

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25

FEDERAL TRADE COMMISSION
COMPETITION AND CONSUMER PROTECTION
IN THE 21ST CENTURY

Tuesday, November 13, 2018
9:00 a.m.

Howard University School of Law
2900 Van Ness Street, NW
Washington, D.C. 20008

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25

FEDERAL TRADE COMMISSION

I N D E X

PAGE :

Welcome and Introductory Remarks	3
Introduction to Algorithms, Artificial Intelligence, and Predictive Analytics	6
Opening Address by Michael Kearns	32
Understanding Algorithms, Artificial Intelligence, and Predictive Analytics through Real World Applications	58
Perspectives on Ethics and Common Principles in Algorithms, Artificial Intelligence, and Predictive Analysis	139
Consumer Protection Implications of Algorithms, Artificial Intelligence, and Predictive Analytics	217

1 WELCOME AND INTRODUCTORY REMARKS

2 MR. GAVIL: Good morning, everyone. My name
3 is Andy Gavil, and I'm a Professor here at the Howard
4 University School of Law. On behalf of Dean Danielle
5 Holley-Walker, my faculty colleagues, and our
6 students, I'd like to welcome the FTC and all of you
7 to Howard for Hearing Number 7 of the FTC's hearings
8 on Competition and Consumer Protection in the 21st
9 Century. We are very happy to cosponsor today's
10 event, and I want to thank the FTC and the many people
11 at the agency and here at Howard who have worked hard
12 over the past few months to organize these hearings.

13 As you all know, today's topic is
14 Algorithms, Artificial Intelligence and Predictive
15 Analytics. As is immediately evident from both the
16 list of questions the FTC has posed and the agenda for
17 today and tomorrow's programs, these hearings have
18 been purposefully designed to take a broader and more
19 interdisciplinary perspective than any of the previous
20 ones.

21 Moving well beyond the usual collection of
22 academic and practicing economists and lawyers, FTC
23 staff have assembled an impressive collection of
24 academics, public servants, technologists, scientists
25 engineers, and industry leaders, but, of course,

1 there's still lots of lawyers and economists.

2 The goal is to educate the agencies and the
3 broader competition and consumer protection policy
4 community so we can all obtain a better understanding
5 of the technologies that are transforming our economy,
6 as well as our political and social environs. We'll
7 hopefully learn more so we can better understand the
8 business models and practices of our time and so we
9 can differentiate myth from reality, promise from near
10 and long-term prospect.

11 The ability to take on this kind of
12 prospective study is a hallmark of the FTC and one of
13 its great institutional strengths. It is especially
14 fitting that such a forward-looking approach is being
15 taken here at Howard. Only two years after Howard
16 University was chartered by Congress in 1867, this law
17 school was founded with the aspiration of producing
18 lawyers who would lead the future fight to realize the
19 full promise of the reconstruction amendments to the
20 Constitution of the United States.

21 Next year, we will celebrate our
22 sesquicentennial, and for that occasion, instead of
23 looking backward, we have selected a theme that looks
24 forward, "The Next 150." As is true for the FTC and
25 for today's hearings, any institution that fails to

1 look forward is bound to fall backward.

2 In closing, please note that the event is
3 being photographed and webcast and will be posted on
4 the FTC's website, and that by participating all
5 attendees consent to those conditions.

6 Please also note that our students will be
7 coming and going throughout the day and are available
8 to answer your questions. Please get to know them
9 while you are here and feel free to seek them out if
10 you have any questions or concerns.

11 Finally, it's my great pleasure to introduce
12 our first presenter. Our scheduled presenter, Michael
13 Kearns, has been slightly delayed, so we're going to
14 start with John Dickerson from the University of
15 Maryland, and hopefully Michael will arrive in time to
16 follow John. Again, welcome, thank you, and enjoy the
17 hearings.

18

19

20

21

22

23

24

25

1 PRESENTATION: INTRODUCTION TO ALGORITHMS,
2 ARTIFICIAL INTELLIGENCE, AND PREDICTIVE ANALYTICS

3 MS. GOLDMAN: Hi, I'm Karen Goldman. I'm an
4 attorney in the Office of Policy Planning at the
5 Federal Trade Commission, and I just want to introduce
6 you to John Dickerson, who is an Assistant Professor
7 in the Department of Computer Science at the
8 University of Maryland, College Park. Welcome.

9 MR. DICKERSON: Thank you, Karen. It's a
10 pleasure to be here. I am John Dickerson, I'm a, I
11 guess, third-year Assistant Professor at the
12 University of Maryland and right up the street in
13 College Park, and today I'll be talking about an
14 introduction briefly introducing the audience to
15 algorithms, AI, and predictive analytics.

16 And so for this talk, I'd like to start with
17 a motivational quote which sounds like it was written
18 a long time ago, and that's because it was. So
19 "although machines can perform certain things as well
20 or perhaps better than any of us can, they infallibly
21 fall short in others," by which means we may deduce
22 that they did not act from knowledge but only from the
23 disposition of their organs.

24 And this sounds old because it was written
25 a long time ago. It was written by Descartes, who was

1 a philosopher and mathematician in the 1600s. So
2 quite a long time ago, folks were already thinking
3 about what does it mean to think, can we mechanize
4 thought?

5 Another famous philosopher from the 1600s,
6 Hobbes, states, "reasoning is nothing but reckoning."
7 So reckoning here is just a reference to mathematics.
8 So reasoning is nothing but mathematics essentially.

9 And some time passed, 1600s, 1700s, 1800s,
10 until the 1900s, when some breakthroughs occurred in
11 logic and mathematics and philosophy. Folks like
12 Boole, folks like Hilbert, made some breakthroughs in
13 the formalizations of mathematical reasoning. So
14 recall, we think reasoning is nothing but reckoning,
15 and now we can reckon perhaps with mathematics.

16 So there were some proofs showing that some
17 hard limits -- there are some hard limits to what
18 mathematical reasoning can do, but subject to those
19 limits, folks like Alan Turing came around, Church
20 came around and said there are certain machines --
21 simple machines -- that for any of these mathematical
22 reasoning problems, subject to these limits, we can
23 create a machine that can do this.

24 So this is nice, this built on now hundreds
25 of years of philosophy and mathematics, but the

1 general pitch here is that if intelligence can be
2 simulated by mathematical reasoning, that is reasoning
3 is just reckoning, and mathematical reasoning can be
4 simulated by a machine, then can a machine simulate
5 intelligence?

6 So AI, artificial intelligence, the word was
7 coined by John McCarthy in either 1955 or 1956,
8 depending on how you count, it's '55 in a proposal, to
9 fund the Dartmouth Summer Research Project on
10 Artificial Intelligence. And you'll hear this called
11 the Dartmouth Conference. This occurred in the summer
12 of 1956.

13 And there are some fun quotes in there
14 saying basically we can solve artificial intelligence
15 in three months or we can solve artificial
16 intelligence in one generation, but the one I'd like
17 to pull out is that every aspect of learning or any
18 other feature of intelligence can be so precisely
19 described that a machine can be made to simulate it.
20 So even in the 1950s, 1960s, folks were making
21 statements like this.

22 So a quick spoiler, this hasn't happened
23 yet. We can just shut this down right now. But,
24 progress has been made. So how does that progress
25 occur? Well, this is a cycle of basically R&D

1 progress that you'll see repeating in the AI world,
2 and this has happened since basically 1956, where some
3 new advance, maybe a new technique, new hardware
4 happens. Fast progress is then made on old, hard
5 problems. So it could be a new mathematical
6 technique, it could be new hardware, GPUs, these
7 graphics processing units, are one of the main drivers
8 in the current sort of fast progress being made on
9 problems that we're seeing now.

10 But eventually you start to hit road blocks.
11 And at this point, the academic community, the
12 industrial community starts to get pessimistic, this
13 bleeds into the press, and at that point, everyone is
14 pessimistic about progress, funding dries up, progress
15 dries up and so on. We wait until the next large
16 advance.

17 And so this is the cycle that occurs in most
18 sorts of verticals. It occurs in AI research as well.
19 In AI, though, we call it a cycle of basically AI
20 summers and AI winters. The winters are when funding
21 dries up and nothing happens; the summers are
22 basically what we're going through right now, where
23 we're seeing large advances driven by sort of recent
24 hardware and mathematical advances.

25 So this is a bit pessimistic, this cycle,

1 but like I said, progress has been made. So this has
2 been cycling for arguably maybe six or seven times
3 since the 1950s, but every time we go through this
4 loop, progress is made, new problems are solved, and
5 new problems are encountered.

6 So what is AI? AI, many definitions, the
7 one I'll use here is the ability to process and act
8 based on information via automation. So we can break
9 this down roughly into four segments. One is
10 perception. I want to be able perceive the world
11 around me. That could be the physical world; that
12 could be the virtual world. I want to be able to
13 learn something about it. So I get some signals about
14 the world, then I learn something about them. Maybe I
15 learn a model.

16 I want to abstract and generalize that model
17 so that I can use it in other situations. And what do
18 I mean by use? Well, maybe I can reason about this
19 information, I can reason using my model and then act
20 within the world. Again, that could be virtual, that
21 could be physical.

22 So if I can create this automated system,
23 roughly, I have created what we would call AI. So
24 let's keep moving through this history of AI until we
25 are where we are today. Roughly we can split AI

1 research into some first-wave AI, second-wave AI, and
2 then maybe 2.5 or third-wave, which is where we are
3 right now.

4 In this first wave, primarily, researchers
5 focused on what is called search. So this is either
6 searching through a potential solution space, some
7 quick examples, chess is a good example here where we
8 had, say, Deep Blue beating Kasparov via a
9 sophisticated algorithm that did search through using
10 domain-specific heuristics, expert knowledge, for
11 instance. Folks who played a lot of chess encoded
12 heuristics into the search algorithm; it would search
13 through the solution space to find, say, the next move
14 to play.

15 Now, another hallmark of first-wave AI is
16 something called expert systems. And this also relies
17 on basically bringing in a lot of expert knowledge to
18 form some sort of large database of rules, of
19 knowledge, of facts about the world, using some sort
20 of inference engine, typically based on logical
21 reasoning, to make new sort of conclusions based on
22 these facts, and then some sort of action, I/O
23 system to interact with the human. So this is
24 basically the world up until maybe the '80s in
25 terms of AI.

1 Now, there were some large successes here,
2 so one example that I used earlier is this chess
3 champion falling to basically a sophisticated search
4 algorithm. And there are many more. And, in fact,
5 techniques from first-wave AI are still used in
6 practice, but they're decidedly brittle and they
7 really don't have any real learning capability. So
8 they're really sort of a function of just the
9 knowledge that you encode into them.

10 There's a huge overhead to encoding that
11 knowledge. Right, I have to ask, say, every member of
12 the audience and everyone watching to tell me all the
13 facts that they know about the world and then I have
14 to store that somehow, and that might be brittle and
15 that might not generalizable. They're very, very
16 brittle systems, but they do allow me to do in-depth
17 specific reasoning. Right, if I ask a bunch of
18 experts for facts on a specific vertical, then I can
19 do a lot of fast automated reasoning about just that
20 vertical. So that can be good, but it's very
21 difficult to generalize.

22 And if you recall back to that earlier
23 slide, we want generalizability, we want abstraction
24 because we want to create some system that's able to
25 encounter new environments and still act in a

1 reasonable way.

2 So in terms of those four boxes, first-wave
3 AI generally does perception reasonably well in the
4 sense that I have asked all audience members to give
5 me facts, and it can do reasoning and acting
6 reasonably well, but it won't learn and it won't
7 generalize.

8 Now, there were some transition points in
9 multiple areas of sort of AI research. One of these
10 is something called natural language processing, which
11 says, can I get a computer to ingest, say, raw text or
12 can I get Alexa to ingest signal from your voice and
13 then have it understand that in some sense. So in
14 natural language processing, up until about the late
15 1980s, most of the rules for doing this sort of
16 translation or understanding were done via hardcoded
17 sort of expert rules.

18 Around the late '80s, probabilistic models
19 started to come into play. Okay, so this is going to
20 sound more like machine learning like folks have maybe
21 heard about in the press. These are models that
22 ingest, in this case, large text corpora and learn
23 patterns in that data.

24 To look at a different vertical in AI, so
25 autonomous vehicles rely heavily on something called

1 computational vision, which says, hey, I have a video
2 image, can I understand what's going on in that image.
3 Say I'm a car and I'm driving along, and I have a
4 still image of the road in front of me, can I
5 understand that there's a stop sign and a pedestrian
6 and dog in front of me and so on. So in autonomous
7 vehicles, in the mid 2000s, ARPA ran what they call
8 a Grand Challenge, in fact their first Grand
9 Challenge, which asks, can I create a vehicle that
10 can drive some hundred-plus miles across the desert
11 autonomously?

12 In 2004, no vehicles completed this task.
13 In fact, I think the longest trip that a vehicle took
14 was something like ten miles. And these vehicles
15 relied heavily on hand-coded rules that say something
16 like, in general, when you're, you know, ten degrees
17 away from the sun and you're driving forward at a
18 particular speed, then a shadow is going to be a
19 shadow instead of a rock with some set of features
20 associated with it. And, again, this is a very
21 brittle system. This is not going to generalize very
22 well.

23 But then in 2005, five teams completed the
24 entire trip, so 100-plus miles. And this is because
25 they started using these probabilistic models. And,

1 in fact, you can see the general manager for the
2 program, Strat at the time, had a fun quote:
3 "[Vehicles] were scared of their own shadow,
4 hallucinating obstacles when they weren't there."
5 And this is for those prior systems. And then
6 probabilistic models allowed them to get around this.

7 So you can see similar transition points
8 throughout all core AI areas, in the late '80s, in the
9 '90s, up and through basically the mid-2000s. And
10 this happened because of three things. One is
11 computational power increased, and this is the story
12 of basically computation since the '40s or '50s. This
13 has played a driving role in AI development as well.

14 Number two, storage costs decreased. I
15 don't have to pay a lot of money to store a lot of
16 data. And, three, everyone in this world now relies
17 on statistical models, maybe with some expert input,
18 but still statistical models.

19 So this takes us into the second wave of AI,
20 and there's no hard date for this because it happened
21 differently in different verticals in this world.
22 Here, we're relying on this assumption now that we've
23 learned the hard way, multiple times, that encoding
24 all knowledge explicitly does not work. It doesn't
25 scale. It's very brittle and it's very difficult to

1 handle uncertainty.

2 The new idea is that we should create a
3 general statistical model for a problem domain. We
4 should create a statistical model for natural language
5 or for a type of natural language or for autonomous
6 driving, a type of autonomous driving. What do we do
7 with that model? Well, we feed in data from the real
8 world or maybe simulated data until it looks right.
9 And this is going to be characterized by statistical
10 learning.

11 So the reason why these models have taken
12 off is because if we input a different data set or,
13 say, set of data sets into these models, we'll learn a
14 different model and then we can deploy that in a
15 different environment. So it's much more
16 generalizable.

17 Now, some examples. In machine translation,
18 for instance, going back to this natural language that
19 we discussed earlier, we can feed in multilingual text
20 corpora to learn relationships between language. So
21 say we want to translate French to English, one of the
22 early multilingual text corpora came from Canada,
23 where there are rules stating that, say, any
24 government ruling has to appear both in English and
25 French. And so now we have a mapping between English

1 and French documents, we can feed that into a model
2 and we can learn a way to translate between the two
3 systems.

4 Autonomous vehicles. We can feed in videos
5 and tests of successful driving into a model and then
6 learn what scenarios are safe or not safe or maybe put
7 some error bars around what scenarios are safe in
8 general.

9 Face detection, face recognition. I can
10 feed in many labeled faces of people. Here is where
11 the face is, or here is where the face is and an idea
12 associated with that, to learn what a face looks like
13 or to learn what, say, your face looks like.

14 So these types of models are very good at
15 perception, and they're very good at learning.
16 Remember, we're training these models, these general
17 models, based on a data set, and if we feed in a
18 different data set, we're going to get a different
19 result, so they're reasonably good at abstraction and
20 generalization as well, so long as your model is
21 general enough and so long as you have enough data.
22 But there is no reasoning or acting. I've made no
23 statements about, say, when one should turn the car in
24 -- turn the wheel in the autonomous vehicle.

25 So a quick example model. Remember, these

1 are systems that rely on statistical learning to train
2 probabilistic models that will tell us something about
3 the world. A quick example is a neural network. So
4 these appear a lot in the news now, which is why I've
5 chosen them, but they're not a new idea. Indeed, that
6 1955 proposal where McCarthy defined AI, used the term
7 AI for the first time, also discusses neural networks.
8 I believe they were called neuron networks at the
9 time. So this is not a new idea.

10 The general idea of neural networks is that
11 one should pass information into this input layer,
12 which you see on the left side of the screen. So that
13 information could be pixels of an image. That
14 information could be something with audio signal. It
15 will cascade through the network, along basically a
16 series of pipes that go through nodes, and these pipes
17 have, say, different widths that can be controlled by
18 a learning algorithm.

19 And then the final layer of this network
20 that has information flowing through it will create
21 some sort of guess. In the case of, say, classifying
22 images, here we have cats and dogs, it's going to
23 create, say, a probabilistic model of whether or not
24 an image is a cat or a dog. And that gives you some
25 signal as to how good or bad your statistical model --

1 in this case a neural network -- is acting.

2 A very general model, so long as we can feed
3 information into it via that input layer and so long
4 as we can judge the output and so long as we can
5 actually learn, so make the network better, using
6 sophisticated optimization techniques, we can use this
7 for many problems and, indeed, that is what we've
8 seen, so long as we can, again, train these models
9 through repetitive sort of optimization algorithms.

10 So another sort of buzzword that one sees in
11 the press a lot is a deep neural network. Again, not
12 a new idea. These existed, I think, since the 1980s,
13 and they're just these neural networks that we had on
14 the last slide but with more, quote, unquote, hidden
15 layers. These are the layers in between that input
16 and that output. So I can add more and more of these.
17 I can create more piping -- intricate piping between
18 these different nodes to learn new patterns in the
19 data.

20 And sometimes, indeed, we can stack many,
21 many, many, many, many more nodes, so order of
22 hundreds of thousands, millions, et cetera. So these
23 are very large models. And, again, this is because we
24 have increased computational power and cheap storage.

25 That idea for deep networks has existed

1 since the '80s, but we've seen them taking off in the
2 last five to ten years because of advances in
3 hardware, because of a huge increase in the amount of
4 data that exists. So we have large firms collecting
5 data; we have the government collecting data; and we
6 can now store it cheaply, access it quickly, and
7 because, indeed, from the R&D community, there have
8 been much better methods developed for learning
9 basically how to make a good one of these.

10 They're hugely successful. They're good at
11 detecting anomalies in data, for instance, credit card
12 fraud. They're good at voice recognition. You've
13 seen Alexa, Siri, Google Assistant, et cetera.
14 They're great at machine translation, language
15 generation, game playing. Some recent high-profile
16 success stories such as AlphaGo playing basically
17 expert-level, Go, DeepStack Plane, expert-level Heads-
18 Up Poker.

19 Self-driving cars are starting to take off.
20 Video search, audio search, finance, et cetera. These
21 are all success stories in part due to deep learning.
22 Not a new idea, driven by advances in hardware and
23 training them.

24 Nobody understands why they work very well,
25 and this is a common story in AI as well and this is

1 something that we're seeing more and more appearing,
2 which is humans have sat down, they've designed the
3 network structure, they've designed what those nodes
4 and what those connections between the nodes look
5 like. Maybe they're encoding some domain expertise.
6 There are some known heuristics that you can rely on.
7 There's a trial-and-error process, and maybe actually
8 other AI is actually coming in and trying to train
9 these models or structure these models in a better
10 way, but nobody knows when or why they don't work in
11 general.

12 So they work well in expectation, which is
13 why we see machine translation systems, which is why
14 we see Alexa and Siri in households now, but when they
15 fail, it can be very confusing, it can be reasonably
16 catastrophic, and it can be very hard to explain.

17 And some recent research pushes funded by
18 the DOD, funded by industry, funded by nonprofits,
19 have started noticing that, hey, an adversary can
20 exploit this kind of behavior. When I have a system I
21 trust most of the time but it can be exploited in very
22 odd ways and I don't understand why or when that
23 happens, then I can wreak some havoc in these systems.

24 So I'd like to take a step back. So now
25 we've talked about deep learning, we've talked about

1 machine learning, and we've talked about AI. And,
2 roughly, AI is this sort of four-pillar approach to
3 perceiving the world, learning about it, building an
4 abstract and general model, and then using that to
5 act and reason. Machine learning is just one way to
6 build these models, where we do not focus on acting
7 and reasoning but we focus on perception, on learning,
8 on abstraction, and on generalization. And deep
9 learning is just a specific form of basically
10 representational learning, so it's a type of machine
11 learning.

12 Right, so every time you hear deep learning
13 in the news, you can replace it with machine learning
14 mentally. It's just a way to solve a machine learning
15 problem.

16 So some present-day movements in AI,
17 understanding bias and methods for debiasing. You'll
18 hear about this I think throughout today and tomorrow,
19 many of the topics on this slide. So this is sort of
20 a teaser. Understanding bias and methods for
21 debiasing. So if I feed skewed training data into
22 these systems -- remember, these are statistical
23 models that are trained on data from somewhere in the
24 world. If I feed skewed data into the system, then
25 I'm going to learn something that represents that

1 skewed data. So how do we understand when that
2 happens and can we create systems that still feed in
3 this biased data which might be the only data that
4 exists but spits out a model that is debiased?

5 As mentioned before, adversarial reasoning
6 in multi-agent systems, learning to act with
7 cooperative actors, learning to act with adversarial
8 actors, so bringing in older fields such as game
9 theory into these new methods for solving those
10 problems. How do I say design -- well, I'll talk
11 about this in a few slides, but how do I design
12 policies as a firm to compete with other, say, firms
13 that are both cooperative and adversarial? Can I do
14 this based on machine learning?

15 Also mentioned on the previous slide,
16 robustness to noise, robustness to adversarial
17 attacks, both in terms of theoretical robustness and
18 empirical robustness. How do I design automated
19 systems that fail less, that are robust to attacks and
20 that fail more predictably, because obviously these
21 systems will always fail at some point.

22 And in that vein, explainable AI, there's a
23 lot of money going into this as well because it's very
24 difficult to interpret the results that come out of
25 these systems from time to time, so can we produce

1 human-understandable models that also work well?

2 And one final move in the AI community has
3 been reinforcement learning. It's a type of machine
4 learning, but it's a type of machine learning that
5 also focuses on learning to act and reason. So now
6 we're getting closer to that initial definition of
7 artificial intelligence. Here we have an agent, maybe
8 physical, maybe virtual, that's going to act within an
9 environment. It's going to receive a reward signal
10 and then maximize total reward. It wants to find the
11 actions to take for any state in the world such that
12 when it takes that action, it is treated well in the
13 future, it receives reward and expectation in the
14 future. And I'll give you some examples of this at
15 the end of the talk.

16 So here again, again, reinforcement
17 learning, not a new idea, but deep networks have been
18 used extensively here to revolutionize their use and
19 practice. So here we have deep networks that are used
20 to, say, reduce the complexity of representing the
21 environment. Remember, I can't actually write
22 everything down, I don't want to represent every
23 single fact in my computer, so now I'm going to learn
24 some abstraction of the world and then act on that.

25 So reinforcement learning is taking us

1 closer to what we want to call AI. We have
2 perception, we have learning. These are just like
3 machine learning. We have abstraction and
4 generalization, again, moving toward that. Again, if
5 we train these models on different data, we get a
6 different trained model, and we're starting to move
7 toward reasoning and acting here.

8 So in the context of this audience, I
9 thought I would do maybe a quick deep dive into a few
10 uses of AI, particularly in something called market
11 design. So markets provide agents the opportunity to
12 gain from trade. Many markets require structure to
13 operate efficiently. Market design is going to tackle
14 this problem via what's called economic engineering.
15 So I put on my economist hat and I put on my
16 engineer's hat and I put on my mathematician's hat.
17 I'm wearing three hats at this point, but I can use
18 these hats to design a market, how do I structure the
19 market, how do I constrain the market such that I
20 achieve some sort of efficiency goals.

21 AI is increasingly helping with the design
22 of these markets. For instance, these automated
23 methods can use data to help designers characterize
24 families of market structures. They can be used
25 obviously for predictive methods that anticipate, say,

1 future supply and demand in electricity markets or
2 finance markets.

3 One example, as a computer scientist, this
4 is close to my heart because a lot of the money in our
5 world comes from this, is using AI in online
6 advertising. So online advertising markets generally
7 match advertisers with consumers. Many billions of
8 dollars, and this is an increasing market, many, many
9 billions of dollars are being used here, and it's a
10 driving force in the technology sector.

11 Machine learning models in this space right
12 now are being used to divide customers into very fine-
13 grained and automatically generated segments. So no
14 longer just male/female but something far, far more
15 fine-grained than that. That's learned automatically.
16 They're being used to set reserve prices and auctions
17 based on user modeling and bidder behavior, again
18 automatically.

19 They're being used to automatically generate
20 the creatives, that is, the artwork that you see pop
21 up on your screen, to automatically generate those,
22 say without human input, to fit a specific customer's
23 predicted wants. All automated.

24 Reinforcement-learning-based tools --
25 remember, this is that form of machine learning that

1 also focuses a bit on acting -- are being used to help
2 advertisers, for instance, bid automatically on these
3 very fine-grained segments. Remember, now we have,
4 say, millions of segments. How do I bid on that, I
5 can use a machine-learning-based model to do this.

6 Another example, AI in electricity markets.
7 Here, matching supply and demand is extremely
8 important. It relies heavily on demand forecasting.
9 Machine-learning-based techniques are going to provide
10 very accurate demand forecasting, which leads to very
11 stable market prices and more efficient power usage.

12 Reinforcement-learning-based techniques --
13 remember ML plus some form of acting -- are going to
14 allow us to activate or deactivate expensive
15 heterogeneous power sources to maintain that
16 stability. So I can predict better demand, I can
17 predict better demand at particular time points
18 further into the future, and then I can make a plan to
19 boot up or boot down particular power sources such
20 that I maintain market stability, such that I reduce
21 brownouts and so on. Again, automated.

22 And my final example is AI and kidney
23 allocation. This is close to my heart. I've done a
24 lot of work in this space. So here, kidney exchanges,
25 for instance, are an organized market where patients

1 with end-stage renal disease enter and are able to
2 swap donors -- willing living donors -- to receive new
3 organs.

4 It's a really interesting paradigm that's
5 been around for, say, 15 years now, and it accounts
6 for something like 10, 11, 12, 13 percent now of all
7 U.S. living donations of kidneys. Hundreds of
8 transplant centers are involved in this organized
9 market, in fact, multiple organized markets.

10 And, here, AI-based tools are also
11 operating. Now, this isn't fully automated, but
12 they are, for instance, semiautomatically and
13 optimally subject to human value judgments, matching
14 donors to patients, both in the U.S. and also
15 worldwide. Here, I've called out the United Kingdom
16 and the Netherlands, but in many countries. They're
17 providing sensitivity analysis at a level that humans
18 cannot for new policies. And they're learning from
19 data the quality of, say, potential matches in this
20 market.

21 Now, let's return to some open questions
22 and some recent pushes which will, I guess, trigger
23 good discussion for the rest of today and tomorrow.
24 So one is, how and why does deep learning work? So
25 I've mentioned not a new idea. Neural networks

1 existed since the '50s; deep learning existed since
2 the '80s. Now we have new hardware and now we have
3 new training techniques, these tend to work very well
4 in expectation, but when they fail, they fail
5 confusingly. Why do they work?

6 How can we handle incentives of competing
7 agents? All those three market examples that I showed
8 you, firms are obviously going to compete against each
9 other in this space. The government, regulatory
10 agencies have their own incentives as well.
11 Individual participants have their own incentives.
12 How do we handle this, how do we encode other aspects
13 such as fairness, accountability, and transparency
14 into these systems?

15 How do I ensure that my automated system
16 doesn't marginalize, say, an already marginalized
17 class in the ever sort of increasing hunt for
18 efficiency? How do I even define this? How do I
19 define fairness? This is a classic question in
20 economics that computer scientists are now starting to
21 struggle with as well. How do I implement this in a
22 scalable way, in an understandable way?

23 Ethical AI, this will be talked about, I
24 believe, later, by folks like Henry Kautz, how do I
25 divide the labor between policymakers, such as those

1 in this audience, who are ethically trained and
2 ethically minded and technically trained, perhaps
3 ethically neutral AI and machine learning experts?
4 So I can implement, say, a very sophisticated system,
5 but I need to then produce some sort of aggregate
6 output that I can pass back to policymakers to ensure
7 that this is reflecting the aggregate human value
8 judgments of those who control the systems. How do I
9 do that? And there are close ties in this sort of
10 exploration to the world of privacy and the world of
11 social norms.

12 So in general, our end goal is to create
13 these systems that perceive the world, learn from it,
14 create some sort of generalizable model and then
15 inevitably learn to act using that model. We're not
16 quite there yet, but there's a lot of hope in this
17 space. But I'm going to say that maybe this isn't
18 even the biggest problem. The biggest problem is
19 going to be the interplay between these systems and
20 society, ethical issues, societal norms, human value
21 judgments. How do we play between, say, these sort of
22 sophisticated machine-learning-based approaches to
23 what I've shown here on this slide and the rest of the
24 real world? So I'll leave it at that.

25 MS. GOLDMAN: Thank you very much, Professor

1 Dickerson, for that excellent introduction to the
2 field and for the questions that will be coming
3 throughout this hearing.

4 (Applause.)

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

1 In particular, following on the last
2 speaker, in the last few years, I've been thinking
3 quite a bit about ethical and social issues in the use
4 of machine learning and algorithmic decision-making
5 more generally. And I also saw that there are some
6 discussion or a panel about sort of competition and
7 marketplace questions introduced by machine learning.
8 I hope to make some less technical remarks about that
9 because I think that's less scientific to say there
10 but a lot of interesting things to discuss, and also
11 relatedly topics related to consumer protection and
12 abuses by machine learning and AI.

13 And so what I just want to do with my time
14 is make some informal remarks, provide some personal
15 opinions on these topics based on my own experiences
16 and research, and, you know, hopefully cue things up
17 for the next couple of days for the rest of the
18 speakers.

19 So as the last speaker mentioned, there has
20 been a lot of discussion not really first in the
21 technical community but first in the mainstream media
22 and society at large, about the many things that can
23 go wrong when applying machine learning and AI and
24 related methods to algorithmic decision-making.

25 And before I describe -- say a little bit

1 more about what can go wrong and what we might hope to
2 do about it scientifically, I thought I would start by
3 sort of just framing how things can go wrong in the
4 first place. And so one thing you might wonder is,
5 you know, if there is a lending model or a consumer
6 credit-scoring model that exhibits racial bias, for
7 instance, or there's some data analysis or machine
8 learning methodology that leaks personal, private
9 consumer data, you might -- it's a reasonable thing to
10 wonder whether this happens through active
11 malfeasance. You know, are there evil programmers at
12 tech companies who, you know, put in a line in their
13 code that says if the person's race is this then do
14 this; if it's some other race, then do something else;
15 or whether they program back doors into their code
16 that permit privacy leaks.

17 And there's good news and bad news here. My
18 strong belief, and I think those people who work in
19 the field would say that, no, there is absolutely no
20 such malfeasance going on by evil programmers at
21 technology and other companies. So that's the good
22 news. The bad news is that the truth might actually
23 be a little worse, which is these sort of collateral
24 damage or consequences are actually the natural
25 byproduct of applying the formal, fundamental,

1 scientific principles of machine learning and AI. And
2 I'll say a little bit more about that.

3 The vast majority of what I think we would
4 call algorithmic decision-making is actually a little
5 bit more specifically almost always driven by machine
6 learning these days. So, in particular, when you
7 think about the algorithms that make things like
8 lending decisions or decide what ads to show you on
9 Facebook or Google, these generally are not what you
10 should think of as hand-coded or programmed
11 algorithms, but rather they're the result of taking
12 data, you know, historical data, whatever that means
13 in a given domain, giving it to an algorithm, and that
14 algorithm, of course, trains a model on the data. And
15 then at the end of the day, it's the model that's
16 actually making the decisions. It's the model that's
17 actually deployed in the field.

18 And, typically, the algorithm that
19 transforms the data into a model is actually
20 tremendously simple and very principled from a
21 scientific standpoint. So if I had slides, one thing
22 I like to do in forums like this is put up the
23 Wikipedia pseudocode for the so-called back
24 propagation algorithm for neural networks which the
25 previous speaker mentioned. And that pseudocode is

1 literally a simple loop with about ten lines of code
2 in it.

3 And a real working version of it wouldn't
4 be that much more complicated. And it's doing the
5 most obvious thing you can possibly imagine, which
6 is essentially going through the training data and
7 adjusting the parameters or nodes of the model in
8 order to minimize some -- you know, usually accuracy
9 or error-based cost function, okay? So that
10 algorithm is not opaque at all. It's entirely
11 transparent.

12 Sometimes, you know, when I talk to people
13 who aren't in the field, they naturally assume that
14 machine learning algorithms -- you know, the code for
15 them might look like something like I imagine the code
16 to a video game like Grand Theft Auto looking, you
17 know, hundreds of thousands of lines of spaghetti code
18 with all these special cases and details, and it's not
19 like that.

20 So, then, the natural question to ask next
21 is if the complexity doesn't lie in the algorithms
22 themselves then where does the complexity creep in?
23 And, of course, it's from the interaction of the data
24 being processed to produce a model mediated by these
25 very, very simple algorithms, okay? And so the

1 problems arise these days not so much from the
2 algorithms themselves, which, again, are very simple
3 and operating on very basic, kind of well-motivated
4 scientific principles, the problem is really when
5 you work in extremely large complicated model spaces,
6 of which, you know, deep learning is just one and
7 perhaps the most recent example, the sort of space of
8 models has a lot of sharp corners in it, as I might
9 put it, which allow to you kind of optimize the thing
10 that you're trying to optimize like minimizing the
11 error on the data at the expense of other things that
12 you didn't explicitly ask for like fairness or
13 privacy.

14 And I think if there's one kind of important
15 adage to understand about machine learning, it's that
16 basically modern machine learning will not give you
17 for free anything that you don't ask for and specify.
18 And in general, you shouldn't expect it to avoid
19 things that you don't want that you didn't tell it you
20 didn't want. Okay?

21 And this is, I think, the source of a lot of
22 the kind of violations of social norms and values that
23 we've seen by machine learning and AI in recent years.
24 So that's a little bit about what can go wrong. Now
25 let me talk a little bit about -- sorry, that's a

1 little bit about how things can go wrong.

2 And, so, with that background, I want to
3 talk about, well, what are the different things that
4 can go wrong, and, most importantly, what can we do
5 about them from a kind of scientific standpoint. So,
6 you know, the things that can wrong are things that
7 I've mentioned already, which is violations of things
8 like privacy or fairness or interpretability and
9 transparency, or even safety or morality, if you like.
10 You know, the sort of logical extreme of this for
11 those of you who've heard of it is, you know, this
12 sort of parlor game or science fiction thought
13 experiment known as the singularity on which AI, you
14 know, sort of -- AI achieves superhuman intelligence
15 to the point that, you know, for lack of a better
16 term, the robots become our overlords.

17 While that's a fun thing to think about, I
18 don't know many sane people in the machine learning
19 community who actually think that that's our sort of
20 gravest technological risk anytime soon. All you need
21 to do is come and see what AI and machine learning can
22 actually achieve right now and compare it to humans or
23 other biological species and you will be deeply
24 underwhelmed by what we can accomplish so far. But
25 violations of social norms are, like, already with us

1 now today and on a very large scale, whether we are --
2 whether we know it or not or whether we're measuring
3 them properly or not.

4 And, you know, I think it's important to say
5 to this audience that I think I and many of my
6 colleagues, you know, we do believe that better laws
7 and better regulations are possible and should be
8 developed. And I'm sure that that's being worked on
9 and is a necessary activity. But I think my opinion
10 is that that will be woefully inefficient in the
11 algorithmic era to actually keep up with the types of
12 violations of social values that we're seeing because
13 it just -- you know, basically human organizations
14 don't scale, and you can't sort of expect to police
15 the sort of violations I'm talking about with sort of
16 regulatory agencies that are pouring over the
17 decisions made by algorithms on a sort of a human time
18 scale and hope to keep up.

19 So an alternative approach, which I'm a
20 great advocate of and as are a growing number of
21 people who do technical work in these areas is to
22 design better-behaved algorithms in the first place
23 and to actually endogenize various notions of social
24 norms inside of our algorithms and asking that our
25 algorithms -- that the actual code obey some

1 definition of privacy, some definition of fairness,
2 some definition of morality, if you like.

3 And, of course, this leads immediately to
4 two very difficult questions. The first difficult
5 question is, you know, how do you define these things
6 as the last speaker said. How do you define
7 algorithmic fairness, how do you define algorithmic
8 privacy? And, then, if and when you can come up with
9 such a definition, it's going to come at some cost,
10 right.

11 So if I have some notion of fairness in
12 models that are used to provide criminal sentencing
13 guidelines, my asking for fairness from that model by
14 gender or race will come at a cost of accuracy. What
15 I'm saying is like a tautology. If I sort of -- if I
16 ask myself to find the model in some space of models
17 which minimizes the error period, and then I ask to
18 find the model that minimizes the error subject to
19 your favorite definition of fairness, the error can
20 only get worse.

21 And so in a model like -- let's say in
22 a setting like criminal sentencing, this means that
23 a cost to accuracy might mean sort of, you know,
24 hard things to think about. It might mean
25 incarcerating more innocent people, or it might mean

1 letting more guilty people go free. So there will
2 be societal and technical costs to imposing these
3 sorts of constraints on our algorithms, but I think my
4 view and the view of many people in the field is that
5 we have to go down the road, we have to decide
6 algorithms that incorporate these values, we have to
7 talk about what the possible definitions for these
8 values are, and we need to study these tradeoffs
9 between the thing that people optimize for in machine
10 learning, which is accuracy, and the tradeoffs to
11 different social norms.

12 Okay. And so what I want to do with most of
13 my remaining time is just tell you a little bit about
14 the sort of very active research that's going on in
15 the computer science and machine learning and related
16 communities on this scientific agenda, sort of picking
17 definitions for different social values or norms and
18 actually encoding those norms inside of our algorithms
19 and then trying to study what the tradeoffs will be
20 with, you know, things like accuracy and other more
21 standard objectives.

22 So let me first talk about the work that
23 goes on in the area of privacy in machine learning,
24 and not just in machine learning but more generally in
25 kind of data analysis and data science. And I think

1 it's helpful to say just a little bit about the
2 distinction between what I'm thinking of as privacy
3 and a closely related and complementary area, which is
4 that of security and cryptography.

5 So security and cryptography, to a first
6 approximation, is a technology about keeping data
7 locked down. It's about controlling access to data
8 and making sure that people who shouldn't have access
9 to data don't get that access by basically hacking
10 into a system that they shouldn't hack into. And this
11 is largely the domain of security and cryptography,
12 and that's one notion of privacy. That's sort of
13 control of your data and making sure it doesn't get --
14 you know, it doesn't get accessed or stolen by people
15 who shouldn't.

16 Here, I'm talking about something a little
17 bit different and more nuanced but in many ways is
18 equally as pervasive and important as notions of
19 security, which is the fact that, you know, you have
20 all of this data that's being collected by different
21 companies and agencies and other organizations. And
22 you might worry about what -- not just sort of, you
23 know, how -- who's accessing that data but what can be
24 inferred about you from that data that isn't directly
25 in the data itself.

1 So the kind of thing that I'm concerned
2 about here is that if your medical record is used as
3 part of a study to build a predictive model, let's
4 say, for some disease based on symptoms, and then
5 that model is published or used in the field, could
6 it be that the use of that model or the publication
7 of the model, perhaps combined with other publicly
8 available data sets, actually reveals a great deal
9 about your own personal medical status and record.
10 Okay?

11 And, you know, if you go down the road of
12 thinking about possible technical definitions of this
13 type of privacy, I believe that most of you would
14 eventually come to two kind of, I think, important
15 conclusions or desiderata from any sort of privacy
16 definition for machine learning or data science.

17 One is that, you know, you need to account
18 for the fact that any particular data set that you
19 want to, you know, keep private in some technical
20 sense, will not be the only data set in the world.
21 And, in particular, that data set might be combined
22 with other data sets that you don't know about or
23 didn't foresee or don't even exist yet but might exist
24 in the future.

25 And one consequence of this that I will

1 state without proof is that this means that any
2 definition of privacy that it involves anonymization
3 is essentially a flawed definition of privacy, right,
4 because anonymization refers to taking the data set
5 that's in front of you and doing things like
6 eradicating personally identifiable information.

7 But the literature and news is, you know,
8 rife with examples where you anonymized one data set,
9 somebody else anonymized a second data set. Those two
10 data sets were combined and then maybe combined with
11 some publicly available information, and your specific
12 data could be backed out of that. You could be, as we
13 like to say, reidentified, or the data set could be,
14 you know, deanonymized as they say.

15 And, you know, I think many people feel
16 strongly enough about this assertion that there is
17 sort of a saying in the field, which is, you know,
18 anonymized data isn't, meaning that, you know,
19 whatever you think you did to deidentify individual
20 identity in a data set, that can often be undone
21 through the unforeseen combination of that data set
22 with other data sets.

23 The other, I think, sort of axiom for any
24 definition of privacy that's important is that in
25 order to have a definition of privacy that still

1 allows to us do anything useful with data, it's
2 important to isolate, you know, the potential harm
3 that comes to somebody as the result of use of
4 their data in some analysis or model-building
5 exercise versus the harm that might come to them
6 just because data analysis reveals some facts about
7 the world.

8 So, for instance, if you were a smoker in
9 the early 1950s before there was discovered a link
10 between smoking and lung cancer, well, when somebody
11 did data analysis and discovered that there was a
12 strong correlation between lung cancer and smoking and
13 you were a smoker, that fact does harm to you, but it
14 doesn't matter whether your data was used in that
15 analysis or not, right?

16 Researchers were going to discover this fact
17 whether your particular data was used or not, and a
18 harm was done to you by the fact that suddenly it's
19 revealed that smoking is bad for your health and you
20 were a smoker and everybody knows it. But that harm
21 was not done to you specifically as a result of the
22 data analysis using your data or not. You were
23 basically -- you know, this harm was going to be done
24 to you whether your particular medical record went
25 into those studies or not.

1 And so there is a very rich definition of
2 data privacy known as differential privacy that was
3 introduced about 15 years ago and has received a
4 great deal of scientific attention in the interim,
5 and now has a very rich theory and a very rich body
6 of algorithms that basically on the one hand meet
7 this sort of very strong notion of data privacy which
8 has to foresee the possibility of triangulation
9 through the combination of multiple data sets on the
10 one hand but still permits sort of powerful use of
11 data.

12 And so, you know, one kind of pseudo-theorem
13 that I will state to you is that everything that we
14 pretty much know to do today in machine learning we
15 know also know how to do in a differentially private
16 way. And it's just a matter of companies adopting
17 this technology and choosing to, you know, do their
18 machine learning and data analysis in a differentially
19 private way. And we're actually starting to see
20 large-scale deployments.

21 Both Google and Apple use differential
22 privacy in meaningful, large-scale ways in some of
23 their services, and maybe more importantly, the 2020
24 U.S. Census, every single statistic or report that
25 they release as the result of the 2020 Census they are

1 planning to do so under the constraint of differential
2 privacy. And so this is an example, I think, of a
3 very promising kind of case study, right? Of course,
4 people have thought about different definitions of
5 privacy and data analysis and machine learning for a
6 long time. There was a struggle to kind of come up
7 with the right definition. Many of us believe that
8 sort of definitions based on anonymization are
9 fundamentally flawed.

10 But then a better definition came along
11 around 15 years ago. There's been a huge amount of
12 research on this particular definition, and, you know,
13 the good news is that in this particular -- for this
14 particular definition of privacy and this particular
15 social norm, it is possible to sort of give these very
16 powerful guarantees at not too great a cost to
17 accuracy or computational efficiency and the like. We
18 can sort of, you know, have the best of both worlds,
19 if you like.

20 So let me say a few words about research in
21 fairness in machine learning and algorithmic decision-
22 making, which is much more recent. It's a much more
23 nascent field than the study of privacy and machine
24 learning and AI, but we already know a fair amount
25 about it. And one of the things we already know about

1 it is that it's going to be a little bit messier than
2 privacy. So my claim is that if you waded into these
3 literatures and you looked at the work that's gone on
4 in differential privacy and looked in particular at the
5 definition of differential privacy, you perhaps, like
6 many people, might sort of agree that this is sort of
7 the right definition of privacy.

8 So we already know that there's not going to
9 be a right definition of fairness in algorithmic
10 decision-making and machine learning. And what do I
11 mean by we know there's not going to be a right
12 definition? So there's already from the last several
13 years several examples, several papers which have
14 results of the following form. They basically say,
15 well, you know, whatever your definition of fairness
16 is, wouldn't we all agree you'd want it to have at
17 least the following three properties. And you kind of
18 look at those three properties and you'd say yes, yes,
19 I would definitely want any definition of fairness to
20 at least meet those three properties and probably
21 other stuff, too.

22 And, then, of course, the punch line of
23 these papers is, well, guess what, here's a theorem
24 proving to you that you cannot simultaneously
25 achieve all three of those properties in any

1 definition of fairness. Okay, so those of you that
2 are -- have kind of an economics or social choice
3 background might know about, like, arrows and
4 possibility theorems for sort of voting systems.
5 These have a very similar flavor.

6 And these -- and so this has very concrete
7 consequences. So in particular, a typical notion of
8 fairness in machine learning would ask for the
9 approximate equality of false positive or false
10 negative rates across different groups. So let me
11 give an example. Suppose you're trying -- you know,
12 suppose you're a mortgage company and you're trying to
13 build a statistical model that tries to predict from
14 people's loan applications and credit history whether
15 they will repay a loan if you give it to them or not,
16 okay? A very natural thing to want to do. And you
17 want this model so that when people apply you can make
18 a prediction about whether they'll repay or not, and
19 then you want to give loans to people that will repay
20 you and not give loans to people that you predict
21 won't repay you.

22 But because this is a statistical model,
23 you're going to make mistakes. You're going to have
24 both false positives and false negatives, right? And
25 we might think of false negatives as really the case

1 that causes harm to the consumer in question, right?
2 False negative being somebody who's creditworthy would
3 have repaid the loan if you didn't give it to them,
4 but you decided to reject them, okay?

5 So we might think of false negatives as a
6 harm inflicted on a consumer, and false positives are
7 sort of the people that got lucky. So a typical
8 definition of fairness would basically say that, look,
9 you're going to make false -- you're going to make
10 false rejections. We're not going to try to prevent
11 that, but across different racial groups, it cannot be
12 the case that the false rejections rates differ
13 wildly. It cannot be the case that the rate at which
14 you falsely reject qualified African American
15 applicants is three times the rate at which you
16 falsely reject qualified white applicants, okay? So
17 this is a very natural constraint. And these
18 impossibility theorems basically say if you ask for
19 that and you also ask for a quality of false
20 positives, i.e., the people got lucky, and one other
21 related condition, you cannot simultaneously achieve
22 all of these.

23 So what this means is that we already know
24 that in fairness we're going to have to simultaneously
25 entertain multiple competing definitions of fairness.

1 And so not only will there be sort of tradeoffs in
2 competition between fairness and accuracy, there is
3 even going to be competition between different notions
4 of fairness. If you optimize for one notion of
5 fairness or constrain for one notion of fairness, you
6 might be damaging or losing on another notion of
7 fairness, okay?

8 But nevertheless, you know, we know this and
9 we have to proceed anyway, and so there's been a great
10 deal of research in the last several years on
11 algorithmic fairness and on different notions of
12 fairness and what the tradeoffs between that
13 particular notion of fairness and accuracy is. And
14 this is an area where, you know, to again echo
15 something the previous speaker said, when you sort of
16 talk about the potential interfaces between technical
17 people and policymakers and other stakeholders, I
18 think there are very, very good starting points.

19 So one thing you can do, for instance, is if
20 you pick a particular definition of fairness like
21 approximate equality of false rejections in a lending
22 application, and you have a data set in front of you,
23 a historical data set of people who did and didn't
24 repay loans, you can trace out empirically -- I can
25 give you -- I would have shown this slide if I'd met

1 the deadline -- I can actually show you an empirical
2 tradeoff where on the X axis would be the error -- the
3 predictive error of your model. On the Y axis would
4 be a numerical measure of the extent to which you
5 violated this fairness notion, so 5 percent would mean
6 -- sort of 5 percent unfairness means that let's say
7 between African Americans and other races there's as
8 much as a 5 percent disparity in the false rejection
9 rates. And I can just trace out a curve for you that
10 shows you the menu of choices you have.

11 So you can get the lowest error, but, you
12 know, if you sort of ignore fairness entirely, that
13 will give you the lowest error but it will give you
14 the worst unfairness. At the other extreme, I can
15 demand that the false rejection rates differ by 0
16 percent across populations. It's a very strong
17 constraint. And I will get the worst error but the
18 most fairness, and in between you'll get something in
19 between.

20 And I think this is the type of, you know,
21 sort of quantification of the tensions that we face as
22 a society in making these kinds of decisions that's
23 the right at least initial interface between, you
24 know, people like me and people like you for lack of a
25 better term, right, because it sort of really shows

1 starkly the choices that you have available.

2 Just to say a little more about fairness,
3 most definitions of fairness, like the ones I've been
4 discussing, actually only hold at the group or
5 aggregate level. So you're only making promises to
6 sort of groups of people in a statistical sense, and
7 you're not making promises to individuals, so, you
8 know, sort of more prosaically, if you are a -- you
9 know, if you're a person of a particular race that was
10 falsely rejected for a loan, you would have repaid
11 that loan, the consolation that you have in these
12 types of definitions is, like, well, we're also
13 rejecting people from other races who would have
14 repaid their loans at the same rate that we're
15 rejecting people from your race, which is sort of cold
16 comfort if you're somebody who was mistreated in this
17 way.

18 And so a lot of recent research, including
19 some of my own, is it trying to move towards
20 definitions of fairness and studying their algorithmic
21 implications that try to make finer-grained promises?
22 Maybe not all the way down at the individual level,
23 but to much finer-grained groups than just things
24 like, you know, race -- you know, top-level race or
25 gender or the like.

1 So these are two social values -- privacy
2 and fairness -- on which in relative terms we know
3 quite a bit already scientifically. And I think we're
4 well on the way to kind of developing both a science
5 and an engineering of designing better algorithms and
6 understanding what the tradeoffs are between accuracy
7 and the various definitions of the social values that
8 we've come up with.

9 My former grad student and colleague,
10 Jen Wortman Vaughan, is giving the keynote tomorrow.
11 And she's done a lot of recent research on
12 interpretability and transparency of machine learning,
13 which is another, of course, important social norm. I
14 think we know a lot less there so far, partly because
15 we just haven't had as much time, but one of the
16 problems with sort of coming up with satisfying
17 definitions of things like transparency and
18 interpretability is that there's fundamentally an
19 observer in kind of the middle of such a phenomenon,
20 right? So when you talk about interpretability, for
21 instance, of a statistical model, you have to talk
22 about interpretability to whom and what reason and in
23 particular the sort of numeracy of the audience you
24 have in mind will matter greatly, right, if we're
25 talking about interpretability to people with like a

1 statistics training, that means one thing. If we're
2 talking interpretability to the average American
3 citizen, well, you know, the average American citizen
4 has not been exposed to linear regression and may find
5 it a little bit bewildering to even talk at all about
6 an abstract mathematical mapping from loan
7 applications to lending decisions.

8 And so I think much of the research that
9 needs to happen on that topic will have to have like a
10 cognitive and behavioral element to it. You'll need
11 to do human subject studies with the type of audience
12 that you're interested in and ask them what they think
13 is interpretable to them or whether you can explain
14 models to them and the like.

15 So I'm almost out of time, but just to sort
16 of quickly touch on a couple of other things that I
17 saw on the agenda, I saw that there was one discussion
18 -- there was one panel title that had a very
19 intriguing name, which was Algorithmic Collusion. And
20 I'm not sure exactly what the context that's meant
21 there is. But, you know, if your concern is that, you
22 know, we might be entering an era where algorithmic
23 decision-making causes in some implicit or explicit
24 kind of large-scale way collusion between different
25 entities, whether it's on things like pricing or

1 decision-making and the like, I definitely think this
2 is already happening.

3 One area in which I'm very familiar with
4 this already is on Wall Street where quantitative
5 trading teams tried to build statistical models to
6 predict the directional movement of stocks and, so to
7 speak, beat the market. And my basic belief there is
8 that there's a huge amount of implicit sort of
9 collusion going on there, and it's really because, you
10 know, when we all use the same or similar data sets,
11 and when we all use the same or similar algorithms to
12 train our models, then even if we think we're clever
13 and independent and creative, we are going to be
14 strongly correlated just through the data, right?

15 If we're trying to predict the same thing
16 and we're using similar data sets and similar methods,
17 then no matter what else we do -- everything else
18 we'll do is a second-order effect from the fact that
19 the data itself will correlate us. And so I think
20 that this is an interesting topic on which there is
21 probably interesting scientific things to say but I
22 haven't thought about it yet, and I don't know of a
23 large body of research on it

24 But I'm out of time, so let me stop there
25 and let the agenda move on.

1 (Applause.)

2 MS. GOLDMAN: Well, thank you so much,
3 Professor Kearns, for that great overview and
4 introduction to all the issues that will be covered in
5 this hearing.

6 DR. KEARNS: Okay, thank you.

7 MS. GOLDMAN: So now it is 10:15, and we're
8 going to be taking a little break until 10:30, at
9 which time we will be back for the first panel.

10 (End of Presentation.)

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

1 UNDERSTANDING ALGORITHMS, ARTIFICIAL
2 INTELLIGENCE, AND PREDICTIVE ANALYTICS THROUGH
3 REAL WORLD APPLICATIONS

4 MS. GOLDMAN: So we're now going to begin
5 the first panel on Understanding Algorithms,
6 Artificial Intelligence, and Predictive Analytics
7 through Real World Applications. As I mentioned, in
8 case someone had just come in, I'm Karen Goldman. I'm
9 an attorney advisor in the Office of Policy Planning
10 at the Federal Trade Commission. And this is my
11 comoderator, Dr. Harry Keeling, who is an Associate
12 Professor in the Department of Computer Science here
13 at Howard University.

14 So we hope that this panel, which will cover
15 applications that are currently in use and on the
16 horizon, will provide a sense of the variety of uses
17 to which these digital tools can be put and highlight
18 that no single application is truly representative of
19 their use.

20 I just want to mention that anyone in the
21 audience who would like to ask questions of panelists
22 should write their questions on the notecards that are
23 being passed out and will be collected later on with.

24 With that, I'd like to introduce the
25 distinguished members of this panel. So we have Dana

1 Rao, who is an Executive Vice President and General
2 Counsel of Adobe. Next, we will have Henry Kautz, who
3 is the Division Director for the Division of
4 Information Intelligence Systems in the Directorate
5 for Computer and Information Science and Engineering
6 at the National Science Foundation.

7 Then we will have Angela Granger, who is
8 Vice President of Analytics at Experian. And then
9 Melissa McSherry, who is Senior Vice President, Global
10 Head of Data Products at Visa. We have Michael
11 Abramoff, who is the Founder and CEO of IDx, and
12 Professor of Ophthalmology and Visual Sciences at the
13 University of Iowa and also Professor of Engineering
14 and Computer -- of Electrical and Computer Engineering
15 and Biomedical Engineering.

16 And then we will have Teresa Zayas Caban,
17 who is the Chief Scientist at the Office of the
18 National Coordinator for Health Information
19 Technology.

20 So with that, Dana, would you like to begin
21 your presentation?

22 MR. RAO: Thank you.

23 Thank you. Thanks for being here. So the
24 first thing I wanted to just sort of get out there,
25 I'm a lawyer, and people are like, why are you talking

1 about AI, and I thought I'd put it out there because
2 there are some very distinguished computer scientists
3 on this panel. So I was actually an electrical
4 engineer undergrad and going to a university, and so
5 when I was at law school, I was going to write a
6 paper, a note for the journal, and this book was on my
7 dad's desk, Understanding Neural Networks. This was
8 back in 1996. So I thought, oh, that would be fun to
9 read, and I read it, and I wrote my paper, which got
10 published, called "Neural Networks -- Here, There and
11 Everywhere," which was a wildly inaccurate
12 characterization of where neural networks were in
13 1996. So don't come to me for your stock advice, but
14 it was -- it's been a fascinating topic for me, and at
15 Adobe, we're really interested in this topic.

16 And for us, AI is special because we have
17 this entire business that's focused on helping people
18 be creative. And creativity is a part of the brain
19 that doesn't follow rules. It's unstructured, and
20 traditional software programming is a very structured
21 form of algorithms. It's predictive. You understand
22 the rules, and you understand how to characterize it,
23 and that's actually not a great fit for creatives who
24 tend to break rules.

25 And so our products have always struggled to

1 bridge that gap between innovation and creativity and
2 the structure form of traditional computer
3 programming. And AI bridges that gap, and it really
4 allows us to create tools that are better for our
5 creative customers.

6 So when we think about how we look at AI and
7 digital creativity, we're really focused on minimizing
8 the mundane, eliminating those repetitive tasks that
9 everybody has in their day. And so for creative
10 professionals, there's a lot of complexity in the
11 tools and in setting up the camera shots or the video
12 shots that are not actually the highest value add that
13 they have, where they're really trying to get their
14 artistic sense across or fulfill the goal of a
15 marketing campaign as they create content for it, the
16 complexity of adjusting each pixel's luminance or the
17 color or the background or the lighting gets in the
18 way of them actually doing the part of the work that
19 they're getting paid to do. So that's where we're
20 really interested in using AI is it would eliminate
21 those mundane tasks.

22 And we also at Adobe, we've noticed there's
23 a huge demand for content now, and that's either
24 because there's social media channels and people are
25 posting content all the time on Instagram and Snapchat

1 and Facebook, or on ad campaigns -- digital media
2 advertisement campaigns where -- digital marketing
3 campaigns where you are personalizing content for each
4 consumer. So there's a huge demand for content, more
5 than ever before, and our creative professionals need
6 to be able to create content at a higher velocity, and
7 that's what AI is helping us do.

8 So when we think about AI, we think about it
9 in the creativity space in two buckets -- content
10 understanding, computational creativity. And Adobe
11 also has an experience intelligence business. I'm not
12 going to talk about that much today, but just for
13 transparency, we also have this other business that
14 also uses AI in a different way.

15 Content understanding is really trying to
16 get behind what's in an image, for example, or a
17 video. So it's easy to look at an image of a cat and
18 say there's a cat, or there's a house and just do sort
19 of basic object recognition. What AI allows you to do
20 is provide that insight into the image and add an
21 abstract layer, a conceptual layer above what you
22 typically can do pre-AI so we can understand things
23 like actions and concepts and styles and sentiments,
24 so just abstract concepts that are in your image that
25 the AI can infer from looking at it.

1 And so we have a couple of demos that we're
2 going to show. We're hopeful they're all going to
3 work correctly. I think this is going -- yeah, it's
4 going. And these -- this deck will be published in
5 the Adobe public policy blog, so anyone who wants to
6 see the full deck and watch the videos through can do
7 that. But we're just going to talk through a few -- a
8 couple seconds of these.

9 So this is a person in the, let me just go
10 back here. Set this up. So this is a person using
11 our stock photography service. And so they wanted to
12 start a creation. And so they wanted to be able to
13 say, I have an ad campaign for Nike, how should I
14 start. And they go to our stock photography site
15 and they just search for things to sort of -- as
16 inspiration for the ad campaign.

17 And so for example, in this example, this
18 person's going to say, you know, I see this image of
19 this woman with a ribbon jumping. That sort of
20 captures the aesthetic of what I want. And here we
21 go, here we go. And so she -- say they choose this
22 picture, and then what Adobe Stock does, it recommends
23 other pictures that are very similar to this picture.
24 So in this case, she says, okay, I like this, this is
25 a good start for me, and then Adobe Stock at the top

1 does sort of a normal picture recommendation. Here
2 are other pictures of people with ribbons, and that
3 may be what you're looking for.

4 But in this case, that's not what we want.
5 Like, Nike actually wants this sort of freedom. And
6 so we select the woman jumping, and our AI understands
7 that what we want is actually the action of jump.
8 Like that's what we want out of this picture, not the
9 color, not the ribbon, not the blue sky. We want the
10 action of jump. And so now we actually recommend
11 pictures that are about jumping.

12 So we can take the concept of that picture
13 and using AI understand, okay, they actually wanted
14 jumping, and so now we can just show them these other
15 pictures.

16 Now, the next level is we say, okay, well,
17 Nike didn't really want a picture of random people
18 jumping. It was actually supposed to be a family
19 picture. So we take family and we use the jump
20 concept from the first image, so you see how they're
21 stacked on the right, and now you have families
22 jumping. And now the creative professional could say,
23 that's where I want to start, I want to choose one of
24 those pictures and start my campaign from there.

25 So how do we do it? So what we do is our AI

1 will analyze these -- in this case, an image -- and
2 look for the concepts behind it. So you can see in
3 the middle, there's concepts, and on the right, there
4 are percentages. The percentages are the confidence
5 that our AI is actually accurately predicting what is
6 going on in there. But what you can see is we've
7 analyzed those faces and we've analyzed the context of
8 the picture, and you can see that where you said, oh,
9 there's happiness there, there's love there, there's
10 joy there, we've understood the abstract concept of
11 those pictures. And so you can go, if you're a
12 creative professional, and say I need pictures, my
13 theme is love, you can type in love as a search term,
14 and you're going to get a wide variety of images, but
15 they're going to have this concept in them.

16 You can also look for families, right? And
17 it will understand that the connection of these three
18 people plus the expression on their face means that
19 they're a family. And you can understand -- and you
20 can search for concepts like family as part of this.
21 And so you can see all the different kinds of
22 categories that you are able to search on using our
23 Adobe AI technology to analyze what is actually going
24 on inside the picture.

25 We also have a PDF and Acrobat service, you

1 know, and that has lots of text, and we've actually
2 run our AI on the text to understand the intelligence
3 behind the words. And we have married that up to
4 images to allow you to do automatic phrasing. And,
5 again, we can do very basic captioning. So you put
6 your photo there, and we can say couple on a bike and
7 that's object recognition. But then we see AI, then
8 there's a little slider you can see that's moving.
9 And you can say I want to see what the AI thinks this
10 is. And it says young couple on a bike, or in this
11 case, it said beautiful peacock, right? So it
12 understands not just the image but also the concepts
13 behind the image. So if you wanted to search for
14 "beautiful," you'd get that peacock, for example.

15 So these are the techniques that are being
16 used when we talk about content understanding, the
17 first part of how we looked at AI and creativity. You
18 know, traditional machine learning, it's traditional
19 deep learning, and when you look at all these things
20 like aesthetics and style and color, as part of the --
21 we train our AI to understand these concepts, and then
22 we are able to provide these services to our creative
23 professional.

24 The second piece of what we do is try to
25 make the creative professional day's faster. And

1 that's what we call computational creativity. And
2 that is trying to help their work flow. How do we
3 help them do those tasks even faster than they used to
4 have to do under traditional software? So here's an
5 example. Let's say somebody wants -- Macy's wants an
6 ad campaign and they told you to go out and shoot a
7 cityscape at night, and you go out and you spend six
8 months getting this shot. It had the right lighting,
9 the right building, the right angle, and you're like,
10 all right, I'm great, I'm happy.

11 And then you turned it in, and Macy's was
12 like, you know what, we've changed our mind, we want a
13 different setting. We want it to be the sunset. And
14 so then traditionally, you'd have to go spend another
15 six months reshooting this picture trying to get the
16 lighting correct.

17 So with our AI, we can automatically segment
18 out the part of the picture that's of interest to you,
19 and then that's the cityscape. And then we let you
20 import another picture that is of the desired lighting
21 and sky that you want. And with one click, you can
22 now take that lighting and put it in your picture.

23 So that's probably not 100 percent of what
24 the creative professional wants for their Macy's
25 campaign, but it's probably 80 percent or 90 percent

1 of what they want, and now they can take this picture
2 and make it into exactly what they want with very
3 little extra effort. So you've just taken six months
4 of extra work, of not exciting work, that was not the
5 fun part of their day. The fun part of their day was
6 setting up that shot to get that image in the first
7 place. And now they can take this and they can go
8 back to Macy's, and if they come back and Macy's says,
9 you know what, we've changed our mind, snowy, blue-sky
10 day, five minutes later, you can just change. And so
11 the AI really helps drive that routine out of your
12 day.

13 Another example is what we call neural
14 stylization. And so, again, this is the idea that
15 we've been able to understand the style of an image.
16 And so we've trained our AI demonstration on the style
17 of different famous paintings. And so if you have
18 your photograph on the left and you said, I want it to
19 look something like the interpretation of these two
20 different paintings, you can do it. All it does is
21 understand the style of whatever painting you put in,
22 and it's just the style of it. So it's not just
23 copying the colors broadly like you might have
24 expected pre-AI. It understands what the style of the
25 image was and applies it to the image.

1 So just understanding that concept of -- I
2 think this is going to play. And so this is not just
3 for creative professionals. This is for hobbyists.
4 You can take your own pictures and you can upload
5 whatever artist you want and it's going to take the
6 style of the artist and apply it to your picture. And
7 it understands that concept.

8 We can also use AI and we do use AI for our
9 video editing products. So this is a project called
10 Cloak, and this is a normal example where you have --
11 you shot a scene and then in post-production, you want
12 to get rid of something you don't like. In this case
13 you don't want that couple there. So in using AI, we
14 segment the image and understand who's in the image
15 and who they are, and we can also fill in the
16 background with copied pixels to make the background
17 look perfect.

18 So on the left is the original footage, and
19 on the right is post-AI, and it looks like they've
20 just vanished, right? And then that used to take
21 months of work to do to edit two people walking out of
22 the video, and now you can do it in minutes.

23 So as I mentioned, we also have an
24 experience intelligence business. This is the other
25 side of our business. This is a digital marketing

1 business that allows you to target advertisements and
2 allows chief marketing officers to understand what the
3 content in the campaign is doing. So we provide that
4 service and use AI there as well. We use it to help
5 you predict the results of a campaign before you even
6 launch it. We may say this is going to be successful
7 in the northeast, or this is going to be successful in
8 California based on our analysis of customer data from
9 interacting with their website. That's another way
10 we're using AI at Adobe.

11 So I think the question is how we get there.
12 How do you actually produce the AI, and I know there's
13 going to be a lot of people talking about the nuts and
14 bolts of the computer science so I'm not going to
15 spend too much time on this, but this is how Adobe
16 does it. Our AI product called Sensei. And this is
17 the architecture.

18 And so what we do is -- what we do typical
19 of any neural network, we have the neural network and
20 then we train it with data, and we train it for an
21 outcome. And using this architecture, we're able to
22 create the neural network; we freeze it in place; and
23 we ship it PhotoShop; we ship in Premiere, and that's
24 the result you see as a consumer.

25 So the principles -- this is my second-to-

1 last slide -- the key principles for training AI that
2 is important to Adobe and just a takeaway for everyone
3 is how do we make this product work as well as we need
4 millions of pieces of data to train it. You need lots
5 of examples of artists; you need lots of examples of
6 images in order to train a neural network to
7 understand the insights that we're able to show you.

8 So when you think about how do we make AI
9 beneficial, how do we get the rewards of AI, you need
10 access to data, you need access to a lot of data and
11 you need access to a variety of data, and that variety
12 of data will make your neural network accurate. And a
13 variety of data will also eliminate bias.

14 You can imagine bias when you're looking for
15 images, that is inherent because you may have trained
16 your AI on that particular kind of a person, and if
17 you go searching for a job or an occupation, you're
18 always going to get that person because that's what
19 you trained it with. So the wider variety of data you
20 put into the AI, the more likely it is your results
21 are going to be unbiased.

22 So thank you for your time. This is our
23 presentation. Creativity in AI, with AI is what Adobe
24 is focused on. It's how we believe AI will help
25 transform the creative professionals for today and

1 tomorrow. Thank you.

2 MS. GOLDMAN: Thank you so much for that
3 colorful and creative presentation.

4 So, next, Henry Kautz will begin his
5 presentation.

6 MR. KAUTZ: Thank you. So I'm going to
7 focus my talk on the work we've been doing at NSF to
8 support AI applications for social good. So when we
9 look at a proposal, we have two major criteria.
10 First, we want to advance science or engineering
11 looking at fundamental advances, but we also consider
12 potential for broader positive impacts on society.

13 Now, the traditional broader impacts that
14 were frequently mentioned in proposals, we're training
15 graduate students and potential future applications of
16 the result. So someone, say, I'm doing this
17 fundamental research, and maybe someone in the future
18 will come along and do something to benefit society
19 with us.

20 But, increasingly, we see that the
21 fundamental science and these broader impacts are
22 entwined, that as you work on an application for
23 social good you discover new questions that require
24 fundamental scientific advances. And from those
25 advances, you discover that there are new

1 opportunities.

2 So AI and broader impacts. So AI methods,
3 taken broadly, that includes machine learning,
4 knowledge representation and reasoning, and what we
5 might call deliberative intelligence, making optimal
6 decisions, are being used by researchers in every
7 discipline that's funded by NSF. I'm from the
8 computer science, and my particular division funds a
9 lot of the fundamental work in AI, but there's really
10 no area of NSF now, including the social sciences,
11 where you don't see people talking about AI. And,
12 increasingly, we're partnering with other agencies,
13 funding or taking advantage of work in these fields.

14 So we've seen -- over the last decade, we've
15 grown up quite a rich portfolio of what we call cross-
16 cutting programs. So these are interdisciplinary
17 funding opportunities that involve multiple
18 directorates within NSF and sometimes with other
19 agencies. Some of the most important are the Smart
20 and Connected Health Program that we run with NIH.
21 And so there, we are looking at AI research that is a
22 bit more applied than traditional work funded by NSF
23 but is not yet ready for the kinds of actual clinical
24 uses that NIH would fund. So we both put money in
25 there, and then we help bridge the gap between those

1 agencies.

2 Smart and Connected Communities looks at
3 applications of AI to all kinds of problems facing
4 urban life from pollution, policing, and violence,
5 transportation, other issues. We've had a program for
6 several years now called Big Data in science and
7 engineering, and that is to support broad
8 collaborations -- or collaborations that can cover a
9 number of fields. So you might have material
10 scientists together with a computer scientist or, you
11 know, electrical engineer together with the computer
12 scientist or even medical people.

13 And through that Big Data program, we've
14 also funded what are called big data hubs, so the idea
15 that these are a set of universities that act as
16 resources to all of the universities in that region
17 for activities such as helping making connections to
18 government agencies. And through that, we've had
19 programs like the Civic Innovation Challenge.

20 One of our most recent programs that is
21 particularly relevant for broader impacts is one
22 called the Future of Work at the Human Technology
23 Frontier. And it's a very interesting combination of
24 directorates -- computer science, engineering,
25 education, and then our social, behavioral and

1 economic sciences.

2 So we're now looking at the future of the
3 workplace and in particular how AI will be impacting
4 that future. So we want to fund both the kind of
5 technology we might see in the future. So, for
6 example, in a recent -- we just completed the first
7 year of the program, and one of the awards was on
8 smart classrooms, so how we might integrate AI as a
9 teacher's assistant, and not replacing a teacher but
10 assisting a teacher.

11 But we also will be looking for work where
12 technologists work with social scientists to look at
13 both the positive and the negative consequences. Will
14 AI throw millions of people out of work? That's
15 absolutely an open question. And if you look back at
16 the history of science and technology, you can make
17 quite good arguments either way that AI will lead to
18 permanent unemployment or that AI will lead to new
19 opportunities for employment.

20 This is another example of the work from
21 this most recent program -- solicitation. So Whole-
22 body Exoskeletons for Advanced Vocational Enhancement.
23 So, here, we're looking, you know, at something a
24 little bit different than your traditional robotics
25 for manufacturing but augmenting the human worker to

1 give the human worker superhuman strength and
2 endurance, or as I mentioned in classroom teaching,
3 where a system that is monitoring a classroom and
4 noticing when students -- those students who have
5 become apparently disengaged are not working or not
6 attending and can perform such tasks as simply
7 alerting the teacher or engaging in a personalized
8 activity with the student.

9 So one of our very largest grant programs
10 is called Expeditions in Computing. These are
11 typically \$10 million over four to five years. So,
12 here, we're really looking for research of the highest
13 intellectual merit. All of our reviewing is a system
14 called peer reviewing, where we get unbiased
15 scientific experts from the community to review. And
16 in Expeditions, we have multiple layers of viewing
17 because we really want to get the best of the best.

18 And in addition, these -- the work we fund
19 should address the nation's greatest needs. So to
20 give just a case study of the synergy between positive
21 broader impacts and scientific merit, I'd like to just
22 mention some of the work going on at the Institute for
23 Computational Sustainability, which is a -- the result
24 of actually two successful Expeditions in Computing
25 that went to a consortium of Cornell, Stanford, and

1 University of Southern California.

2 So the problem here is looking at
3 sustainability problems, and by sustainability, we're
4 looking at environmental sustainability, economic
5 sustainability, resources, social sustainability, very
6 broadly, as complex problems that are really too
7 difficult to solve with human intelligence alone. So
8 we want to employ AI techniques and large amounts of
9 data to solve optimization -- essentially resource
10 optimization problems that are far beyond the kinds of
11 linear optimization that most of the people in this
12 audience would be familiar with.

13 These are highly nonlinear problems where we
14 must model uncertainty. So we can't -- we just can't
15 ignore the fact that many -- there are many variables
16 that are unobserved. Okay.

17 Now, you might think that, well, these are
18 all different problems, but what has been so
19 fascinating by this Expeditions is that problems that
20 seem to be quite different often have very -- have
21 shared technical solutions, okay? So this is a subway
22 map that the research group created. And as we see,
23 each of the tracks of the subway, the six tracks --
24 the six tracks are scientific themes. So pattern
25 decomposition, crowdsourcing, mechanism design, so

1 social choice theory, and economics, spacio-temporal
2 modeling probabilistic inference, and sequential
3 decision-making. And then each of those tracks is
4 going through the stops, where the stops are the
5 particular application.

6 So in each application you had domain
7 experts. So let's say there's one there on landscape
8 scale conversation and rural communities. That
9 included, you know, people who knew a lot about that
10 topic and had been studying and working with
11 communities in Ecuador, but it made use there of
12 temporal modeling, probabilistic inference, and
13 sequential decision-making. So you see it's quite a
14 variety here -- flight call detection, and I'll
15 mention that again, wind and solar forecasting, all
16 the way over to microbial fuel cells.

17 Now, but one thing I should point out is AI
18 covers many things. There's sometimes a tendency
19 because of the great success of what are called
20 artificial neural networks to say that that is AI.
21 And as we just saw from the previous speaker,
22 artificial neural networks are wonderful when you're
23 dealing with patterns, doing pattern recognition, and
24 essentially trying to emulate those parts of
25 intelligence that don't involve essentially logical

1 thinking but are more based on pattern recognition and
2 intuition, the kinds of problems we don't think about
3 when we solve them -- recognizing your friend's face,
4 right? We don't think consciously about it.

5 By and large, the work in this particular
6 set of projects, though, involves what we may call
7 your Type 2 intelligence, your deliberative rational
8 intelligence where you have many alternatives to
9 consider. In fact, there is such a large number of
10 alternatives, you can't simply enumerate them all one
11 after the other. You have to have very clever ways of
12 essentially searching through an enormous, sometimes
13 infinite space of possibilities and narrowing in on
14 those points that are near optimal.

15 So just going down a little bit deeper, the
16 problem of data -- of decomposition in big data. So
17 this is -- so a core technical problem. You have some
18 kind of very complex signal, and you want to reduce it
19 to something simpler, right, to a small -- the one
20 measurement or a small number of measurements. So
21 this is also called dimensionality reduction, source
22 separation, sometimes called segmentation. But it
23 makes use of a body of algorithms that have come up in
24 computer science, electrical engineering, and
25 particularly more and more in work in AI.

1 So we had a -- there were a series of
2 projects, one on detecting gunshots. And you can
3 imagine security applications in a city. Another one
4 detecting elephant calls. So you can put out audio
5 monitors in the jungle and use that to conduct a
6 census of elephants, right, based on their calls.
7 That same work was then used to detect birdcalls of
8 actually birds in flight for a project with the School
9 of Ornithology at Cornell. And perhaps, surprisingly,
10 is with very few changes, that same algorithm was used
11 in a project on crystal phase mapping, which is in
12 material discovery, so a problem where you're coming
13 up with a mix of new materials, you hope they have
14 some property, and you're analyzing the results of
15 shooting x-rays at those new materials.

16 Another example -- my last example here --
17 is dealing with hydropower in the Amazon Basin. So
18 there are a great potential for getting more
19 hydropower from the Amazon Basin. And, in fact 170
20 dams have already been built or under construction,
21 about 300 dams are planned or proposed.

22 Now, there's obviously a big problem
23 here. If all of these damns are built, not only will
24 there be quite a lot of devastation to wildlife, but
25 they will become less effective because one damn is

1 going to affect the water flow to another damn.
2 So you have to look at this as a multi-objective
3 optimization problem to balance off energy,
4 fisheries, transportation, and navigation. Obviously,
5 as you put in more dams, you're going to make river
6 transportation more expensive, and finally looking at
7 the long-term effects, how will all these dams affect
8 the natural flow of sediment and nutrients and how
9 that affect farming. So this becomes a multi-
10 objective optimization problem.

11 And then the goal is to look at the
12 tradeoffs between these different factors and have a
13 new algorithm that can present, well, here is the
14 possible best tradeoffs. There's no single best
15 tradeoff, but you can look at that any solutions that
16 don't fall along this line are provably worse, so
17 they're worse in some respect and no better in any
18 other respect. So this tremendously reduces this sort
19 of infinite space of the number of dams and the
20 placement of dams to one that now can be decided by
21 humans. Yeah, that's showing where they're the dams.

22 And interesting that this same effort has
23 led to startups. For example, ATLAS AI, that is
24 basically a for-profit AI for social good company.
25 This also received funding from the Rockefeller

1 Foundation, looking at providing -- helping developing
2 nations be more sustainable in their agricultural
3 practices. Networks of CompSustNet, a larger network
4 that includes this group of these three universities
5 with others to address these problems.

6 And with that, I'll conclude. Thank you.

7 MS. GOLDMAN: And thank you so much for
8 showing us the diverse portfolio that NSF is
9 supporting.

10 And, now, Angela Granger will begin her
11 presentation.

12 MS. GRANGER: Thanks. Sorry, it's a little
13 tight up here, so we thought that would be the better
14 route to get around.

15 I lead analytics for Experian, and one of
16 those areas that I'm responsible for is credit scoring
17 product development, and for those of you that don't
18 know, Experian is a global leader in consumer and
19 business credit reporting and marketing services. We
20 support clients in over 80 countries, and we have
21 approximately 17,000 people in 37 different countries.

22 We believe it's our responsibility to assist
23 lenders in managing consumer credit risk and
24 empowering consumers to understand and responsibly use
25 credit in their financial lives. We're committed to

1 being the consumers' credit bureau, and I thank you
2 guys for having me here today.

3 To set the context for today, there's a lot
4 of different areas of application for credit scoring,
5 so we're going to -- I'm going to specifically talk to
6 scores used to assess eligibility for credit where
7 adverse action may be taken. The example was used a
8 couple of times earlier today specifically of
9 application of credit for an example where you could
10 be approved or declined, your application for credit.
11 That would be the credit scoring context we're talking
12 about today.

13 Benefits of AI or machine learning, for both
14 lenders and consumers in our industry, are ultimately
15 better lending decisions. If you have greater
16 insights into the data that you're using, better
17 accuracy in the scores, you're going to have better
18 decisions being made.

19 And, secondarily, financial inclusion.
20 Where we're really finding the power of AI and machine
21 learning techniques is our ability to evaluate new
22 data sources more quickly and incorporate that new
23 data into credit scores, thus broadening the access
24 for credit for people who maybe have thin credit or
25 are new to credit and don't have a credit file with us

1 today.

2 Where we like to start is with the data. If
3 you think about predictive modeling, and any kind of
4 modeling for that matter, it's important to understand
5 the data that's feeding into the model. For us, we
6 talk about traditional credit data. And when you
7 think about traditional credit data, what we refer to
8 is what you typically find on the core credit
9 databases at the major credit reporting agencies. And
10 this includes information around what we call trade
11 lines or account-level information where you get
12 access to a consumer's payment history on a certain
13 type of account, their outstanding balances, that sort
14 of thing.

15 We also have information on inquiries that
16 are made into the credit bureau for applications for
17 credit for example. And we have public record
18 information, particularly on bankruptcies. We also
19 maintain some additional information that you might
20 think of as being part of a credit application, such
21 as income and employment.

22 We also like to talk about alternative
23 credit data. So this goes by many terms. In our
24 industry, when we say "alternative credit data," we
25 really mean data that is not on that core credit

1 database that I talked about a minute ago. So types
2 of alternative credit data that aren't reported to the
3 core credit database today can include rental
4 payments, asset ownership, alternative financing such
5 as payday loans, short-term loans, rent-to-own-type
6 loans.

7 There's additional public record information
8 out there that's not on the core credit database.
9 And, most recently, we've incorporated consumer
10 permission data.

11 Both alternative data and traditional credit
12 data have been found to be very predictive of a
13 consumer's creditworthiness. And, particularly, the
14 alternative data comes into play in those cases of
15 thin file and no-hit-type consumers that I mentioned a
16 minute ago.

17 The Fair Credit Reporting Act regulates the
18 collection, dissemination, and use of consumer credit
19 information, and so all data used in credit scores are
20 what we would call FCRA-complaint. What does that
21 mean? That means the data needs to be accurate, so
22 the credit reporting agencies must do their best to
23 ensure their data is accurate. The data is
24 disclosable, so consumers can see that information.
25 Consumers can get one free credit report every 12

1 months, and they can see their credit information if
2 they're denied credit as an example.

3 The data furnishers also play a role in the
4 process. They have to confirm information when
5 disputes happen, and they're held to certain
6 turnaround times as well as part of the dispute
7 process. And, lastly, we were set up pretty nice
8 earlier around fairness. Fairness is another part of
9 the FCRA. So scores are -- they cannot discriminate
10 based on these different ECOA factors such as gender,
11 marital status, race, and religion.

12 So for about 30 years, we've been using
13 scores kind of in their current form, which means
14 they're using this core credit information that I
15 talked about earlier. And so between that and our
16 experience over time, we've come up with things that
17 are generally acceptable in our space, data that
18 complies with those FCRA rules that I mentioned
19 earlier, proven payment information, rental data,
20 account transactions from your demand deposit accounts
21 are generally deemed acceptable. Generally not
22 acceptable are things like social media data, you
23 know, who your Facebook friends are sort of things,
24 and any data that could discriminate in decisions or
25 that could be discriminatory, I should say.

1 Under consideration right now, we're looking
2 at education level, again, something to help us in
3 that new-to-credit space. Think of students
4 graduating from universities and having that
5 information available so that they can more easily get
6 credit and join the credit ecosystem.

7 So one of the things about our industry is
8 not only is the data itself, which we just went
9 through, regulated but the scores or the models are
10 regulated as well. There's regulatory guidelines
11 around accuracy and fairness that have been put out by
12 the OCC. Those documents or those guidelines, I
13 should say, are pretty extensive. They cover the
14 model development process, they cover model use, they
15 cover model monitoring, when to redevelop. And they
16 create quite an extensive amount of documentation.

17 And in order to meet these model governance
18 guidelines, many of our clients -- so think of, you
19 know, big banks, big lenders -- have had to create
20 entire staffs just to take on this model governance
21 requirements.

22 We talked about the controls around
23 discrimination which lead to the need for
24 transparency. And then in the FCRA, we are also
25 required to provide your top four reasons for your

1 score being what it is as well. And so the need for
2 transparency, or what we call explainability in
3 scores, is very high.

4 Some key considerations when developing
5 credit scores to meet all these needs, I won't go
6 through all of these in particular, but they really
7 cover the full life cycle. We talked about, at one of
8 the earlier sessions, generalization. So our models
9 need to essentially replicate. They can't just work
10 really well on the training sample. They have to work
11 well in production. If you think about credit scores
12 in use today -- think about mortgage in particular --
13 the credit scores being used there are about 20 years
14 old, right? So these models need to continue to
15 replicate and still rank-order consumers in terms of
16 their creditworthiness.

17 Today, models have an average shelf life of
18 about three years, so we're looking at AI to help us
19 get models to market faster. Some research that we
20 did, we tested several different techniques around
21 machine learning. I won't go into each of them. You
22 can see that here. But suffice it to say the gradient
23 boosting models are the ones for credit scoring and
24 credit risk in particular that seem to be rising to
25 the top.

1 When we let the machine run by itself, these
2 are the type of results we get. We see anywhere
3 between a 5 percent to 10 percent lift depending on
4 the situation. This is a more generic sample for auto
5 and bank card, so we see about a 5 percent lift if you
6 were to do the math here. But our clients report
7 anywhere up to a 15 percent lift as they start to
8 really look at specific portfolios or specific
9 lenders.

10 This, however, is when you just let the
11 machine run itself and you don't take into
12 consideration some of those things we talked about
13 earlier.

14 We do something that we call model
15 refinement, and this is where you have to go in and
16 ensure your model is working as expected, that you can
17 explain what's happening. You want to make sure that
18 a credit score doesn't go down if a consumer has made
19 some impact to their credit such as paying off some of
20 their debt or lowering their utilization. And if you
21 don't do this refinement and you don't understand
22 what's happening under the covers, that can happen.

23 So when you go in and you refine the model
24 through the requirements that we talked about before,
25 you'll see that the lift in performance from the -- in

1 this case, extreme gradient boosting methodology, is
2 lessened. So in our particular example, the lift went
3 from 5 percent to 2 percent. In other examples, we've
4 see that 15 percent or 10 percent lift come down to 5
5 to 8 percent, right? So on average, we're seeing
6 about a 5 percent lift in accuracy from applying some
7 of these techniques outside of our traditional
8 regression methods.

9 This is just another example of addressing
10 overfitting, which tends to be a problem with some of
11 these new methodologies that aren't -- haven't been
12 used in practice as long. What you tend to do if you
13 throw all of the data into the machine and let it do
14 its work, we have over 2,000 attributes, variables,
15 characteristics that we will throw into a model, and
16 it will use almost all of them if it can, right.

17 And that tends to overfit and the model
18 doesn't generalize. And so you do have to go in and
19 manually intervene and not let the machine do all the
20 work.

21 Some of the advantages for AI in credit
22 scoring go beyond just the modeling. You know, I
23 mentioned a 5 percent improvement, and I'm sure you
24 guys are all sitting there, going, whoo, 5 percent, 5
25 percent, right? But in the credit risk world and

1 creditworthiness world, we have very predictive models
2 today. And so a 5 percent improvement is actually a
3 big improvement. The data that we use in the models
4 is very accurate, and so we get very good models. So
5 5 percent improvement is significant, but we're
6 looking to use machine learning and AI methodologies
7 across the model development life cycle and not just
8 in the model development itself.

9 Probably most importantly to take away from
10 today is in credit scoring. Credit scores are static
11 models. So most of us when we think of AI think of
12 realtime updating, self-learning type models. Those
13 are not in use in our industry today. These are
14 static models. So while we're looking at these
15 additional techniques outside of regression, we're
16 still talking about static models. I mentioned the
17 turnaround time or the shelf life of a model is about
18 three years right now. With these new techniques,
19 that's going to come down, but we have to have the
20 ability to go back in time and replicate our models.

21 So, lastly, there's some future policy
22 regarding credit scoring that we wanted to make sure
23 you were aware of. Today, unlike what people think,
24 your telephone bill, utility payments are not reported
25 to the credit bureau. Those are very powerful

1 predictors just like other payment methods of future
2 payment behavior and so of creditworthiness. And
3 there's been several studies that show that today.

4 And so with that, I would like to thank you
5 for giving me this opportunity and hopefully this gave
6 you a quick glimpse into the status of AI and how it's
7 being applied in credit scoring. Thank you.

8 (Applause.)

9 MS. GOLDMAN: And thank you, Angela, for
10 that very interesting presentation on credit scoring
11 and bringing in the related legal and policy issues.

12 So, next, Melissa McSherry will begin her
13 presentation.

14 MS. MCSHERRY: Thank you very much, and
15 thank you so much for having me today. I work with
16 Visa. Visa is the world's largest payment network,
17 and what that means is basically when you use a Visa
18 card your -- the merchant where you use the Visa card
19 basically calls their bank and says can I authorize
20 this transaction. And then Visa connects the
21 merchant's bank with your bank, who says yes or no,
22 that's a good transaction to authorize. And then that
23 signal goes back to the merchant, and all of that
24 happens if everything goes according to plan. All of
25 that happens almost instantaneously.

1 In that -- in that context, Visa is very --
2 we work very, very hard to make sure that the
3 transactions that are going through are legitimate
4 transactions or not fraudulent transactions. I think
5 fraud worldwide today is something like \$600 billion,
6 so it's a lot of money, and we want to make sure that
7 we do as much as we can to help banks prevent any of
8 those fraudulent transactions from going through while
9 still making sure that all of the good transactions go
10 through. Basically, when you are actually the one
11 using your card, if you try to use it, that it
12 actually works.

13 So what I'm going to talk about today is one
14 way in which Visa is using AI, specifically computer
15 vision, to help us do that work of looking after and
16 preventing fraud on the Visa system.

17 So you might be asking what do puppy dogs
18 and blueberry muffins have to do with preventing
19 fraud. And I put this up just to sort of illustrate
20 both the challenges and the opportunity in computer
21 vision. So all of you could look at these pictures
22 and very easily discern what's a blueberry muffin and
23 what's a puppy dog. But using the techniques that
24 were available up until, you know, call it 2012, 2013,
25 this was actually a pretty hard problem for most

1 computers to solve. They would get it right about 75
2 percent of the time.

3 And in I think it was 2013 -- there's a
4 competition that is run every year. And new
5 techniques, specifically things called convolutional
6 neural networks, started coming into play and started
7 dramatically improving the ability of computers to
8 correctly differentiate the muffin from the dog. And
9 so we're now at the point where these techniques can
10 generally differentiate not just muffins and dogs but
11 can differentiate different images about 97 percent of
12 time as opposed to 75 percent of the time, which is
13 really quite good.

14 If you think about human beings -- although
15 if you're sitting there concentrating, you know, you
16 would always be accurate since most people don't
17 concentrate all the time and they do sometimes make
18 careless errors, human beings run at about 95 percent
19 of the time, right, when you give them a lot of
20 images. So this ability to look at a picture and so
21 to say this picture looks like this one, and this
22 other picture looks like this other one, this is one
23 of the applications of AI that has dramatically
24 improved.

25 And so now I'm going to talk a little bit

1 about how we use that application of that computer
2 vision application of AI in the context of fraud. So
3 just a couple of terms before we get started with this
4 particular example. First of all, what is a fraud
5 score? Like I said, whenever you use a card, Visa
6 basically attaches a score to the transaction that
7 goes to your bank that says how likely is it that we
8 think that this is actually you using your card versus
9 someone who's trying to commit fraud using your card.
10 We provide that information to the bank so the bank
11 can make a decision about whether or not they want to
12 authorize the transaction.

13 And as you can imagine, we process a lot of
14 transactions, right? So that first thing we do in
15 every transaction is we attach a score from zero to
16 99. But then if we look across all of the
17 transactions, we can actually say, for instance, all
18 of the transactions in an hour, how many of them were
19 at, like, the highest score, got a score of 99? How
20 many of them were at the lowest score, got a score of
21 zero. And it's helpful to us to look at the
22 percentage of scores that are in each of those bands.

23 And the reason why is if you -- if we're
24 running along and 1 percent of the population is
25 getting the highest score, that 99, and it's nice and

1 steady and then all of a sudden like 10 percent of the
2 population is getting a 99, that means that probably
3 one of two things is happening. Either there's a
4 giant fraud attack, and there are fraudsters that are
5 trying to, in a very coordinated way, steal a lot
6 money, and this does happen sometimes, right, in which
7 case we need to intervene. And we typically intervene
8 by calling the banks that this is happening to.

9 Or there is something wrong with our models
10 or system or how we're processing things. And, again,
11 that's a situation in which we need to intervene and
12 we need to make sure that everything is actually
13 working as we expect. So not only do we look at the
14 fraud scores, we also look at the distribution of
15 those scores.

16 And so the next page, this is just -- this
17 is a made-up example, but I think it sort of
18 illustrates what's going on. So you can imagine that
19 this is a graph looking at the percentage of
20 transactions in a particular score band. And in this
21 particular case, I just did it over days, and it goes
22 up and down, and it goes up and down because, for
23 instance, the kinds of -- the mix of transactions that
24 you see on like a Friday night can be pretty different
25 than the mix of transactions you see on a Sunday

1 morning. And so the mix of transactions in a
2 particular score band can go up and down.

3 Now, if you look at this, it's pretty easy,
4 again like the puppy dogs and the muffins. It's
5 pretty easy to see that at the end there's something
6 that looks a little bit different, right, that doesn't
7 -- that pattern doesn't look like all of the other
8 patterns that came before it.

9 And this is, again, pretty easy for everyone
10 in the audience to see that that's different, but it's
11 actually kind of hard for the tools that we had prior
12 to those computer vision tools to pick this up, like
13 you can't -- like a traditional control chart, it's
14 hard to write a rule that will get this because the
15 actual numbers are sort of -- they're inside the range
16 of the historical range, they're going up, they're
17 going down. They're not -- it's just -- it's hard to
18 write the rules. But, again, it's easy to see it
19 using computer vision tools.

20 And so what the computer vision tools let us
21 do is basically do what a person would do in terms of
22 looking at this and seeing a pattern that's different.
23 But the computer vision tools let us do that every
24 hour of every day. I mean, the computer doesn't get
25 tired and people do, like, they need to go do

1 something else other than look at charts all day.

2 It lets us look at hundreds of metrics, not
3 just one, right? And if you think about this, this is
4 a pretty simple chart that I put up here, right? This
5 is basically one-dimensional, right? We sort of look
6 at the scores, versus one-dimension. And so it's easy
7 to see the variation. If I had put a chart up here
8 that had multiple dimensions, like we were varying a
9 couple things at the same time, that very quickly gets
10 really hard, even for people, to see the differences.
11 But, again, the computer vision techniques that we've
12 been talking about can pick those variations up pretty
13 quickly and can identify those. So we can not only
14 monitor what's going on versus one dimension, we can
15 monitor what's going on versus multiple dimensions,
16 and it makes our monitoring that much better and that
17 much faster.

18 So just a quick explanation of how we've
19 applied this in our particular situation. Basically,
20 we built a model that looks at the distribution of
21 each of those score bands that we just talked about,
22 so, you know, for instance, scores of 10 to 19, right,
23 so it does this for each score band. And it looks at
24 those distributions for a five-hour period over each
25 of the last 120 days. Right, so this is lots of data

1 that's coming in. Think of the computer as looking at
2 a chart, an hourly chart over the 120 days.

3 And from that, it forms an expectation of
4 what the current five-hour period is going to be,
5 right? Is the score -- is the distribution going to
6 be going up and then down? Is it going to be going
7 down -- you know, down and then up? Is it going to be
8 going, you know, one direction -- it forms an
9 expectation. And, then, and this is the part that
10 relates back to the puppy dogs and the muffins, it
11 looks at the actual picture and it compares it to its
12 expectation that it created based on the last 120
13 days, right?

14 And so on the top row, we see on the right
15 is sort of what we would expect, right, for this time
16 period from the data that's come in over the last 120
17 days. And what we see on the left is what actually
18 came in. In those two pictures, the computer would
19 say, yep, those two things -- they look similar,
20 they're both blueberry muffins or they're both puppy
21 dogs, right?

22 But on the lower band, what we see is the
23 expectation for the particular time period that we're
24 looking at is just that the scores will be going up
25 during the time period. But what we actually see is

1 that they're going up and then coming back down. And
2 the computer at that point says, no, no, no, these do
3 not look like they're the same. This is not --
4 something is not matching here.

5 And that, then, causes the system to
6 generate an alert and say, hey, a person, a human
7 being, needs to go look at this, right? It might be
8 that it's fine. It might be that it's just, I don't
9 know, Black Friday, right, and so all kinds of things
10 are different. Or it might be that there is an actual
11 problem and we need to get engaged and figure out what
12 the problem is, and we need to figure that out
13 promptly.

14 So in this particular case, what's going on
15 is the computer is basically taking a lot of work that
16 might have been kind of boring and tedious for the
17 people and doing the boring and tedious part and then
18 just pulling out the things that are interesting and
19 require human intervention so that human can then go
20 and figure out what we actually need to do
21 differently.

22 One thing I just want to call out about this
23 particular example is, you know, so Visa is using a
24 lot of different AI techniques across a lot of
25 different places in our system. These particular

1 can you replicate that or what alternatives should
2 there be?

3 MR. CALO: Fred is a deep expert on notice
4 and choice, one of the leading experts on notice and
5 choice in America. But I will hazard something which
6 is that what is interesting about artificial
7 intelligence, at least when we come to embody it in an
8 agent, which is something that somebody asked about,
9 is that it can be awfully contextual and dynamic.

10 So I think that we ought to be encouraging
11 -- you know, the possibility of having a conversation
12 with Alexa about Amazon's privacy practices is, I
13 think, quite exciting, you know, and the idea -- maybe
14 you are anti, but the idea being that you could ask
15 specific questions rather than have some stupid thing
16 that was like really long and you are never going to
17 read it. But you could say, hey, Alexa, can Amazon
18 turn on you remotely to listen in on a conversation,
19 and then get an answer about that. I think that is
20 actually pretty powerful, personally.

21 MR. GILLULA: I am not anti, I just think
22 maybe only the people on this panel would find it
23 super exciting to have a conversation about Alexa
24 about --

25 (Laughter.)

1 MR. GILLULA: Which is not to say it would
2 be me, I agree. I just do not think the vast majority
3 of consumers would get a ton out of it.

4 MR. CALO: I mean, I think it is critical
5 when you are thinking about emerging technology
6 generally not just think focus on what is loss, but
7 what new affordances might be there or what you might
8 gain. I think that these things are quite powerful.
9 I think we are getting to a place where natural
10 conversations are becoming more viable and I think
11 that we should therefore -- I mean, if you think about
12 it, notice and choice, we have been operating under
13 basically Gutenberg technology all this time, right?

14 We just publish a long thing whether it is a
15 digital or a print, just a bunch of words on a page.
16 Yet, you know, here we have companies that are doing
17 these amazing things about organizing information and
18 gauging you and so on. Anyway, I think there is a lot
19 of innovation that could be occurring with notice.
20 And part of it would be to contextualize and actually
21 answer questions about this consumer instead of just
22 having something that no one reads.

23 MR. CATE: I would echo everything Ryan
24 said. I would just like to make two points. One is
25 we put in the record a paper that I did with some

1 colleagues at the Center for Information Policy
2 Leadership about AI, how it is used today and some of
3 the issues it raises, and one of the things we talk
4 about in there is the way AI is already being used to
5 enhance privacy protections, not just to make them
6 more easily understood or explainable, but to actually
7 activate them. So in other words, you can identify
8 somebody's privacy preferences as they start
9 expressing them and then you can start predicting what
10 they will be so that you offer them the default they
11 are more likely to care about. Rather than the
12 default that you want, you try to give them the
13 default that they want.

14 I would say just, in general, though, back
15 to the original question on notice and choice. As I
16 said earlier, we have relied on this largely because
17 we have not known what else to rely on for 50 years
18 now, with not a lot to show for it. And so I think we
19 should recognize that notice should be used and choice
20 only where there is something meaningful to tell the
21 individual and only where there is something they can
22 do about it. So I think it is terrific when my iPhone
23 says, did you know this app is using your contacts, do
24 you want to permit that? That is meaningful notice
25 and I can do something. I can say yes or no, I can

1 alter it.

2 But making my doctor add another paragraph
3 to the 65 paragraphs of the HIPAA notice saying, by
4 the way, your scans are going to be read by AI and, by
5 the way, you have no choice about that whatsoever
6 because it is far more accurate than humans, I am not
7 sure that is overly meaningful. I think we have to be
8 very contextual with notice because the effect when we
9 do not is that we just teach people to ignore all of
10 it. We get people in the habit of knowing that notice
11 is meaningless and so they do not read it, whereas if
12 we would use notice when there actually is something
13 worth telling them and something they can do about it,
14 we might resurrect notice as a meaningful data
15 protection tool.

16 Now, having said that, I am not disagreeing
17 with Irene. The law requires, both in Europe and in
18 some industries in the United States, notice and
19 choice, it is just bad law. In other words, it is
20 causing people to ignore these notices by providing
21 them when you cannot do anything about them and nobody
22 cares.

23 So one of the things we often talk about at
24 universities is, you know, a teachable moment. You
25 know, you can only teach someone when they are

1 interested in learning. Similarly, you can only give
2 meaningful notice when there is something that is
3 going to cause them to care about it. And that cannot
4 be they woke up in the morning or they went to a
5 doctor's office. It might very well be where there is
6 an event, there is something happening, there has been
7 some effect on them, there is some reason that they
8 would care, and then using the tools that Ryan was
9 talking about would be fabulous to really make notice
10 meaningful and interactive.

11 MS. LIU: There is always a conflict within
12 companies with product design when you are trying to
13 design products that is easy to use and that is easy
14 to understand. When you are throwing in all sorts of
15 consents and notices, it can make it really difficult.
16 And so there is often a conflict between the lawyers
17 and the product design teams about how can we make it
18 look beautiful without all your verbiage. So that is
19 something that we struggle with.

20 And I completely agree with Fred that
21 meaningful consent is ultimately more beneficial to
22 society and to consumers for how their information is
23 being used and how the company is using it versus just
24 providing our lengthy privacy policies that most
25 companies have.

1 MS. GEORGE: And as a corollary to that,
2 does the notion of opt-out work in an AI context and
3 does that vary based on I think the stage of the
4 product life cycle, be that data collection, you know,
5 product design when it is rolled out to market and
6 being used or other instances?

7 MS. LIU: Jeremy and I were talking about
8 this earlier. So from a GDPR standpoint, you do have
9 a right to erase your data. So there is an obligation
10 for companies to be able to remove that data. And
11 depending on how you configure that information, it
12 can be difficult. So that is something that you need
13 to think about from the beginning in the design phase
14 to ensure that companies, especially with the
15 California Privacy Act as well, it is important to
16 design these products in such a way that there is an
17 opt-out notion.

18 To opt out of AI, typically if a company --
19 if someone wants to opt out of AI completely, that is
20 like let's say if I am using Netflix and I want to opt
21 out of using the choices, the different types of
22 videos or shows that they are showing to me, it is
23 basically opting out of using Netflix completely. So
24 you have to think about, like, are you trying to opt
25 out of the product or are you trying to opt out of the

1 database use as well? So there are different ways of
2 viewing opt out, and I think Jeremy can probably talk
3 more about the technological ways of opting out.

4 MR. GILLULA: Yeah, there has been some
5 recent papers that show that for neural networks you
6 can actually reconstruct what the training data was if
7 you are given sufficient time and access to -- and
8 able to run test data through the neural network,
9 which basically means that if I am a service and I
10 used your data to train my neural network, I cannot
11 remove your -- the fact that you are -- the imprint
12 your data has left on my neural network basically
13 without retraining it from scratch and retraining it,
14 once again, without your data. So it is technically
15 -- is it technically possible? Yes. Is that
16 potentially a huge burden on the company?
17 Potentially, yes.

18 Then there is the other question of, how
19 much benefit do you get from having your imprint
20 removed from whatever model was generated? Because it
21 does take quite a bit of effort to reconstruct all of
22 the training data, and so that is in the unlikely but
23 feasible attack. So there we do have to get sort of
24 into this balancing act a little bit.

25 MS. WORTHMAN: Another question from the

1 audience. In cases where autonomous systems result in
2 consumer harm, who should be held liable and to what
3 degree?

4 MR. GILLULA: Just send the robots to jail.
5 (Laughter.)

6 MR. CALO: Well, I mean, I think that is a
7 genuine puzzle. I mean, so you have -- in criminal
8 law and in tort law, we generally require that you do
9 something either on purpose or that you -- a
10 reasonable person would be able to foresee the
11 category of harm that occurred, right? And so when
12 you, for example, have a bot, which this really
13 happened, that is supposed to buy things randomly on
14 the web and buys methamphetamine and the police come
15 and say, you know, you bought methamphetamine, and you
16 say, no, no, it was the bot, right?

17 Or in another instance, where a company made
18 a bot that was arguably hacked into, but at least was
19 subverted by trolls that wound up denying the
20 holocaust which is not lawful in some jurisdictions
21 where -- that had access to this bot. You know, you
22 would be sort of hard pressed to bring a criminal case
23 to it. And certainly in many categories where --
24 something happens where the system just behaves in a
25 way that was not anticipated, you do not have what is

1 called proximate causation for purpose of bringing a
2 tort lawsuit, which is what I teach.

3 And that is not a great place to be because
4 you wind up in a situation where you have victims, but
5 not perpetrators. And I do not know how much that
6 would really matter to FTC enforcement, specifically,
7 because I think you could get around it just by
8 saying, look, you created these conditions that were
9 deceptive or unreasonable and these unexpected things
10 happened, but something was going to go wrong. But I
11 think it is pretty serious in tort and criminal law.
12 I think it is hard.

13 MS. WORTHMAN: We have also had a question
14 from the audience about retail price discrimination at
15 the individual consumer level and what is the material
16 harm to the consumer in price discrimination and maybe
17 price discrimination can be sort of whether or not it
18 is advertising different things, not on a prohibited
19 basis under ECOA, but just because you are using a
20 different type of computer, because you are purchasing
21 tickets on your mobile rather than on a laptop. What
22 is the harm, what is the cost-benefit analysis in that
23 particular instance?

24 MR. CATE: So this is a place where actually
25 notice would be quite useful. This would be much

1 more, in other words, to say if you visit this
2 website, we are going to use pricing based on
3 information about where you are coming from, the
4 computer you are using, whatever because it would then
5 empower you to say, well, I am going to go have my
6 friend check and see what the price is to see if I can
7 get a better price. In other words, that would be
8 actionable notice, you could really do it. And by the
9 way, having to disclose it would probably slow people
10 down -- companies down actually wanting to do that.

11 I mean, remember, we have discriminated on
12 price for forever, I mean, for generations. Every
13 time you fly, there could not be -- there is more
14 discrimination for all sorts of reasons, how long you
15 are willing to stay, what nights you will stay, and so
16 forth. We discriminate based on zip code, we
17 discriminate based on all sorts of other information
18 that have been imperfect. Now, we are going to be
19 able to discriminate better. I mean, we are going to
20 have both better technology and better data and the
21 two together are going to make much more precise
22 discrimination. You know exactly what I will pay.
23 eBay knows exactly what I will pay for something
24 because it has watched me pay that for years.

25 So this is actually a place where you could

1 say, first of all, we need to figure out is that a
2 harm. Is it something we are going to say is unfair?
3 Is it something that we are going to say causes
4 injury? And if not, maybe disclosure is sufficient.
5 To say, look, we are not willing to say we are going
6 to prohibit it, but we are going to say you get
7 notice. So now, you can figure out if you want to try
8 to come back at the system the other way. They are
9 doing it to you, can you do it to them?

10 But this is why we have to remember, again,
11 it is going to be very contextual and it is not
12 something new. It is not something AI is going to
13 create. AI is going to make it better in the sense of
14 potentially more precise or more tailored.

15 MR. CALO: I will give you my two favorite
16 examples of price discrimination after -- I mean, and
17 by favorite, I do not mean I like them. One of them
18 was a couple of years ago a marketing firm was using
19 this tool to figure out when women felt worse about
20 themselves and they labeled these "prime vulnerability
21 moments." And they suggested that perhaps you should
22 advertise or charge people more during those moments,
23 you know what I mean? That strikes me as not a very
24 good use of price discrimination.

25 Another one of my favorites, although they

1 claim they never did this, is when Uber experimented
2 with figuring out whether you would be more willing to
3 pay surge prices when your battery was low on your
4 phone because maybe you would get stranded there.
5 Lovely, also. They say they have never done this and
6 I believe them about that.

7 The issue is not price discrimination. The
8 issue is taking advantage of people, which happens, it
9 happens a lot. And, yes, from an economic
10 perspective, better information is better. Maybe we
11 would worry at one level about all the social surplus
12 going to the firm. You know, they know your
13 reservation price. There is no windfall for you
14 because they charge you more if you would be willing
15 to pay more, so they get social surplus. We have
16 seemed to have moved away from the original
17 understanding of how consumer protection worked, which
18 was that it was immoral for firms to extricate all of
19 the social -- we seem to have moved away from that
20 model, and that is fine.

21 But I think it is the advantage-taking that
22 I really would worry about, and that is the kind of
23 thing I want there to be hard questions asked about.

24 MS. LOPEZ-GALDOS: Yeah. What I think is
25 that the questions we are addressing here, like from

1 the liability question and the answer from a total
2 perspective to this question right now is that there
3 are no new issues. Discrimination, price
4 discrimination has existed forever. It does not
5 matter whether a machine makes the decision or not,
6 the debate is the debate. We should analyze whether
7 we still -- whether price discrimination, for example,
8 is procompetitive or not or on the other side whether
9 consumers are being harmed or not, which approach we
10 want to take. But it is a debate that we should be
11 having and we have been having even without machines.

12 So I think we just need to continue talking
13 about these things, but I do not think it makes a
14 difference whether a machine makes a decision or a
15 human being makes that decision.

16 MR. GILLULA: So actually, I want to
17 disagree. There is something fundamentally different.
18 And if you lump in AI and big data and predictive
19 analytics altogether, then I agree there is nothing
20 new separate on AI. But a major difference is that
21 now there is a -- you are making a decision based on a
22 tremendous amount of data that has been collected as
23 opposed to just like, say, one data point that you
24 happen to notice or one data point you got like the
25 zip code or how many nights you want to stay for the

1 flight or something like that, something that is very
2 clear.

3 Now, you can potentially make price
4 discrimination decisions based on what websites the
5 person was visiting. Were they looking at budget
6 travel websites versus high-end travel websites? And
7 then there is the question of what happens if -- how
8 were the price discrimination decisions made if you do
9 not have any data on the person? And do they suffer a
10 penalty for preserving their privacy?

11 If I use a tracker blocker app on my phone
12 and I go to your website and I try to buy a plane
13 ticket and you do not have any history, am I
14 automatically categorized as I have to pay the highest
15 price or not as a punishment for not giving you data
16 about what level I might be willing to pay? So I
17 think that is a difference as opposed to say, you
18 know, what we have been doing for generations. It is
19 not different versus what we have been doing for the
20 last 10 or 15 years.

21 MS. WORTHMAN: Following up on that, though,
22 is there -- even though these are problems that we
23 have faced before, are there any particular harms that
24 are new based on price discrimination from AI or that
25 is a result of AI? Any new types of harms or is this

1 just the same thing that we have seen before?

2 MR. CALO: Well, I think there is a huge
3 difference. I think that -- again, I do think you
4 have to group together a bunch of different
5 phenomenon. It is not AI particularly. But, you
6 know, look, for a long, long time, companies have
7 noticed that -- and not just companies like mom and
8 pop shops, everybody has noticed, that there are just
9 cognitive limitations that we all have, right? We
10 just have these limitations to our rationality and
11 that is why everything costs \$9.99, right? I mean,
12 obviously, okay?

13 There is a set of cognitive limitations that
14 behavioral law and economists, Ariely, Kahneman, and
15 so on have -- Christy Jolls at Yale -- have been
16 surfacing over a long period of time. And these are
17 things like prospect theory and status quo bias. And
18 sometimes the FTC actually intervenes and says, you
19 seem to be using status quo bias here with these
20 rebates. We are going to intervene because it does
21 not seem to be fair and you do not seem to understand
22 what is going on.

23 The issue is that even with all these
24 behavioral economists thinking about how we have
25 cognitive limitations, the list of cognitive

1 limitations is about 45 long, okay?. What artificial
2 intelligence permits you to do because it is so good
3 at pattern matching is to model what rational consumer
4 behavior would look like in a particular environment
5 and then look for deviations that are particular to
6 you, even if they are explicable. Turns out when you
7 are watching "Buffy the Vampire Slayer" on Tuesday
8 night, you are going to pay more for ice cream. I
9 know I am. But the point of the matter is that there
10 will be situations that are very, very specific to you
11 and perhaps not even have a theory behind them.

12 But what it allows is the mass production of
13 bias. That is what it allows. It allows these
14 systems to figure out where you are specifically
15 susceptible. And, indeed, we see early signs of this
16 already. I mean, you heard earlier a presentation
17 about how Netflix is showing different people posters
18 for shows based on guesses about their demographics or
19 their qualities. You know, that is part of the
20 phenomenon that in the literature is referred to as
21 persuasion profiling, the idea that not just that you
22 be matched to your interests, but that the messages to
23 you to sell you things would be matched to your unique
24 vulnerability.

25 So, for example, for some reason in your

1 life you are really worried about scarcity, well, that
2 advertisement will say, "while supplies last," right?
3 And this is the kind of move that marketers are making
4 and it is only possible because of the way that we are
5 mediated by digital technology and we have these
6 intense analytic capabilities and, respectfully, I
7 think that is an enormous distinction from what has
8 come before.

9 MS. LOPEZ-GALDOS: So obviously, before, we
10 did not have self-driving cars and now, apparently, we
11 are going to have self-driving cars. So we are going
12 to see new things happening. Now, a self-driving car
13 might just cross over a person. What I was trying --
14 the point that I was trying to make is that the
15 thought process of analyzing the problems and
16 analyzing who is at fault, what was the causality, I
17 mean, the thought process is the same. The same that
18 exists without human beings is just applied to the new
19 setting.

20 I think the theories and thought process
21 should remain -- we should not think in the abstract.
22 We should think like we have a lot of analysis in tort
23 law, for example, and we want to say who is
24 responsible, who is not. In a self-driving car, there
25 is software, hardware, there are apps, there might be

1 somebody inside the car that was doing something as
2 well. And what I mean is that in the thought -- when
3 we are analyzing who is at fault and who is liable for
4 crossing over two people, the thought process of, for
5 example, causality should be the same as without AI.
6 That is an example -- for example, of the point that I
7 was trying to make.

8 MS. GEORGE: So I am going to ask one final
9 question and then I think we are going to wrap up. It
10 is going to be a compound question. Because I like
11 that.

12 So are there ways in which the FTC should
13 expand or rethink the notions of unfairness and
14 deception when it comes to AI and what educational
15 role should the FTC play with these new technologies,
16 both for consumers and businesses?

17 Marianela, do you want to start?

18 MS. LOPEZ-GALDOS: I think it is a very good
19 final question. I think the FTC is doing a great job
20 in putting together these hearings, as I said in the
21 beginning. I think AI is just a machine learning --
22 it is at a nascent moment. I think it is very
23 important to keep having a dialogue with businesses,
24 with the community, with the consumers, with experts,
25 and see where we are going to and see whether there is

1 anything that needs to be refined, for example, of
2 existing laws or not.

3 But what is very, very important is not to
4 think in the abstract of AI. We talk about AI as if
5 -- you know, at this moment, there are marvelous
6 things that can be done. I think there is a lot of
7 potential, but I really think that before stepping and
8 regulating or saying, oh, this is going to be a
9 disaster, everything is going to be mass-biased, et
10 cetera, we really need to understand where we stand,
11 what engineers can do, what companies are working on.

12 I think companies, at least the ones that
13 CCIA represents, are willing to cooperate with the
14 authorities, are willing to engage in adopting
15 principles. And I think having an open and frank
16 dialogue about what is going on is key to make sure we
17 get the right approach. So society can really profit
18 from AI.

19 MS. GEORGE: Irene, you just want to
20 continue down the line?

21 MS. LIU: Sure. Again from the beginning, I
22 feel that the FTC framework and the existing laws are
23 sufficient and the fact that it is broad enough that
24 it can capture AI, I think that is great. I think FTC
25 has withstood the test of time because it is broad.

1 But at the same time, I do think -- I agree with
2 Marianela that it is important for the FTC to continue
3 talking to the industry, also with other regulators
4 and academics to make sure that they are abreast of
5 this nascent technology.

6 There is also movements across the
7 globe, it is not just a U.S. phenomenon, but just
8 globally. Again, there is a recent universal
9 guideline for AI that was launched in 2018 by a
10 number of data protection officers recently. The
11 World Economic Forum is working on this issue.
12 Regulators in Europe, China, have taken a deep
13 interest in AI and so there is a lot of cross-country
14 developments within AI as well that the FTC can also
15 engage in to make sure that it stays ahead in terms of
16 the policy developments around the world so that we
17 are not hindering innovation, but fostering it as
18 well. So from that perspective, I think the FTC Act
19 is moving in the right direction with these types of
20 hearings as well.

21 From an education standpoint, the FTC can
22 also play a role in educating consumers to understand
23 what is AI. Again, because it is a new technology,
24 people hear about it. We talk about it all the time
25 in Silicon Valley, but it may not be known to the rest

1 of the country. So just educating people about what
2 chatbots are, what it means when you are choosing
3 Netflix on a Tuesday night and watching "Buffy the
4 Vampire Slayer," what the impact might be. It might
5 be that your ice cream prices might go up or it may be
6 that your Netflix fee might go up if you are a more
7 avid watcher than others.

8 So just understanding the impact of the data
9 would be helpful to consumers and also encouraging
10 companies to implement AI not just to exploit data,
11 but to think about it holistically is really important
12 and encouraging companies to do that from that
13 framework of advancing society versus exploiting the
14 data is something that FTC can take on, too.

15 MR. GILLULA: So I am actually going to
16 answer the question in reverse order. In terms of
17 consumer education, I think to accomplish that
18 mission, the FTC needs a much more robust staff of
19 technologists. They have only somewhat recently
20 started having technologists on staff. I feel like
21 the FTC should have as many technologists as lawyers
22 at this point. And, obviously, that is not where we
23 are.

24 I also realize that is not in the FTC's
25 ability to change. So if you are a Congressman or a

1 Congresswoman sitting in the audience, this is my plea
2 to you is increase the funding for technologists at
3 the FTC because those technologists can help with
4 explaining AI and what to expect in a consumer
5 standpoint to consumers. They can also help explain
6 it to the lawyers at the FTC when they are doing
7 enforcement actions or they are doing investigations.
8 They can help explain it to policymakers. So I think
9 there is a real need for a really robust technical
10 staff there.

11 In terms of whether or not the FTC Act
12 sufficiently captures everything that we might worry
13 about with regards to AI, I still worry a little bit
14 about the fact that -- I mean, I guess there are two
15 parts. One is whether -- I mean, at least -- and,
16 again, you got the only nonlawyer I think on the panel
17 talking. The FTC Act -- when you are talking about
18 harms and unfair and deceptive, you are talking about
19 what is the cost-benefit analysis. And I worry a
20 little bit that when we are talking about privacy, in
21 particular -- so, again, this comes back to rolling AI
22 and big data and predictive analytics into the same
23 thing.

24 But when you are talking about privacy, what
25 may be good for society is not necessarily good for

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25

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.

s/Linda Metcalf
LINDA METCALF, CER
Court Reporter