust debate going on in this area, and I think we've got the right group of folks to cover that with you.

There are more detailed bios online, but just very briefly, starting next to Ellen, we have Maurice Stucke, who is a Professor at the University of Tennessee College of Law and Cofounder of the law firm the Konkurrenz Group. He's also a senior fellow at the American Antitrust Institute and on the board of the Institute for Consumer Antitrust Studies. Maurice advices governments, law firms, consumer groups, and multinationals on competition and privacy issues.

Next, we have Ai Deng. Dr. Deng is a Principal at Bates White, an adjunct faculty member at Johns Hopkins University, and an invited expert for the Romanian National Council for Scientific Research. He has over a decade of experience in litigation, business counseling, and academic research, and he has worked on some of the largest price fixing and market manipulation cases of the past decade. His current research interest focuses on the intersection between technologies and antitrust.

Then we have Kai-Uwe Kuhn, who is a Senior Consultant to the competition practice of Charles River Associates. He’s also a Professor of Economics and Deputy Director of the Center for Competition Policy at the University of East Anglia School of Economics. Previously, he was Chief Economist at DG Comp, where he worked extensively on antitrust issues in financial markets and the internet economy.

And after that we have Rosa Abrantes-Metz, who is a managing director in the antitrust, securities, data mining, and financial regulation practices of the Global Economics Group. She’s also an Adjunct Professor at NYU’s Stern School of Business. She works on matters involving collusion, manipulation, and fraud in a variety of industries and has published many articles on econometric methods, screens, conspiracies, and manipulations.

After that, we have Sonia Pfaffenroth, who is a Partner at Arnold & Porter, where her practice focuses on complex antitrust investigations, litigation, and client counseling. She recently coauthored an advisory paper on the antitrust implications of pricing algorithms. Previously, she served as the deputy assistant attorney general for the civil and criminal operations at the Department of Justice's Antitrust Division, where she oversaw some of the DOJ's most significant antitrust matters.

And, finally, definitely last but not least, we have Joseph Harrington, who is the Patrick T. Harker Professor of Business Economics and Public Policy at the University of Pennsylvania's Wharton School, and is Department Chair in the Business, Economics, and Public Policy Group. His research is widely published and currently focuses on collusion and cartels, with the objectives of understanding observed collusive practices, developing observable markers of collusion, and designing competition policy to deter collusion.

Each of our panelists will have between five and ten minutes to make brief opening statements, and we then move on to moderated Q&A. As we did yesterday, we will take questions from the audience. If anybody in the audience has a question, please flag down one of our conference staff for a comments card; they'll collect them and pass them to us.

And so with that, we'll start off with Maurice.

MR. STUCKE: All right, well, thank you very much for this invitation. A few years ago, Ariel and I, we were thinking about the migration to online and online pricing, and we thought what would be the implications then that might have on price fixing. Can computers collude? And so what we came up with were four possible scenarios of collusion. And the first one, messenger, is the easiest. And, there, humans collude and they use then algorithms to help perfect their collusion.

And this is really for antitrust a no-brainer. You have evidence of an anticompetitive agreement, the illegality inheres in the agreement, and intent evidence plays a lesser role. And we already have a couple of cases along these lines. First is the Topkins case in the U.S., and in the U.K. it was against Trod and GBE.

The second scenario is hub and spoke. And, here, you have a series of competitors that are using the same common algorithm. And one way to think of this would be platforms such as Uber, whereby the users, the consumers, as well as the drivers, the pricing was all determined by a single algorithm.

And then the second would be when multiple competitors are outsourcing their pricing to the same third-party vendor. So here you have a series of vertical agreements, and the issue is when do those vertical agreements become a hub-and-spoke cartel? And, here, we could see that you have evidence of an agreement, it’s really how you classify the agreement, and you can look at possibly intent evidence to then determine what the likely anticompetitive effects might be.

The third scenario, predictable agent, is trickier. Here, you don’t have evidence of any agreement. There’s no meeting of the minds. But there’s strong evidence of anticompetitive intent. Each firm unilaterally decides to use, let’s say, a price optimization algorithm. And the industry-wide adoption of this algorithm helps foster what we call tacit algorithmic collusion. And this presents various policy changes that I’ll address at the end.

And then the final scenario, which is probably more in the future, is digital eye. Here, there’s no evidence of agreement, nor is there any evidence of anticompetitive intent. Each company utilizes a price optimization algorithm, let’s say through machine learning. The algorithms then all determine that the profitable outcome is tacit collusion.

So we don’t -- the owners of these algorithms don’t know necessarily if and when their algorithms are colluding, but nonetheless, it has the same effect. So what, then, are some of the policy implications of this? Well, for messenger, the first one, there really isn’t any concern. Our tools are well equipped to address that.

Second, for hub and spoke, we still have the tools to address that. It’s going to be trickier than how you characterize that agreement and what sort of guidance can the agency give market participants of when a series of vertical mergers -- vertical agreements, rather, raise antitrust concerns.

But the last two, and I think that’s what we’re going to largely talk about today, will likely then raise more significant policy issues. So does our current policy towards conscious parallelism apply when price optimization algorithms can enhance firms' ability to tacitly collude? And we're not saying that tacit collusion will occur in every industry, but in industries where tacit collusion might be on the margin, will algorithms help then push it over the edge? And so you might have industries where four to three, five to four mergers, in industries characterized with algorithms may be more acceptable to tacit collusion.

Second is our legal concept of agreement outdated for computer algorithms? Are current laws sufficient to deter and prevent tacit algorithmic collusion? Third, how can the agencies identify when algorithmic collusion occurs, especially when pricing is dynamic. It's very difficult to detect express collusion. Are the tools up to snuff to detect tacit collusion?

Next, what additional measures should be considered to reduce the additional risks associated with the use of price optimization algorithms? So our book really wasn't based on Terminator; it was based on discussions with computer scientists who raised these concerns. And, moreover, when you look online, what do they promote? They promote avoiding price wars. They promote enabling companies to maximize profits. They talk about how pricing is maybe good for the consumer but bad for the business. And they can help companies avoid these price wars.

Now, is this just puffery, or is this actuality? And I think we're going to talk about what other agencies are doing. So I think it's very important for the FTC not to discount this as Terminator, but rather to take this seriously like many of the European officials and start devoting resources to this. That's why I very much as encouraged that Bruce and others at FTC held this important policy hearing today.

And then, finally, in what ways should firms be obligated to integrate ethics and legality into a computer program? And to what extent are companies going to face liability for their algorithms? To what extent will independent software developers face liability? One of the interesting things in Trod, I don't know to what extent, but it seems that the companies were going to the software developers and saying, this is not working, we need to tweak this in such a way.

If the software developer was aware that these algorithms were being used to help a cartel, should they be liable? And to what extent are companies, should they have an affirmative duty to program their computers so as to not tacitly collude? And is that even possible? Those are other policy issues that I would encourage the FTC to explore. Thank you.

MR. RHILINGER: Next to Dr. Deng.

DR. DENG: Thanks, Maurice, and thanks, Bruce for setting the stage for the discussion. I also want to thank the FTC for inviting me here. It's an honor to be here today and to speak to you all this morning. For me, it's always fun to join a conference where my name is on every single slide or in caps, so very happy to be here.

As Bruce and Maurice just summarized, we really have seen a great deal of interest in and concerns with algorithmic collusion. What appears to be particularly troubling is the type of algorithms that are capable of collusion, tacit or explicit, all by themselves without human interference.

There are at least two interesting questions in this discussion. The first is obviously just how close we are to having colluding robots that are production-ready, ready to be deployed by businesses. And, secondly, if so, what can we do about them? What can we do about potential antitrust risks?

I'm going to argue that we can go a long way in answering those questions by taking a close look at the literature of economics and artificial intelligence. Now, the existing literature has already a lot of insights to offer. Now, I'm not saying we have all the answers yet, which is why the discussion that the one, like the one we're having today, is still so relevant and important.

Okay, so what do I see as some of the most important lessons we can learn? First of all, there is clear experimental evidence that an algorithm or a robot could be designed to tacitly cooperate with opponents in environments such as, you know, social dilemmas, such as prisoner's dilemma, which is kind of a protocol -- in prototype models that economists study competition.

So in these experimental settings, I would say colluding robots are no a longer science fiction. Secondly, I guess fortunately for us, designing an algorithm to tacitly collude turns out to be a very challenging technical problem. Now, I'm not going to list all the technical challenges here, but I just want to give out one example based on my recent AI research that is published just earlier this year.

So the researchers pointed out that a good algorithm must be flexible in that it needs to learn to cooperate with others without necessarily having prior knowledge of their behaviors. But to do so, the algorithm must be able to deter potentially exploitative behavior from others. And I quote, “when beneficial, determine how to elicit tacit coordination -- cooperation from a potentially distrustful opponent who might be disinclined to cooperate.”

The researchers of the study went on to say that these challenges often cause AI algorithms to deter -- defect, I should say, rather than to cooperate. And I quote, “even when doing so would be beneficial to the algorithm’s long-term payoffs.” Now, there are several reasons why the fact that there are, you know, a lot of technical challenges in designing such an algorithm is relevant to us in the antitrust community.

First, I would argue that, you know, they show that there’s perhaps a lack of support for a popular belief that just any learning algorithm, any kind of machine learning algorithm that tries to maximize a firm’s individual profits would necessarily and eventually lead to tacit collusion.

This also tells us that to design an algorithm, then, has some degree of guaranteed success in eliciting tacit coordination from opponents or competitors. This capability to collude most likely needs to be an explicit design feature. Now, this observation itself has further implications. First, it suggests that at least from an antitrust policy perspective we ought to consider the possibility of prohibiting the development and incorporation of certain inclusive or problematic features while balancing the pro and -- you know, potentially pro and then anticompetitive effects of algorithms. And Joe here actually wrote a recent article in which he explored some of the issues, including this one.

Second, as a result of the challenges, there may very well be important leads in the records that antitrust agencies and even private parties could look for in an investigation or in a discovery process and all without technical expertise. Several documents are going to be of particular interest. For example, documents that shed light on the design goals of the algorithm. Documents -- any documents or any document behavior of the algorithm, any documents that suggest that the developers may have modified or revised the algorithm to further the goal of tacit coordination. Those are going to be very, very helpful.

Now, another type of document I think really should raise red flags is any marketing or promotional materials that suggest that the developers may have promoted their algorithm’s ability to elicit tacit coordination from competitors to their customers. Now, what’s interesting here is that I hope you can see that it’s not necessary for the investigators to have any sort of intimate understanding of the AI technology to look -- number one, look for such evidence and even interpret some of those evidence.

Another important lesson I think we can learn from the AI research is that at least if you look at academic literature, the algorithms being designed are not necessarily what economists call equilibrium strategies. Equilibrium strategies are intuitively stable in the sense that, you know, I’m going to define this loosely, we have economists, you know, on the panel here, so I’m going to define this loosely.

Equilibrium strategies are stable in the sense that, you know, if you and your competitors know that all of you are adopting certain strategy you will have no incentive to change, right? This is known as Nash equilibrium and game theory. As two recent -- as two AI researchers put it in a recent article, the question of designing a good agent for social dilemmas, kind of like the competition environment, can be sometimes very different from computing equilibrium strategies.

Similarly, in another recent AI study, despite the promising experimental findings, the researchers acknowledge that unless their learning algorithm is an equilibrium strategy, it can be exploited by others, meaning that the players who started out using their algorithm may have incentive to deviate, to move away from their algorithm. This means that, you know, if a firm happens to adopt an algorithm that is a nonequilibrium strategy, they may have the incentive to move away from that and, as a result, potentially disrupt the potential inclusive environment.

I'll just talk very briefly on economics literature, and I'm sure my copanelists are going to have a lot to say on this. So there is one literature in economics that studies the interplay between information flow and cartel stability. One early and seminal paper shows that in an environment where firms have very flexible production technology, so you can change a production level very, very quickly, and if the information arrives continuously, it turns out that the cartel becomes very difficult to sustain.

Okay, and further study even shows that in that environment one way to sustain the cartel is actually to intentionally delay the information flow. Now, to me, this is a very relevant line of research, because presumably if you think about algorithms, robots, they are potentially much more capable in processing and collecting information potentially in real time and really, really quick.

In a recent article of mine titled “Four Reasons why We Won’t See Colluding Robots anytime Soon,” I made two more points. I have time to just talk about one. That is, despite the fact that algorithms, which are, you know, computer codes, right, are undoubtedly hard to interpret, especially for many of us in the antitrust community, I do want to note that cartels may affect themselves in other ways that are observable and interpretable.

In fact, economists and courts have long been well aware of what’s known as plus factors, right? To quote a paper, plus factors are economic actions and outcomes, above and beyond parallel conduct, but are largely inconsistent with unilateral conduct, but rather, largely consistent with explicitly coordinated action.

So I won’t give an example here in my

opening remarks, but we can get into some of the examples. With that, I'm going to close my remarks and look forward to the discussion. Thank you.

MS. CONNELLY: Thank you, Dr. Kuhn.

DR. KUHN: Well, thank you very much as well for the invitation. It's very nice to be here and participate in this discussion. And some of the things that I have to say really come from some of the research on collusion, especially the experimental research that I've been doing in recent years.

I think in order to think about policies in this area, it's really important to understand what issues we're exactly addressing. And one of the things that I'm concerned about in this debate is that that sometimes gets mixed up. That is of particular import in terms of the ways that collusion theory is being used because they're two really very separate, and different parts of collusion theory that are both important but where we know a lot more about one than about the other. Or what about the other we now know a lot more, but that's not generally very well known.

One aspect, and that is what enforcement really targets, is how do we actually come to a common understanding of what we should be doing and what are the consequences of if we're not doing it or if we're actually sticking to the agreement. That's what we usually call the coordination problem in that context. And that in theory doesn't play very much of a role because it's very, very hard to model in a polite way what coordination activities are, how they work and how their effectiveness changes in different market environments. So there's basically very little kind of theoretical work on that aspect.

The other aspect is what I call the stability of cartels, do I have an incentive to deviate, because I always have? If I raise the prices, I have an incentive to deviate; therefore, there needs to be some punishment on the market. If it's tacit collusion, that has to be implicitly learned or intuited.

But we have the literature that says if we can coordinate on an outcome, can we sustain it, and under what circumstances are there more outcomes that we can sustain, but it doesn't tell us really anything about the likelihood that in a particular market situation we are going to see collusion. So that's what's really the question to understand, when do we actually see coordination. Is something that's coordination activity usually talking about it, something that's essential or not? And that leads to the question with coordination, how likely is tacit collusion actually?

And what you want to do in the policy area really depends on whether you think the coordination problem is relatively easy to solve in AI or algorithmic acting is going to make tacit collusion a lot easier so that coordination is less of a problem, or whether you think, well, maybe the rapid interaction is good for stability, but it doesn't really affect coordination all that much, because in the first case, you want to just use the existing and maybe expand and adapt instruments on enforcing against coordination activity. In the other case, you have a real problem, and those are the kind of things that Joe, I think, has been thinking about.

Now, I believe, and this is something that's very important, is that out of the research in the last 15 years, we've actually learned that coordination is actually much harder than we always thought, especially in situations that are relatively complicated. There's an experimental literature on coordination games that has shown already in the early 1990s, even if you have ranked equilibria, you might actually go to the worst one if people are doing it experimentally.

And the reason is if you're trying to achieve something that's very good for everybody, if someone isn't coordinated, that's really bad. And just the fact that you want to ensure against that, then under those circumstance kind of leads to very bad outcomes. And I've argued many years ago in a policy article on collusion that the reason why you want to enforce against coordination activity is precisely the fact that if we don't see that, we're going to have a reversion to very competitive behavior because collusion models have that structure that it's actually very risky to collude at high prices, because if someone else doesn't understand it and get it and we don't have a fully common understanding, then that's very risky and you want to ensure against it and that brings the prices down.

That's what we kind of see in those things. We do see in a lot of situations that there's collusion but very much from what you've heard about algorithms, people have run these things in the past, on simply two-by-two games -- two strategies, two players. And, there, you've got a lot of experimentation between people because people do experiment, and you see a lot of what happens with contingents.

Now, the interesting thing is if you're going into the experimental literature and have three players, usually you don't get the coordination without communication and it just all collapses. We've even seen this a lot in two-player situations, as soon as the games get a bit more complex, you have price setting with capacity constraints, you have a larger set of strategies. Kind of in the first place we tried to write an experimental paper on coordinated effects of mergers, and I couldn't get the guys to tacitly collude, it just wouldn't work. As soon as they communicated, the theory worked out perfectly.

And we see in all of that literature, at least from a minimum of three players onwards, if you can't communicate, collusion just basically is very rare. And the same thing happens if, even if you just announce prices, right? That's not enough because what the coordination really involves is learning how one should be thinking about contingent strategies, which are very complicated coordination to do, okay?

So the question here is, if individuals can't do this very well, would algorithms do this a lot better? And one of the arguments are that they're -- you know, they're profit-maximizing, uncompromising on profit-maximizing. They're really good. We're just a bit more boundedly rational and so they're going to get there much better.

Now, the reason why that is not right is that the coordination problem as such is something that you can't solve by rationality. You cannot reason through by knowing that you're rational that everybody knows that everybody else is rational. You can't reason through how you should be playing something that in principle has two equilibria.

So what we're consistently seeing in those types of situation is that the thing that brings you out is actually talking about it. And basically making sure that you come to a common understanding. That's been the subject of a paper -- of an experimental paper we've written where we've analyzed the communication, and the really effective thing was to communicate about contingent strategies and say, if you don't, then I'm going to punish. And the other guy says, Why would you do that? And they have a long conversation until they understand why that makes sense, and then they implement it. When they don't do this, they basically don't get to collusion in the long run.

Now, if you're taking that to the algorithms, you're kind of asking your question, do we have anything else that might tell us that if it’s just an algorithm we might have the similar problems. There’s an interesting literature out there from the early 1990s where people were doing dynamic evolutionary games, not evolutionary stability, but it has the same thing where you say what’s an equilibrium, does someone deviate?

All the questions we’re asking with algorithms is how do you get to the agreement, how do you get to equilibrium, right? And, there again, there is a very strong result out there that says if you have this type of evolutionary games as they were specified then, which I think you could think about as a genetic algorithm as well, you will get something that’s called a risk-dominant equilibrium that is this problem of going very high to a high price but then having bad payoff if someone is not coordinated is actually a very large one, and you’re selecting these -- but the push in the collusion games would be going towards lower prices.

So I think the question that is -- you know, is there anything that we would know from the AI literature -- from the artificial intelligence algorithm literature that would tell us that algorithms would have less coordination problems.

There are specific situations in which algorithms are very good at that.

And I haven’t quite seen that, and I was thinking I would be telling you that there’s all this literature out there where this might actually be done, and I’ve seen literature on algorithms that do get to collusion, but again, they’re in the context of very, very simple gains, and the complexity of this with as soon as you’re getting to something with realistic markets, it gets much, much higher. And dimensionality is there kind of a curse in all situations.

So I think once you start thinking about it in this way, there’s kind of the question, well, there are a lot of things that you can do with the current instruments. There’s literature that would suggest that, yes, if you’re exchanging your algorithms, both sides know what it is, you might get to collusion, even if you’re not explicitly talking about it. Well, that’s like information exchange where you’re telling others what your proposed price is. Actually, it’s even more than that. You’re telling them what your contingent price is for all eventualities in the future, right? I would think that would come under the typical prohibitions of information exchange on prices that we already have.

I think that the way to think about some of these things is, you know, can we think about how coordination, the mechanism, work. Can we give obligations on transparency on those types of things were that is necessary? And do we have to kind of come to some kind of transparency, for example, on issues where we would have AIs, like, communicating and what would be meaningful for regulation. But I think that's more the issue and that's what I'm much more concerned about than rampant tacit collusion.

MR. RHILINGER: Thank you.

Next up, we have Dr. Abrantes-Metz.

DR. ABRANTES-METZ: Good morning, let me start by thanking the invitation to be here. It's a pleasure to be here. I would like to take a step back and think about algorithms in study in a little bit of a different way. If as economists we think about the situation where we have many competitors, we have homogeneous products and cost prediction functions, we have perfect competition and no entry -- perfect competition means full transparency about everything -- then we have perfect competition. Price is equal to marginal cost. That's the socially desirable outcome, and that's what economists take as the benchmark and compare real market outcomes against.

So then the question becomes actually whether pricing algorithms, given that they are associated with higher transparency and through them there's a higher chance and normally it happens that you can more quickly respond to changing market conditions and competitors, including aren't they actually fermenting more -- the likelihood that we will see more perfect-competition-like outcomes then instead of collusion.

So I think we need to start by thinking about taking this as the benchmark and then start thinking about as we deviate from it, is it really more likely that we're going to see tacit collusion coming out of these algorithms or not. I think that there is, even given the limited empirical evidence to date, a high chance that we're talking -- that we're going to see higher and more fierce competition coming out of these algorithms than necessarily a lot of evidence of additional tacit collusion. That doesn't mean that that has not already occurred and that it won't occur. The question is whether the likelihood is higher or if those are more isolated events.

So I think what we have to understand really also is that both situations will lead to similar prices among competitors. Perfect competition will lead to completely identical prices, but low prices, and the tacit collusion will lead to equal prices at a higher level. And so we need to be able to distinguish the two situations if we're saying that algorithms tacitly collude and they are leading to equal prices, well, are those prices necessarily too high? Is that a necessarily highly undesirable social outcome?

So we know from theoretically that it is possible that particular market structures will enable the enabling factors of collusion when pricing algorithms are used. But I think what is really important to understand is whether the empirical evidence backs that up and also how do pricing algorithms actually change what's called the plus factors in a way that make it hard to provide the general rule as to whether tacit collusion is more likely to occur or not.

Of course, we always start with thinking of the situation where we have just a small number of players. We have high barriers of entry, some high product homogeneity, and then because pricing algorithms are usually going to work in high transparency worlds and they enable more interaction, they can even replace the direct communication among competitors, then it is possible that they will facilitate tacit collusion in theory because they facilitate signaling potentially, they facilitate the monitoring of prices, and they facilitate the punishment of deviations from a potential collusive agreement.

But as it has been mentioned earlier, what we are worried is that these kinds of concerns that are typically in the oligopolistic situation will extend to situations were markets are less concentrated. But let’s start by thinking also how do price algorithms and the availability of so much data and market transparency actually affect some of the components, some of the market structure, and the maintenance supply factors that would normally tell us that if X exists, then collusion is more likely or not.

Let’s think, for example, just to give a couple examples in terms of demand. Everything else the same, typically the availability of these pricing algorithms in retail internet trading is going to reduce -- is going to increase, I’m sorry, the elasticity of demand by consumers simply because it’s much more easy -- it’s easier. The search cost is low, it’s easier to search across different webpages, my elasticity of demand is higher and, therefore, market power is lower.

We can think the same way about barriers to entry. We know that large data in highly concentrated markets may provide an additional barrier to entry. On the other hand, the digital economy is full of examples where those situations were overcome by entrance and in which that level of high transparency actually enabled a reduction of entry costs to the potential entrant.

Also, markets where there’s a lot of innovation tend to be markets that are typically markets in which a lot of these pricing algorithms are applied, tend to be markets that are more difficult to collude upon. So there’s a lot of structural components that do get changed in these situations that make it hard to have that general rule and assessment in terms of the typical plus factors that we tend to use in collusion matters as to whether we should expect, even theoretically, for tacit collusion to be more likely in these situations.

I would now like to talk just a little bit about whatever empirical evidence exist out there that may give us some more information as to whether tacit collusion may be more likely. For example, the S&P 500 releases every year industry-specific returns on equity and profit margins. And every year, systematically, the retail sector has the lowest profit margins of all industries, between .5 and 3.5 percent, and that's particularly true for web-only-based retailers.

So are the prices probably converging to the same level? Probably. Are they monitoring each other? Yes. But they don't seem to be making that much money compared to others. So, again, how likely is it that these pricing algorithms are really going to lead us under certain circumstances to more competitive rather than less competitive outcomes?

And so another example that is particularly more familiar to me because those are the type of cases that I tend to focus on the last couple of decades are cases involving, for example, commodities trading cases and financial markets in general. Over the last two decades, particularly the last decade, there has been a large effort to move trading from over-the-counter to exchanges.

Now, what is just in a couple of words the main difference between the two? Over-the-counter trading, you typically -- the information is not available to every market player. You don't really know what are all of the offers to buy and sell at any moment in time. You have no visibility, no transparency to where the market is, aside from some average value that somebody provides to you. Highly opaque markets.

When these products get moved into exchanges, where at any moment in time you know where all of the market is, you know, what everybody's willing to buy and sell, you don't know who you're buying and selling with until you actually trade and execute the trade, but you have transparency which has enabled a lot of pricing algorithms to emerge and be more widely applied.

What have we observed in terms of market efficiency with this move? We have observed that the bid-ask spreads, which are actually the dealer profit margin, the difference between that which they buy and they sell, have shrank drastically. So we have observed lower prices, even in situations where the exchanges that are more expensive to operate than over-the-counter trading, there's a lot of fees that go into operating an exchange, we actually see that prices are going down.

Now, do we see collusion situations happening? Absolutely. But, actually, we see a whole lot less collusion happening in these exchanges where pricing algorithms are so enabled due to high transparency. Prices are more correlated because everybody is training their algorithm in the same data set, but the episodes of collusion in exchanges that are exchange-specific are actually a whole lot lower. We know we have seen so much collusion and manipulation lately, but those situations -- 90 percent of them -- were related to deficient structures such as benchmarks-rigging, auction rigging, that were themselves deficient, which led and facilitated rigging.

With respect to actual trading that occurs naturally in exchange and in over-the-counter, there is no comparison between the incidence of collusion in these very highly transparent market-based on exchanges and the over-the-counter. So I think that even though the empirical evidence is limited, I think we need to sort out through what is already available out there and think about whether if we are to regulate a problem that we may potentially be misdiagnosing if we're actually going to undercut all the potential benefits that we may have from these techniques. Thank you.

MS. CONNELLY: Thank you.

MS. PFAFFENROTH: Thank you. And I'd like to thank the FTC for the invitation to be here today. It's a pleasure to be here. And I'd just like to start by saying that the views I express today are my own, not those of Arnold & Porter or any of our clients.

So I'd like to shift gears slightly and talk a little bit about enforcement currently. You know, in the current time where algorithmic-enabled collusion still requires human input at some point in the process. And Bruce mentioned the OECD paper that the agencies drafted last year. And that paper drew the distinction between interdependent behavior and collusive behavior. And collusion requires an agreement between two parties.

The enforcers have said that algorithms are a tool, and you have people determining the goals and designing the algorithm to meet the goals of that tool. And as a tool, the algorithm can be a mechanism to implement a collusive agreement. It could be a technology that assists in policing, an agreement that's already in place to deter cheating. But as a tool, the algorithm in that context is sort of the technological equivalent of the stereotypical meeting in the smoke-filled room, where the agreement is reached and facilitated.

So in that context, you have a person, a human being, putting the algorithm in motion and directing it to perform a set of actions in the context of a collusive agreement that is in violation of the antitrust laws. And even if once that’s set in motion it becomes self-executing, there’s still predicate communication. There’s still a predicate agreement between parties that led to that action.

Maurice referenced the Topkins-Trod-Kik. So this was a case prosecuted by the DOJ in which Topkins and his coconspirators were accused of fixing the prices of art, of posters that were sold online through the Amazon marketplace. And in that case, the DOJ was alleging that the coconspirators had used commercially available algorithmic-based pricing software that operated by collecting competitor pricing information and then applying certain pricing rules to that data to set pricing.

And in that case, the way DOJ described the conduct was that specific pricing software was adopted with the goal of coordinating pricing changes. So one conspirator would program its algorithm to look at the price of a nonconspiring competitor and set the price slightly below that, and then other conspirators would set their pricing software to look at the price of the first conspirator, and therefore, through the use of that software, it was executing on an agreement to coordinate pricing changes, to control price.

And the way it was described, after that initial agreement, it was largely self-executing, but there was an agreement at the beginning. And so that enforcement action is an example of competitors agreeing directly within the traditional framework to use that algorithmic software to execute an anticompetitive agreement. It’s an electronic tool. It’s not the first time that electronic tools have been pointed to by enforcement agencies as a tool to enable collusion.

Back in the ‘90s, the DOJ settled charges that airlines that had a jointly owned computerized online booking system were using that as a tool to fix prices. There was also a reference to Uber, and so on the side of the private litigation, there was a case pending in the Southern District of New York, and not commenting on any merits of the case, but just with respect to the framework in which the court looked at that, and the case ultimately went to arbitration instead, but there was a consideration of the merits of the arguments and a motion to dismiss before that happened.

And in that case, you had the court looking at it, as Maurice referenced, a hub-and-spoke framework, where there was allegations that drivers that joined Uber are agreeing with each other to use the same algorithm to set prices. So that that -- that there was a rim and a hub, again within the traditional framework of considering collusive agreements.

If there isn't an agreement between competitors, then algorithms have the capacity to allow competitors to observe more quickly, match prices more quickly and maybe more effective than other types of observation capabilities that companies have had available to them in the past. But without the underlying agreement, it's still parallel conduct. It's still parallel pricing, which is not illegal under antitrust frameworks. And something enforcers have made clear is that independent action -- independent action is still parallel.

So for example, if two competitors independently, without communication, go out and adopt the same pricing software, and that increases the likelihood of interdependent pricing and may even act to stabilize pricing, there’s still no agreement. There’s still no collusive conduct that forms the basis of an antitrust violation.

And so you have had historically the agencies articulating this as focusing on the behavior, focusing on the anticompetitive behavior between parties, not the outcomes of the consequences of certain actions that are taken independently. And so, you know, thinking about it from a business perspective, from the practical counseling perspective, if that bright line weren’t there, that agreements between competitors to collude with respect to price setting is unlawful, independent action that may result in price stabilization but does not involve any communication between competitors is not unlawful.

If that bright line is taken away, it would make it very complex and difficult for a business to determine where the line is, where is market transparency no longer procompetitive and when does it become anticompetitive? You know, when is the threshold for when conscious parallelism, which is lawful, when does that come off? Well, that would be very difficult to define and very difficult to counsel with respect to.

All of that said, I think that even in the current environment, and this is something that others have alluded to and Maurice talked about at the beginning, there is still the opportunity for risk for companies even if they are not engaged in collusive agreements, that certain behavior or business strategies or the adoption of the same pricing software or the use of a common platform could give rise to inferences that there is, in fact, an underlying agreement.

And that's something from a business risk perspective that businesses have to focus on to make sure that conduct which is, in fact, lawful under the antitrust laws doesn't give rise to an inference, potential investigation or litigation risk, that it is, in fact, the product of an underlying agreement. And I'll stop there.

MR. RHILINGER: Thanks very much. And I think that leaves us with Joe.

MR. HARRINGTON: Okay, thank you. And thank you to the FTC for putting together this panel.

Suppose managers at competing companies independently decided to let AI determine the prices they charge. Due to the complexity of AI, these managers are unable to foresee what will result.

Further suppose that these AI programs have learned to collude as reflected in prices above competitive levels. Algorithm collusion has emerged and it is harming consumers.

Now, the legal challenge in prosecuting those companies is that the law is rooted in conspiracy, but there is no conspiracy here. To be more specific, what is unlawful is an agreement between competitors where an agreement is, according to the U.S. Supreme Court, a meeting of minds in an unlawful arrangement, or a conscious commitment to a common scheme.

This legal perspective is also present in European Union jurisprudence where an agreement means that companies have joint intention and a concurrence of wills. In other words, companies have an unlawful agreement when they have mutual understanding to restrict competition.

Now, the courts have laid out various paths towards proving that there is an unlawful agreement. Common to them is an overt act of communication between companies intended to coordinate their conduct. There must be evidence of communication. However, neither mutual understanding to limit competition, nor communication to facilitate that mutual understanding, is present with algorithmic collusion.

The AI programs are simply setting prices, recording prices and sales and other relevant data, and adapting the pricing rule in a manner to yield higher profits. There is no overt act of communication between the managers, nor between the AI programs. There is no mutual understanding to restrain competition between the managers as they acted independently and did not foresee the collusion that would emerge. And there is no mutual understanding among the AI programs unless one is prepared to attribute to understanding to AI.

According to the law, algorithmic collusion is legal because there is no agreement; still, prices are above competitive levels.

Now, in developing a legal approach to prosecuting algorithm collusion, it will prove useful to first ask, why is it that the courts have made communication to limit competition unlawful rather then limiting competition? It is the practice that facilitates collusive pricing which is unlawful, rather than collusive pricing itself.

To elaborate on this point, suppose Company A verbally expresses to Company B that Company A will raise price and goes on to say that it will keep price at that high level only if Company B matches it. Otherwise, Company A will return price to its original low level.

After Company A conveys this message to Company B, suppose Company A raises price and Company B matches it. Based on their communications and their pricing conduct, Companies A and B would be convicted of violating Section 1 of the Sherman Act.

Now suppose Companies A and B use those same pricing rules, whereby Company A raises price and keeps it there if Company B matches the price, and otherwise drops the price back down. Well, Company B’s pricing rule hasn’t matched Company A’s price increase. If the companies use those pricing rules but did not communicate, the result is collusive prices, but they will not have violated the law. There is collusion, by which I mean the use of pricing rules to support supercompetitive prices, but no communication.

Now, the reason that collusion without communication is lawful is because of an evidentiary hurdle. Collusion is about the use of a reward-punishment scheme. If you price high, then I will reward you by pricing high. And if you price low, then I will punish you by pricing low.

One can think of it as a contractual arrangement among competitors for sustaining prices above competitive levels. The evidentiary challenge is that we observe prices but not the reward-punishment scheme that may be sustaining them. The reward-punishment scheme resides in the heads of the colluding managers. If we see one company raise price and the other match it, we cannot be sure that it's a collusive deal or that these price increases are driven by, say, a common rise in cost.

We cannot get inside the heads of the managers to know what is underlying their conduct. Did a manager raise price with the intent that its competitors match that price increase and put in an end to price competition? Or is there a legitimate competitive rationale for companies that raise their prices?

Now, returning to discussing the algorithms collusion, here's the critical observation. While we cannot get inside a manager's head, we can get inside the head of an AI program. At any moment, the program's code includes a pricing rule, which it uses to set price. We can engage in testing to learn the properties of that pricing rule, and, in particular, whether those properties are collusive.

Is the pricing rule designed to punish competitors with low prices? Should they seek to undercut price? It is a pricing rule designed to raise price but maintain it there only if rival companies match that price increase. More generally, is the pricing rule collusive in the sense of using a reward-punishment scheme to sustain higher prices and eliminate price competition?

The realization that we can in principle determine the pricing rule that an AI program is using is the basis for a different legal approach designed to deal with algorithm collusion. This approach makes limiting competition illegal rather than communicating to limit competition. My proposal is to have a per se prohibition on pricing algorithms that limit price competition. Liability would be determined by dynamic testing, which means entry and data into the pricing algorithm, and monitoring the output in terms of prices to determine whether the algorithm is unlawful.

Having established this set of prohibitive pricing algorithms, the burden would be on companies to monitor their AI programs to ensure that their pricing algorithms comply with the law.

Implementation of this legal approach will require extensive research by economists and computer scientists to identify a set of prohibitive pricing algorithms. This set should include pricing algorithms that promote collusion while at the same time not including pricing algorithms that promote efficiency, for example, algorithms that adjust prices in response to demand information.

I believe this is feasible because the properties that enhance efficiency seem quite distinct from those that promote collusion. Towards identifying a class of prohibitive pricing algorithms, I would propose the following three-step research program. In the first step, create a simulated market setting with AI programs that produce both competitive and collusive prices as outcomes. And, in fact, that is currently ongoing.

In step two, investigate the resulting pricing algorithms in order to identify those properties that are present when collusive prices emerge but are not present when competitive prices emerge. Those properties serve to define a candidate set of prohibitive pricing algorithms.

Step three, test the candidate set of prohibitive pricing algorithms by assessing the impact on market outcomes from restricting those pricing algorithms to not lie in the prohibited set.

Now, let me conclude with a kind of cautionary comment. Should at some future time algorithmic collusion occur and should it become ubiquitous, existing jurisprudence would offer no legal recourse of stopping it. Consumers are currently unprotected from algorithmic collusion. To my knowledge, a per se prohibition on collusive pricing algorithms is the only available approach to preventing algorithmic collusion.

While implementation of this legal approach faces some significant technical challenges, they are not insurmountable. But more daunting than those technical challenges is the alternative, which is leaving a massive loophole in the law that would allow companies to limit competition through algorithmic collusion. Thank you.

MR. RHILINGER: All right, I want to thank all of our panelists for interesting opening remarks there. I would like to spend the rest of our time with a moderated question and answer. And to kick things off, we've heard a lot of references, both in the opening remarks of the panelists and in Bruce's introduction about the debate that's going on. There have been some interesting comments here about the ways that we can potentially identify and deal with any collusion that's going on today.

I'm curious to get the panel's reaction on just the sufficiency of the tools that are available to enforcement agencies today. And really you can focus on tools to detect, tools to deal with whatever we find, policy proposals for us to think about. And I thought maybe we could start with Maurice.

MR. STUCKE: All right, well, thank you very much. We have a new paper that we just put up on SSRN, "Sustainable and Unchallenged Algorithmic Tacit Collusion," in which we address some of the concerns, and what we first find is that express collusion is often more durable than what we identify.

Second, what we find is that in the legal world, there is the assumption that tacit collusion can occur without communications. But, third, and I think which is particularly interesting here is recent experimental evidence that justifies some of the concerns that Joe has raised, whereby you have algorithms that then collude when playing with a human. And, in fact, they reach a collusive outcome earlier than when humans -- human and human experiment.

And then also they see tacit collusion among algorithms. They first tried it with 2Q learning algorithms and then they went to 3Q algorithms. They then had 30 price levels. They went up to 100 price levels, and then what they found was that tacit collusion occurred and was very stable.

And, then, finally, we have some real-world evidence, although indirectly, with RPM. There was the recent case that the European Commission brought against Pioneer and other electronic developers. And what was interesting here is because the industry relied on these pricing algorithms, Pioneer only had to go and target, let's say, the one discounter. And then once it did so, once that discounter then increased its price, all the others then followed rather quickly thereafter.

And you see this in some of the literature for the software vendors, how do you identify leaders, how do you identify followers. And if you can identify the leaders, then you can avoid these price wars.

So what should the agencies do? Well, let's look at some of the things that are happening now. First is research projects, and I think that would be key. I mean, the Germans and the French announced in 2018 that they're going to engage in extensive research projects; the European Commission as well.

Second is to have a dedicated team within the agency. The ACCC has a data analytics commission. Third would be looking at some of the policy proposals already on the table. So Germany’s Monopolies Commission had some recent proposals on algorithmic collusion, including systematically investigating these markets to see what risk will likely emerge, because as Joe points out, this can be quite pernicious and detecting actual collusion is already difficult enough, detecting tacit collusion can be really difficult.

And then, finally, what I think here -- one of the things that we recommended in our OECD paper was creating these tacit collusion incubators. And we’re already starting to see scholars doing that. That’s the two studies that we cite in our paper were based on that. But I think this would be an excellent opportunity for the agencies, particularly to better understand under what circumstances will this tacit collusion occur and then prevent it through merger policy.

I mean, I remember when I was at the DOJ. You know, we were told, well, with collusion, stuff happens. We don’t really know when it happens, when it doesn’t happen. We had very good tools for unilateral effects, but not so much for collusion. And these tacit solution incubators or these algorithmic collusion incubators can really give us insights into what conditions may emerge or substantially lesson competition along this dimension.

DR. DENG: I would just echo what Maurice just said. I think he gave a lot of good advice. And to me, I mean, although I said that I do believe that there is a lot we could do even without expert -- you know, technical expertise on AI to uncover and interpret evidence, I do think that having technical expertise within the agency or at least have easy access to that type of expertise I think it’s going to be very helpful.

As Joe pointed out, I mean, if you look at the algorithms, you know, it’s basically saying a piece of computer program and you can read, you can, you know, try them out in different environments. And I do want to caution that, you know, right now, if you look at the literature, a lot of studies, of course, they are largely experimental studies, meaning the researchers really need to specify the market environment, you know, the demand, the supply, the pricing options, the strategies available to the AI agents. You know, as in any simulation studies, the limitation is that there is always a concern that when you get out of that environment, that controlled environment, do you still see the same kind of phenomenon.

I think that's always something to keep in mind when we interpret experimental studies. And I do think that there is a lot we can learn from just keeping a close eye on the technical side, the AI literature, as I said. I think we as the antitrust community can benefit a lot by simply keeping a close eye on those because there is a lot of interest in the AI field to develop those algorithms.

Now, of course, their goal is not to develop colluding robots, right, just to be clear. Their goal is to develop algorithms that could, you know, work with humans and make our life easier, even in social dilemmas. Even when the algorithm's subjectives kind of, you know, conflict with human objectives and how they can learn to work with each other in particular with humans. So I just want to be clear, it's not, you know, the AI fields are, you know, evil colluders trying to design things to hurt us.

But the research that they have done, you know, we can learn a lot in terms of the limitations, the challenges of designing collusive algorithms. Thank you.

MR. RHILINGER: I don’t mean to interrupt, but just one quick question. You mentioned earlier a lot of evidence that as someone that manages merger investments I see a lot of, you know, documents and that sort of thing. Do you still see a role for technologists in helping to interpret that sort of thing, because, again, as you were describing it, the material sounded familiar, but I was just thinking as this field is changing so fast, do you still see a role for technologists in that process?

DR. DENG: Yeah, that’s a good question. I do think that at least in the initial stage I don’t see that you need a lot of technical expertise. I mean, I can give you a couple papers in the AI field, and, you know, if you just read the abstract and the conclusion section, you know exactly what they’re trying to do, you know exactly how their algorithms performed in kind of a controlled environment, you know, that simulates competition and how they were able to collude or not able to conclude.

So I do think that in the first pass, you know, people with experience in antitrust and understanding the markets already can go a long way.

And I think, you know, eventually, if you go into the program, that's where absolutely I think you do need experts to interrupt.

MR. RHILINGER: Thanks. Sorry, Kai-Uwe.

DR. KUHN: No, that's fine. I do think we have a lot more possibilities with traditional tools even in this field than we're kind of admitting in this context. And I think this is a little bit underestimating also the coordination activities that are just necessary in order to get there. And I found that very revealing with one of the comments that Joe made when he was talking about the algorithm can be designed in a way to collude.

And that's essentially what otherwise the coordination activity would be. I mean, there's a great difficulty, and I talked about this, which is in principle, if you don't know what the other guy's algorithm is you're playing against lots of algorithms, and that becomes a really complex problem in how you're getting the other algorithm to converge to common behavior, and how to induce that, I'm not quite sure what anybody knows.

But even if you're trying to do something like this, I think the activity of trying to put a mechanism into the algorithm, that would lead to collusion. It’s much more detectable than actually looking at the algorithm and asking the question, is if it reacts by saying cut the price if the other guy cuts the price, is that part of a collusive strategy, because we see lots of markets in which there’s sequential price setting, under virtually all markets where there’s sequential price setting, and those tend to be very competitive markets in which prices sequentially are lowered.

So I’m not convinced that we’re going to be very good at identifying collusive strategies from very complicated algorithms or maybe not so complicated algorithms but basically saying this is a collusive strategy because we only know that if we know what they had in mind, what the strategies were of the algorithms that they were trying to play against and that they were trying to coordinate with.

So on the other hand, if there is an attempt to do this actively, then there are people around who know that we were trying to design an algorithm like this. And you will be generating the same information as you’re getting now from kind of someone spilling the beans internally. And so in that sense, well, maybe that wouldn’t be the typical communication or coordination behavior and one might want to increase that scope a little.

But that’s what I said before, you actually want to look at the coordination behavior, the sharing of a price, the clear intention of having a rule in the algorithm that is trying to lead to collusion, that you would want to target, because you’re much more likely that you’re going to get evidence about that while price setting and price movements and even strategies are really, really hard to interrupt, because, you know, how you were going to test the algorithm, what did they have in mind, what the algorithms were on the other side. That’s kind of the unknown in this.

And that’s why I’m much more circumspect about what Joe is suggesting, but certainly I think if one is thinking much more about what are the activities to kind of get there, you’re getting much more step-by-step increments in the direction of dealing with the issue that you can actually understand and that fit into the current framework.

DR. ABRANTES-METZ: I would like to just make a small comment on I think that it would benefit the business community if there were general principles, general rules not necessarily forbidding per se. It doesn’t mean that it can’t be, as Joe suggested, but having general rules, guidelines on what should we desire in a pricing algorithm and what we should not and the conditions under which we should be more concerned about certain features than others.

We have that for communications among competitors. And I think that if we are to build structures that are better from the start, we are then less likely to find ourselves in bigger problems later on. You know, I always think about what happened with the financial benchmark situation where for years I said that these structures were easy to wreak and pretty much everywhere we did we found rigging, extensively and massively. But somehow the authorities were distracted, I believe, because only after LIBOR broke we started to come up with guidelines on what are the good principles for financial benchmarks.

So I think we should have a more proactive role in this case and start by conducting more research and having more of these type of discussions and come up with good principles on which to base on this pricing algorithms that the business community knows and to Sonia’s point that don’t suddenly get shocked, that something that they did had no clue, they were now liable at some level, and then start from then on and see whether the guidelines that we come up with do need some sort of an extension or a little bit from a broader view of what an agreement actually is.

MS. PFAFFENROTH: And I just wanted to build quickly on something that Kai-Uwe mentioned a minute ago. So something else that’s important to consider in the context of the increasing use of algorithmic pricing for businesses is not just a situation where, you know, you have two competitors agreeing that they’re going to adopt certain pricing software, but also thinking about where information sharing, the sharing of information itself regarding what specific algorithm has been adopted, what software has been adopted, or certain aspects about technologically how it functions, that that type of information sharing between competitors, even if there is no explicit agreement that they are going to set the parameters to a certain set of actions or to take a certain set of outcomes still gives rise to antitrust risks because sharing the algorithm, the existence of the algorithm, the choice of a certain algorithm or the mechanisms by which it function could conceivably be closely akin to sharing pricing information, which itself can be risky or violative behavior, even in the absence of the explicit agreement.

MR. HARRINGTON: Let’s see. Let me kind of respond to a couple of remarks made and then kind of address the question. So to be very clear, my remarks had nothing to say about the likelihood that I would assign to algorithmic collusion. It was saying that if it were to occur what would be the legal response. Right now, the legal response would be we couldn’t do anything; we need to develop something else.

You know, I’m also kind of sympathetic with the challenges that Kai-Uwe mentioned with regards to the approach that I’m proposing. It’s not going to be easy but I do think collusion is a discrete phenomenon. That’s not just something that’s a little bit less competitive. We know in practice, we know in simulations, and I would say practice in actual conduct by humans, that there is a discrete change in conduct, and it’s all rooted in this idea of reward-punishment. Quite different from competition. And so it’s starting from that principle that I think that, you know, it is -- it offers enough potential to be able to try to identify properties of collusive pricing rules, that this, I think, is a viable approach.

How exactly that will workout? You know, we really won’t know until the research is conducted, but there’s going to have to be lots of problems solved. You know, in terms of the original question, I’m going to respond in a much broader way in terms of, you know, what we can learn from other jurisdictions, which is one of the things that is going to become more common in the midst of collusion by algorithms. Well, there’s algorithmic collusion or it’s just pricing algorithms being used to kind of supplement kind of existing modes of collusion, is detection, because what we’re imagining here is that these pricing algorithms, however they’re being used, is conditioned on easily available prices of rivals. So we’re not thinking about intermediate goods markets here; we’re thinking about retail markets on the whole.

So we’re looking at a setting in which a competition authority or any third party could, in principle, engage in screening that is looking at that same data to try to find patterns that are consistent with collusion. So the idea of screening for cartels as looking at market data to try to identify them, is something that’s being done in a number of jurisdictions but is not being done in the U.S.

I was recently at a meeting with about 25 to 30 chief economists from various jurisdictions. About two-thirds of them said that their agency was engaging in some form of screening -- some just kind of experimenting with it, some putting lots of resources into it, such as in the case of Brazil. The U.S. DOJ was there. They were part of that minority that was not engaging in screening.

So I would say, you know, what we can learn and what we can do is to try to make screening a kind of a -- more of a standard practice for competition authorities because I think that's going to become more and more useful if, in fact, pricing algorithms become a more important component of collusion.

DR. ABRANTES-METZ: Let me just add one point on that. Competition authorities are also, some of them, starting to be interested in developing these types of AI techniques to detect. So beyond the typical screening, many of them have very large data sets of actual bid rigging. They have collected for a very long time.

And I, for example, am working on one of those projects where we are starting to develop a model to detect potential bid rigging, apply it to a different data set, but training it on a particular data set. So some of the agencies are actually going much beyond the typical screening that we have been doing for, some of them, for some years to getting more up-to-speed into AI techniques. So I do agree with Joe. This is something that should definitely be done.

MS. CONNELLY: Any other comments?

Yes, of course.

DR. KUHN: Yeah, just to rejoinder on two of the remarks that were done in your information exchange. So I think in developing rules, it’s always important, if you want to have a per se rule, which is really good for incentives and for firms to have clarity, you want to make sure that the costs are relatively low. And I think some of the suggestions that come here in order to say certain -- basically, any information exchange about what your algorithm is, you can make illegal because it’s very hard to think of any good reason why you should be sharing your algorithm with your competitor, or information about your algorithm to the competitor.

So this is kind of one of the examples where I would say we basically have the legal framework on information exchange. It falls very much into the same similar category of exchanging prices that you want to set in the future. Why not do that if you need an extension there to make it clear that that falls under it legally, well, do it. But that's a very traditional approach that I think would already go very, very far, even in addressing Joe's concerns because it then makes it unclear what I'm actually competing against, and that makes it much, much harder to get through.

Just on the screening, I think one has to be very cautious about thinking that you can screen everywhere. There are a couple of markets, and especially with bid rigging and so on and so forth, where the structure of the price setting in the market is very, very clear. Now, in a lot of other markets it's very, very hard to do screening of that type, and I think even in some of the retail markets that you're looking at.

So as a general proposal of doing it everywhere, I'm not really convinced. And when the European Commission tried it, it really failed because you couldn't make an inference that was good. So you need secondary information for the inference that very often comes from the price-setting structure. Now, you have that in financial markets, you have that in bid rigging, but in other commercial markets, I think

I'd be -- I'd be very, very cautious and would ask myself what would actually be the criteria for knowing that you should be starting to intervene.

DR. DENG: Can I quickly follow up on the screening and monitoring? Joe and -- Bill, Joe and Romi (phonetic) have done a lot of work on this. And I think I made a similar point in an article called "Cartel Detection and Monitoring: A Look Forward," making the point that there's almost an interesting paradox here because AI, we're talking about AI being these evil colluders, but at the same time, I do think that there's a lot of potential for the AI technology to help us detect and monitor the markets.

And, you know, subject to Kai-Uwe released comments on, you know, it's not always you can apply those techniques.

MS. CONNELLY: I'd like to move on to a few questions from the audience. We've actually gotten quite a few. I think this one actually plays nicely off the comments that I just made. The question asks, at what point or how should the agencies think about setting the balance between antitrust enforcement in this area and not deterring innovation or additional sort of innovative competition?

Would anyone like to start us off? Maurice.

MR. STUCKE: Yeah, one thing. I really think there's four prongs to respond to that. And the first thing that I think came out from -- I think everyone on this panel would agree, is to better understand the risks. And that's why I think these market studies and the like are really helpful. And also speaking with the people that are promoting this.

I mean, for example, the Italian competition authority observed, "a number of specialized software developers offer solutions that allow even small companies to implement strategic dynamic pricing strategies, offering tools to autodetect pricing wars as well as to help drive prices back up across all competition. So I think that's one.

Second is improvements in tools to detect collusion. You already heard one proposal here. Other proposals include auditing the algorithm. There are pros and cons involved with that. We promote the algorithm collusion incubator, but then there's also the market studies.

The third thing, and I think this is key, is refining the tools for merger enforcement. Bruce mentioned that that's going to be one of the primary mechanisms to target tacit collusion and to get a better handle on this. And, then, I mean, the other thing that’s coming out through this hearing is that the United States has a market power problem. And we’re seeing increased concentration in many industries, market power and the like. Some dispute the evidence, but all the evidence seems to be pointing in that direction.

And to the extent that’s true, to what extent does it not only affect then algorithmic collusion but also maybe perhaps switching the presumption in mergers. For example, that if you have highly concentrated industries, there’s already legislation now on the Hill that the presumption would be changed. And we’d propose that as well in our effective competition standard paper.

And then the final way, so far, we’ve been talking about ways to deter and detect collusion. Another way to think about this is are there other mechanisms to destabilize tacit collusion. For example, you know, industries that have high entry barriers because of regulatory restraints and the like, and other jurisdictions are now experimenting, for example, with the speed in which companies can change pricing. There may be pros and cons. That’s why I think the algorithmic collusion incubator could be helpful. But then also what about on the consumer side? Is there ways that you can reduce price transparency to the buyer’s advantage? So for example, offering reverse bids and giving buyers call options on multiple sellers to help destabilize tacit collusion.

So the thing is I’m driving for a gas station, I could then put in an app to the multiple gas stations, what’s the best price you can offer me. And now I will know the price but not necessarily my rivals.

MS. CONNELLY: Would anyone else like to comment?

We’ll move to another set of questions just in the remaining few minutes that we have from the audience. We’ve gotten a couple questions on this point and I think it relates nicely to some of the conversations yesterday on the consumer protection side and also to, Ai, your comments about the level of technical expertise or understanding that might be necessary to address these issues.

So yesterday, on the consumer protection side, it was suggested that the FTC should consider hiring as many technologists as lawyers and that we really do need a much more robust technical understanding to be able to address these issues.

We've gotten a couple of similar questions from the audience asking about the impact of the fact that many of the algorithms are proprietary, what the impact of that might be on our ability at the antitrust agencies to address the types of conduct that we've been discussing on this panel, and also the impact of the extent to which some of the more complex technologies are actually explainable or understandable to us at the agencies and also to even the companies who are using them.

I'd like to see if the panelists have any comments on any of those topics. Anyone like to start? Sure, Maurice.

MR. STUCKE: I would -- I mean, the first thing I would do is I would go to the ACCC and ask them their experience because they are now hiring data specialists on this. And I think it's -- you know, look, we want to find out what the other agencies are doing, to what extent are they using data technology, and then -- data technologists, and then to what extent can you use them then effectively, both for behavioral discrimination, price discrimination, as well as collusion and other issues that may arise as well. I think you definitely need that expertise going into a data-driven economy.

MS. CONNELLY: Anyone else? Rosa.

DR. ABRANTES-METZ: My experience in these financial and commodities markets have been telling me that often -- and a lot of these include -- relate to spoofing schemes, also to pricing algorithms that regulators are very, very much behind everything else that is ongoing. And it is hard to keep up with somebody who just does that every day, every single minute of the day and invents new ways of adjusting prices all of the time.

So I don't think I would have ever the expectation that the agencies would be able to be monitoring all of these aspects from everybody all of the time and know all of the technologies. I do think, though, that they should have some of that knowledge in-house, and wherever the suspicion does come from whatever source that happens, that a particular pricing algorithm may be causing problems, anticompetitive effects. Then I do think the agencies need to have that knowledge to get into there and even if it is proprietary obviously having the authority to go review and have their own experts with them.

I don't think, though, that this would be something, again, that would be feasible to do or even desirable. The amount of costs at the firm level to be able to keep up with this kind of regulatory oversight would be large. But I think that occasionally that may well be justified and so that expertise would be needed.

MS. CONNELLY: Anyone have any comments on that?

DR. DENG: So maybe just a quick comment. So I do think that the first line of defense -- the line of really information source should be the developers themselves, the companies who adopt those technologies. You know, being in a research community myself, I mean, every time I could write a very technical article with all the mathematics, you know, simulation behind, but I always want to make it easy to read, have a very easy-to-read abstract and conclusion. So I do think that's the first place that agencies and anybody without technical training should go to.

And after that, I echo what Maurice and Joe's proposal. I think after that, you know, to really understand how the algorithm behaves, you probably will need to have, you know, the simulations, experiments, and research after that.

DR. KUHN: I actually think there is another aspect to this which is very important to actually have some people with expertise, which is really a checks-and-balances issue. You very often get, if you are -- you know, if you're a competition expert but not an expert in the other things, everything you see you interpret as a competition problem. And that's often not appropriate to the things that you're seeing, but the reason why you interpret it in that way is that you're not understanding the rest of the framework.

And so everywhere where we've seen economists come in, patent lawyers come into the agencies and so on, I think we've had a much more differentiated and broader view. In the end, I think that also enhances enforcement because it enhances a distinction between something that's problematic and something that's unproblematic, and especially something like collusion where the important thing of policy is giving the right incentives, right? It's really important that you punish things that are for sure bad because if you're punishing things that might not be bad, you're actually reducing the incentive effects of what you're doing.

So I think just from that perspective of kind of distinguishing and having the perspective of saying, oh, but this is also relevant for X, which has nothing to do for competition, just that big-picture item is something that's, I think, of critical importance if one is engaging, even if it's not replicating the algorithms that one is looking.

MR. RHILINGER: With that, we are over time, so I'll ask you to please join me in thanking our panelists for an interesting session.

(Applause.)

MS. CONNELLY: Now we have a short break.

(End of Panel 1.)

FRAMING PRESENTATION

MS. GOLDMAN: Okay, so I'm Karen Goldman. I'm an attorney adviser in the Office of Policy Planning at the Federal Trade Commission. So I would like to introduce our next speaker, Michael I. Jordan. Professor Jordan is the Pehung Chen Distinguished Professor in the Department of Electrical Engineering and Computer Science and also in the Department of Statistics at the University of California, Berkeley. He is a leading figure in the field of machine learning. We will now begin his prerecorded presentation.

MR. JORDAN: Hi, I'm Mike Jordan from the University of California, Berkeley. I'm glad to be joining you. I'm going to be talking about emerging challenges in AI, taking a perspective that brings machine learning together with economics, which is a relatively new way to think. So I've been working in AI for over 30 years now.

I should say I don't think of myself as an AI researcher. I'm really a statistician, sometimes a computer scientist, sometimes a control theorist, increasingly somewhat of an economics person. And part of the message here is going to be don't take this AI buzzword too seriously. It's not the buzzword that most of us use who've actually been working on machine learning for all these years. It's an aspiration perhaps even for the future but it's also a unhelpful buzzword for many of the situations it's being used in.

So let me get started here with a little bit of historical background at least from my prospective. So first of all, this field really is just statistical data analysis. Around 1980, it started to become called machine learning, at least by people in computer science, and it already had a large number of applications in industry that have changed the world, going back already to the 1990s.

So the back end in many companies, such as Amazon, was formed on machine learning algorithms, meaning really statistical data analysis with large amounts of data at scale and done in relatively close to real time. So fraud-detection systems to bring fraud rates down so you could do online commerce were critical in the development of those companies. Search algorithms are based on statistical data analysis and machine learning, and, critically, supply chain management. So a company like Amazon that serves billions of products has got to know where every piece of every product is in the supply chain at every moment, so they model things like storms in the Indian Ocean, and that's critical already in the 1990s.

And, in fact, the algorithms being used now are not so different from the ones being used in that period of time. Having built those systems, it was natural for companies to think about the human side, turn this towards -- away from the back end because a lot of the data was foreign about humans. And so systems like recommendation systems started to emerge, where you would take in data -- do data analysis on one person's buying patterns and use that to recommend products to other people.

Now, if you do this at scale of tens of millions of people, or even hundreds of millions as we're seeing in China and, you know, interesting new issues start to come up, and those were already being faced, you know, 20 years ago. And, now, we've moved to the third generation. This is often called the deep learning era or the AI era, but really it's not that different.

The applications have kind of sort of focused more on human-imitative things -- speech recognition, computer vision and so on, but I think of these really as end-to-end era. It's that we've been able to commoditize something like computer vision or speech recognition. So that end-to-end is specifically used for new purposes and used in creative ways.

But there's really not been a qualitative transition in the ideas, per se. The algorithms have not changed that much. There's lots more data and lots more machines but sort of those are just really quantitative changes.

So what's new to my view of what's happening now is not really this imitative -- human-imitative AI. It's the emergence of new markets based on data analysis and producers and consumers all coming together. So I'm going to be focusing on that, all the challenges there.

So in thinking about what AI is today and how it might be regulated and what are the meanings of that and consequences, I don't think you need to think too much about the history of AI. You really want to know what's happening, and it really is something changing in, in fact, I think exciting new ways.

So let's go back a little bit in history. How did people make money off of the web using machine learning, and now I kind of have Google in mind, or Facebook. So their argument has been that they provide a service to humans -- search or social networks -- but they need to provide better and better services somehow, and they're sort of stuck in the virtual world, so all they know about humans is the data they get, and so they have to analyze that data to learn more about the preferences and needs of humans. So with all the attendant issues about privacy and data analysis and all that we're seeing play out, kind of the problem is they don't know what to do with that data in terms of providing better services.

So what have they done? Well, they've advertised -- they've made their money off advertising. So they created a market, but it's not between the consumers or the producers of the data. It's between themselves and advertisers. And they're trying to figure out what humans want, but the data are pretty weak really. People talk about all the data search engine companies we have, but, you know, at the scale of tens of millions of people or more, that data is not that good an indication of any individual human's preferences or needs. So the service gets a little bit better but not hugely better, and they're kind of embracing AI in the hope that it will lead to even, you know, more impressive service. But, still, people are not going to be willing to pay for that service, so it’s not really yet an economically new model, and advertising remains the corn in the realm.

So I think what’s new right now, one of the big trends, is that there are companies that have different kinds of data, not just clicks data and, you know, browsing data. So the e-commerce payment companies have transactional data, and I think it’s a better place to start. So it allows already a notion of a two-way market to arise. It’s a transaction not between Google and the person but between a producer and consumer both who are on some platform.

So Uber is actually an example in one particular vertical. They have producers and consumers, and they don’t provide any extra value themselves beyond linking the producers and the consumers really. I believe that this is actually a better starting place for starting to think about data analysis and algorithms and people altogether because there’s going to be economic value associated with data now, and that’s actually better. Economic value is something that humans can build on and start talking about issues such as fairness and what’s the value of my data. It makes sense that the data already has some value.

So let me actually step back for a moment and think about this buzzword “intelligent.” Again, I think a lot of us think of ourselves as statistics and machine learning people, and we don’t think that we’re really working on human intelligence, AI. And, in fact, as someone who was in a neuroscience department and had a background in psychology, frankly, I don’t think there’s been that much progress. We don’t understand intelligence, certainly human intelligence. We have a very long ways to go.

And we haven’t, over the last 40 years, really deeply understood intelligence. Our learning systems mimic human intelligence. They take data out of an intelligent system and they mimic that. That’s very far from actually getting at the core of intelligence. And I don’t think that’s the future, actually. I don’t think at least in my lifetime that we’re going to deeply understand the intelligence of a five-year-old boy or girl. And we don’t really need to is the point. It’s not necessary to build the kind of intelligent systems that we need to have our life be better.

So if you think about intelligence, there’s another kind of intelligence on the planet. It’s not just human brains and minds. A market is an intelligent entity. And if you're looking down at the earth from Mars and you say what's intelligent down there, you notice that every city has food coming into it every day, every restaurant has the right number of items for all of its menu, every household has the right amount of food and every store and so on, and that's done by a huge network of, you know, millions of local decisions not really coordinated. So it's the usual perspective of microeconomics, but the point is that that's an intelligent system. And it's -- arguably it's intelligent in its own way as a brain or a mind. It's adaptive, it's robust, and so on.

And perhaps oddly, that perspective has not really been part of the dialogue on AI, and I think it should be. I think we should be thinking of creating artificial markets, artificial intelligent markets, and not just old kinds of markets, new kinds of markets will emerge as we bring statistics and data together with market principles.

And so new consequences will emerge, and I think they're actually more favorable than some of the ones we've seen in the current dialogue over just classical AI.

So here's a little formula, AI should be thought of, if we're going to use that buzzword, as data plus algorithms but also plus markets. So we're not simply trying to imitate humans and find out about their needs by looking at data. There's a lot of guessing in that, and I think that will be true for the foreseeable future.

Rather, we're trying to use market design and have data flows being created between producers and consumers, not just between companies and users. And that will provide better services that people will be more interested in and be willing to pay for. And, moreover, if you're going to talk about a concept like fairness, it's not just the data analysis and the way the data were collected that leads to fairness. You need economic concepts like utility. You should not give the same service to everyone. That's not fair. Rather, I should have my own utilities be expressed in some way in the system.

Let me begin with a concrete example of this. So music is arguably a domain in which there has not been a real living market. More people are making music than ever before. People drive a taxi during the week and put their music up on a SoundCloud during the weekend, but they're not making any money off of that, and they're engaged in no market. They put their product out there and it disappears from their life.

More people are now listening to that than ever before, however, but there’s no connection between the producer and consumer. So sites such as Spotify or Pandora stream the music to people; however, they don’t -- how do they monetize that? They’re not creating a market. What do they do? Well, they do what you think they do. They use advertising to make money.

So I think that’s broken. I think we’re missing a market here, and so a lot of human happiness is being left on the table. People who might like to make -- have their career be play music for other people can’t because there isn’t a market in which they can participate. There’s the record companies, but that’s a tiny and mostly broken market.

All right, so how do you create this? It’s in some sense not that hard. It’s just data analysis, so it’s not fancy, schmancy AI, but it’s really an important way to think about how to use the data. Just take the data of who listens to who -- maybe YouTube provides it, maybe Spotify, make a dashboard for someone who’s been putting their music on SoundCloud. They can now look at a map of the United

States, say, and see that they were being listened to this past week in Fort Lauderdale, Florida by 10,000 people. Not that they know that, that's economic value. They can give a show there and make maybe a few tens of thousands of dollars. And if they do that a few times during the year, there's a salary for that person. They can leave their taxi job.

Moreover, a market is creative, so they can -- now they're connected to their fans they can make other kinds of offers like I'll play at your wedding for $10,000 and so on. And I could imagine like a million people in any given country doing this. So there's AI being used to create new jobs, not to take away jobs because when you link customers and producers, you've created a market that creates new kinds of value.

Of course, the company that provides this is going to make money as well. They simply take a cut from the transactions because these are real economic value transactions. But they're not the one who are having to create the value and you worry about the -- their use of the data, okay? They have to be careful with privacy, certainly, but it's somehow easier.

There is a company doing this in the United States. It's called United Masters. If you are curious, go have a look at what they are doing. It's actually real musicians and real tech people doing something of this form. But I think this is actually far broader than music and far broader than this one company. I think that is going to happen not just in music but more broadly in entertainment. You have all kinds of producers and consumers who could meet up and provide value to each other, information services, personal services, people who want to cook for others, people who make haircuts and so on and so forth.

Now, part of this is that you want to make recommendations. You want to have people have data being brought into play here. It's not just a classical old market on a new platform. It's actually new kinds of markets, all right?

So let's think a little bit about that. So a classical recommendation system makes independent recommendations to people who come on their site. No economics is involved because there's no scarcity and there's no interactions of the decisions. So that's not going to be true in real world markets. There's going to be interactions and scarcity.

So think about a classical recommendation system. You all know what these are. A record is kept of a customer's purchases. Similar customers are recommended similar purchases. And, you know, Amazon pioneered this. Right, but these recommendations are done independently, and it's quite plausible that we could make the same recommendations to two people, three, hundreds of thousands of people. And is that a problem? So if I recommend the same movie to everyone, it's not at all a problem. I can copy the bits. It's classical. I'm in the virtual world, not in the real world, and so there's no scarcity.

What if I recommend the same book to everyone or to hundreds of thousands of people? Still not such a problem because there's something called print on demand. I can copy it quickly and have it out in three days to everybody.

But if I recommend the same restaurant to everyone, I'm really trying to provide economic value to people, tell them that you've arrived in a city, here's -- you push a button like an Uber person would push to get a ride. The restaurants around me see that I'm now ready to eat, and they make offers to me, maybe discounts, and so on. And I look at the offer, I say that restaurant, that's for me, and accept. There's now a transaction being made. So it's not just an advertising of restaurant service or, you know, kind of classical push service; it's actually a transactional service.

But now if I recommend the same restaurant to everyone, they'll all go there and there will be congestion. If I recommend the same street to every driver, I build a system that independently recommends routes to the airport, I'm going to create congestion. And if I recommend the same stock purchase to everyone, I'm going to create instability in the market.

All right, so these are the kind of problems that arise when you think of an economic perspective, and the solution really is straightforward in some sense. Just set up markets between restaurants and diners or even between streets and drivers, between financial consultants and people who want to invest their money.

So I hope you see that there's many challenges of this kind. This is one actually in creating a different kind of AI that's not just the kind that focuses on imitating humans but is broader than that. Here's a list of some of the things I work on in my own group, and you can see things like realtime, fairness, diversity, providence. These aren't the classical robot vision, you know, sort of style machine learning. They're broader, they're sort of reflecting a broader goal in terms of economic networks.

I’m going to skip the next two or three slides of my slides here. You can look at them afterwards, but just to say multiple decisions is not just economics, it’s also statistics. We are starting to make decisions under uncertainty. You have to worry about hypothesis testing and multiple decisions, and so a lot of our systems have to make not just one decision but huge numbers of decisions. And when you do that, you start getting false positives becoming a big concern. And classic statisticians worry about this and scale maybe a few decisions, but now a system like Uber or a medical system or a commerce system is making hundreds of thousands or millions of decisions per day. You really have to worry about all the interactions.

And there are schemes called false discovery rate schemes which worry about controlling those errors. And I’m going to skip over the slides that talk about this. I just want to say there has now been some work on any time control of false discovery rates, where you can have a person make or a group making decisions over time and you can stop them at any time in their error rate up until that time it’s under control. So it has more of a control or almost economic perspective, but it’s statistics now being brought to bear. So I’m going to skip over the slides that talk about that.

And let me move to my final slide. So some parting comments on this buzzword “AI.” I do have an op-ed called “Artificial Intelligence, the Revolution Hasn’t Happened Yet” that provides some background to what I’ve been talking about today. It’s not the same material but starts to give a little bit of a breakdown of what AI refers to.

And the one that you mostly see in the newspapers is human-imitative. I don’t think that is the right goal. I also don’t think autonomy should be the right goal, but really what I think is emerging is a new engineering discipline, and it blends economic ideas, computer science, statistics, and related fields to build networked, large-scale social decision systems with a wide range of applications.

So in thinking about what you’re doing in this meeting and what you want to write about, I hope you’ll at least have a nod in the direction of something new is emerging that isn’t just data analysis and the replacement of human beings by computers, but it’s really this broader engineering context. So thank you very much.

MS. GOLDMAN: Please join me in thanking Professor Jordan for his excellent presentation.

(Applause.)

EMERGING COMPETITION, INNOVATION, AND
MARKET STRUCTURE QUESTIONS AROUND ALGORITHMS,
ARTIFICIAL INTELLIGENCE, AND PREDICTIVE ANALYSIS

MR. WILSON: Good morning. My name is Nathan Wilson. I’m an antitrust economist at the FTC, and I’ll be one of the moderators of this panel. The other moderator is my colleague, Brian O’Dea, who is an attorney in the Bureau of Competition at the FTC and is seated to my right.

Before we get to our panel, however, I’d like to begin by doubling down on what Karen said and thanking Dr. Jordan for his helpful and interesting remarks on the various challenges and prospects that AI practitioners may face in the coming years.

Now, our panel is going to keep that focus on what lies ahead in terms of algorithms and AI but shift the emphasis to how those technologies may affect competition and market structure throughout the U.S. economy. Now, we are fortunate to have a great panel to discuss these issues with us today. And I’m going to turn their introductions over to my colleague, Mr. O’Dea.

MR. O’DEA: Thanks, Nathan. Our first panelist is Robin Feldman, who is the Arthur J. Goldberg Distinguished Professor of Law and Director of the Innovation Law Institute at the University of California, Hastings. She has published four books and more than 50 scholarly articles. Professor Feldman testifies frequently before Congress and federal and state agencies. Her empirical work has been cited by the White House, along with numerous courts and agencies.

Professor Feldman participated in the GAO's report to Congress on AI; the Army Cyber Institute's threatcasting exercise on weaponization of data; and the National Academies of Sciences roundtable on AI and life science.

In addition to her scholarship, Professor Feldman runs the startup Legal Garage in which students provide free legal work for 60 early-stage technology and life science companies each year.

Our second panelist is Joshua Gans, who is a Professor at the Rotman School of Management at the University of Toronto and Chief Economist of the Creative Destruction Lab. His most recent book is Prediction Machines: The Simple Economics of Artificial Intelligence, which was published earlier this year.

Our third panelist is Preston McAfee. McAfee is a former Professor of Economics at the

California Institute of Technology and the University of Texas. He has written extensively on auctions, pricing, antitrust, business strategy, and the intersection of computer science and economics. Previously, he was a researcher and executive at Microsoft, Google, and Yahoo.

Our fourth and final panelist if Nicolas Petit, who is Professor of Law at the University of Liege, Belgium, a Research Professor at the School of Law at the University of South Australia in Adelaide, and a visiting fellow at the Hoover Institution at Stanford University.

His current research focuses on three areas: antitrust in digital economy firms, patent protection as an engine of innovation, and law creation in a context of technological evolution. His recent written work deals with the limits of antitrust economics in relation to technology giants and the legal frictions created by the introduction of artificial intelligence in society.

So last is a disclaimer before we get started. Any questions or statements by Nathan and myself are our own and do not necessarily reflect the views of FTC.

So with that, I’d like to start out with a definitional question, and this may be a bit of a recap to folks who've been with us over the last two days here, but I think it's helpful to set up some of the discussion that we'll be having today. And that is sort of at a high level, what are core futures that define an algorithm? How do those differ from the core characteristics of AI, and are there antitrust or competition contexts in which differences between algorithms and AI are likely to matter?

So, Robin, why don't we start with you on that on.

MS. FELDMAN: Sure. So an algorithm is a relatively simple beast. In the broadest sense, the computer context, an algorithm is just any series of steps performed by a computer on input data. In contrast, when we talk about AI, most people are talking about machine learning, which these days, generally means using past data to train a model to by itself make predictions on future data and direct choices based on those predictions. For example, is it a stop sign or a speed limit sign? So should the computer apply brakes to the car?

It's important to understand that AI and that machine learning is not just predictive analytics. We've had that for a while. Rather, AI, on its own, can make assumptions, test, learn, reiterate, do all of those things by itself. So from a competition perspective, one might think of three distinctions that might matter in the algorithm AI context. The first is the evil you specifically programmed as opposed to the evil that a reasonable programmer or a reasonable user could have predicted as opposed to the evil that is entirely unpredictable.

So with a simple algorithm, we're probably talking about the first category, that is, the evil that you programmed. And in that case, the blame and the sanctions are relatively easier. But with AI, maybe you didn't task the computer to behave in a manner that is anticompetitive or discriminatory, but that's where you've ended up.

So when bad things happen that a reasonable programmer or a reasonable user could have predicted, competition authorities might want to react in a manner that's similar to misconduct that was specifically programmed. However, when bad things happen that were entirely unpredictable, one might want to react differently. We may not want to hold you liable, or at least not to the same degree, although, we certainly would want to hold you responsible for fixing the problem.

Of course, a reasonable framework still leaves extraordinarily difficult questions. How are we going to determine what a reasonable programmer or a reasonable user could have predicted, and for that matter, how will we determine what the AI did and why? Thank you.

MR. O'DEA: Thanks.

Nicolas.

MR. PETIT: Sure. So the reason there is a difference between the two sets of technologies insofar as antitrust is concerned, so on the one hand, simple algorithms, which follow given rules for pricing purposes represent, I'd say, a moderately interesting problem for antitrust policy. On the other end, sophisticated AI systems like, you know, deep learning, neural networks and reinforcement technologies that turn data inputs into outputs have much bigger implications for antitrust policy, and most of that is tied to the so-called black box problem.

So the black box problem basically means that neural networks and deep learning cannot really tell you -- tell the programmer or manager or shareholders or regulatory agencies how the linkage between an input and an output has operated and what decision-making process has been going on there. And that, I think, has three implications for antitrust policy. The first one is the liability problem. Is it fair to impute liability to the firm, its managers, or its shareholders when it's not possible to tell, you know, what happened.

Is it better to think of other allocation of liability regimes? Sharing that between technology manufacturers and firms operating in markets, all the more so when the technology is actually not owned by the company on the markets? Should we think about absolute strict liability regimes like product liability or move to joint liability regimes? I mean, there's a ton of questions here.

The second, I think, implication which we'll face in antitrust, if ever we have these technologies deployed at scaling markets, is whether we need to abandon decision-making rules which seek to elicit with their firm conduct is competition on the merits by inference of anticompetitive intents or by reliance on plus factors and whether we should not actually move to an antitrust regime, which basically instead of prohibiting selective types of bad conduct should actually prohibit bad outcomes in themselves.

So, you know, you could think about an

antitrust regime based on pure levels of harm, our type of prohibitions, that would bring antitrust policy very close to regulation, actually.

And the third, I think, implication for antitrust is one of remedy. So Computer Scientist Gary Marcus, he’s quite famous in the AI field, he talks of a debuggability problem. So when you have a black box, there is actually no clear way to diagnose or design data defects that led the AI system to predict or command an anticompetitive outcome and, therefore, the points where we can actually remedy those effects are very obscure and opaque.

Now, I just don’t want to suggest that we should actually change the antitrust policy and enforcement regime today based on the three problems because there is an ongoing discussion in the AI field today that actually AI may be hitting a wall. The deep learning, you know, type of conjectures that we are sold by the press are far from real and certainly not at scale, so we should be very careful here.

MR. MCAFEE: I want to make a relatively simple point that old AI, that is AI from the 1970s and ‘80s was actually designed by humans and we could understand what it did and why. And the same thing is true if you run a giant regression. So if you --

regressions have been run with a billion right-hand-side variables. But even so, if I ask why is it making this prediction, well, that's the sum of the predictions from all of these coefficients, and we understand at some level where those coefficients come from. Deep neural nets, on the other hand, don't work like that. They have extraordinarily complicated interactions.

And they have what's a very entertaining feature of them -- and let me apologize for my voice -- is just like humans have optical illusions, right, you've seen optical illusions where you look at a printed picture and it appears to be moving, or there's two gray bars that you would swear one is twice as long as the other and they're, in effect, actually exactly the same length, as you can verify with a ruler.

Well, AI has -- at least deep neural nets have optical illusions as well. And some of these are quite scary. So there's been attempts to trick -- to fool automated driving programs with a minimum number of pixels. And it turns out not to take very many pixels to convince an automated driving program that a stop sign is, in fact, a speed limit sign. And when I say not very many, you still have an octagonal red sign with the word "stop" written on it and two little one-inch by eight-inch stickers that are gray, and it comes out saying, oh, yeah, that's a speed limit sign.

There are also some pretty entertaining optical illusions for AI, and I want to emphasize, no human is fooled by these. We're fooled by other things, but we're not fooled by these. We're fooled by squiggly letters that shows you a picture of a giraffe -- or shows the AI a picture of a giraffe and convinces the AI that it's a house cat. So -- and this may be the wall that you're referring to, is that we are running -- so there are things that are just not understood about them.

And, then, finally, I think Google did something of a disservice to say -- to distinguish between algorithms and data because all of the modern examples, the algorithms are typically quite simple, and it's the data, you know, that's voluminous and complicated.

MR. WILSON: Thanks a lot, Preston. I think those comments actually tee up where I was going to take this conversation next. So we've been thinking a bit about the difficulties perhaps of really implementing AI and algorithms at scale and some of the factors that could affect that.

How should we rank order in these different elements that are necessary inputs for firms looking to deploy AI and algorithms at a kind of substantial level? So how should we weigh data versus, you know, the labor element, finding people with the talent and expertise to appropriately deploy these technologies versus other types of physical and technological capital that may be required?

So, Joshua, do you want to take first crack at this?

MR. GANS: Sure. So just to preface, I've been listening to the discussion here and the discussion in the previous panel, and there's a kind of, well, I wish artificial intelligence was anywhere near as intelligent as has been out thus far. You know, I come here to make this session as boring as I can possibly do. And I'm going to do that in two ways.

First, I'm going to tell you that artificial intelligence is currently no more than an improvement -- a really big improvement -- in statistics. It's as good -- you know, it's as intelligent as multivariate regression. It is basically a prediction machine. It can take data you don't have and convert it into information that you need at a much greater rate than previous.

Secondly, I don’t think it’s hitting any sort of wall. It may be hitting a wall in terms of its ability to do traditional tasks that it’s been -- we have been benchmarking on with a number of applications of AI are quite voluminous in the economy, so we should realize that, which brings us to the issue here, which is the thing that I think we need to focus on is, is artificial intelligence representing something new that we have to worry differently about market power and also barriers to entry and those traditional antitrust things.

And for want of, again, pouring water on what is going to surely be an otherwise interesting session, I’m going to suggest that, in fact, when you think about it, there’s nothing currently indicated that suggests we need to do anything to change our approach to antitrust whatsoever, at least in regard to structural elements or abuse of monopoly power and things like that.

And that’s because of the inputs. The inputs to AI -- there’s numerous ones that we’re going to talk about, but let me talk first about data since that gets a lot of note. Data is used in AI in two respects. One, it is used in order to generate algorithms that can serve predictions and then be embedded in other things and improve productivity, product quality, et cetera. So data is used essentially for the same purpose we would use it in scientific tests or anything like that, to innovate.

The second part is data is used in order to personalize products. It's used in interaction to learn things about consumers, to come up with more tailored -- more product variety, if you will, in that respect. The two roles of data are very distinct. Occasionally, they all happen within the one firm, but data needed to train algorithms, to train machine learning, invariably can exist in a lot of places.

It's no more an issue for barriers to entry or anything like that as, you know, someone having patents, key patents or key scientific personnel or specialized research equipment or anything like that in terms of giving them some leverage in the market for innovation.

In terms of the personalization and the ability to have data that really learns about the customer very well and can tailor products to them, well, that's where traditional market power comes in. You have to have access to the continual interactions with those customers in order to generate the improvements and generate the advantage.

So Google gets an advantage because its customers are continually searching and, therefore, it can -- because of its share can generate some advantages that way. Facebook gets it because individual customers happen to use Facebook a lot, and it starts to learn about them. So both of those are very traditional market power things. One is about advantages in innovation markets and other things, which may be a technology side. The other side is simple advantages in market share that can give firms potentially a leg up over others.

Either things, we’re very familiar with dealing with it. We’ve done it before. We’ve done it with other technologies. It’s just a relabeling of the -- what’s going on.

MR. WILSON: Thanks a lot.

Preston, would you like to extend that?

MR. MCAFEE: Yes. Actually, I disagree a little bit with Joshua. Not maybe fundamentally, but -- so most of the technological innovations that have come about over the past 300 years actually substituted more for human brawn than they substituted for human thinking. There were some, the cotton gin, that actually was a descaling one. And the first one

-- but the first really big one of these was the adding machine.

And all of a sudden, now, you could work in a restaurant as a clerk and not be able to do math. And that was -- that was very different than the bulldozer, which substitutes for lots of people with shovels in the sense that it was substituting for thinking rather than for physical exertion. And this is on that scale except much larger. We already have news stories, sports stories are written by machines. Corporate earnings reports are written by machines. Why? Because there you're in a race who is first to market, and so that's really important.

Where I completely agree with Joshua is that I don't see much of a constraint in processing power. We're in a terrible situation with respect to talent. That is to say, you could double the number of people who are classified as data scientists and machine-learning experts and employ all of them tomorrow. Wages are rising sharply. So we have a significant talent gap.

And we have a data gap that I think -- I have the sense that the data gap will likely go away but is significant today, and partly it's significant just because we haven't taken advantage of all the data that corporations have.

Let me say that one thing that's very -- you know, if you think about electrification as a major technological shift, electrification presented the United States with a serious problem, which is that there was a giant economy of scale in turbines. So you wanted to have a big turbine, and that tended to create monopoly. And we addressed that problem by having either a municipally owned electric utility or regulating the electric utility.

In contrast with artificial intelligence, we have a lot of suppliers and a lot of automated tools. There are tools that are, you know, attempting to make AI accessible to people who are not technical at all and are attempting to commoditize artificial intelligence.

MR. WILSON: Thank you.

Nicolas?

MR. PETIT: Yes. There's two things I want to say. So the first one is about disputing empirical antitrust economics topic, you know, whether data -- whether there are increasing returns to data. And I think it's properly more right than wrong to say that there's -- with scaling data, you have, like, you know positive demand effects, network externalities, and so on and so forth, meaning increasing economic returns to scale. But when I talk to engineers, often I hear that scaling data displays diminishing technological returns. And I think that was said by Sue Lacey some time ago at a conference, and especially when used in AI systems.

So the hard and forgotten truth there that it’s not cost less to scale up and firms need to incrementally invest in fixed and variable assets when they analyze collection rates, you know, more voluminous amounts of data systems, especially with combined -- in combination with AI.

And, of course, the rates of diminishing technological returns to data in AI systems is probably dependent on the class of application that we’re talking about. So there might be differences across families of AI applications. But, again, I think we can’t just proceed on the assumption that there is the -- there are increasing returns to scaling data insofar as the technology is concerned.

And, in fact, again, another famous AI scientist the other day referred to the risk of exponential inefficiencies in relation to convolutional deep learning, noting that the reliance on large numbers of labeled examples in deep learning systems may actually lead to their demise because it's just too costly to actually scale up.

The second thing I wanted to say is do not underestimate the barriers to entry that will be generated by regulatory initiatives, maybe not in this country but in other regions of the world. There is a lot of demand in the number of regions in the world, in particular in the European Union, for regulators to step in and impose all sorts of compliance systems on AI companies, AI development companies, ethical concerns and so on and so forth. And we may move that field of the economy and technology developments towards, you know, regimes which look more like, you know, maybe pharma, where, you know, there's sort of sunk investments to comply with the regulatory structure are actually absolutely incommensurate.

And so if you think about that, you know, you can build on top of that the fact that most countries advance on that journey in a way which is completely uncoordinated. And that, again, will actually probably increase the, say, returns to compliance to big firms and decrease them for smaller firms.

MR. WILSON: Thank you. Anyone else want to chime in before we move on?

All right. Well, let's turn now to something that Preston teed up, which was the market to supply AI technologies themselves. Do we think that that market is competitive today? How do we see it developing in the future? And is there anything that we as antitrust agencies should be thinking about? Preston, do you want to start us back off?

MR. MCAFEE: Sure. So Google, Microsoft, IBM, Amazon, and at least 100 small companies that you've never -- mostly you've never heard of like Noodle, a variety of Chinese companies, all offer what amounts to off-the-shelf AI. And while they're different, they have two big things going for them. So if you look at, for example, the Google and Microsoft systems, they have a variety of data. They can already translate languages. They have a variety of data that they begin life with.

So you as a, let's say, a lipstick manufacturer don't have to put in language translation because that's already built into the AI systems. And if you want to build smart apps, actually, which is a thing that we're going to see a lot of competition over the next half decade as AI chips start to roll out in our phones, you want to build apps to take advantage of that, these systems give you a way to do this -- it’s not literally one button, one touch but it’s really simple.

Often they’re set up in such a way that you don’t need to know what the data is. Now, there’s a famous computer science saying -- garbage in, garbage out. If your data is all messed up, what comes out of this is not going to work very well. But nonetheless, they have really commoditized the provision of AI services. By the way, they also recognize photographs, they can tell you what’s in video and so on.

And we’re in a really fortunate position that we have large, very deep-pocket, well-funded firms who have all convinced themselves that AI is the future. And so they made giant investments to become vendors of AI. And so this looks to me like quite a competitive market in the sense that there are four very general purpose, large American firms and then there are dozens of more specialized firms selling this technology. And so I make this to be a market that’s supplied quite competitively.

MR. WILSON: Thank you.

Nicolas?

MR. PETIT: Yes. So I have no particular view on the evolution of industry structure insofar as these technologies are concerned, but I was sort of recently struck by the sort of movement that we're seeing in the industry where large tech companies acquire open source companies, so I'm sort of thinking here about, you know, Microsoft buying GitHub and IBM buying Red Hats. And I was sort of, you know, trying to make sense whether there was an AI angle to that.

Now, I don't want to sort of, you know, push that idea too far because, you know, I'm not a business analyst. I have very, very low skills in that area and in many others actually. But when you think about AI, there's likely two things that spring to mind that could probably, you know, sort of explain in the background also part of the transactions from a strategic standpoint.

So one of them is that AI is sort of understood and seen as a general purpose technology. And, you know, you said general purpose earlier. I think that's quite -- that it's assumption. So general purpose technology is not like electricity or the steam engine. They have a lot of complementarities which are horizontal across the technology and economic sectors but also vertical across the sort of value chain.

And with general purpose technologies, we know there’s always -- there’s two faces. I mean there’s literature which say there’s two faces. There’s a face-off pushing adoption, trying to make sure that, you know, a lot of sectors horizontally and vertically embrace and adopt the technology, and the second one is basically investing and appropriating the returns of the technology.

And maybe what we are seeing here, since maybe 2010, 2012 when massive advances have been made in deep learning is basically we are in the adoption phase, and those large tech companies are basically trying to sort of force adoption also by the open source community in terms of all of those technologies, so bringing the open source community to adopt the AI source, which have been developed with like, you know, billion-dollar investments in the past -- in the past years.

The second thing that I want to say about those movements and those transactional movements in open source of large tech companies is that as I said before, has scaling increases and has you moving AI technologies across technology applications. Problems of defects and the fact that AI is very brittle. So when you move an AI sort of natural language processing system to, say, pricing, there’s a lot of fragility in that. And the AI might be subject to an optical illusion.

So having, like, many people onboard from different industry and especially from the open source community, people who are used to think about removing problems, solving problems, is probably a clever move insofar as working towards better AI is concerned.

MR. WILSON: Thank you very much.

Robin, I think you come at this question from a slightly different perspective, or your focus was different.

MS. FELDMAN: Sure. So although I largely agree with what has been said about access to all kinds of things, including access to data processing with one exception. And that is the very early end of the startup market. So right now you can access data processing for about $4 an hour from any of the big three major services. That doesn’t sound like a big deal. But it can be for an early-stage company because of how it adds up. So I talked to one company yesterday who’s doing biophysics. It’s a spinoff out of the university setting.

And at the university setting, the founder had access to federally funded networks that had 1,000 GPUs in them. Outside of the university context, it took this company, and they're looking for nonaddictive pain-relieving substances, which is important in society. So it took them 48 hours to train one agent and then they've got to test that. So coming up with one decent agent cost about $10,000. And that's going to add up very, very fast if you're an early-stage startup.

Now, if you think that disruption and innovation are going to come largely from later stages, not a problem in development. But if you think about past systems such as the programming cost it took for Facebook or two guys in a garage for Hewlett Packard, that's a bit of a barrier for the early end of the market.

MR. WILSON: Thank you very much.

Joshua?

MR. GANS: So I just wanted to -- so my experience with regard to the early-stage startups has been a bit different, and quite obviously it's coming from Canada, which is -- potentially has a different environment regarding resources for artificial intelligence, but at the University of Toronto and now elsewhere, we run this program called the Creative Destruction Lab. And over the past three or four years, I've seen maybe 300 early-stage startups in the artificial intelligence, machine-learning space which form the basis for the book that I wrote.

And I must admit that while talent is a huge problem, getting the data sciences, machine-learning experts and people who can understand how to optimize training of algorithms with respect to the CPU power and GPU power and other things like that, it hasn't been my experience that the startups have found themselves wanting when they've had the talent there. There has been -- they have been able to train their algorithms, they have been able to innovate, they have been able to launch products and do things.

Now, invariably, like with every startup, they have to make choices a bit different. And one of the things about our program is, you know, people coming out of university settings tend to get advice from one or two people and things like that. The problem with that is, you know, that largely depends on their experience of those advisers and which direction you should go.

Invariably startup choice is a lot wider than that. So if there was a constraint in sort of pushing the technology in one direction, there are substitute options, different customers and other things from where to start in order to sort of sensibly build your startup. And we've found that startups have been quite able to take advantage of those options.

Now, you'll never know if that -- it certainly wouldn't lead to the same outcome as if they made other choices. But from the overall perspective of thinking about antitrust, I don't see them as constrained from being able to innovate, enter, and provide some competitive pressure in that way.

MR. O'DEA: Preston, I wanted to follow up on a point that you had made about algorithms and off-the-shelf solutions. And I think you talked about translation, artificial intelligence, and the fact that you can take maybe a business report and put together some language around it.

Do you see certain applications that would be less commoditized such as in pricing applications where some of the off-the-shelf solutions being offered by, you know, some of the big folks out there might not work as well and that there might be specialization? Or do you think that the competition to provide AI will sort of be robust to whatever those applications are?

MR. MCAFEE: So, wow, that's a great question. We're in the snake oil phase at the moment.

So there's lots of stuff being sold that just is like nonsense.

Pricing, I worked on building a pricing engine for sale at Microsoft, and one of the big challenges you run into immediately, I'll just put in terms of Microsoft Surface. When does Microsoft run sales? At the back-to-school and holidays. That is to say they run sales when demand is at its highest. So if you just look at the data and run a regression or, you know, build a machine-learn solution, it actually doesn't work.

It gives you -- and there is actually a solution to this problem. And the form of the solution is called MML, first you build a model of what the people were doing, that is, you build a model of what generated the data, which is to say what were they responding to with prices. And then you use the errors from that model to identify the -- treat the errors from that model as experiments, and that gives you data. And that actually works pretty well.

But the point is, and this is why I say it's snake oil, that we're in the snake oil phase, is that if you just run the data, the data wasn't generated by a random process and it does not measure what you want to measure. So with pricing in particular, if you just try to take the data and run with it, it just doesn't work. And I can tell you that from personal experience.

I think more broadly, you know, there's a lot of data that wasn't stored very well. People created what they called data lakes. And they just dumped the data in, and actually any economist who's worked with government data finds out that, wow, stuff -- there's something just wrong here. And it will turn out, you know, in 1981, they changed the definition of the unemployment rate.

And so industry data is full of those sorts of problems. Actually, there's -- Gartner has the hype cycle. This is a really smart thing because we see it just happened over and over again where we see this peak of enthusiasm. You know, everything is red hot. If you used an Excel spreadsheet, you can call yourself a data scientist and get a great job, buy a house in San Francisco.

And we're in that -- this, you know, peak of the hype cycle. What happens next is the trough of disillusionment. And then it starts taking off. And I think we're going to see that, that is to say, I think we're going to see a lot of the things that we thought were going to work about AI just fail because

I gave you the example of the data, but there's also the optical illusions, and polluted data is going to be a big one. Or just -- you know, there's a certain amount of skill needed. If it's implemented without adequate skill, it's not going to work very well. And so there's going to be a lot of -- yeah, we spent a bunch of money on this and it was all wasted. I think you're going to hear that over the -- you know, as we go into the next recession.

And then sometime after that, it's going to turn out that all of our lives are affected by this everywhere. I'll give one example. If you use a Microsoft computer and you go chat to get help with your computer, you're actually chatting with a robot. That's a robot. That's a chatbot. It's a nice test, actually of how well this technology works.

Now, that's a situation where it works great because you've got very structured data, you had answers to questions, you know, frequently asked questions, and so on that they could draw on. And we're going to see a lot more of that, though, just that you're going to chat with a machine to get answers to questions, and you're going to be happy with it, I think.

MR. O'DEA: Anyone else?

MS. FELDMAN: I would just say I agree with Preston. I think it was Preston earlier who said that computers are easily misled or can be misled. Humans are misled all the time by data. Just throw some data in front of a human being, tell them there's a sophisticated algorithm behind it, they'll follow you off a cliff.

MR. WILSON: That prompts me to want to follow up a bit on something that's come up a couple of times, which is that finding qualified talent seems to be a real problem potentially for firms looking to adopt AI and algorithms. Is there something idiosyncratic about this technology that makes the labor market harder to understand, or this is just this is a new technology and eventually hiring managers will learn the signals to look for?

MR. GANS: I think it's just a training gap. I think it's taking a while for people to be appropriately trained. It's not only just being trained in machine learning and being able to do something off the shelf. There's still a considerable amount of artisanal or artistic-type characteristics to it, the sort of thing that only comes from experience. And so I think we are likely to have this sort of talent issue for some time, I mean, especially if the goal is -- you know, and we're going to realize this when the goal is to make AI deploy without errors or cause massive reductions in product quality or worse or harm. And I think that's going to show up.

And so I think it's going to slow the diffusion of AI throughout the economy unless, you know, it turns out that some applications can be very easily scaled and all of a sudden you have an AI solution that can just be deployed without the customer fully having to develop, personalize, or understand it. But I think we're still -- it seems like we're a ways off that yet.

MR. MCAFEE: Yeah, let me add to that, that traditionally the skill sets that you needed, which are things like building pipelines that move data around and process it, using like scaled cloud computing, those often didn't come in the same -- like if you got a statistics degree, you wouldn't necessarily get either of those two things, and yet you would get the other part that you need, which is understanding statistical data.

And so we haven't historically taught the skills that are needed in the same program. And we instead got them by hiring physicists who had had to learn some of those skills in order to do the research. That has changed completely. And now we're like generating people with exactly the right skill sets and so on. And so I think that will speed up the process of providing enough data scientists.

MR. O'DEA: Thanks. And I should mention -- I should have said this at the beginning of the panel, but there are colleagues of ours who are walking up and down with cards for questions. We have reserved time at the end of the panel. So if you have any questions, write them down and it will be delivered up to us to ask at the end.

So I'd like to move the discussion now to what effect we think that AI and algorithms may have on market structure for various industries across the U.S. economy. And, you know, I think there's three possible options that we talked about on the precall before this panel, and one is to what extent do we expect that it will create entirely new markets, to what extent do we think that it will sort of allow challenges to companies who have been entrenched in a dominant position for some period of time, and, lastly, do we see certain markets where it may be likely to lead to increased consolidation? And sort of what factors might lead to each of those three outcomes and which of those outcomes do you think are most likely?

So, Robin, why don't we start with you.

MS. FELDMAN: So on a simple level, we will see the emergence of new markets for creation, production, and implementation of AI. You think about the market we've been talking about on the market for AI processing power with its three key players that are Amazon, Microsoft, and Google. Those three players existed and they competed with each other in the past, but this market didn't.

You're also going to see what are new markets for new societal activities -- so driverless cars or what I call implantable nurses. And we aren't just going to see new markets but also adaptation markets. That is, as AI spreads throughout industry, some existing players will try to bring in AI expertise in-house, and others are going to turn to third parties to develop the AI for them and to use it externally.

It's these middle-level players, I think, that are important to watch because they reach across competitors and across industries. Anyone who reaches across competitors has the potential to operate as a hub-and-spokes, that is, connecting the competitors for the purpose of collusion through those third parties. But I think there’s a much trickier issue as well. And that is with mid-level players who reach across industries, we may have to adapt our notions of market definitions.

So right now, current market definitions tend to be grounded in the idea of a specific product market, but when you have key players that are working across market and across industries, we have to worry about multiplicity effects. So when can a wide-market player, using interactions across those markets, impact price and supply in those markets without having power on all of those markets or maybe even in any of those markets? Now, I can’t predict for you where that will happen. I’m not in AI, but I can tell you it’s happened in other contexts and it will be important to watch.

And, finally, in a period of disruption and creation, competition authorities want to keep an eye on big players. And I don’t just mean tech. So think about the transportation industry where trucking and delivery is going to be completely changed. So big players are unlikely to disappear quietly into the night. And they may go to great lengths to try to hold onto their power. So it’s going to be a tricky time.

Perhaps one of the most important things competition authorities can do during this period of time is not get dragged into what is essentially big players trying to rev up government forces to protect them.

MR. O'DEA: Thank you.

Joshua?

MR. GANS: So I think that is a largely correct view. I imagine that companies that were born just before AI or a decade before Amazon and Facebook and so AI has been a gift to them to be able to improve what they were doing and in the process increase their shares of the market and continue to grow.

What's interesting is that especially when we've got a new technology like this coming in, there's so much that is unpredictable about where it's going to hit and who's going to be favored, and other things like that. You know, to the extent that AI is statistical tools, improving product quality, improving productivity, you know, we don't necessarily expect much impact on sort of a general competitive landscape except that things just get better.

Where we might get some bigger effects is that there are times in which these new technologies manage to completely transform and surprisingly kill incumbents that were previously the darling of antitrust focus. And, you know, we saw that with Blockbuster. That was always listed as that. And, you know, it disappeared quicker than any antitrust case could be build against them.

And I suspect, and I just want to give you an example, and I'm just going to preface this by it's pure speculation, is I wouldn't be surprised if a company like Google might be particularly susceptible to some startup applying AI in an innovative way. I know that everybody looks at Google and says, wait, that's a quintessential monopoly. That's the company that we want to focus on. But it's hard. It's got a search engine.

And the search engine, while certainly when it first appeared and you know, depending on who you talked to, is at the frontier right now in terms of being able to search for stuff, is not perfect. It's not perfect. And I'll tell you why it's not perfect. Just think to yourself when you've done a search for a thing that you know is there and you're just trying to search for its location on the web, and Google doesn't serve up that result, and you have to modify the search and other things like that to properly communicate with Google as to what you want.

Well, that's the kind of thing that AI could come in and provide a different way of sorting the information, aggregating it, trained on it, that could do a much better job than that. And if that appeared tomorrow, subject to, you know, the ability to roll it out and other things like that, Google could lose market share very, very quickly. It's entirely possible. You know, while there's default behaviors and other things like that, those things are possible.

So I wrote a previous book, a few books ago, called The Disruption Dilemma, which was about this. And there's no doubt that contrary to sort of the management theorists who talk about disruption is everywhere and we're all whatever, it's all competitive and business is hell, blah, blah, blah, you know, having key assets, having various entries still can soften the effects of that and give you time to regroup.

But there are other cases in which the way of doing production in the industry so changes that your incumbent firms are actually at a serious disadvantage because they both -- not only do they have to build a new system, but prior to doing that, they have to dismantle their currently profitable system. And so that's two things, whereas a startup can just do one. And so I think that sort of thing might happen here.

Now, that's not a suggestion to be anything less than vigilant on antitrust, but it's something to just give us some pause as to which way this is all going to go.

MR. O'DEA: Thank you.

Nicolas?

MR. PETIT: Yes, sure. So in your initial question, you were referring to the effect of algorithms and AI on market structure, and one aspect which is slightly distinct that I want to address is whether the research that we're having today on algorithms, AI, and markets is too much focused on the supply side, sellers using AI to price products and whether we have been thinking enough about the effect on the buyer side.

And so while there's been some discussion and thinking about, you know, whether AI technologies could actually capacitate and enable buyer power for consumers and, you know, there's been reports, OECD, CMA, talking about that. Now, what I want to talk about very briefly is about sort of distinct thinking about buyers in those markets. And the question is whether agents on the demand side can deploy AI systems to subvert the use and employment of algorithms by strategic sellers. And the optical illusions that you were talking about before, in the field, we talk of adversarial examples are a case in point. So we know that AI systems are extremely brittle, that deep learning algorithms are very vulnerable to small perturbations of the inputs, imperceptible to humans.

So you change a pixel in a panda picture, and you're going to see a lion, right? The AI is going to see a lion, where, you know, no human would make that mistake. And so we are seeing today some technology developers develop technology which uses adversarial examples and other sorts of technologies to entitle buyers to actually undermine the working of algorithms on the selling side.

So to give you a bunch of examples of those bot-management or bot-mitigation technologies, we talk here about the use of Captcha. So you know those boring -- those boring tests that you have to go through to prove that you are a human, they're actually named -- the Captcha acronym is named after the Turing test, automated Turing system for -- to detect humans from machines.

Software developers are selling technology to manager whether visitors click on certain areas of buttons on websites because algorithms always click, say, on, you know, the right corner, whereas we humans would sort of randomly touch, you know, whatever area on a button.

Technology providers also sell software which entitle buyers or, you know, rival companies to detect whether a certain query is issued from a mobile phone. And so for instance, they managed to do that by retrieving information on the phone through the accelerometer or gyroscopical information. So, you know, when a human touches a phone, there is slight movements, and the technology can detect whether that's human or whether that's a bot.

So what's interesting about those technologies that we are seeing and I was discovering that a few months ago, a middleware market segment where technology companies are developing such technologies to develop defenses for buyers and rival sellers to undermine the working of algorithms on the selling side. And so, for instance, a company called Akamai Technologies develops defenses for firms which want to avoid scraping bots. Another company called Luminati, they have developed technology to mask bots.

And the end equilibrium of those technological interactions is not a given. And so I would say if antitrust enforcers want to be on the lookout, maybe they want to make sure that there is competition and innovation in this middleware segment, which will provide solutions -- technological solutions to market players willing to get good bargains in transactions.

MR. O'DEA: Nicolas, do you see some of these tools being used by sort of individual consumers, or would this primarily be by firms and actors who are on the buy side in markets?

MR. PETIT: That's a very good question. So most of the evidence that I have gone through is analytical evidence, right? There's a huge fact-finding exercise that needs to be made in relation to the technologies. What I understand, that sophisticated buyers and sellers use those technologies, but we should not -- I mean, competition is all about that, actually. It's about, you know, making sure that markets expand and that consumers from all sides -- sophisticated and less sophisticated -- can avail themselves of them.

I want to add something to your point earlier. In this middleware market, you're seeing a lot of, say, small companies. I'm not sure if, you know, $2 billion turnover per year is a small turnover, but you're seeing that kind of companies, but you're also seeing companies like Amazon, for instance, which provide such tools as part of its available U.S. offerings. So, you know, large tech platforms, smaller middleware companies.

MR. O'DEA: Preston?

MR. MCAFEE: So, first, I just wanted to follow up on both Nicolas' and Joshua's point is that AI assisting consumers doing things like, let's say, looking for airplane fares, so this is you set it to go and it monitors the fares, I don't know if you know this, but airplane fares change multiple times a day. And so if you don't need the fare right now, it's actually optimal to search, but it's kind of costly. And so there are companies monitoring airplane fares.

And this is the kind of thing that is a threat to Google. In fact, there was a period of time where people thought Google might fall just because it was having trouble making the transition to the phone. Actually, the same thing was said about Facebook. Now, they both succeeded in making the transition, but when you get these new technologies that change the way we behave, and it's pretty interesting thought experiment, but what comes after the phone? What's the next one? And then the companies having trouble with that.

I want to make a very different point, though, which is AI generally is going to -- well, related to this, it’s going to facilitate lots of new business models. So just the way that companies deal with their customers, so can now change because they can have smart -- especially smart interactions on the phone as a way of dealing with customers. And when you gets new business models, will the existing firms respond to that by trying to either incorporate those business models or change their business model to survive?

And then -- so that actually -- when we get new technologies, we often get a wave of entry into many different businesses, so we get the -- you know, if you think about electricity, we got the creation of lots of new industries that didn’t exist at all before, and we got new ways of doing old businesses that created more competition, at least maybe temporarily, but it created more competition in those industries.

Another thing that you get is a merger wave. And, in fact, all of the merger waves except one -- I think there’s six or seven of them -- all of them but the 1980s merger wave were brought about by new technologies. And so AI could easily create that kind of merger wave. And that comes about because as firms try to evolve their business model, they realize if I'm going to make this business model work, I need a new capability I didn't have and they turn around and try to buy that so that they can get that capability.

And so I expect to see that -- another merger wave set off by AI over the next ten years.

MR. O'DEA: Does anyone have any thoughts if that merger wave comes? Should the agencies approach it the same way that they are currently, or are there any special sort of rules or techniques that we should be applying in this setting?

MR. MCAFEE: Well, I have a lot of thoughts on this. But, first -- well, overall I think the antitrust laws, they have the right focus and they are up to the job. That is I'm not one of the people that say, oh, everything has changed, we need new antitrust laws. No, I think the antitrust laws have been remarkably good.

The one thing that I would point to, though, is that you often see -- now, let me use the defense consolidation as an example. You often see one merger spawning another. And so that is -- well, actually, the example -- a good example of that is the cable companies buying content. And that seemed really like approving the first merger causes additional ones. And that's one thing our antitrust laws can't handle, is that they -- you know, this merger is either anticompetitive or it's not.

And I like the defense example because we let Lockheed and -- or, excuse me, we let Boeing and McDonnell Douglas merge, and then we let Raytheon and TI Electronics merge. And what that did was create one company that was dominant in air frames and another company that was dominate in defense electronics.

Had we done it the other way, that is to say, rejected the Boeing-McDonnell Douglas and maybe gotten Boeing-TI and Lockheed-Raytheon, we'd have had two firms that had much more similar capabilities and hence would have produced a more competitive environment. And so that's one place where the merger guidelines -- or, excuse me, the merger precedent don't -- can't accommodate.

MR. O'DEA: Anyone else?

MR. GANS: I'd just second that as well. That seems something that would be a good place to have some sort of process that allowed the broader review of sort of these industry knock-on effects going on.

I would also -- you know, I don’t know how you would do this, but it’s clear from numerous examples, and it’s not just here, it’s around the world, that this is a sort of blind spot for legislative-based antitrust.

MR. O’DEA: Okay, so to focus in specifically on AI and algorithms and some of these technologies, are there any general rules that you can think of to help identify when the technologies are likely to facilitate entry and disruption versus restricting entry? Are there any market factors we should be looking to? And does anyone see any rules of thumb or screens for identifying when AI tools or data are likely to make markets less contestable or when we may be reaching tipping points?

MR. MCAFEE: So I’ll just mention, I would look whether a merger seems to be locking up data. So, for example, I probably would not want to approve a merger between any of the credit rating agencies just because that’s going to limit the competition and the supply of data.

MR. GANS: I thought they have open data policies?

MR. MCAFEE: Well, no, they only give it to the Russians. So I'd be looking at does -- is this, you know, creating controlling interests in sources of data that don't have substitutes for rivals? And I think, you know, in some sense, the standard way that we do merger analyses is going to catch this, because we're going to talk to the rivals and they're going to be screaming about the data. We'll talk to the rivals and they'll be screaming about the data.

So I don't think that that's -- it's not that we wouldn't catch it, but that would be the -- I would be looking specifically for is this really locking up, you know, merging two similar sources of data and leaving us with no competitors or one weaker competitor.

MR. O'DEA: Thanks.

Does anyone else see any market factors or screens that we should be looking for?

(No response.)

MR. WILSON: Well, let me, then, shift the conversation slightly to concerns related to intellectual property and the defenses and mechanisms to encourage people to continue developing new IP. I guess in particular I'm interested in thinking about how do various IP regimes fit with AI and does the intersection raise particular competition concerns.

Robin, do you want to start us off?

MS. FELDMAN: Sure. So when we talk about intellectual property rights in AI, we're really talking on two levels. One is rights in the AI program itself, and the other is rights in those things created by the AI program. So let me talk for a moment about rights and those things created by the AI program. And those creations could be data aggregations, software, or processes like the advice to give a loan applicant or the direction to send a car in or a disease treatment.

So protection for things created by AI under U.S. law is very uncertain at this point. Copyright Office language casts doubts on your ability to copyright things created by AI. And with patents, things created by AI are likely to fall into the baskets of software or business method patents. And the Supreme Court has drastically cut back on your ability to protect those things with patent. Forget about the obstacles you have related to something created by AI. The U.S. courts haven't ruled, however, on any of this stuff. And I think it's going to be somewhat of a slog for protection.

But the real issue is the following, and that is whether we're taking about protection for the AI program or protection for those things created with the AI program, copyright and patent systems are not a good fit. So think about transparency. Patents are supposed to teach anyone skilled in the art how to do something, but that's not how it plays out in the fields in which artificial intelligence is likely to interact with patents.

So specifically with software and business method patents, you only have to disclose in your patent application the outcome. You do not have to show very much about how you got there or anything you're doing, if at all. In contrast, consumers and regulators are going to want to have confidence in AI's trustworthiness. So nontransparent protections like copyright and patent, not to mention trade secret, are in tension with this.

Second, consider the issue of contributions to creativity. If AI programs are deriving their creative results in part through the collective decisions of many people, should that creativity be solely attributable to the program, or do we have concerns when those who are first to large amounts of data or bottlenecks, do we really want to give them the ability to exclude everybody when a lot of "everybodies" may have contributed in some way to the development?

And, then, finally, patent and copyright systems operate on a timeline that is entirely foreign to AI. It just doesn't fit the shelf life. Patent protection lasts 20 years, which is an eternity in the AI field right now. Forget about copyright where for something created by an institutional author protection lasts 120 years. The point is simply that patent and copyright may not be the best fit for protecting AI systems, and certainly not if we're worrying about international competitiveness.

MR. WILSON: Thank you very much.

Preston, did you want to pick things up?

MR. MCAFEE: Absolutely. I can summarize my remarks with nothing is obvious to a patent examiner. I think I agree with Robin on many different things, on all aspects of this is that we've issued patents -- well, it will be interesting to see whether the Patent Office allows the following kind of patent. I take something everybody -- you know, that has been around for 20 years or 30 years and I stick a little box in it that says AI and they say that's novel.

MS. FELDMAN: I'll invest in that patent.

MR. MCAFEE: So, yeah, we could just

actually go issue -- we could make 9 million applications of those right now, just stick AI in the existing patents. So in some sense the software patenting has really been broken. And that's been a -- we have lots of overlapping patents. You know, if you look at, like, mapping program -- so the statistic on cell phones is you need access to 250,000 patents to make a cell phone. There's too many. They can't all have been novel.

In fact, probably 249,950 weren't novel. I have 11 patents. You can go look at them. They're public. I'm not going to remark on whether they should have been issued or not. I want to make two other points, though. One -- actually, I want to make three other points. The Supreme Court has actually been pretty hostile to software patents, and I think rightly so. And they may fix what the Patent Office didn't fix. And so that -- it's unfortunate that the way that they're fixing it is kind of expensive because we have to litigate it as opposed to just doing it right in the first place, but at least going forward, it may be better.

I think they made a mistake when they said that you can patent a life form. And I am kind of worried about -- you know, one thing about AI is is that a lot of AI is quite generic until you stick data in it. You can't patent the generic thing. That's been around too long. So that won't -- you won't get patent protection on that. And it's pretty hard to protect the specific numbers that come out because they change all the time. So it would have to be the process of applying AI to some field is what's getting protection.

And so I have some hope that having been down this way with this path with software patents that we won't do it with AI, but I'm certainly worried about it. And I think there is an analogy to patenting life forms as I think we called that one wrong. We should have said you can't patent a life form. It's a living thing independent of the person that created it. But I'll ask Robin afterward whether she agrees.

And, then, finally -- are we still talking about privacy actually, or have I gone too early?

MR. WILSON: No, no, by all means.

MR. MCAFEE: All right. So the EU with its General Data Protection Rule has run a grand experiment. And this is a giant benefit to the United States because we get the what, did this work or not. The EU is big enough to be relevant to us in scale.

And that is to say, people will redo their business models in order to serve the EU because it’s valuable enough, whereas as if, say, North Dakota did it, probably not. And it’s -- we’ll learn a lot. Like, this is either going to cause lots of problems or it won’t. If it doesn’t cause lots of problems, we should probably just adopt it. If it does cause lots of problems, then we at least -- okay, but it caused them problems and not us. And so I’m really glad they did that. And I think it’s going to be of great benefit to the U.S. as we learn how well it works.

MR. GANS: You better put the word "potential" benefit.

MR. MCAFEE: Potential benefit.

MR. GANS: Yeah. You have to learn from it.

MR. MCAFEE: An unexpected value.

MR. WILSON: Thank you. And does anyone want to chime in?

MR. PETIT: Yeah, I just want to remark that the European Patent Office recently issued guidelines on whether AI and algorithms are able themselves and made very clear that computational models and mathematical formulas were not in themselves subject to patents and that the patent applicant had to prove that this came with a technical purpose, which has a state-of-the-art, you know, set definition and, therefore, we should not sort of, you know, create a strawman that, you know, algorithms and AI systems will in themselves -- by in that generate form be subject to patentability. I just want to make that clear, and so, you know, I sort of refer people to the guidelines of the European Patent Office.

MS. FELDMAN: So I would comment that I heartily agree with Preston. My concern is that even though the Supreme Court has cut back drastically in the last 18 months to two years, the Federal Circuit, which is the appeals court right below that hears all patent cases, has swung the pendulum entirely in the other direction, reading the Supreme Court decisions to give lots of room.

The U.S. Patent and Trademark Offices has jumped on this and said, grand, and is handing out patents hand over fist, particularly in the AI field. So, you know, it may be a little soon to declare victory and brings the troops home.

And I would also just push again on the international competitiveness point. If we make a mistake and we tie up things too early and we intern some early market players and we slow down our innovation that way, there are other countries like

China that are poised to just eat our lunch in this field, and we really have to keep an eye on the context, not just internally but externally.

MR. O'DEA: So I have a quick -- oh, sorry, go ahead.

MR. WILSON: I'll go first. So my question is I think the divergence in IP regimes between the U.S. and EU provides us with an interesting natural experiment, but, you know, how long do we give it before we either adopt or start gloating?

Nicolas?

MR. PETIT: Yeah, I want to say two things again. So on GDPR, one often mistaken element of GDPR regulation across the world is that GDPR is there for competitive reasons or for to address market failures of the kind we've discussed in the antitrust field like, you know, problems with monopoly power and so on and so forth.

Now, the rationale for GDPR is almost exclusively moral. Right? And I'm not too sure that a piece of legislation which stands on the basis of a moral choice unrelated to market outcomes lends itself to impact assessment of the kind we're running in terms of competitiveness, whether it's going to be good for firms, bad for firms, good for industry, bad for industry, and so on and so forth, of course, is a relevant concerns, but insofar as GDPR has basically been predicated on the basis of very strong moral choice by the European Union rulemakers, I'm not too sure, you know, we should read too much into that.

Now, of course, others systems of flow, other jurisdictions that may have a different feel about those moral values at the heart of GDPR and whether they can be compromised with more economic objectives such as industry performance and so on and so forth, but that's not how GDPR was conceived in the EU.

The second thing I want to say is before we sort of try to draw the lessons of the GDPR natural experiments, I think we should need to wait a little more because enforcements of the regulation has not yet started. So we are yet to see which firms will be fine for infringements, whether the large players are the massive infringers, whether small players are on the receiving end of enforcement.

MR. MCAFEE: Okay?

Mr. WILSON: Please.

MR. MCAFEE: I take the second point as a complete answer to the question of when should we consider this experiment done. We have to see the experiment through first. One thing about GDPR is it says you can't keep someone's data -- you can't use it for a purpose other than a purpose that they supplied it for directly without permission. So it flips the -- like, you have to give your address to Amazon for them to send you stuff. Otherwise, how would they know where to send it?

So what this says is Amazon can use your address to send you stuff, that's the service that you signed, but they can't use it for anything else. That's what GDPR would say about addresses. This is a pretty -- this flips the ownership rights of the data from the companies to the individual with some limitations because companies had this data -- these data in the first place because they needed -- you know, again, you can't get a Google search query if you don't give Google the query. But what it says is Google can only use that to answer your query and not use, you know, to offer you advertising, for example.

And so I think as an experiment, it's a pretty interesting one, and we can learn a lot from it.

MR. O'DEA: Thank you. Yes, I wanted to go to a couple of questions we got from the audience. One, I think, Joshua, this is primarily to you, but

I'd be interested in the reactions of all the panelists.

Following up on the point that you had made about AI and its potential capacity to allow a new entrant to challenge Google and search, and the question is how do we square that point with some of the conversation that we had earlier around the importance of data and how data can act as a barrier to entry, given that there are, you know, millions of searches going on with Google sort of instantaneously, to what extent will that data be relevant?

And I don't want to make it just a question about Google. So are there, you know, situations where that balance between AI as a challenge versus the data that an incumbent are sitting on will be particularly relevant, or how should we look at that?

MR. GANS: So just to put this in a historical context, we've had already a situation of significant entry by a startup into the search space starting from no data or market share, and that was Google. Google did it. And it did it because it scraped the web itself for information and was able to, you know, through page rank and other means, contextualize it. It only more recently evolved into a situation where the leading way of doing search engines was to wrest it off what humans were doing essentially in trying some sort of artificial intelligence for it.

Now, it is entirely possible that a startup could -- the web is still out there. It's still visible. That is there for startups to use. So the answer would be, it would not use that same data that Google currently has an advantage on. It would find some other way, and that's precisely why that's vulnerable because Google at the moment is probably -- well, if Google were like my other companies historically in this situation, they're probably not -- don't have a team out there saying, I wonder if we do just as well if we don't look at our own data? Why would you do that? They've got their own data and they do very well with that. There's no real thesis for it.

The chances are that thesis will develop elsewhere and moreover because that is in a constraint that people will be able to enter. In other words, what might have been a barrier to entry in the past if the new sort of technology is reconstituting things is not a barrier to entry in the future.

Now, that doesn't happen very often. Let me preface that, it doesn't happen very often, but it did happen once in recent memory, and that is when we expelled all of the incumbent mobile handset makers from the industry -- Nokia, Blackberry, Motorola. These were firms that had been very successful, pretty much dominated the industry, all gone because the way a phone -- what a phone was and did was just reconstituted.

And, you know, did it -- you know, so that just happened. And that's happened in recent memories as well. So, you know, there is some vulnerability there. If you've got network effects like Facebook, if you've got a massive real infrastructure like Amazon, you've got your traditional barriers to entry. And Google have some of that as well, again, but I just wanted to put in the thought that they may not be invincible.

MS. FELDMAN: So here's a concrete example. Right now, data is king. Machine learning, systems need large amounts of training data, past data. But imagine if in the foreseeable future, AI systems develop so that they can create their own training data. And that's not something that's just a pie-in-the-sky idea. In that case, having massive amounts of past data becomes less important and is more subject to disruption.

MR. O'DEA: Preston?

MR. MCAFEE: So I agree with Joshua, but actually, you can look at Google itself and see where Google thinks this is about to happen, and it's the smart speaker. And they think, you know, the idea of the smart speaker or for that matter talking to your phone is -- it will understand you better. In fact, there are something like 50 million Chinese use this product called Xiaoice, which is a chatbot, mostly teenagers. And they chat with it. It's like almost 30 million people chat with it an hour a day. And 2 million Japanese as well.

So that opens a new opportunity to handle search. A chatbot that you've been chatting with for an hour a day for many years understands you way better than Google can. And so that's a threat to Google. And, of course, they -- as a result, they're doing everything they can to have the best smart speaker in the market because they think you're chatting with them and ordering things and so on is a threat.

The other thing I would say is, is that the kind of data that you want -- you know, they have a lot of one kind of data, but Amazon's got way more data about what I buy than Google does -- much more.

Even though I might search for some of those things, Amazon knows whether I actually bought it or not. And for that matter, my credit card company knows all that stuff, too.

And so this ability -- you know, it's true that you need data, but it's not necessarily -- you can't assume that Google's data is like the perfect data. They do everything they can, of course, to have as much as they can. They are extreme in that regard, and I think -- but Facebook has a lot of data, too.

MR. O'DEA: Thank you.

Nicolas?

MR. PETIT: Sure. So, again, you know, the semantics of the discussion are sometimes a little disconcerting because we talk a lot about data and barriers to entry, but the question may be what are the instruments that entitle companies to harvest data. And the better your instruments, you know, the higher the barrier to entry.

So, for instance, you know, Google has, you know, the search engine as the sort of massive harvester of data, but, you know, when mobile came, Google was very concerned that, you know, people spend more time on their mobile phone than on a search engine, and so, you know, it took like, oh, so many attempts to be on the mobile phone, which actually generated antitrust proceedings in the European Union in the Android case.

Now, the next question is, of course, what will be the next user interface which will harvest more data and be the barrier to entry. And so, you know, Google invests in driverless cars because it thinks people spend a ton of time in their cars. Maybe, you know, we'll have the shower or whatever. I mean, there's an example in my family at some point, like when broadcast TV was introduced in the 1950s, the grandfather of my wife, you know, was telling his wife, you know, shut up, they should not know what we're doing. You know, so there was this idea that the people in the broadcasting channel were actually observing what people were doing.

And so I think this battle is more of this kind than the battle for data in itself. The instruments, the entry points where you harvest data are really what matters and where you can see markets reconstituting around new technologies and disruption.

MR. WILSON: Thank you very much. And though I have no doubt that we could keep going for solidly another 90 minutes, I'm afraid that our time has all but elapsed. So if you wouldn't mind joining me in thanking our panel for their interesting remarks, that would be greatly appreciated.

(Applause.)

PRESENTATION

MS. CONNELLY: It is my great pleasure to introduce Joy Buolamwini, who will speak about her work on facial analysis technology. Joy is the founder of the Algorithmic Justice League, which researches the social implications and technical capabilities of artificial intelligence and increases the public's understanding of bias in technology.

She is a Rhodes Scholar and a Fulbright Fellow and holds two master's degrees -- one from Oxford and one from MIT. Her bachelor's is from the Georgia Institute of Technology, and she is currently completing a Ph.D. focused on participatory AI at MIT.

Joy?

MS. BUOLAMWINI: Thank you for the introduction. Well, today, it is my pleasure to share with you some of the research that we've been doing with the Algorithmic Justice League that shows how facial analysis systems being developed by leading tech companies have concerning issues. So here on this intro slide, you see Amazon mislabeling Oprah's face as male, and why might this matter? Well, Amazon currently sells facial recognition technology to law enforcement departments.

You have IBM misclassifying Serena Williams' face here. And with image captioning, you see that Microsoft is struggling on Michelle Obama, describing her as a young man. So these examples I show to remind us that technology is not infallible, and even the largest companies that are making billion-dollar investments into this space run into issues.

So I want to go over facial analysis technology major tasks just so we are clear on the type of technology that's being discussed. So if you look at facial analysis technology, it's broadly about pattern recognition. Machine-learning techniques are used to come up with these patterns for various tasks using large training sets. So the most fundamental task for facial analysis technology is face detection: Is there a face or not?

Once you pass that in the pipeline, you might ask different types of questions, like what kind of face are you seeing in the first place? What's the gender of the face? What's the age of the face? Then you have another set of questions you can ask, which is really about do you know the identity of the face, have you seen this face before? So this is what's generally referred to as facial recognition.

And you have facial identification, which is a one-to-many matching. So think of searching for a missing person or a criminal suspect. And then you also have face verification, which is looking at a one-to-one matching. So think about unlocking your iPhone or paying for something with your face.

So all of these tasks are based on data, and they’re also susceptible to something that I call the coded gaze. So let me go back here. So the coded gaze is my term for algorithmic bias that can lead to exclusionary experiences or discriminatory practices. And in this video, which I hope we’ll play in a while, it shows that actually coating in a white mask to have my face detected by the system, whereas my lighter-skin colleague in this particular video just has her face detected without needing to put on a white mask.

And so this personal experience is what led me to start exploring issues within facial analysis technology. And I decided to look beyond face detection because there were some systems that detected my face, and there were other systems that didn’t detect my fact, but those that did ended up labeling me male or getting my age off. So I wanted to see if this was just my unique facial structure or something more systematic.

And these might seem like innocuous mistakes, but when I came across the perpetual lineup report from Georgetown Law that showed over one in two adults in the U.S., that's more than 117 million people, has their face in a face recognition network that can be searched unwarranted using algorithms that haven't been audited for accuracy, I realized these types of errors could have real-world consequences.

And if you look across the pond in the U.K. where they are reporting real-world performance metrics on these systems as deployed, you're getting false positive match rates of over 90 percent. So in the U.K., you've had more than 2,400 innocent people with their faces misidentified as criminal suspects. And you even have a case where two innocent women were misidentified as men. So some of those misclassifications that I've shown earlier do make an impact.

And when we're thinking about facial analysis technology, we're not just talking about its application for law enforcement. You also have systems that are being used in hiring. So Hirevue is a company that purports to do video analytics, and in these videos, they apply AI to pick up verbal and nonverbal cues to help inform predictions about a potential candidate's performance.

So in this case of predictive analytics, the face is being analyzed, but they say the way that they analyze the face is they compare it to the top performers at an existing firm. So if you have a largely homogenous group of top performers, it could be the case that it's picking up on mannerisms that are more to the demographic and less to the actual task.

Beyond facial analysis technology, AI is being used in a host of decision-making areas, which makes it even more pertinent to make sure we're understanding how these systems function across a diverse range of individuals. And so this is what my dissertation work, my MIT master's thesis, focused on, which was saying for commercially available AI systems that do gender classification, how accurate are they across different genders, and does the skin type also matter?

But before I could really investigate this question, I ran into a problem, which is that the existing standards, the existing gold standard measures for success in the field are actually largely flawed in that they're overwhelmingly male and predominantly lighter skin. So if we are in a case where we have pell-mell data sets setting the benchmark we're destined to fail the rest of society for technologies where data is destiny, and that is where we see ourselves now.

And to bring this point home, if you look at Facebook back in 2014, they released a paper called DeepFace. And there was much rejoicing in the computer vision world. Why? Because they improved the state-of-the-art performance on the task of face verification by almost 20 percent, which was great news because it showed that there were effective techniques being employed using deep learning.

However, if you look at that gold standard benchmark, right, you'll see that it is 78 percent male and 84 percent white. So if this is the gold standard we're using, we're giving ourselves a false sense of progress which can lead to misleading technology. And it's not just the industry benchmarks that are vulnerable. Even if you look at the benchmarks from the National Institute for Standards and Technology, you'll also see that they reflect some of these large skews.

So in the case of the IJB-A benchmark, you'll see that it is about 76 percent male. Now, if you do an intersectional breakdown of this benchmark where you're looking at skin type as well as gender, you'll see there's an over-representation of lighter skin men, here, 60 percent, and less than 5 percent of that particular benchmark are of darker-skin women. So it became a bit more evident to me why some of the issues I was encountering might not have surfaced in the industry or in the research.

So given these skews, I developed a more inclusive benchmark so we could see the performance of these systems across a range of skin types and again with this benchmark that was better balanced on gender parity. And so I was able to test commercially available AI systems that are being sold right now. And I chose IBM and Microsoft, given their huge investment within AI cloud services and also Face++ in that Face++ has access to one of the largest databases of Chinese faces, and we’re often hearing that China will have the data advantage when it comes to AI, so did that play out?

Well, when we look at the results, we’ll see that the overall accuracy of these systems on our particular benchmark seems all right -- 88 percent to 94 percent. But once you start to break down the performance by gender or skin type or the intersection, that’s where disparities begin to emerge. So if you look at the breakdown by gender, you’ll see that there is an air gap, right?

So this doesn't depend on the skin type at all, just one gender or the other. And if we do a breakdown by skin type, we also see that there's a substantial gap in terms of the performance with much better performance on lighter skin than darker skin.

Now, once we start to do an intersectional breakdown, we really start to see interesting patterns emerging. So in this case, the best performance group -- performing group are lighter-skin males, and in the worst performing group, we have darker-skin females. This was the best-case scenario with Microsoft.

When we moved to China with Face++, right, here we see the best performance is on darker males, showing the importance of an intersectional analysis, but we also see that it's failing in one of three women of color, right, 65 percent accuracy. And similarly for IBM, you see that the worst-performing group, darker-skin females, and IBM also doesn't perform as well on darker males compared to its peers. And even if you look at the lighter-skin section here, right, there's again a difference between male performance and female performance.

Now, if we just aggregate these numbers, we get performance results that I found quite surprising for commercially sold products for binary classification, where you have a 50/50 shot of getting it right by just guessing. So you see here, for type skin, women, dark-skinned women, you have error rates as high as 47 percent on a binary classification task, real-world commercially sold products.

So I decided to share these results with the companies to see what they would say, and IBM and Microsoft got back to the research group, and all of the companies released new APIs after this external audit, so new systems that were reportedly improved.

And if we look at the self-reported improvement, right, we see that there is a significant jump in accuracy for their worst-performing group, but when we did our external evaluation, we did see an improvement, but the improvement was not quite as high as they reported because the type of data that they're using and also the thresholds they're going to set it to will, of course, put the companies in the best light.

But even if we have more accurate systems, accuracy does not mitigate abuse, and you have a case, for example, where IBM was reported to have equipped the New York City Police Department with facial analysis technologies that could search video footage by skin color, by facial hair, and even the clothing people were wearing, so essentially providing tools for racial profiling that could violate civil liberties.

So, here, the question isn't about accuracy; it's about abuse and use, which is why I'm here speaking to the FTC because it's up to regulators to protect us and within the face space, our research shows there are specific steps that can be taken to make sure these systems are not abused or weaponized.

One is making sure that companies actually know the performance of their system so they're not misleading us by presenting software that supposedly works well for everybody but is truly just optimized for a small subset of the population. We also need the results to be published in terms of how they're performing on the benchmarks that exist. And they need to support independent research evaluation. Otherwise, the self-reported results we'll get will not give us the true picture.

But we also need to make sure that when we are doing these national benchmarks we're also making sure these benchmarks are representative. So an immediate step that can happen right now is requiring NIST to publish the demographic and phenotypic breakdown of the existing benchmarks, and then also making sure that these numbers are just aggregated in a way where we can see if there are intersectional performance differences.

Beyond the research, we also need to be thinking about consent. Do we have a choice in whether or not our faces are being used? Facebook right now has over a billion face prints of biometric data that many people don't know they are collecting. Could there be an option to purge that information? Transparency is often crucial but not just in terms of how systems are performing based on these benchmarks but what they're doing in the real world. And we saw the importance of that when we see the results from the U.K.

And I see that time's up, so I'll go quickly on these last parts. We need due process. If you have a company like Hirevue using face-based analytics to predict your potential job performance, is there any way to contest that kind of prediction and what mechanisms can regulators put in place so that there is more due process?

And given the rapid adoption of facial analysis technology, we really have to think about its implications on privacy. You can change your password; you can't necessarily change your face as easily. So I'll leave it there for regulators to think about how to safeguard our faces in this new frontier of algorithmic justice.

Thank you.

(Applause.)

MS. CONNELLY: Thank you, Joy, for that very interesting talk. Now, we will take a lunch break until 2:15, and then we will reconvene for the last sessions. Thank you.

(End of presentation.)

(Lunch recess.)

KEYNOTE

MS. CONNELLY: Hello, welcome back from lunch. We are delighted to have Jennifer Wortman Vaughan here to speak about fairness and intelligibility in machine learning. Jennifer is a Senior Researcher at Microsoft Research in part of Microsoft's Fairness, Accountability, Transparency and Ethics Group. She is especially interested in the interaction between people and AI and has often studied this interaction in the context of prediction markets and other crowd sourcing systems. She completed her Ph.D. at the University of Pennsylvania in 2009 and subsequently spent a year as a Computing Innovation Fellow at Harvard.

She is the recipient of Penn's 2009 Rubinoff Dissertation Award for Innovative Applications of Computer Technology, a National Science Foundation Career Award, and a Presidential Early Career Award for Scientists and Engineers. She is also involved in a variety of efforts to provide support for women in computer science. Most notably, she cofounded the Annual Workshop for Women in Machine Learning, which has been held each year since 2006.

Please join me in welcoming Jen. Jen?

(Applause.)

MS. VAUGHAN: Thanks. I am supposed to have some way of changing slides here, right? Is it this? Okay. Perfect.

Yes, so thank you so much for the introduction. I am really excited to be here today, talking to all of you. I am going to be talking about fairness and intelligibility in machine learning, which are topics that have come up a lot over the past couple of days. But I hope that this talk will provide a different and maybe a little bit broader perspective on these issues.

This may sound like a cliche by now, but we are living in the age of AI. Artificial intelligence is everywhere and that is why we are all gathered in this room today. We are at the point where AI systems can recognize individual people and images and translate speech on the fly.

The plot that I am showing on the right here has registration numbers for NIPS, the top academic conference on machine learning over the year. Last year, the conference sold out with 8,000 participants registered, and this year, we do not have final numbers yet, but the first round of registration sold out in less than 12 minutes. All of this means that there are some amazing opportunities and it is a really exciting time to work in machine learning.

But at the same time, we are seeing that these new opportunities also raise new challenges and these challenges tend to receive a lot of attention in the media usually when things go wrong. We are hearing more and more stories about algorithmic bias or algorithmic discrimination. And these high-profile stories have really highlighted how important it is to get AI right and to make sure that our AI systems do not discriminate or further disadvantage already disadvantaged groups.

Our CEO at Microsoft, Satya Nadella, takes seriously both the value of AI and also the importance of addressing all of these challenges that come with it. Satya published a great slate piece in 2019 that outlined his principles of artificial intelligence. These later evolved into the six principles laid out in The Future Computed, four core principles of fairness, reliability and safety, privacy and security, and inclusiveness underpinned by two foundational principles of transparency and accountability.

These principles are at the heart of the research that my colleagues and I do within the FATE Research Group at Microsoft. The four pillars of the

FATE Group are fairness, accountability, transparency, and ethics.

Of course, we are not the only group within Microsoft thinking about these issues. Microsoft's AETHER Committee was established in 2016 in order to discuss and recommend programs, policies, procedures and best practices on issues to do with AI, people, and society. The AETHER Committee now is working groups focused on seven topics, including fairness and bias and intelligibility and explainability. And Microsoft is part of larger efforts, such as the Partnership on AI, which is a multi-stakeholder organization with around 70 companies and other partners involved that is dedicated to studying and promoting best practices in AI.

So before I jump into fairness and intelligibility, let me just take a step back for a few minutes and say a few words about what AI and machine learning are. I know that you have been hearing a lot about these topics over the last few days, so I will keep this short. But I just want to make sure that we are all on the same page here.

There are many different ways of defining artificial intelligence. Nobody really agrees on one, but my view is that, roughly speaking, AI is about computers doing things that we would normally think of as intelligent. Now, in some cases, this means mimicking human intelligence, as is the case with computer vision or speech recognition, but in other cases, it might mean performing tasks that humans are not any good at at all, things like making quick decisions about which link a user visiting a website is going to click on.

Machine learning is a subfield of AI that is focused on systems that learn from data and experience as opposed to being explicitly programmed to behave in some way. Machine learning algorithms search for patterns in data and use these patterns to make predictions about the future. Examples include spam filtering, music recommendation systems, and targeted advertising.

Now, a neural network is one specific type of machine learning model. In the ‘80s and ‘90s, relatively few people were working on neural networks and they made up only a small part of the machine-learning landscape. These days this picture has changed a bit. Because of increases in computational power and the availability of huge amounts of data that enable neural networks to perform well, there is a lot more emphasis on them these days. This is often under the name “deep learning,” which I am sure all of you have heard. Deep learning is most often used for tasks like speech and vision where there is a lot of structure in the data.

Finally, I want to mention that machine learning can, loosely speaking, be broken down into three categories. First, in supervised learning, we use labeled data instances, such as medical scans labeled as containing a tumor or not containing a tumor, to learn a general rule mapping inputs to outputs, so mapping a new scan to either tumor or not tumor.

In unsupervised learning, the goal is to uncover hidden structure or patterns in the data, perhaps by clustering similar data points together. Finally, in reinforcement learning, the goal is to perform a task, such as driving a vehicle or playing a game in a dynamic environment and learning takes place over time through trial and error.

Now that I have said what machine learning is, I want to spend the next few minutes giving some intuition for why it is that machine learning can be biased or unfair. To do this, it is useful to consider the machine-learning pipeline. So a typical machine-learning pipeline looks something like what I have here. We start by defining the task or problem that we would like to solve.

We next construct the data set. Data set construction involves selecting a data source, acquiring the data that we want to use, preprocessing the data, and perhaps labeling the data.

Third, we define a model. Are we going to use a linear model or a decision tree or a neural network? What is our objective function? Each of these choices is associated with its own set of implicit assumptions.

Fourth, we train the model on the data. We next test and validate the model before deploying the model in the real world. Finally, we gather feedback about the performance of the model in practice and use that to improve the system. We will see that decisions made at every point in this pipeline can introduce bias into a system.

So let’s start with the definition of the task itself. What problem is it that you are trying to solve? In 2016, a research paper came out by a group in China who were training a face recognition system to predict who is going to commit a crime based on images of people’s faces. This is extremely concerning for a whole suite of reasons and could lead to substantial harms for the people who are misclassified. I would argue that this is just not a task that people should be trying to solve with machine learning. It is that simple.

But there are more subtle examples, too. Consider the problem of gender classification that Joy discussed earlier, so predicting someone's gender from a photo. On the surface, it might be less clear what the harms are here, but there are a couple of potential issues. For example, if a gender classifier only predicts binary gender, it is not going to work on people whose gender is nonbinary and likely will not work well for transgender people either. There are other issues as well. And in this case, it might be worth rethinking the task definition or, at the very least, talking it over with diverse stakeholders who can share their own opinion.

Let's move on to data set construction. So there are several different ways bias can arise at this stage of the pipeline. One is that the data source may reflect societal biases, right? The world has a lot of bias in it and our data sets reflect the world. This is what happened when Amazon tried to build a machine-learning-based recruiting tool. If your data source contains mostly male resumes and you have historically hired mostly men, your machine-learning system is going to pick up on this.

Linguistic bias is also a problem. Researchers at Princeton found that translating he as a nurse and she as a doctor into Turkish, a gender neutral language, and then back into English yields the stereotypical she is a nurse and he is a doctor. I want to emphasize here that these translations were not explicitly programmed, but were a result of the data that the translation systems were trained on. Loosely speaking, people are more likely to say she is a nurse than he is a nurse. So a translation system trained on speech generated by people is going to prefer that translation.

And to show that I am not just picking on Google here, I will point out here that the same thing happens with Microsoft’s Translator for exactly the same reason.

Bias can also arise if data is collected from a skewed source. As one example that we also saw in Joy’s talk, if we train a face recognition system on images that are mostly white men, then it will work well for white men, but maybe less well in other populations.

Yet another way that bias can arise in data set construction is through the use of human labelers. For example, there is a lot of research out there showing that human biases come into play when people are grading essays. But some states are still using automated essay grading systems that are trained on essays that are graded by humans, treating the human scores as if they are ground truth.

Okay. Let's move on to the model definition. So a model is a mathematical abstraction of some part of the world. For example, we might assume that the price of a house is a linear function of the number of bedrooms, the number of bathrooms, and the number of square feet with a little bit of random noise or variation. By its very nature, a model is simpler than the world, and so choosing a model necessarily means making some assumptions. What should be included in the model and what should not? How should we include the things that we do? And sometimes these assumptions privilege some groups over others.

Consider predictive policing. A predictive policing system may make predictions about where crimes will be committed based on historic arrest data. One implicit assumption here is that the number of arrests in an area is an accurate measure of the amount of crime. This does not take into account that policing practices can be racially biased or that there may be historic overpolicing in less affluent neighborhoods.

Moving on to the training process, this is the stage where we optimize the parameters of a model, so the weights, W1, W2, and W3, in the example that I showed earlier, based on your training data and whatever optimization criteria you have decided on.

There is some good news here. Once you have actually settled on your data set and your model and objective, the actual training algorithm is probably not going to introduce any additional bias. We see this as a common misconception. You generally do not have a biased algorithm, at least not a biased training algorithm. The problem usually really stems from the data or the model or the objective that you are trying to optimize or any of these other issues that I brought up earlier.

The testing phase of the pipeline is your opportunity to check for biases and potential harms and problems can come into play if you do not have the right metrics in mind here. There are a lot of different fairness metrics out there that are more or less appropriate in different contexts. And there is actually a great tutorial on this from last year's FAT* conference by Arvind Narayanan, who I think was supposed to be here today.

So to define these metrics, it is useful to start with the idea of a confusion matrix. Suppose an AI system is making a binary decision, such as whether to reject or hire a candidate. We can take any population that the algorithm is run on, say all the men, and divide them into four groups. The unqualified candidates who are rejected, these are true negatives; the unqualified candidates who are hired, these are the false positives; the qualified candidates who are rejected, the false negatives; and the qualified candidates who are hired, the true positives.

Most of the fairness metrics that people discuss can be defined in terms of the number of candidates who fall into each of these buckets. For example, we could ask what is the probability that a woman is qualified given that you choose to hire her? What about a man? Predictive parity requires that these probabilities, which can be calculated separately for each group, men and women, by looking at the number of true positives divided by the number of true positives plus the number of false positives should be almost equal for the two groups. You can think of this metric as assessing a form of calibration of the system.

Instead, we could choose to ask what is the probability of hiring a woman if she is unqualified? What about a man? False positive rate balance requires that these probabilities be just about equal for both groups. And, again, we can calculate these probabilities by looking at entries of this confusion matrix.

Similarly, we could ask what is the probability of rejecting a woman if she is qualified? What about a man? And false negative rate balance requires that these probabilities be almost equal.

Now, you may have heard about some of the controversy around the ProPublica investigation a couple of years ago which showed that COMPAS, a widely used recidivism prediction tool, was, according to some metrics, racially biased. In their audit of the COMPAS system, ProPublica considered some metrics, which basically boiled down to a false positive rate balance and a false negative rate balance, which I just showed you.

In other words, they asked whether COMPAS makes similar errors in terms of both type and quantity for black and white defendants. Indeed, they found that it does not. Because of this, they said the system was racially biased. In response, Northpointe, the company behind COMPAS, argued that COMPAS does satisfy predictive parity and so, therefore, it is fair. There was a lot of back and forth between people about this and about why the system did not satisfy all of these metrics.

However, the situation here is more complicated than it might appear on the surface. It turns out that it is actually mathematically impossible for a system to simultaneously satisfy these three properties at once, predictive parity, false positive rate balance and false negative rate balance. Any system that satisfies two out of three of these properties necessarily must fail to satisfy the third.

I will not go into more detail, but the takeaway here is that there are always going to be tradeoffs that we need to consider when thinking about fairness and we should choose our metrics carefully with these various tradeoffs in mind.

Moving on to deployment, the most common issue here is that the deployment population is somehow different from the population that you assumed that you would have. That is, your deployment population is different from the population from which your training and test data were generated, or the population you had in mind when you were defining your model.

So a common example here is collecting training data from people in one country and deploying a system in other parts of the world. There is actually some interesting research way back in 2011 that looked at available face recognition tools and showed that the location where the face recognition system was developed had a significant impact on its performance on different populations. Specifically, systems were substantially more accurate on faces from the same geographical region that they were developed in.

Finally, there is the feedback stage. And this is something that is discussed a lot in the context of predictive policing and hot spots. As we have already discussed, predictive policing systems operate under the assumption that more arrests in an area equals more crime. This can create a feedback loop or self-fulfilling prophesy. More officers are deployed to the neighborhoods where more crime is predicted. This leads to more arrests in these neighborhoods which leads to higher crime being predicted and even more officers being deployed there.

All right. So I have shown you how bias and unfairness can creep into AI systems. What can we do about it? Unfortunately, there is no silver bullet or one-size-fits-all solution to bias. But there are strategies that we can take to mitigate possible harms.

First and foremost, fairness needs to be prioritized at every stage of the machine-learning pipeline. We simply cannot hope to address the problem if it is not. Second, we must think critically about the implicit assumptions that we are making at each stage. How might the model that we choose introduce bias? What about the metrics that we use to train the system?

Third, we should pay special attention to potential biases in the data source and data preparation process since we have seen that so many of the biases in machine-learning systems are introduced through the data. This has been a point that I have heard several times this morning. The data is really what matters here.

Next, we should ensure that the population whose data is used for training, matches the population where the system will be deployed. We should involve diverse stakeholders in discussions at every stage of the pipeline and gather multiple perspectives. Diverse teams have an advantage here -- and this is something that we should consider in hiring as well.

And, finally, we should acknowledge our mistakes and learn from them. When it comes to bias and fairness, perfection is not possible. So we need to be willing to learn when we make a mistake and do better next time.

For the last few minutes of my talk, I want to move on from fairness and talk about transparency and its relationship to intelligibility. Within policy circles, it is common for people to use the term “transparency” in two somewhat different ways. First, it represents the idea that people should be able to understand and monitor how AI systems work.

Second, it is used to suggest that those who deploy AI systems should be honest and forthcoming about how and when AI is being used. In machine-learning circles, the first idea is usually referred to as intelligibility or interpretability. One important thing to realize here is that literal transparency, that is, providing information about model internals, can actually work against it.

In particular, one way of being transparent would be to expose the source code used to train a machine-learning model. However, the source code really would not tell us much about why an AI system behaves the way it does, especially if we do not have access to the training data or modeled parameters. If I just tell you that my source code is optimizing a linear model, this does not give me a lot of insight into how the model works.

Another form of transparency might involve exposing the internals of a model, such as the learned parameters or weights. However, several research studies, including a recent study that I ran with colleagues at Microsoft, show that at least in some situations exposing model internals can overwhelm people with information and actually make them less likely to notice instances where a model is making a mistake.

In our study, we found that this information overload effect could happen even with the simple linear model with only two features in it. I would argue that in most cases it is intelligibility and not literal transparency that we want. To give you a few examples of why we might need intelligibility in an AI system, suppose we have an applicant who wants to know why she was denied a loan. In this case, intelligibility helps with accountability, allowing consumers to understand why a system is treating them in a certain way.

Suppose instead we have a model that is deployed in a school that predicts that a student is likely to drop out. Knowing which factors are relevant for this prediction could help this teacher decide whether to believe the prediction and how to best intervene. In this example, intelligibility can lead to greater trust in a system’s predictions.

Third, suppose we have a model that matches candidates to jobs. By understanding characteristics of the training data, an employer may see that female candidates are underrepresented, leading to some potential bias. This is an example about the assessment of bias and relates back to the first half of my talk.

I want to point out that, in this example, intelligibility is coming from understanding the training data rather than understanding the machine-learning model or the full AI system. As with fairness, we can think about intelligibility in different parts of the machine-learning pipeline like this.

Finally, suppose a data scientist sees an unexpected prediction from a model that she has trained. Knowing why this prediction was made could help her debug the model. In this example, intelligibility leads to greater robustness in the system.

Now that I have argued for intelligibility, let me mention a few different approaches that have been proposed in the literature. One approach is to design and deploy models that are intuitively simple. Simple might mean something like a small decision tree or sparse linear model. For example, my colleague and collaborator, Dan Goldstein, has some nice recent work on simple point systems that assigns scores based on a small number of features, resulting in models that can be easily understood and even memorized.

Of course, as I hinted at several slides back, simplicity does not always lead to intelligibility. And in some cases, simplicity is just not possible; for example, when designing an AI system for a highly complex task like a web search or when the goal to provide intelligibility for an existing complex system rather than starting over from scratch.

Because of this, a second common approach is to design simple post hoc explanations for potentially complex models or systems. One thread of research in this discussion -- in this direction looks at how to explain individual predictions by learning simple local approximations of a model around a point. Another focuses on learning simple approximations of a full model. These approaches can be useful, though there is a danger that simple explanations they provide may not be perfectly capturing what the true complex system is doing and may, therefore, be misleading if we take them too seriously.

Given the importance of the data used to train a model, we may also be interested in providing intelligibility around the data source. In the world of electronics, every component, ranging from the simplest resistor all the way up to the complex microprocessor has a corresponding data sheet that details the operating characteristics, test results, recommended usage, and other information about that component.

Inspired by data sheets for electronic components, some colleagues of mine and I put forth a proposal that data sets, models and APIs be accompanied by a data sheet that documents the creation, intended uses, limitations, and so on.

To help teams construct data sheets for their own data sets, we put together a set of questions that cover different types of information that we think belong in a data sheet. These questions are divided into categories listed here, motivation, composition, the data collection process, preprocessing, distribution, maintenance, legal concerns, and ethical considerations. Each category has about five to ten questions.

There are several possible use cases for data sheets. First, they could be posted alongside public data sets to inform potential users about the makeup and origin of the data. Second, they could be included with a company's internal use data sets to provide information to future internal users. This is something that we are starting to pilot on a small scale within Microsoft in the near future.

Just as with fairness, none of these approaches are a silver bullet that will solve every need. The right approach to intelligibility is always going to depend on the context. The approach that works best for a CEO making a strategic decision is likely to be very different from the approach that works best for a regulator who wants to understand why an individual was denied a loan.

There are, therefore, a number of questions that people should ask when trying to decide on method of achieving intelligibility. We have already touched on a few of these. Why is the explanation needed or what is the goal of the explanation? Do we need to explain a single prediction or a whole system? What is it that we want to understand here or who is it that we want to understand the system?

But there are a whole host of other questions that go into determining which solution is right for a particular scenario and understanding the space is an active area of research that a lot of people are working on, including myself.

So in my last minute, I would like to conclude by reviewing a few key points that I hope you will remember after you walk away from this talk. First, as I have tried to stress throughout this talk, there is no one-size-fits-all solution to fairness, transparency, or intelligibility.

Second, fairness, transparency, and intelligibility cannot be treated as afterthoughts. These principles must be considered at every stage of the machine-learning pipeline, right from the very beginning.

Third, there are countless opportunities for technology to play a role in the solution. I mentioned a variety of intelligibility methods that we are starting to explore and there is lot of active research going on in fairness, too, around algorithmic solutions. We just need to use the technology with care and also understand its limitations.

Fourth, it is important to involve diverse stakeholders and gather multiple perspectives. These diverse stakeholders are likely to notice our own blind spots that we might miss.

And, finally, since there is no perfect solution to fairness or bias or intelligibility, we are all going to make mistakes in this process. The way forward is to acknowledge these mistakes and learn from them so that we can build better AI systems that benefit all. Thanks.

(Applause.)

MS. CONNELLY: Thank you very much, Jen. We will take a minute and assemble our panelists for the last panel, it is wrap-up panel. If the panelists could come up to the stage, we will get started in a minute.

W
RAPPING UP AND LOOKING AHEAD: ROUNDTABLE DISCUSSION OF KEY LEGAL AND REGULATORY QUESTIONS IN THE FIELD

MS. CONNELLY: Good afternoon, everyone. I am Ellen Connelly. Some of you saw me earlier today. I am an Attorney Advisor in the Office of Policy Planning at the FTC. My co-moderator today is Ben Rossen. He is an Attorney in the Bureau of Consumer Protection's Division of Privacy and Identity Protection.

We want to welcome you to our final panel for this series of hearings about algorithms, AI, and predictive analytics. That is our wrap-up panel and we are hoping to have a good conversation about some of the ideas that have been discussed over the past few days as well as to look a bit ahead and highlight some things that policymakers and enforcers might want to be thinking about going forward.

We have a very impressive group of panelists here with us today. There are detailed bios online, but just very briefly, we have Pam Dixon, who is the cofounder and -- sorry, the Founder and Executive Director of the World Privacy Forum, a public interest research group focused on consumer data privacy issues. She was also the lead author of the Scoring of America: A Substantive Report on Predictive

Analytics and Privacy Issues Associated with Consumer Scoring.

Next, we have Justin Brookman, who serves as the Director of Consumer Privacy and Technology Policy for Consumers Union. He works there to shape the digital marketplace in a way that empowers consumers and prioritizes consumer data privacy and security. And he was previously Policy Director at the FTC's Office of Technology Research and Investigation.

After Justin, we have Salil Mehra, who is the Charles Klein Professor of Law and Government at Temple University's James E. Beasley's School of Law where he teaches courses in antitrust, contracts, and law and economics.

Next, we have Joshua New, who is a Senior Policy Analyst at the Center for Data Innovation, a nonprofit, nonpartisan public policy think tank affiliated with the Information Technology and Innovation Foundation. Josh leads the Center's work on issues related to AI, the Internet of Things, and open data.

And, finally, we have Nicole Turner-Lee, who is a fellow at the Brookings Institution's Center for Technology Innovation. She researches public policy designed to enable equitable access to technology, as well as global and domestic broadband deployment, regulatory and governance issues. She is also a visiting scholar at the Center for Gender Equity in Science and Technology at Arizona State University, and she is an appointee with the FCC's Advisory Committee on Diversity and Digital Empowerment.

Arvind Narayanan was supposed to join us, but unfortunately he was unexpectedly unable to come. We will hope to get his views on these important issues at another time.

So just a few procedural points, as I said, we are not having presentations, we are just going to launch into a moderated conversation. As we did with all of the previous panel discussions, we will be collecting comments and questions from the audience. So please look for conference staff should you have a question, they will collect the comment cards and bring them to us.

With that, I would like to get the conversation started by asking a somewhat open-ended question of our panelists. I know that many of you have been able to attend, perhaps not all of the prior sessions, but at least some of the discussions over the past day and a half, and I would like to just open the discussion by going down the line and asking, what are your views on particular items that have been discussed in prior sessions which might merit more elaboration or which might merit additional highlighting for policymakers or, alternatively, are there things that have been missed in the prior conversations?

We will start with you, Pam.

MS. DIXON: Okay, thank you. And thanks to the FTC for holding this important conversation.

So I am just going to launch in quickly. I did not see the sessions yesterday. I was flying home from the OECD meeting in Paris on the development of the AI global recommendations. I am a delegate on that group and I am going to be incorporating some things from that here.

Let me launch. The state of machine learning and AI, it is really important as we think about these policy issues to understand that there is a really bright line. AI is moving in two different directions toward a more opaque direction with the machine-learning side and toward, very clear, the older statistical models. Those two may well require different approaches and it is a good idea to disambiguate those approaches.

I want to specifically talk about deep

convolutional neural networks and some very significant recent advances in that area. We heard a presentation on facial recognition algorithms. They are very important to consider. So let me give you an example here -- and I think it is just really important to draw this out. In the past year, there have been meaningful advancements in facial recognition analytics. The NIST tests, the most recent facial vendor recognition tests, are completed. I have seen the results and the advances in accuracy are remarkable. They are now at 99.8 percent and the tests were really robust across a lot of meaningful parameters.

There is also something called sublinear search, which means that really for the first time, we have the possibility of very accurate biometrics that can also be searched very rapidly. So it is really at the first capacity for accurate mass surveillance.

So a lot of times when we hear examples in fora like these, it is a lot of self-driving cars. But we need to remember that there are other examples. And what I really like to think of is, is this a voluntary use of AI or is this a mandatory use of AI, and we really need to think about those things. And I have not really heard that discussed here today.

I will give you a great example of voluntary/nonvoluntary. Self-driving cars are right now highly voluntary, right? What about scores? Consumer lifetime value scores, something that we are given by businesses, that is not voluntary. What about if you live in India and you are trying to just simply pay your taxes, use of biometrics in that case will be nonvoluntary. It will be mandatory. We need to think about that, what is the voluntary nature or nonvoluntary nature.

In terms of the dispersion of AI and machine learning, I really have not heard about the global dispersion of that today. I hope that there has been discussion of it in prior days. I would just bring forward that AI and machine learning is advancing in different rates, in different locales. But it is pretty much advancing everywhere. And under different jurisdictional regimes -- so in India, you have the massive case study of the Aadhaar biometric ecosystem. In China, you have social scoring. In the United States, we have all manner of consumer scores, including the credit score.

Then in terms of framework, someone today mentioned GDPR, which is great. I would also say that we need to consider in our analysis credit scoring frameworks because credit scores are a form of AI. We have to consider soft law frameworks -- the OECD framework is in process and it is soft law in the countries that adopt -- and then, of course, the self-regulatory frameworks. The self-regulatory frameworks and the soft law frameworks and the law frameworks are all quite different that are in place.

And we are seeing huge differences coming in from Asia and from the developed nation and from the global south. What I can say is that so far Japan wins the prize because they have a very advanced look at what the framework looks like and they have incorporated the best of the west and of the east. They have published -- and there is an English version. They have published ten principles.

Something I am extremely concerned about, and I hope this was mentioned yesterday, but it is incredibly important to understand something about gender and AI. So all of us in this room right now here today are tremendously privileged. We live in a country where when statistics are gathered by the U.S. Census Bureau they are gender disaggregated. This is actually a privilege. It is not so in all parts of the world, particularly in the global south.

And, unfortunately if there is, for example, murder rates and only the murder rates are collected for all genders, it can create a lot of problems over time in telling the story of that particular jurisdiction or that particular culture. And when analyses is done and you do not have gender disaggregated statistics it can be a huge, huge, long-term problem. This exact same issue applies to poverty statistics. And poverty statistics are somewhat controversial, but again they are not adequately collected in all jurisdictions.

In order to really think about AI and machine learning, we have to think globally and we have to think about these fundamental disparities that exist in other jurisdictions.

And then without taking any more time, inputs data has been mentioned, I want to highlight that. Fairness of purpose has to be mentioned. I am so glad that people have been mentioning this. How to ensure uses, back-end uses is something that is going to be very careful and redress has been mentioned. But something that has not been discussed here today is what I call governance.

So after we have all the principles in place, how do we, on a day-to-day basis, govern AI and machine-learning system. So we have to have a cognitive context that is going to fit actual reality. There has to be governance that actually works for these systems.

So just drawing from Elinor Ostrom's principles of governing shared pooled resources, I am just going to draw out three very important things to think about, which is all stakeholders in these processes need to have an appropriate voice. Whatever process is in place needs to be ongoing and iterative. In other words, you cannot make a rule for AI and then it is static for a year, that will never work. Then there needs to be collaborative governance frameworks, not command and control governance frameworks. If there is, for example, a self-regulatory model and it is a command and control where it is disbursed but it is not collaborative, it is not going to work in the long run. So these are just some initial comments.

MS. CONNELLY: Thank you.

Justin?

MR. BROOKMAN: Thank you for inviting me. I am going to pick up on a couple of the themes I heard. I was not able to watch this morning, but I was here yesterday. So I am going to talk about a couple consumer protection themes and then tie it to some of the legal policy issues.

So, first, I think there is broad agreement that there is need for more, I do not want to say transparency because a previous speaker said that and that is a wrong word, but more information available, more accountability out there. And I think it is important to think about what the role that policy can play there is. I think that we absolutely -- there probably should be some more mandates to make information available and, again, for different stakeholders, different sorts of things might be relevant.

In addition to information, maybe there should be some obligation to make these systems testable by outside people, make APIs available for folks like the FTC, folks like Consumer Reports, I think there should be legal obligations to test themselves and to make sure that they are working as intended. But there needs to be more external accountability, too.

I think it is hard to get there with existing law. I think it is hard to make argument under Section 5. I think we may need be to explore some other things. I think one thing Section 5 could be useful for is efforts to defeat transparency. So one example that came up yesterday was Uber’s use of the Greyball program, which is when Uber was trying to get a sense of when someone like a regulator or a tester was trying to evaluate their systems, they would change the protocols or how it operated in order to defeat that. Is that deceptive? Can you make an argument that that violates Section 5?

The deception policy statement today talks about deceiving consumers. But with the advent of AI, I think we may need to think about maybe broadening that somewhat. So, one, use of AI to deceive testers or potentially regulators in that example, I think, maybe should be expanded. Alternatively, an attacker trying to confuse AI, I mean, should that be considered a deceptive practice? Say my operating system is using AI to protect me from someone, should that be considered deceptive even though it is not deceiving the consumer?

I think we should probably expand the policy statement to address that. The FTC has gotten close to that in a couple of areas like the Volkswagen case when Volkswagen was trying to figure out when a regulator was revving the engine and not maybe using AI, but was using some sort of algorithm to change the processing. But there the behavior itself was not, per se, deceptive; it was like the false statements to regulators.

Similarly, Google, there was a case against Google for dropping cookies on Safari when there should not have been. You can make the argument that Google was tricking Safari by doing that, instead -- and, actually, state AGs made that argument. The FTC relied kind of more narrowly on FAQs on the Google page to bring a case. But I think expanding our deception concept to address AI I think is important.

The other theme that I heard a lot yesterday and I think is actually a little bit harder is how to forestall adverse for consumers' uses of AI. So one example that came up a few times is price discrimination and price discrimination is obviously not always bad. But in some cases when there is lot of imbalanced information and perhaps there is corporate concentration, then, yeah, I think it kind of is. I think this was a theme a little bit yesterday, but also when Professor Stiglitz talked to the FTC at one of the first couple of hearings, he mentioned this is his like primary concern with AI.

Is that harmful? That was not listed in the FTC's harms roundtable, but it does -- it is bad for consumer welfare. So do we need a more expansive idea of harm to get to issues like that?

And then, finally, you know, manipulation. Obviously, commercial human interactions, are they a little bit manipulative, are they trying to get you to do something, to buy something? But with AI, you know, they can iterate through a thousand things or pick up on signals to maybe make it like super-manipulative and does it ever cross a line there? I am not sure.

An example that Ryan Calo brought up yesterday was using AI to figure out if someone is like depressed in order to kind of get them to binge purchase. Is that so exploitative that that is going to be prohibited? Addiction, like these devices are designed to get us pressing buttons over and over again. Can that kind of harm be included in a -- again, AI makes it a lot more better, a lot more efficient at addicting us. Should that be included as well? Should there be broader tech mandates around ethics, which is something that a lot of folks have talked about, too. I think privacy legislation can address some of that, but not all of it. So I think there are important questions to consider.

Thanks.

MS. CONNELLY: Salil?

MR. MEHRA: Thank you for having me here

today.

So the recurring theme I would like to address from especially today's presentations is to think about the implications of these technologies from the sort of historical view. This has been a theme, this sort of focus on kind of march towards AI, right? Starting from sort of ex ante trying to program rules to, you know, what we might think of as predictive analytics, which is essentially massively applied data to what we see developing, which is essentially AI or deep learning.

You can think about it from the examples of language, right? Thinking about predictive analytics or data analysis. Right now, your digital assistant, whether it is Siri or Alexa or something is comparing what it hears to a large data archive of audio. It is essentially brute force crunching of data matching the sound files. But technologists are working on sort of deep learning technologies that are closer to something like semantically understanding language.

So if we think about this from the competition perspective of pricing and markets, the sort of programming of a generation ago, setting forth pricing rules ex ante for all occasions, that is really hard to do, right. The world is very complex place. But as you move towards predictive analytics, this massively applied statistical analysis, it draws on some of the technologies that came out of fintech where there is a lot of observable pricing, the crunching of a lot of data, much more open data, basically hugely applied statistics, maybe some human machine collaboration.

So we have seen -- and there is already literature on this -- that this would be relevant to things like tacit collusion, right? The possibility that it is increasingly possible to anticipate your competitor's pricing and moves. This would be relevant to explicit collusion. We often say competition is a click away, but if we think about cases like the posters or wall decor case, right, we get the idea that maybe price fixing is also a click away, which has implications for the sort of norms that ordinary people or ordinary firms bring to the table when they think about antitrust and antitrust violations.

We might be concerned, in particular, if you think about the history with copyright and unauthorized consumption of copyrighted goods, you might be worried about that kind of breakdown of norms against, for example, price fixing.

I also think there is sort of a longer term, sort of more future-looking implication here with AI and deep learning. So this is computers that have the ability to draw and software that has ability to draw in patterns and actually shape their own rules of engagement with the world. That is one way to think of it.

In conjunction with this, we have seen the greater reliance on what we might think of as sort of captive data. So when you think about -- and we saw this in the last panel before lunch -- when you think about digital assistants, when you think about the spread of these technologies to cars, you are not just sort of learning a language or a dialect or series of words, you are actually focusing on an individual's own particular patterns, for example, patterns of speech in a closed environment, their home or their car, an idiolect, if you will. This is not necessarily observable -- this data that is gathered, it is not observable to your competitors in the way that, for example, the internet was or the web was when Google was launching its search product, right?

So where data on the internet, for example, seemed open and accessible, this type of data collection may be turning more proprietary. So I would like to leave you with sort of a bigger question about competition laws, which is -- or a series of questions, which is how are we going to fit this into our current competition law, structure, right?

You could see this as a barrier to entry, but I think it will be difficult to deal with as a barrier to entry, this type of specific individualized idiosyncratic data collection. You might wonder about the degree to which we should empower as a remedy or as a solution, empower user control over data. When people think of the GDPR and the idea of it seems to enshrine this principle of owning your data, you know, should there be some sort of fostering of user choice to multi homes so that you do not see as much captive individualized data.

But these questions I think are sort of the tip of the iceberg and the sort of things that sort of start us rolling.

MS. CONNELLY: Thank you.

Josh?

MR. NEW: So again, thanks for having me. I think this has been a great discussion from what I have been able to see so far.

I want to echo what Pam touched on about the need for governance in this space. I think this room is probably much more in the know than most people having these kind of conversations, but AI and its impact on society has become a pop culture issue and I think that is very beneficial in certain ways, but also very frustrating when you are trying to have nuanced policy discussions about how you can actually approach governance of these technologies, because most popular ideas we have seen so far about how to address a lot of the harms that we talked about today, like broad mandates for algorithmic transparency or algorithmic explainability or the creation of an AI regulatory authority, you know, an AI regulator or a robotic commission that we have heard similar proposals for. You know, Elon Musk had said something like that.

People who are technically savvy, they understand AI's value, but proposing some really short-sighted solutions. I mean, the presentation we just saw earlier, Jennifer -- and I think she just walked out, but I wanted to thank her -- that was fantastic. That demonstrated that these are really complex technical challenges. How we approach governance needs to be equally nuanced. There has been so little discussion about how you actually focus on implementing these approaches to governance.

We see companies do this in like their statement around AI use and ethics. We see policymakers do this. Theresa May made a speech at the beginning of this year that was particularly egregious that basically said, you know, AI is valuable, but we need to make sure it is safe and ethical, and then the conversation ended there. And, like, of course that happens. But that is vapid. That is a truism. No one is going to disagree, but that does not actually help. That is not a model for governance.

So, of course, I am biased here. We published a paper early this year titled, "How Policymakers Can Foster Algorithmic Accountability," that takes a stab at making an actual implementable model for regulators to approach these issues. I am definitely open to debating those ideas. It might not be right; I think it is. But those conversations are -- have not been happening so far. I think this event today, in going forward, we are going to start to seeing them more often.

But I really just want to reiterate the need for kind of issuing -- devoting all this political capital just to saying, oh, we need to do something, then actually focus on doing something because that just has not been happening yet. Well, other countries are being more proactive about it. The EU had GDPR, and I think that is actually really detrimental to AI in a lot of ways, but they are recognizing the need for action here. Don’t interpret that as praise for GDPR. My boss would be very mad to hear me say that.

But I would really hope the FTC and policymakers, in general, work on this quite a bit going forward.

MS. CONNELLY: Thank you.

Nicol?

MS. TURNER-LEE: Thank you. So last, but certainly not least, I will add a little bit more value to this conversation, particularly focusing on an issue, an area that I am most concerned with which is the application of these technologies to historically disadvantaged populations and vulnerable communities.

So first and foremost, I think generally what I gleaned from the presentations that have taken place over the last couple of days is that we have some definitional concerns when it comes to what is AI. And those definitional concerns sort of create some problems when it comes to what is the appropriate regulatory structure and policy structure for it, as well as the use cases that will be more ethical and appropriate for the application of AI.

And in the body of research that I do at Brookings, part of my concern is, if we are still sort of debating these definitional concerns and many of the use cases will actually further disadvantage groups that are already on the margins of society, then how do we begin to sort of make sure we build in equity and fairness and inclusivity from the onset.

I would say from what I have heard from the conversations there are probably three critical areas that are related to this. The first -- and I am looking at Joy, who I am a fan girl of her work, you know, clearly starting with the right training data set is one that is particularly of interest to myself because that inclusivity of the data set will actually help us to come out with outcomes that are much more fair and accurate when it comes to representation.

And I would even argue -- and this is something that we will have a paper coming out at Brookings on algorithmic bias detection and mitigation with the University of Michigan and the Better Business Bureau Institute, that we have to look at this diversity and design structure that not only pushes for when we put these products to market, do we have the proper coloring of those folks that are going to be the subject or the targeted focus of what those algorithms are? For example, that goes to facial analysis software, that goes to search query software.

Any application that has to be representative in diversity and design starts with that as a presumption rather than an aftereffect of the application, the procedure, and potentially more diversity in those work forces would probably help as well, and ensuring that you have a check and balance that gives some context to whether or not that algorithm or AI application will oppress versus, you know, advance the needs of particular populations.

I would say in this nascent technology as well, it is very important for us to understand and perhaps do -- and this is something I gleaned from the hearings as well -- an exploration of the statutory guardrails that are in place. There are simply things that we cannot do in the U.S. when it comes to credit, housing, and other civil rights laws. And we have not had, I think, a really thorough conversation and exploratory conversation on whether or not those statutory guardrails actually apply to this space and in what way and in at what point and what type of retribution do consumers have when these things happen to them.

I think that conversation, particularly we look at the human consequence of credit worthiness, applications for credit worthiness, bail and sentencing, housing and surveillance, it is particularly important that we actually have that conversation up-front. One of the things that we are going to be proposing in our paper is this framework of a bias impact statement and template. You know, are companies in a self-regulatory mode or operators of algorithms doing good scrubbing and house cleaning of the purpose of that algorithm and the potential unintended consequences on protected classes, and if not on protected classes, on other vulnerable populations where that training data may eventually end up further oppressing or discriminating against those groups.

Those are very dangerous alleys to go through because they generate disparate impact, disparate treatment and disparate error, and sometimes those are irreversible when it comes to historically disadvantaged and vulnerable populations. They cannot come out of it. In my research on digital divide, when we look at populations of color, the most valuable asset that they have if you look at the settled research on wealth, is their Social Security. We already know what happens when people are foreclosed on their personal identity. As we look at these emerging technologies, the question becomes the degree to which they will foreclose on other opportunities that limit people's access to social and economic mobility.

I would say on that piece, one thing that also struck me, I want to say in the hearing was a statement by one of the panelists that as AI gets more precise in its ability to discriminate; it gets more precise in its ability to discriminate. To me, that is a problem. As a sociologist what that says is that we also need more interdisciplinary connections between technologists and social scientists to sort of match the settled research on what happens when you actually look at online proxies of zip code and you match that with employment applications.

How does it look when you look at chronic disease and how it affects certain populations and you create scores or AI applications that further keep people within that box that may actually limit or restrict them from getting healthcare? So I think having more of those cross-functional dialogues will be something that is particularly important at this time as we see -- and it is so most relevant that the FTC is doing this -- more of these applications go into civil society and touch upon public interest.

I would end with this, that clearly -- and having just returned from China, who has proposed that they will be the number one in AI -- part of this conversation, too, at Brookings, we are concerned about AI from variety of verticals, whether it is autonomous weapons, whether it is the commercial applications or public interest applications. But common to all of these are conversations around privacy, conversations around ethics, conversations around innovation and consumer protection.

What I think is missing, if I may add to the conversation when we look at regulatory and legal frameworks, is how do we create this Venn diagram that pulls all of that together? Across all of these use cases, are there principles that we should be standardizing that apply to the ethical use of an autonomous weapon to the ethical use of an application that is going to predict or impact one’s ability to get into a school of their choice for higher education?

So I think, going forward, that would be a very interesting exercise in terms of again more multi-stakeholder engagement, more interdisciplinary cooperation, more global and domestic governance structures to really think about where are their commonalities when we look at AI applications and emergent technologies where we want to pay attention. And how does that diagram look where there may be some deviance from that model, but there are key structures that apply to all of these use cases that are important for the public good of this launch of AI.

MR. ROSSEN: Well, thank you to all of you. There has been a lot to unpack already. I want to follow up on an issue that a couple of you mentioned, which is about how other jurisdictions are approaching some of the issues of balancing policy goals with respect to these technologies while furthering innovation. I know a couple of you mentioned GDPR already, as well as some other jurisdictions. We have had six months or so now of the GDPR in effect. Maybe that is enough to start measuring what is working and what is not or what the U.S. might learn from some of those jurisdictions or might want to avoid.

So, Josh, I will start on your end of the table this time and then maybe Pam and we will see if others want to weigh in.

MR. NEW: Sure. I think this would be a good opportunity to do kind of a study in contrast versus what the European Union is doing -- a region that very, very highly prioritizes consumer protection, in our view, at the expense of innovation in many cases versus what China is doing, which is very, very invested in advancing AI with pretty much no regard to consumer protection.

So we put out a report early this year about analyzing the impact of GDPR on AI development and adoption. We found some pretty concerning things because the EU has stated that it wants to be competitive in AI; it wants to foster advanced technology industries, use AI in areas like manufacturing and healthcare to capture all the benefits, which is all well and good, but they have really kind of shot themselves in the foot in certain areas.

There are two provisions, in particular, that I want to mention. There is the right to explanation of significant decisions or a right to meaningful information. And then there is the right to erasure. So the first one -- and the wording is a little bit vague and I think that was by design because they were waiting for the court system to figure out enforcement and implementation issues when they arose. But it basically says that if an algorithmic decision is used to make -- or an algorithmic system is used to make a significant decision about a person, they have a right to meaningful information about that system, which sounds good and the concept of, you know, right to explanation is not uncommon in law, it is very common in consumer finance. If you are denied a credit card, you are owed an explanation why whether or not an algorithm is involved.

But the GDPR's wording on this is so vague that it does not really -- it very likely applies that standard of a right to explanation to all decisions whether or not -- to all algorithm decisions that could be significant, but not to the same decisions when a human makes them. And that is a regulatory burden. If a company is concerned about that regulatory burden, they will just use humans to make those kind of significant decisions that do not have preexisting statute for explainability, which comes at the direct expense of productivity and does not actually protect consumers any more. Companies will just simply avoid doing that because that is the pragmatic approach to doing this.

And if you think that all of those decisions could cause harm, you should pass a law that says, all these decisions need to be explainable whether or not an algorithmic system is involved. It is kind of really short-sighted to only target a decision when an algorithm makes it, even though that does not make it inherently more dangerous or risky.

The second is the right to erasure, the right to remove your personal data from a database that could eventually be used in algorithmic systems. When you are training a machine-learning system on massive amounts of data and then you take away a portion of that data that was used in that training data set, there are lot of concerns that could significantly impact the performance of that algorithmic system, potentially making it unsafe or unusable or less viable a product, cause consumer harms in other areas. It is not even clear that that is necessarily technically possible in all situations. But that is a pretty broad mandate that does not actually provide immediate benefit to consumers.

The reason that these are problematic, which tie into our argument about why we should focus on accountability on outcomes rather than processes, is that explanation or erasure are not ins and of themselves, they are means to consumer protection. But they focus on process rather than outcomes and I think that is a really flawed approach that Europe has kind of adopted in many areas.

So, in stark contrast to that, this will be much quicker because there is a lot less to talk about, China just simply does not prioritize consumer protection like Europe, like Canada, like the United States, like many countries do that are also competing in AI. They have access to massive amounts of personal data about their citizens. There are not really any concerns about how that data is used in potentially very invasive ways. That could be because, you know, dissent is not really permissible in the same way in these countries -- in the United States and other countries.

But they are racing, as Nicol mentioned, to be the world leader in AI. They are putting all their chips on AI. By 2030, they want to be the global innovation hub I think is the way they describe it.

So if all this concern about consumer protection is good, these are good discussions to be having. But if we are not also having conversations about how to support AI, how we can accelerate its growth and adoption so we can actually compete for global market share with Chinese-developed AI where they do not embed those kind of values in their systems, then all of these conversations are going to be moot.

If we are not investing in accelerating AI that abides by values that we care about, then it simply will not exist in the world more broadly once China beats us to the punch. And that is something that Europe really missed the boat with, and as the U.S. kind of figures this out, I hope we kind of shoot the middle effectively to address that problem.

MR. ROSSEN: Pam?

MS. DIXON: All right, thank you. So, I am going to draw examples that are different. Thank you for covering that. I am not going to repeat.

I want to talk about two examples. I am going to talk about India and I am going to talk about the U.S. So I am going to make the examples as close as possible. So I think most of you who know me know that I spent a year in India doing research on the Aadhaar biometric ID system. I tracked it from 2010, from the very first person who was enrolled in the biometric ID when it was completely voluntary to 2016 when over a billion people had the ID and it had been made retroactively mandatory.

So what I want to say about India is basically they had the installation of biometric technology AI, very sophisticated AI technology, before there was any policy put in place and before there was any governance put in place. This went on for years. It was made mandatory. Unfortunately, people literally died as a result of the failure to authenticate. For example, in the State in Jharkhand in India, there was approximately a 50 percent failure to authenticate rate. That means that 50 percent of the people could not get their food when they lived below the poverty line. They could not get it because their biometric ID did not work.

So this is a big problem. Additionally, women and children who were trying to flee and be rescued from human trafficking were denied healthcare. That is in contravention to UN policy and to EU convention where victims of human trafficking are not supposed to have to become identified to folks who will require them to be a witness for the prosecution. So big, big problems.

Now, what happened in India that solved these problems happened very recently with the Supreme Court ruling in India called the Puttaswamy Aadhaar, most of the mandatory uses of the ADAR were overturned, and in what is now a very famous dissent, there was the do no harm principle that was discussed in the ruling. And this do no harm principle talked about if you are going to use these technologies, you must ensure that they create a public good and do no harm. This was very, very new in India, and we will see where it goes from there.

Now, in the U.S., we have a much different situation. We have so many more laws. We do not have a biometric being installed in the country where there is technology before policy. But we do have semi-mandatory system which is the U.S. biometric entry and exit. We are going to have biometric entry and exit. It is something that is coming, it is already being pilot tested.

So here is my question for the U.S. What is the specific governance for that system? Is it going to be command and control where we do not have a choice? These are very, very sophisticated AI systems. So you see certain parallels and certain differences. But in all of them we have to ask ourselves, is this a mandatory system or is this a voluntary system or a mix of the two? And how we determine policy is going to make a really big difference on whether that happens.

In terms of another nonvoluntary thing that I want to mention -- and this is really across jurisdictions. I have not found a difference. I found it in China, I found it in Europe, I found it in the U.S., and I found it in almost all global south jurisdictions, which is an issue of scoring using various -- it is typically machine learning.

When individuals are scored or classified or given an output of machine learning, the number matters, because as humans we just love to score. It is a shorthand and we are ultimately going to use something that is a shorthand, more than a long table that we have to actually analyze, this is just human nature. What are we going to do with this? What are the policies that we have about things that we do not know about?

So the GDPR attempts to address this, but I have not seen specific governance that would actually solve the problem. In the United States, we have the Fair Credit Reporting Act, which effectively regulates credit scores that are derived from consumer credit bureau reports. But when you have credit scores that are derived from other data points and used for the same -- well, almost the same purposes, they are not regulated.

So what do we do about this issue? It is so nuanced, it is so subtle, but it is already here, it is already in use, we do not have lot of choices here. So I just leave you with these thoughts. I think that we have a lot of work to do.

MS. CONNELLY: Justin and then Salil.

MR. BROOKMAN: Yeah, I just have one minute. I just wanted to respond briefly to Joshua’s point. One, on GDPR, we do not really know what it does, right. GDPR is a very high level, vague document. On the privacy side, the primary effect has been a bunch of companies emailing you their privacy policy and then putting really obnoxious consent flows up there. I am not entirely sure how companies are responding to the profiling elements. So I think there is a lot of vagueness there and I think we are not entirely sure how it will play out in practice.

On the outcome side, I hear what you are saying, but I think that trusting entirely to outcomes means you trust companies to always get it right. And it is really hard to test here. It is hard for the FTC to test, it is hard for consumer reports to test. It is certainly hard for any ordinary consumer to test. I can certainly see a consumer rationally saying, you know what, I do not really trust you with my data, I understand that you have a privacy program in place and theoretically accountability, I am just going to go ahead and take my data back. I hear what you are saying, that there is a cost there, though, I mean, all data is messy. So I am not entirely convinced it will be that deleterious to the learning algorithms. But certainly giving consumers some degree of agency or autonomy over their information does provide a meaningful check on company's power over them.

MS. CONNELLY: Salil?

MR. MEHRA: This is sort of a brief comparative point that relates to the FTC's competition mission and also sort of a big picture view on a need for competition law. Joshua brought up the issue of AI development in China. Some of you may have seen the recent book by Kai-Fu Lee that talks about the development of AI in China and there is sort of an argument about thinking about algorithms as the -- and data as sort of the two big factors in developing AI, sort of the recipes and the ingredients and whether the ingredients or the data is actually maybe more important than we think. China makes available a lot of this data, right, big gaps of data to some Chinese firms in the AI space.

What I would suggest is that might highlight, you know, thinking about this in perspective, the potential need to preserve and promote competition, first of all, to stimulate innovation in the space for development of algorithms, but also second to maintain access to the flow of data if that is also very important to this kind of competition.

MS. CONNELLY: Nicol?

MS. TURNER-LEE: May I add one thing?

MS. CONNELLY: Sure.

MS. TURNER-LEE: Yeah, I was going to add in one thing with regard to the GDPR. So I think it is interesting. You know, I agree for the most part with what the other panelists have said on the GDPR and China and their handling of data and how that ties into AI applications. But I think one thing that is interesting that the GDPR has done is it has informed the public around how our data sort of flows through the internet ecology. And it has given some framework, even though I think the United States -- you know, it would be impossibly -- somewhat hard to actually apply that here because of different things -- and Josh and I have debated this.

But I think that one thing the GDPR does do, it sort of unpacks the opacity of the internet to a certain extent, right, because people have to opt in to various applications. The question for GDPR is where in the onion do I get to peel back some of these applications that may be producing a disproportionate output.

And I think that is where the GDPR will really struggle to figure out, is it at the beginning, the middle or the end. For those of us that study algorithms, it sort of begins to look at the black box framework and maybe white boxes it a little bit, but not completely. I think that, again, as the internet has evolved, it will become much more difficult for regulatory frameworks to figure out those pinpoints for consumers to sort of jump in and correct, which is sort of the intent of the GDPR going forward.

MS. DIXON: Can I just jump in very briefly?

MR. ROSSEN: Sure. I have a short followup and then we can move forward.

MS. DIXON: I want to just touch on your white box analytics point. That is the other thing I did not hear about is white box analytics.

MS. TURNER-LEE: That is right.

MS. DIXON: So we are hearing a lot about the black box. But there is such a thing as white box analytical process, and I actually just submitted extensive comments to the NTIA about this and about the need for doing this. So, look, it is very, very possible for even the most complex machine-learning process to be done in a way that is deidentified and it is using deidentified data.

I am not saying this is a perfect privacy protection, by no means. However, it can really help preserve a lot of privacy in certain use cases and situations, and as a general rule of thumb, using raw data should be kind of like walking naked down the street. It is not necessary in every instance. If you decide to do it, great, but you better have some very good reasons for doing it and you better know what you are doing. That is really kind of the white box analytics methodology.

There have been some major -- talking about economics, there have been some very major acquisitions in this area. Lexis Nexis -- or, excuse me, RELX just made a massive over $1 billion purchase of a company that is doing white box analytics and my understanding is that one of the impetus of this purchase acquisition was because competing financial institutions needed data analytics, needed machine-learning analytics, but they did not want their competitors to know what they were getting analyzed and the exact nature of their data. They were not going to hand that over to a third party for both compliance and other competitive reasons. White box analytics solved that problem. Thank you.

MS. CONNELLY: Thank you. I would like to follow up on sort of down a path that Salil, I think, started us on in his opening comments. This relates to further exploration of how we, at the agencies, as well as other policymakers who might be looking at these issues, can better prepare ourselves to handle any competition or consumer protection issues that might be raised by these technologies going forward.

For instance, is there a set of key questions on the antitrust side, Salil, or on the consumer protection side to some of my other panelists, that we should be asking? Is there a set of study or additional resources that we should be looking to build up to sort of better position ourselves looking a bit ahead.

Salil?

MR. MEHRA: So I think one way to think about this is, actually, to think about the way that our current legal framework is essentially our model, right, thinking about the way people develop technology in this area. And so if we think about current legal framework, I know there is debate about consumer welfare and whether we should maintain that as a traditional touchstone, but let's start off with that. These technologies can really still, I think, even if we do not change our legal framework, it can impact how we apply the decisional rules that we have developed over the history of antitrust law and practice.

I will give you a couple of examples. One would be, you know, think about HHI and merger analysis. We have used this for decades, you know, as an indicator of likely loss of competition due to concentration even in the absence of, for example, explicit cartel behavior. Predictive analytics or further into the future AI or deep learning make these anticompetitive effects likely at a lower threshold, then even without changing our legal standards, we might want to apply these standards differently, more stringently. This is ultimately an empirical question.

But it is one that I think the FTC is actually well positioned to consider, for example. In the longer term, right, just like you test a model and you reconsider a model, it feeds into whether you would want to reconsider your legal or regulatory framework down the road. Another example of our existing legal framework and how these technologies might affect how we think about it is to think about price discrimination.

So antitrust law in this area has, over the past couple of generations, has moved towards thinking about this price discrimination as not a problem, essentially, or not a problem from a consumer welfare perspective. Or more specifically that it is only a problem when it impacts competition and thereby consumer welfare, which the Chicago School would tell us never happens or almost never happens, right?

But even if our legal rule does not change, we might be concerned that the increased ability to use machine learning or AI to price discriminate based on the collection of big data could actually change the results, right, change the results of what happened. So what do I mean?

Here is what I mean. Here is an example. It could have negative social welfare effects if -- and this is a big if -- if big data operates as a sort of input entry barrier in some markets, you could see situations where cost rises because big data comes at a cost, so cost rises. The average price to consumers rise through price discrimination, but ex post versus ex ante, the profit to the price discriminator actually increases.

So this would be negative on the whole, but there would be a privately optimal reason to do it, right? So we already have legal authority right now to prohibit price fixing where it lessens competition or tends to create a monopoly. So the issue here would not be about some new law; this would be about applying existing law. It is not necessarily the case that the scenario that I sketch out will always happen. But it is worth being aware that it could happen. If you apply existing law and you start to find the model not tracking what you are finding, then you can reevaluate and think about, well, what needs to change. That is a couple of ways to think about that, how to deal with technology.

MS. CONNELLY: Thank you.

Josh?

MR. NEW: In terms of questions policymakers should be asking or regulators should be asking in the space. Great, thank you for asking that. I get to talk about algorithmic accountability more. When --

the model we developed that we think regulators should be considering when evaluating harm to consumers from an algorithmic systems, they are going to have two really important questions that they should be asking when deciding when they are investigating this case, whether or not the operator of the algorithms or the person who deployed it, the company, should be punished.

The first is whether or not the algorithmic system had mechanisms in place, either technical or procedural mechanisms in place to verify if a system was acting the way they intended it to. So they can verify that they are not acting with malicious intent, they are not actively trying to harm consumers, which is an important part of determining how you would sanction a company. And there are a couple ways you can do that.

The reason that we think this is an outcomes or ends-focused approach is that it could involve transparency, it could involve explainability, it could involve confidence measures. There are bunch of different tools you can use to achieve that, but they are all going to be contextually specific. So algorithmic transparency, as some describe it, does not add a whole lot of value when you are using really advanced deep learning applications when you cannot interpret that code. Even the people who are developing it, cannot explain its decision-making process. But in certain more static algorithms where it is very clear, transparency could add lot of value.

The second question regulators should be asking is whether or not the system had a mechanism in place that the operator could identify and rectify harmful outcomes and that can demonstrate whether or not they were acting responsibly to prevent harm from coming to consumers. And, there again, a series of different kind of mechanisms you could use to accomplish that, both technical and procedural, you could do impact assessments, you could do error analysis. However -- and the -- I think the AI side of the room can tell you about all the different ways you can actually go about doing that.

Then you can -- once you ask those two kind of questions, it gives you kind of a flow chart. We called it a regulator’s neural network, which is kind of corny, I know. But so there is a significant harm that occurs, a harm that is significant enough to warrant regulatory scrutiny. It is not just an inconvenience or a really poorly designed product. It is something that actually caused consumer harm.

So if it passes the first check, they did demonstrate that they could -- that system was acting the way it was intended to, yes or no. If no, then they are already subject to a modest penalty. If they -- if yes and you go to the second point -- or you go to the second point regardless, if you can identify and rectify harmful outcomes, if you answered yes to both of those questions, you are left in kind of this weird area where you were acting in good faith, a bad thing happened that might not necessarily be illegal and harm occurred, there are different ways you can approach incentivizing that kind of thing not to happen again.

But if you answered no to at least one of those questions, you get sanctioned moderately. If you answered no to both of those questions, you get sanctioned very heavily. That creates a kind of -- a pretty clear process about how you can actually go about enforcing the company’s acting in ways designed to -- you know, they are actively invested in ensuring that their algorithms do not cause harm.

Again, this is our stab at the model, I am sure there are other ones. I would love to debate them. But, right now, I think that is the best idea that we have had about it.

MS. CONNELLY: Pam?

MS. DIXON: Thank you. So I love talking about the governance. I like talking about it because it is practical and it means that you are down there in the nitty-gritty where it is actually happening.

So the model we have been working on is really the Elinor Ostrom model, which was -- she has eight principles and they have been extensively ground-truthed and tested over and over in the environmental context, but they really work, also in the data protection, privacy, human rights context.

So let’s just talk about -- basically, the idea is you end up with a broad framework of things you want to accomplish, bad things you do not want to happen, good things you do want to happen. You develop a risk mitigation -- iterative, ongoing risk mitigation process so you can identify the bad things you do not want and make sure they are not happening. And then, of course, you have the ethical guidelines that articulate what you do want.

But within that, what Elinor Ostrom found through her work over decades is that if you have these systems be macrocosms it is extremely ineffective. Rather, she ends up with microcosms. So smaller slices of data ecosystems and machine-learning ecosystems are going to work more effectively than taking some gigantic slice of the pie.

And then identifying the stakeholders that are impacted by those machine-learning algorithms, perhaps bisect or even making it smaller slices. So, for example, in the healthcare environment, what do the stakeholders have to say there about, for example, a frailty score that someone gets or the use of medical diagnostics, et cetera, et cetera.

You have to take small slices, work through that in an ongoing, iterative analysis of the risks and the specific mitigations for those risks and it is a collaborative model of the shared resource of data and the data outputs and the data inputs, the entire spectrum, not just one chunk, the entire spectrum. But it has to be collaboration. If it is command and control, it will not work because you still then end up with disenfranchisement.

MS. CONNELLY: Anyone else on this? Justin?

MR. BROOKMAN: Yeah, sure. So, first, I want to echo Salil’s point. He made a point that I made in my earlier comments, but in a far more informed and articulate manner, on price discrimination. So I appreciate that.

I am going to answer in a slightly different way, but also it is like a theme that I have heard throughout a couple of days, which is the need for technology staff at the FTC. So having been in OTEC, I think OTEC plays a tremendously helpful role there, but it is like a handful of people. You can make a compelling argument they should expand ten-fold. I know I heard Commissioner Slaughter and other folks talk about the need for a bureau of technology to address these issues.

I do not think I would go quite as far as Jeremy from EFF when he said there should be 50-50 split between technologists and attorneys at the FTC. Rather, I think actually they need lot more of both to address these issues. The FTC is, what, half the staff it was in the ‘80s. The economy has grown three times as much and there are a lot of very challenging consumer protection issues that did not exist back then.

Also, at the same time, more technologists is not a panacea. Even if it was 70 people in a bureau of technology, the FTC is going to have less people than -- less technologists than any Silicon Valley company of moderate size. They are going to be generalists, right? They are going to be working on AI; they are going to be working on security; they are going to be working on ad tracking. I mean, you are always going to be outgunned. I think that imbalance of tech expertise cannot be an excuse for inaction. The FTC cannot wait until it is like 99.999 percent sure that it has the right approach.

I know that Chairman Ohlhausen used to speak about regulatory humility, which is fine, but I think there is also -- that cannot turn into regulatory timidity. It cannot be excuse for inaction in this area.

MS. CONNELLY: Nicol?

MS. TURNER-LEE: Yes, I was just going to add -- so Justin kind of stole my thunder. I think there definitely needs to be some technologists at the FTC and perhaps one social scientist would do to add to the team. But I also want to say the FTC should really look at -- you know, the FTC has done really a great job I think prior to this discussion on artificial intelligence when it came to big data.

Very rich, robust reports have come out of the FTC with regards to algorithmic bias that was something that FTC took on last year or the year before. It has continued to talk about it. The Obama Administration came -- conversations around equal opportunity frameworks when it came to algorithmic design.

The FTC could play a role and I think regulators, in general, should play a role in leveraging their pulpit for more algorithmic hygiene. You know, how do you create a set of criteria or triggers for even companies to, you know, first look at what are they doing in terms of their hygiene when it comes to the purpose or the intent of the algorithm, the feedback mechanisms that are embedded in the systems, the involvement of civil society on those applications that will have potential unintended consequences or predictions that may be wrong.

You know, having that conversation and using the regulator to sort of advance that discussion would be equally helpful because what we see in Washington oftentimes is, again -- and I want to go back to the black box -- a lot of the discussion has been on the output of the black box versus understanding what is actually the input. And when you are in Washington doing policy, your concern is really for the output. It is for what is at the end of the spectrum not necessarily for what is going into the recipe.

And having that disconnect with the FTC and other regulators, raising awareness of what that looks like, advancing consumer algorithmic literacy is also, I think, a role of a regulator so that we can get to a place where we can all sit at the table and have this conversation. Because I think in many of the conversations that I am personally in, when we convene various stakeholders, they are talking on two ends of the table. When you place a regulator in the middle, they are trying to figure out which side to pick.

So I think, again, in addition to what has already been said about consumer welfare standards and some of the tools that the agency and other regulators have at their disposal, the real question is, are we raising the level of awareness of, again, what are those use cases and the extent to which we all have a basic understanding of what we are trying to regulate. I think that definitional hiccup will sort of stand in the way of us making a lot of progress.

MR. ROSSEN: So following up on a couple of things that you all have mentioned -- and maybe Justin and Nicol, I will sort of direct this first to the both of you. You know, we have heard over the last couple of days a lot of discussion about fairness and ethics being baked into AI and tools that might be available to make a difference in that.

One of the things we heard about a bit yesterday was this idea of differential privacy and I do not know if we got a sort of full picture as to exactly what that is and what it means, but there was discussion about how technology has improved to the point that differential privacy might be a bigger player than it has been. Is that something that more companies should be looking to? Are there incentives that are needed in order to sort of push folks to do that? Are there things needed to encourage companies to bake fairness and ethics in sort of from the outset?

MR. BROOKMAN: Yeah. So I think differential privacy has a lot of positive applications and it was cool to hear that the 2020 Census will be using that for all their early results and that some folks like Google and Apple, who have some external brand name pressures, are adopting those. Is there enough pressure for the industry to be doing this, to do robust de-identification e-type things? I would argue not. I think there really do need to be some more bright-line rules in this space.

I think the wait-and-see approach, which I heard also mentioned a couple of times here, I think -- I do not know that they have done enough. I think that is kind of the reason we are having all these hearings. The wait-and-see approach has not really been good enough. I think Chairman Simons basically said that when he kicked off the initial approach. There needs to be more rules in place.

I think one way to do it would be mandating, limiting inputs in some ways around things like background checks and credit scores. Did I pay a bill, does that go in there, maybe that is fine. Was I arrested, sure. What I got at grocery store, you know, maybe not, right. What I do in social media, maybe we should just say that is out of scope for this sort of thing.

FTC has said that if FCRA applies to those sort of things that you got to let them know. Maybe we can go a step farther and just say, you know, the social cost of those sorts of things, even if they are right, the chilling effect on free expression extended to autonomy just is not worth it. I mean, more broadly, I think we do need privacy law to help, again, arm consumers against potentially adversarial AI. Technologically, everything about us is collectible now.

There was a paper out last week about how people can use WiFi signals to kind of see through walls to see when you are walking around your apartment. You know, we have this concept and the

Fourth Amendment that there are some things that are just off limits. Even if it is collectible, it is just not reasonable to collect it, like that sort of thing.

I think we need to transport some of those ideas over to commercial privacy as well and it needs to include things like collection limitation and data minimization. These were, I think, relatively more controversial ideas maybe five years ago. I think now even like Google's privacy principles recognize, you know what, some things should just be off limits.

MS. TURNER-LEE: Mm-hmm. Yeah, I want to echo what Justin is talking about in terms of things being off limits, and I was not here to hear the conversation of differential privacy, but understanding that companies are trying to create these larger tents so that they actually do not find themselves creating these discriminatory effects, I think is important.

But, you know, one of the things that I think is a technical limitation of where we are with this harvesting of this new data is the fact that the connections that happen on the web -- and this was Michael Kerns' piece on the inferences that are actually adopted -- they do not have a start or stop and there is no causality to it, which is something that we used to see in the harvesting of big data, right, this relational database.

Now, what could start as me liking red shoes and ending up with me receiving a predatory credit card or loan because the red shoes somehow got associated with the fact that I am a single parent and, you know, I search certain things because I am limited in income. I think that is, again, going back to Justin’s point, where there might be areas that are off limits when you actually look at that.

I was also going to say, too, I have been pushing -- and, again, as sociologist who looks at the social science aspects of AI application -- you know, where is the strict scrutiny where it comes to these data sets and the checks and balances that are associated with that. When I want to study human subjects, I have to go through IRB. There are certain things that I have to actually check off that I am not harming individuals when it comes to the harvesting of the information that I am collecting on a simple research study.

Because what we are seeing today with AI is a rush to market and a rush to innovation, I think goes back to Justin’s point, even if companies like

Apple apply differential privacy the question is, it is still not necessarily giving you discrete variables as to whether or not I am an African American woman, my direct address. It is inferring that which, again, goes back to making uneducated guesses around my behavior, which then can have an outcome.

So I think, again, having good comprehensive privacy law at least starts the process, but like many people who I think we heard throughout couple of days, we are all baffled on what do we do next and the extent to which we apply strict scrutiny to certain things. I think having use cases that are off limits may actually do that or creating regulatory safe harbors or sandboxes where we can experiment in those cases, where people are very much aware that they are being experimented upon, versus finding out later that because of something that they did online, they were denied a credit or a loan and cannot take that back.

MS. DIXON: We really need to mention data brokers here in these contexts.

MS. TURNER-LEE: Yes.

MS. DIXON: And I do not know if it came up yesterday, but it did not come up today until now. Look, please go back and look at all the testimony I have given since 2009 on data brokers. Look, we have a big problem, especially regarding transactional -- financial transactions. When our financial transactions are largely digital, either debit cards or credit cards, it leaves a juicy trail that is just beautiful analytic material. Imagine this over the course of maybe 30 years, 40 years.

And you know what, it is really difficult to get away from that trail and to get away from the enormous predictive qualities that that trail allows for. And then there are generational issues there as well where you can also have entire families' transactional histories. We have actually been working on analyzing some of these data sets and the data sets are available in the U.S. and the U.K. and Canada right now. They are absolutely profound data sets and they are a little bit terrifying as well.

So what do you do? So, you know, one of the questions that I have been having in regards to some of this research is what is human subject research in the context of machine learning and AI. Do we need to take a new look at that? And I think the answer is yes. A lot of what I see that is characterized as A/B testing is not actually A/B testing, where an academic institution covered under the common rule was conducting the research, they would have to go through an IRB and the IRB would not approve the study. So we have to look at that.

The other thing I would say is this, you have to look at every single step and micro step along the entire continuum of the AI process. I appreciate the constraint on uses on the back end, but I really do believe that looking at an ethical impact assessment of the data collection, the data quality, is it disaggregated gender data, is it aggregated data, what has been aggregated with the data, what is the context of the data, there are a lot of pieces of the puzzle that could be added, and I do believe it is highly context specific, which means a lot more work for a regulatory agency.

But I think even laying out a series of like a dozen very specific sector-based use cases would be very, very helpful.

MS. CONNELLY: Anyone else on that point?

(No response.)

MS. CONNELLY: I would like to circle back to something that I believe was said on the very first day of hearings, so way back in September. I would like to get this panel’s views on this idea. It also connects to a number of the presentations and discussions we have had over the past day and a half about this concept of intelligibility and the extent to which some of the more complex, perhaps machine-learning technologies or more complex algorithms are or are not intelligible.

So in the first day of hearings, I believe that one of the panelists, towards the end of that day, made a comment along the lines of consumer protection is a much harder task for the FTC without clear visibility into what is going on. I would like to ask that question. Perhaps Salil could comment on that same concept from the competition side. Is antitrust also a much harder task for the FTC without clear visibility? Is it true that we do not have clear visibility or that there is not a way to get clear visibility into what is going on and then also come at it from the consumer protection side? Maybe we will start with Salil.

MR. MEHRA: Yeah, I have thought about this a little bit and I think it is going to be a problem for you potentially. I do not think it is an insoluble problem, thankfully. You are talking about this idea without clear visibility, without intelligibility, without sort of transparent prices and outputs, right. So one of the thing these technologies help you do -- it is not the only thing they help you do -- but one of the things that these technologies help you do is to match, right, match buyers and sellers, match whatever, people on a transactional platform or other platforms.

And they are matching in what is, as people say, a black box so you do not have as easily observable prices and outputs without some sort of compelled data disclosure, right, through litigation or otherwise. I think that there is a potential danger to that. You sometimes will see people worried about, for example, Amazon with the analogy as a trader or a broker with a broker system with a frontrunner inside the broker, someone who can see the orders as they come in and price in advance of them.

Where I am going with this is there is an analogy to some of the things I think the SEC has been dealing with in terms of market fragmentation and trying to deal with the possibility that fragmentation is not necessarily to the benefit of the consumer. You know, they have been dealing with this for I think almost 20 years at this point. So I think it is something to think about as these technologies develop.

MS. CONNELLY: Thank you. Anyone else? Justin?

MR. BROOKMAN: Yeah. So I think in some ways -- sometimes explainability is mandated and I think that should remain the case. FCRA says you have to be able to explain it. You cannot say, I do not know, machine learning. That is prohibited. I think that should probably remain the case for especially essential decisions.

I already talked about the role that transparency plays and I think there should be greater obligations there.

Substantiation is an interesting area when it comes to AI. So I really enjoyed Professor Dickerson's intro yesterday when he described neural networks as they kind of throw together a model and they run it. They step back and are like, hmm, that does not look right, and they are going to rejigger stuff and kind of back into it, it sounds like. That may be a lofty distillation of it. But I do feel that in AI there often is like false promises of precision and dodgy accuracy. You know, we are testing your saliva, we will tell you you are 38.742 percent Irish. You know, at what level -- and the FTC requires substantiation around advertising claims. At what level does an AI system have to be substantiated?

Like they kind of got there a little bit in the Spokeo case. Like Spokeo was an online data broker and they were like five people, but they had like records on everyone in the country and they had some algorithm, but it was deeply stupid. I mean, it was comic. I was listed as Hispanic Jewish, who made a lot of money, but I had a lot of debt. But they made like very precise determinations about everyone in America. And the FTC ended up bringing a case, but it was limited to FCRA claims. They were saying, hey, use this for employment purposes and they were not following the Fair Credit Reporting Act.

There was an element in there about like accuracy under the Fair Credit Reporting Act. But I think there are interesting questions more broadly about the FTC could be doing more to kind of come in and say, you know, you have to have some basis for making these very precise claims other than I do not know, the machine said it.

MS. DIXON: I am just going to pick up on just a few things. I really -- I really agree with that.

So in terms -- there is a continuum of explainability on AI. Some of it is incredibly explainable and transparent and then it goes to the other end as well. I want to focus on two things, explainability and interpretability. So explainability being are the results explicable and defensible? And there is so much research being done on this now. So I do think that there is a lot of hope there, even for very opaque systems.

Interpretability, though, is something I do not hear a lot about. How do you interpret the ultimate output? So I really like to always talk about the credit score in regards to interpretability. Why do we care about our credit score? The reason we care is because if we are going to buy a home, it matters; if we are going to buy a car, it matters. In large credit decisions, it matters. It has a meaningful impact on what we are going to pay, what interest rates and whatnot.

Well, if you have a score of 100, it is so substantially different than having a credit score of 700. How do we know that? It is because there is a limit. We know that the top perfect score is 800. So we have a very clear idea of what is not so good, good, really good, and just perfect.

So a key to interpretability is to have that kind of very specific boundary and definitional boundary of what that particular output means no matter what form it is in, whether it be a score or some other categorization.

MS. TURNER-LEE: Can I say something? I think those are really good points, but you also have to do regular audits and have imbedded feedback mechanisms to continue to see if the algorithm is still learning and training itself in the way that you actually designed it.

What I found to be interesting, in Allegheny County, Pennsylvania, governments have actually, you know, had the pulse on this because they have had no choice to do so. They developed -- an algorithm they developed about vetting child abuse cases in Allegheny County, Pennsylvania. They decided, okay, we are going to develop an algorithm, cut down on the number of calls. They tested for one thing and had a researcher come in only to find out that there was bias imbedded in it and that African American kids were most likely to be removed out of the home compared to white kids just based on the algorithm alone. But what was interesting about them and responsible was the fact that they did that check.

So I think that, again, as you look at the intelligibility of the algorithm, it is important, I think, to Pam’s point, you have to have the explainability, you have to have the interpretability, but you also have to have these mechanisms built in throughout the process.

That was Joy's work, right? In developing facial analysis software or doing her research on that, she said, hey, companies, guess what is happening here. And those are things that companies will not predict or may not seem intelligible at the time or they may seem intelligible at the time, but the data may actually output a different result.

So I think, again, there are subsets to everything that we are talking about that will move it from a big tent to smaller tents and potentially into smaller areas of concern, which I think goes back to the earlier point that Justin made, which is what is off limits. Once you figure out in that feedback loop that, hey, this is discriminating against kids of color who are going into foster care at a much higher rate because of the AI, then what do we need to do to take this off limits and maybe not use or apply this?

MR. ROSSEN: So we have just ten minutes left and we are going to try to get to some of the questions we have received from the audience. I will start with this one. So we have heard about multiple jurisdictions that are developing AI governance models. Should regulators build up consensus in this process? Are there risks that disconnect in regulatory approaches from one jurisdiction to another that could result in AI being developed or deployed in one country but unable to be extended elsewhere? Are there are other risks posed from these different frameworks as they evolve?

Josh, do you want to take it?

MR. NEW: Sure. So there are risks. A lot of the discussions about how we can approach governance is, you know, encouraging ethics by design or encouraging fair and responsible systems that reflects our values to society. But Pew just came out with a study the other week about kind of surveying different cultural attitudes about the trolley problem, which is like the worst conversation you could have in AI. But, you know, whether or not a vehicle will -- you know, if you leave it going and you do not stop it, it will kill one person or it will kill five people or you could switch the tracks and kill one person, that is an ethical debate.

So with autonomous vehicles, you are going to have to, at some point, make decisions about who to save in an accident. I think that is a preposterous discussion that influences this so much. But their survey found that from country to country, across different demographic and social economic groups, people will choose to save -- there was a pretty wide divergence in who people would choose to save.

In Europe and the United States, we would prioritize younger people over older people. That is just not true in China and Japan where the value of like an elder is held in much, much higher regard than it is in the United States and they would opt to choose -- they would save an elderly person over a child if they had control over that car.

And the same conversations -- there is a lot of effort on global consensus here, about how we actually enforce this kind of ethical human rights by design thing. But I think that study demonstrates that that is an unworkable approach. What ethics and values are are going to vary so much from country to country, and in some countries, their social values are disenfranchising minority groups or women, or sacrificing the lives of some to save other groups that we would just not do in the United states.

So I think we really need to kind of avoid those approaches, these really broad global governance style things that rely on a really subjective notion of ethics and values.

MS. DIXON: I would just say very briefly there is not going -- it is unlikely that China is going to reach a consensus with Europe.

(Laughter.)

MS. DIXON: So given that, where does that put the rest of the major jurisdictions that are working with AI, and I think that different frameworks will be possible. I really agreed with the person from Microsoft who talked about there is no one silver bullet anymore. We are going to end up with layered ecosystems. It is going to be a layered approach.

MS. TURNER-LEE: Although, I mean, I would just add, having just got back from China and having this conversation, I think there is concern, though, when you start to go up on the scale of the severity of the AI application, particularly when you are looking at autonomous weapons, that there is a need for some type of conversation on global governance.

We do not want AI innovation used I think across the globe in ways that can be detrimental and harmful to countries in weaponry, and I think it is important that those conversations happen. I know that OECD has been having this conversation. But that global conversation needs to happen and potentially that will find itself in the financial sector and other sectors, which have also become weaponized in many respects that will have to look at it.

MS. CONNELLY: Salil?

MR. MEHRA: Just really quick, we see a lot of divergence in terms of institutions for making decisions generally and you can think of AI as another tool of making decisions. We see some convergence in certain areas, corporate governance, et cetera. You might find some areas of commonality where you can pursue that as well with AI.

MS. CONNELLY: Thank you. We have about five minutes left, so I think I would just like to ask one wrap-up question and go right down the line. I would like to know from each of the panelists, is there one application or use or sort of one particular policy issue that you think we really should focus on going forward? Where should the debate go from here? Whoever would like to start and we will just --

MS. TURNER-LEE: Ah, are you going to start with me?

MS. CONNELLY: Sure.

MS. TURNER-LEE: You know, without picking one because I think the area in which I study has become very interesting because historically disadvantaged populations in vulnerable groups have already been disenfranchised and marginalized, so I think any of these applications could be one of focus.

I would like to actually shift it -- and this is something that we are going to be presenting in our paper to the FTC focusing on the output, whether it is the disparate impact or disparate treatment of populations caused by the particular application. Impact could be or treatment could be applicable in the bail and sentencing examples that we see using the COMPAS algorithm. Impact could be something -- and I know that the company has sort of retracted the algorithm, but, you know, Amazon and its gender bias in their recent algorithm could have led to reduced wages for women and the lack of representation in their workforce, which could have other impacts generally.

For me, I think we should move away from a conversation of just which application and really prioritize on what are the disparate effects of those particular applications and have more of that view whether it is surveillance being another one that we need to pay closer attention to.

MS. CONNELLY: Josh?

MR. NEW: So I think particularly as it relates to issues around consumer protection and discrimination, what gets left out of these conversations is that, for the most part, companies have a pragmatic interest in ensuring that their algorithms do not discriminate. You can argue that that market force is very imperfect and I would agree with you and they do not always do a good job of fulfilling their own pragmatic ends.

I think the presentation we heard earlier about facial recognition demonstrated that quite significantly. Microsoft or IBM, if they are selling facial recognition, they want to say it is accurate as possible for all demographic groups, but they are not there yet. But recognizing that an incentive exists for them to get it right because, you know, if you are a bank and you implement an AI-alone granting system, you lose money in the long run if you are denying loans to people who deserve it or issuing loans to people who cannot pay it back. There is a force pushing you in the right direction. There is definitely a need for insistence.

What I think the biggest priority for policymakers should be is identifying areas where those market forces do not exist. So it is when the cost of a faulty decision from an algorithmic system are not borne by the person -- by the operator, the person who makes that decision.

So the most obvious example is in the criminal justice system where if a court uses a sentencing decision support system for issuing parole and they are wildly discriminatory, they are not going to lose customers. That is not how the court system works. A judge might be reprimanded maybe, but the court will still be there doing its thing. They do not really have a strong incentive to get it right, other than social value. But, you know, we have seen that not work out before.

So the public sector, more broadly, the market forces are not nearly as significant as they are in the private sector because the really entrenched relationship with contractors, it is not a widely competitive market, those market forces are muted. But there are other areas -- and I am still struggling to identify what they are -- where those market forces are either not present or not significant enough to actually have an impact of encouraging good behavior. I would be really, really fascinated to see what regulators or policymakers can come up with by surveying what kind of potential applications for those market forces would be relevant because that is exactly where we need new laws, regulations, and a lot more insight.

MS. CONNELLY: Salil?

MR. MEHRA: Sure. There has been this tendency so far -- it is not universal -- but to see or promote big data, algorithmic processing, and AI as almost a new form of IP that justifies a kind of hands-off competition law approach in some lines. But I would point out that unlike other forms of IP or things like IP, they have the longer-term potential to impact not just what is in a market, but what a market is. And I think what I would like to see going forward is for the FTC to continue to foster competition, promote consumer welfare and further innovation, and I think that may require some outside-the-box thinking so to speak.

MS. CONNELLY: Justin?

MR. BROOKMAN: I have a slightly different issue that has come up a little bit -- it came up in Professor Dickerson’s intro -- which is gameability, how attackers can exploit AI. AIs tend to be really good at very narrow tasks. They will start out okay and then they will surpass human cognition, but then you will change a rule slightly and it will become terrible.

I think this is a problem for attackers on

AI, that these systems are designed kind of assuming everyone is a good actor, but everyone is not a good actor. So I think we saw around like the 2016 election, like, you know, how bad actors can weaponize algorithms. And if we are going to be relying on AI systems to protect us, you know, are the incentives sufficient for companies to deploy them at scale? Are they workable to protect against these sorts of bad actors? Because, again, this seems like something AI is not necessarily well designed for. So I think there is a lot of -- I mean, we can have a whole other panel on like, you know -- there are a lot of issues there that are important to consider.

MS. CONNELLY: Pam?

MS. DIXON: A few brief things because I cannot just choose one. So first, in terms of privacy, privacy is so much broader than the right to be left alone. I think pretty much everyone recognizes that. Privacy is the core set of rights that really enable human autonomy. In light of that, just acknowledging that as a baseline rule, I mean, something very important that can be done particularly by the FTC is what are the rules regarding de-identification of data and can we please make it so that raw data use as a, you know, just automatic default is literally like running around naked in the streets. I think that that is doable. There are so many entities that are like, oh, we anonymize data. No, no, no, you might be de-identifying it, you might be aggregating it, but, you know, really tackling that issue.

And then something that is a big picture, but I think that it is absolutely central to all of the principles and ethics and all of these things is how is it that the Federal Trade Commission could allow all stakeholders along the entire continuum of AI and machine learning to have an appropriate voice and stake in the process so that all parties have a voice. Because, right now, I think a lot of what we are hearing is parties who do not have an appropriate voice, and I do think that could be remedied with good governance and really a focus on governance.

MS. CONNELLY: Thank you.

Please join me in thanking our panelists from the last panel. A really interesting discussion.

(Applause.)

MS. CONNELLY: If you would indulge me for just a moment, I want to note that we got a number of questions related to privacy topics and I will use that as a plug to note that we will be coming back around to some of these issues in future hearings in 2019.

I would also like to just take a moment to give our sincere thanks to Howard Law School for hosting this event.

(Applause.)

MS. CONNELLY: And, also, just to note that there is a lot of work that goes into this behind the scenes and, in particularly, to thank our AV team and also all of my colleagues in OPP and, in particular, the Office of the Executive Director. Without all of these people helping out, we would not be able to put this together. So thank you.

(Applause.)

MS. CONNELLY: And with that, I would like to have our panelists maybe step down and I will introduce our closing remarks.

CLOSING REMARKS

MS. CONNELLY: So we are very privileged to have the Dean of Howard Law School, Dean Danielle Holley-Walker, here to deliver our closing remarks.

Thank you, Dean.

(Applause.)

MS. HOLLEY-WALKER: I just want to say what an honor and a thrill it has been for Howard University School of Law to host these FTC hearings and to cosponsor this event. I really want to thank all of the organizers with the FTC and also our law school staff who have worked so hard.

I particularly want to think Professor Andy Gavil, who is here in the audience, who gave welcoming remarks on my behalf, and also had the idea -- we loaned him to the FTC, I like to say, for several years and he has been just an outstanding antitrust expert here for almost 30 years. So his guidance and ability to really provide antitrust knowledge to our students here at Howard has really culminated I think in this moment with us having the FTC hearings.

I am actually right next door in room 2 teaching introduction to administrative law to our students. And so it is such a -- and some of them have had the opportunity to come over the last few days and hear this remarkable set of hearing. And I think for us to be able to host the hearings on competition and consumer protection, particularly as related to algorithms, artificial intelligence, and predictive analytics has been a special treat.

I sat through one of the panels earlier today and learned a tremendous amount from the panelists, and all of the expertise of the academics, public servants, scientists, engineers, industry leaders, and lawyers and economists who have been here to present has been a tremendous value to the law school and I hope to the FTC.

I hope before you leave the law school -- this is our 150th year. In 2019, we will be celebrating it. I hope you have had the opportunity to walk around the grounds of this incredible institution, see the history on the walls, and all of the people we are influenced by who have made such a big difference in the profession.

And my second hope is that this will not be your last visit to Howard and your last visit to Howard University School of Law. I hope that you will be back many times over and come back and share your expertise and your ideas with us, help us create the next generation of outstanding antitrust lawyers and outstanding people who work in all of your fields.

So thank you so much for being here.

(Applause.)

(Hearing adjourned.)

CERTIFICATE OF REPORTER

I, Linda Metcalf, do hereby certify that the foregoing proceedings were digitally recorded by me and reduced to typewriting under my supervision; that I am neither counsel for, related to, nor employed by any of the parties to the action in which these proceedings were transcribed; that I am not a relative or employee of any attorney or counsel employed by the parties hereto, not financially or otherwise interested in the outcome in the action.

LINDA METCALF, CER
Court Reporter