
Docket No. FTC-2018-0101

IN THE
FEDERAL TRADE COMMISSION

Hearings on Competition and
Consumer Protection in the 21st Century

COMMENT OF
WASHINGTON LEGAL FOUNDATION

ON
ALGORITHMS, ARTIFICIAL INTELLIGENCE,
AND PREDICTIVE ANALYTICS

Corbin K. Barthold
Washington Legal Foundation
2009 Massachusetts Ave., NW
Washington, D.C. 20036
(202) 588-0302

November 14, 2018

WASHINGTON LEGAL FOUNDATION
2009 Massachusetts Avenue, NW
Washington, D.C. 20036
(202) 588-0302

November 14, 2018

Federal Trade Commission
600 Pennsylvania Avenue, NW
Washington, D.C. 20580

By the time you read this, something in it will be obsolete. That’s an unavoidable consequence of writing about artificial intelligence. There’s no keeping up with it. But this is not cause for alarm, or, for that matter, for regulation. That AI is evolving every day does not mean that replicants will soon roam among us.¹ AI is changing our lives—but incrementally and for the better.

Washington Legal Foundation is a nonprofit, public-interest law firm and policy center that promotes free enterprise, limited government, and the rule of law. We are pleased to contribute to the Federal Trade Commission’s *Hearings on Competition and Consumer Protection in the 21st Century*. Here we respond to the FTC’s request for comment on Hearing #7, an investigation of the “consumer protection issues associated with the use of algorithms, artificial intelligence, and predictive analytics in business decisions and conduct.”²

In the first half of this comment we will review the current state of AI. We present this discussion of deep neural networks, driverless cars, Go-playing computers, etc., for two reasons. *First*, the FTC has asked to hear about AI fundamentals.³ We hope to help shed some light on that subject. *Second*, regulation and misinformation should not walk together. AI is beset

¹ Philip K. Dick, *Do Androids Dream of Electric Sheep?* (1968).

² FTC Hearing #7: Competition and Consumer Protection in the 21st Century, <https://www.ftc.gov/news-events/events-calendar/ftc-hearing-7-competition-consumer-protection-21st-century>.

³ *Id.* (seeking information about “the current and potential uses” of AI, and asking questions such as, “How are these technologies affecting . . . industries[?]” and “How might they do so in the future?”).

by hyperbole. The recent spate of sweeping statements made about it—including by prominent figures such as Elon Musk—should not inspire heavy-handed regulation. A close look at the state of AI reveals that the techno-panic is, for the foreseeable future, quite overblown.

The benefits of AI are clear, while the potential harms are speculative. The near-term impact of AI is positive, while the medium- and long-term impacts are impossible to know. The FTC should therefore proceed with caution. In the second half of this comment, we will provide some guidelines that can, we believe, help the FTC avoid doing more harm than good. The FTC should:

- *Be Predictable.* Let innovators innovate without fear of capricious unfair-practices prosecutions under the FTC Act.
- *Adhere to FTC Act § 5(n).* Take seriously the need to prosecute only practices that cause an unavoidable and unjustifiable “substantial injury.” Never bar an innovative use of AI based on speculative harms.
- *Understand Incentives.* Let companies create the reasonable solutions that consumers demand. Do not short-circuit free-market experimentation with strict rules that are hard to amend.
- *Respect Revealed Preferences.* Do not treat consumers like dimwits. Do not disrupt their free choices.
- *Take Stock of History.* Do not participate in manias. Do not overreact to discrete scandals or misadventures.
- *Favor American Interests.* Do not be the cause of innovation arbitrage.

The core theme of these guidelines is restraint. In the coming years AI will benefit us immeasurably—if we don’t get in the way.

I. The Uncertain—But Promising—State of AI.

A. AI Progress: Frontiers or Boundaries?

The debate over the state of AI is clouded by a few difficulties:

- AI is a complex set of technologies. Progress in different areas is and will remain uneven. One area can languish while another thrives. In recent years neural-network-based AI has flourished as rule-based “Good Old-Fashioned AI” has hit obstacles. In coming years the momentum might reverse—or some other, previously obscure area of research might leap ahead.
- As a result of this uneven progress—and of external factors such as the state of investment, regulation, and public opinion—AI will affect different areas of life unevenly. Cancer-diagnosis programs might become universal before driverless cars do—or vice versa.
- AI attracts a great deal of hype. Remote or implausible breakthroughs—e.g., general AI, human immortality—are often treated as imminent and inevitable events.
- The hype surrounding AI is exacerbated by the broad and doom-laden but speculative pronouncements of certain eminent individuals.

The second two factors—unwarranted hype and prophecies of doom—ensure that the first two factors—uneven progress and uneven effect—get short shrift in public debate.

How rapidly is AI progressing? Is it leaping across frontiers, or is it crawling over boundaries? A review of some of the most important AI topics, the most successful AI products, and the most prominent AI pundits reveals that, although AI is indeed moving forward, the pace of progress is uncertain and often overstated.

1. AI Frontiers.

Watson. IBM’s Watson system became famous in 2011, when it trounced two of the all-time great champions of the quiz show *Jeopardy!* Watson’s *Jeopardy!* victory was a remarkable achievement in natural-language processing, data processing, and pattern matching.⁴

The first challenge for IBM’s researchers was simply getting Watson to understand the clues. They programmed Watson to diagram each clue’s sentences and then dissect them in search of useful information. Breaking down a clue this way helped Watson to engage in “focus recognition”—identifying which clause in which sentence of the clue asks a question—and “answer-type recognition”—identifying what the clue is seeking (a name? a date?). Watson received many rules for doing this. If a clue had a noun phrase such as “This man,” for instance, Watson knew the clue was probably seeking the name of a man. If a clue was just one word—“Coleslaw”—Watson knew to look to the title of the entire category of clues—“The Main Ingredient”—for information about what type of answer was sought.

Watson treated each clue as a massive research project. Entering terms derived from the clue, Watson used a custom search engine to comb a 200-million-page database—Wikipedia was a cornerstone—for relevant entries. The searches would generate several hundred possible answers to the clue. Watson applied a light filter to these—e.g., retain only those results that contain the clue’s “answer type”—and then conducted new searches on the remaining hundred or so candidates by using elements of the clue combined with elements of each candidate as search terms.

Watson then applied a hundred or so simple metric functions to the second round of search results. Applied to each sentence in each article that appeared in the search results, the metrics combed the data for things like sentence diagrams or rare words that appeared in both a search-result sentence and the clue. Running the metric functions and then summing the results produced a score for each sentence. The sentence with the highest score was used as Watson’s answer.

⁴ The following discussion of Watson is derived primarily from Sean Gerrish, *How Smart Machines Think*, pp. 171-206 (2018), and Jo Best, *IBM Watson: The Inside Story of How The Jeopardy-Winning Supercomputer Was Born, And What it Wants to do Next*, TechRepublic, <https://www.techrepublic.com/article/ibm-watson-the-inside-story-of-how-the-jeopardy-winning-supercomputer-was-born-and-what-it-wants-to-do-next/>.

Jeopardy! is, of course, a first-to-the-buzzer quiz show, so for every clue this process had to occur in a matter of seconds. And it did. Watson won almost twice as much money as the two human contestants combined. Ken Jennings, one of Watson’s human opponents, wrote, below his Final Jeopardy response, “I for one welcome our new computer overlords.”

To triumph, however, Watson needed ten racks of servers. Watson relied heavily on a “brute force” or “GOFAI” (“Good Old-Fashioned AI”) approach, a computing method that draws on pre-programmed rules and raw computing power. All Watson could do, moreover, was play a precisely jury-rigged form of *Jeopardy!*

Deep Neural Networks. Many tasks a human completes each day are too complex and too context-specific for pre-programmed rules. Humans tend, in fact, to be quite bad at breaking down and explaining how they know what they know.⁵ Good luck explaining, in a set of rules, how to recognize whether a picture contains a chair.

A deep neural network enables a computer to recognize complex patterns without having to rely on a complex body of pre-set rules.⁶ The trick is to use an array of algorithms that loosely mimics, mathematically instead of chemically, the way a human brain uses inputs to tune neural pathways. With enough neural pathways and enough data—enough inputs; enough examples—a deep neural network can “learn” to recognize sounds or images or even optimal game moves.

The neurons in a deep neural network are layered. The first layer is an input layer. It breaks the input down into its constituent parts. In a black-and-white image-recognition network, for example, each of the first layer’s

⁵ Irving Wladawsky-Berger, *What Machine Learning Can And Can't Do*, Wall St. J., <https://blogs.wsj.com/cio/2018/07/27/what-machine-learning-can-and-cannot-do/> (Jul. 27, 2018) (“‘We know more than we can tell,’ said [Michael] Polanyi in what’s become known as Polanyi’s paradox. This common sense phrase succinctly captures the fact that we tacitly know a lot about the way the world works, yet aren’t able to explicitly describe this knowledge. Tacit knowledge is best transmitted through personal interaction and practical experiences. Everyday examples include speaking a language, riding a bike, and easily recognizing many different people, animals, and objects.”).

⁶ The following discussion of deep neural networks is derived primarily from Gerrish, note 4 above, at 109-13, 129-31, and Max Tegmark, *Life 3.0: Being Human in The Age of Artificial Intelligence*, 72-80 (2017). For a tutorial, see Brandon Rohrer, *How Deep Neural Networks Work*, YouTube, <https://www.youtube.com/watch?v=ILsA4nyG7I0> (Mar. 2, 2017).

neurons might attach only to a single pixel. Each input neuron will generate a numeric value—say, a number between zero and one—based on the grayscale brightness of the pixel to which it attaches.

In the next layer, each neuron generates a weighted sum of the values of some set of the input neurons. And the next layer does the same thing with the previous layer. Because the neurons are connected by a range of different weighting functions, the input values will, as they course through the network, cause some neurons to produce higher weighted-sum output values—cause them, that is, to “fire” brighter—than others.

As the weighted sums become larger and larger aggregates, they come to represent more and more of the image’s pixels. An early-layer neuron might “fire” for a pattern as simple as a black pixel above a white pixel. A middle-layer neuron will “fire” if presented a pattern of patterns—a line or a curve. By the end-layers, a neuron might “fire” only if presented a pattern so complex as to loosely resemble a chair.

At the end of the web of neurons lies an output layer. Each neuron in this layer matches with a potential result. Each neuron might represent a different kind of chair. If the weighted sum of a set of neurons in the preceding layer causes one of these output neurons to “fire,” the network concludes that the image contains a chair in it.

The deep neural network just described would not come into being already “firing” accurately when presented various kinds of chairs. The neuron weights in a deep neural network are, at the beginning, entirely random. The network gains accuracy from being fed many already-labeled images. Each time the network offers a random result from being fed one of these images, the correct numeric value is “back-propagated” through the network. An algorithm is used, in other words, to adjust the weights within the network closer to what’s needed to produce a right answer. The more samples the network is fed, the more finely tuned and accurate its neuron weights become.

The concept of a deep neural network has been around for decades, but progress was long hampered by technical and mathematical hurdles. In recent years, however, researchers have made great progress. In 2010 the first ImageNet image-recognition competition was held. Programs were presented images from a 1.4 million-image database and asked to identify a

thousand categories of objects. The error rate of the top entrant was 28 percent. Deep neural networks first appeared in the competition two years later, in 2012, and they dropped the error rate to 16 percent. Today the error rate is down to 2.3 percent.⁷

AlphaGo. Like Watson, Deep Blue, the computer that beat Garry Kasparov in chess in 1997, used a “brute force” or “GOFAI” protocol. It systematically mapped out future moves and, through raw processing power, examined them—200 million of them per second—and picked the most promising one.

That approach won’t work for the ancient board game of Go. Go players take turns placing white and black stones on a grid. The object of the game is to occupy the most territory on the board. The grid is 19-by-19. This produces around 2×10^{170} playable game positions. “If every atom in the universe were itself an entire universe full of atoms, there would still be more possible Go games than atoms.”⁸ The game is so complex that “top players are often at a loss to explain their own strategies.”⁹

DeepMind, a company acquired by Google in 2014, designed AlphaGo to overcome this problem. AlphaGo was fed 30 million Go positions to study. This data, along with self-play, was fed through AlphaGo’s deep neural networks. This enabled AlphaGo to “learn,” through a mind-boggling amount of trial-and-error, what good moves look like. This knowledge, in turn, enabled AlphaGo to get good at narrowing down options and examining only the potential moves most likely to bring advantage.¹⁰

In March 2016 AlphaGo played Lee Sedol, a Korean Go master, in a best-of-five match. AlphaGo won four games to one. As a result of its capacity to learn, AlphaGo was more than just a powerful calculator with a repository of human Go tactics. At one point AlphaGo made a move so unexpected that Lee left the room in shock. AlphaGo appeared to have placed a stone far from

⁷ Gerrish, note 4 above, at 133-142.

⁸ Andrew McAfee and Erik Brynjolfsson, *Machine, Platform, Crowd: Harnessing Our Digital Future*, p. 2 (2017).

⁹ *Id.*

¹⁰ *Id.* at 4; Gerrish, note 4 above, at 229-48.

anywhere of strategic importance. Fifty moves later, however, the game had shifted to this area, and AlphaGo's move gave it a decisive advantage.

AlphaGo's victory over Lee (and others) is by now well known. Less well known is DeepMind's development of AlphaGo Zero, a version of AlphaGo that beat the old version—the one that defeated Lee—100 games to zero. Unlike its predecessor, AlphaGo Zero is entirely self-taught; it learned Go solely by playing 29 million games against itself. Here is a form of AI that is not constrained by the amount of data that humans can feed it.¹¹

Recently DeepMind had AlphaZero teach itself to play chess. It did so in four hours. It then played Stockfish 8, a top computer-chess program—one trained primarily with data from past human games—a hundred times. AlphaZero never lost; it won 28 games and drew the other 72. AlphaZero won in part by making innovative and unexpected moves that were, at first blush, inexplicable to human chess masters. “Having learned the game without human input—apart from the rules—AlphaZero produced strategic masterpieces that stunned both the chess and AI worlds.”¹² What's more, Stockfish 8 lost while examining 70 million positions per second. AlphaZero, with its ability to focus on promising moves, won while examining only 80,000 positions per second.¹³

Driverless Cars. The promise of driverless cars is immense. Switching to them could save a million lives a year and generate trillions of dollars in economic growth.¹⁴

¹¹ Tom Simonite, *This More Powerful Version of AlphaGo Learns on its Own*, Wired, <https://www.wired.com/story/this-more-powerful-version-of-alphago-learns-on-its-own/> (Oct. 18, 2017). AlphaGo Zero even uses less computing power than its predecessor. Sam Shead, *DeepMind's Human-Bashing AlphaGo AI is Now Even Stronger*, Business Insider, <https://www.businessinsider.com/deepminds-alphago-ai-gets-alphago-zero-upgrade-2017-10> (Oct. 18, 2017).

¹² David W. Smith, *Google Deep Mind's "Alien" Chess Computer Reveals Game's Deeper Truths*, Eureka, <https://eureka.eu.com/innovation/deep-mind-chess/> (Mar. 20, 2018). Chess fans interested in a breakdown of one of AlphaZero's most brilliant moves are invited to see Chess.com, *AlphaZero vs Stockfish Chess Match: Game 10*, YouTube, <https://www.youtube.com/watch?v=DiJeP9oifsc> (Dec. 13, 2017), starting at 14:15.

¹³ Richard Waters, *Techmate: How AI Rewrote The Rules of Chess*, Financial Times, <https://www.ft.com/content/ea707a24-f6b7-11e7-8715-e94187b3017e> (Jan. 12, 2018).

¹⁴ Adrienne LaFrance, *Self-Driving Cars Could Save 300,000 Lives Per Decade in America*, The Atlantic, <https://www.theatlantic.com/technology/archive/2015/09/self-driving-cars-could-save-300000-lives-per-decade-in-america/407956/> (Sep. 29, 2015); Aarian Marshall, *After Peak Hype, Self-*

In the last fifteen years researchers have made great strides. Progress was ignited by a series of challenges put on by DARPA (the Defense Advanced Research Projects Agency) starting in 2004. In the 2004 contest—a “race” on a closed Mojave-Desert highway—the most successful entrant was a Humvee that ricocheted off this and that off-course object before grinding to a halt, on a rock, after 7.4 miles.¹⁵ In the 2007 challenge, however, (some) driverless vehicles could stay on the road, navigate through intersections and around other drivers, and even park.¹⁶

The software in self-driving cars is becoming ever-more sophisticated. The cars use fast and accurate laser scanners and cameras to see the world.¹⁷ They use data to tune their predictions about road conditions and other actors. They can even engage in higher-level reasoning. Confronted with an obstacle, they can ask, “Can I safely get around it?” and then, if not, “Can I back up and go a different route?”¹⁸ The goal is to ensure that driverless cars can deal with unexpected situations, such as occurred a few years ago when a self-driving car stopped for a woman in a wheelchair chasing a duck in the middle of the road.¹⁹

AI Apocalypse. Some smart people are alarmed at these advances. Elon Musk, for example, believes AI is humanity’s biggest existential threat. He frets about the creation of “a fleet of artificial intelligence-enhanced robots

Driving Cars Enter The Trough of Disillusionment, Wired, <https://www.wired.com/story/self-driving-cars-challenges/> (Dec. 29, 2017) (“Intel estimates self-driving cars could add \$7 trillion to the economy by 2050, \$2 trillion in the US alone—and that’s not counting the impact the tech could have on trucking and other fields.”).

¹⁵ Gerrish, note 4 above, at 19-22.

¹⁶ *Id.* at 41-42.

¹⁷ Jamie Condliffe, *A New Sensor Gives Driverless Cars a Human-Like View of The World*, MIT Tech. Review, <https://www.technologyreview.com/s/609718/a-new-sensor-gives-driverless-cars-a-human-like-view-of-the-world/> (Dec. 11, 2017).

¹⁸ Gerrish, note 4 above, at 42-50.

¹⁹ Google Self-Driving Car Project Monthly Report, <http://static.googleusercontent.com/media/www.google.com/en//selfdrivingcar/files/reports/report-0815.pdf> (Aug. 2015).

capable of destroying mankind.”²⁰ The only way to save humanity, Musk thinks, is “a merger of biological intelligence and machine intelligence.”²¹

Shortly before his death, Stephen Hawking, too, expressed concern. He thought we might be heading toward an “intelligence explosion,” a runaway AI evolution that leads to computers that exceed our mental capacity “by more than ours exceeds that of snails.”²²

It hardly needs saying that the opinions of luminaries such as Musk and Hawking are likely to influence the public’s perception of AI. Their opinions are, indeed, likely to have a much greater influence than the opinions of AI computer scientists.²³

2. AI Boundaries.

Watson Revisited. Soon after Watson’s *Jeopardy!* win, IBM turned the technology toward the problem of diagnosing medical ailments. The move made sense. Watson’s strength, after all, is digesting heaps of data and then recalling pertinent information in response to an array of triggers. Doctors struggle to remain informed about advances in the medical field. Watson, meanwhile, can in theory know and apply the content of every medical journal article. The hope—as a relentlessly positive 2016 *60 Minutes* piece noted—is that Watson “can help diagnose diseases and catch things the doctors might miss.”²⁴

²⁰ Maureen Dowd, *Elon Musk’s Billion-Dollar Crusade to Stop The A.I. Apocalypse*, Vanity Fair Hive, <https://www.vanityfair.com/news/2017/03/elon-musk-billion-dollar-crusade-to-stop-ai-space-x> (April 2017).

²¹ *Id.*

²² Abigail Higgins, *Stephen Hawking’s Final Warning For Humanity: AI is Coming For Us*, Vox, <https://www.vox.com/future-perfect/2018/10/16/17978596/stephen-hawking-ai-climate-change-robots-future-universe-earth> (Oct. 16, 2018).

²³ See also Henry A. Kissinger, *How The Enlightenment Ends*, The Atlantic, <https://www.theatlantic.com/magazine/archive/2018/06/henry-kissinger-ai-could-mean-the-end-of-human-history/559124/> (June 2018).

²⁴ See *Artificial Intelligence Positioned to be a Game-Changer*, 60 Minutes, <https://www.cbsnews.com/news/60-minutes-artificial-intelligence-charlie-rose-robot-sophia/> (Oct. 9, 2016).

IBM has spent more than \$15 billion developing Watson. Yet recently “more than a dozen IBM partners and clients have halted or shrunk Watson’s oncology-related projects.”²⁵ Sometimes not enough data exists for Watson: it struggles to spot rare diseases. Sometimes, conversely, there’s too much data: in some areas medical knowledge is developing too fast even for Watson—or, at least, too fast for the humans trying to keep Watson up to date. Sometimes there is the right amount of data—in treatment histories, for example—but not in formats that are easy to collect and feed to a computer. And, finally, sometimes Watson simply proves less insightful and less accurate than hoped.²⁶

Although AI might one day diagnose and treat diseases, monitor chronic conditions, and even help arrest the aging process, “Watson’s struggles suggest that [this] revolution remains some way off.”²⁷

Deep Neural Networks and AlphaGo Revisited. Even in their best domain—pattern recognition—deep neural networks can still get turned around. They remain bound to a form of literalism that humans easily avoid. Present a deep neural network an abstract pattern of black, gray, and orange, and it is liable to tell you, with great confidence, that it is looking at a remote control.²⁸

Deep neural networks can do jaw-dropping things, such as self-learn complete mastery of Atari games like *Breakout* and *Pinball*.²⁹ But these games provide a simple objective—score points—in a straightforward, rule-bound environment.³⁰ Deep neural networks cannot yet aspire to sufficiency (never mind mastery) of games that require memory, exploration, or higher strategy.³¹ The decision-trees in such games remain too vast to be understood

²⁵ Daniela Hernandez and Ted Greenwald, *IBM Has a Watson Dilemma*, Wall St. J., <https://www.wsj.com/articles/ibm-bet-billions-that-watson-could-improve-cancer-treatment-it-hasnt-worked-1533961147> (Aug. 11, 2018).

²⁶ *Id.*

²⁷ *Id.*

²⁸ Gerrish, note 4 above, at 151.

²⁹ See YouTube, *Google DeepMind’s Deep Q-Learning Playing Atari Breakout*, <https://www.youtube.com/watch?v=V1eYniJ0Rnk&v1=en> (Mar. 7, 2015).

³⁰ Gerrish, note 4 above, at 121.

through current methods of backpropagation and reinforcement learning. The games' context-richness remains beyond AI's grasp.

Even AlphaZero's chess triumph is, as one computer scientist put it, "still a discrete task."³² Deep neural networks cannot yet plan or improvise. They cannot yet adapt to change or ambiguity.³³ In this regard they are at present not much further along than "GOFAI" systems such as Deep Blue or Watson.

Driverless Cars Revisited. Although driverless cars are no doubt improving, it is important "to look beyond the extravagant claims of the press releases pouring out of tech companies and auto manufacturers."³⁴ Some companies are predicting that commercially available driverless cars will be released as early as 2021.³⁵ "The reality," however, "is that there are no driverless cars anywhere today, not even in Silicon Valley, in the sense of a vehicle that could be left to its own devices to navigate city streets or even a road."³⁶ Even on roads they "know" very well, "driverless" test cars continue to require occasional human intervention.³⁷

³¹ *Id.* at 122, 250-51.

³² Samuel Gibbs, *AlphaZero AI Beats Champion Chess Program After Teaching Itself in Four Hours*, The Guardian, <https://www.theguardian.com/technology/2017/dec/07/alphazero-google-deepmind-ai-beats-champion-program-teaching-itself-to-play-four-hours> (Dec. 7, 2017).

³³ Shead, note 11 above ("AI agents today can typically excel at one task (such as a game) but they'd struggle to do multiple tasks at the same time, especially if those tasks are in different domains."); Simonite, note 11 above ("For computers, looking into the future of a board game defined by fixed rules is relatively easy. Engineers have made little progress in having them make sense of messier, everyday scenarios. When taking on a many-faceted challenge such as assembling an Ikea sofa or planning a vacation, humans draw on powers of reasoning and abstraction to plot a path forward that so far elude AI software.").

³⁴ Christian Wolmar, *The Dream of Driverless Cars is Dying*, The Spectator, <https://www.spectator.co.uk/2018/07/the-dream-of-driverless-cars-is-dying/> (Jul. 7, 2018).

³⁵ Ford, *Looking Further: Ford Will Have a Fully Autonomous Vehicle in Operation by 2021*, <https://corporate.ford.com/innovation/autonomous-2021.html>.

³⁶ Wolmar, note 34 above.

³⁷ As of 2017, Uber's self-driving cars could drive, on average, only 0.8 miles without needing human intervention. Gerrish, note 4 above, at 56. Google's Waymo is far ahead of this; their cars can, in the right conditions, go thousands of miles per intervention. *Id.* But even Google has yet to master the tricky interplay between smart cars and humans. See note 38 below.

The ongoing wave of investment in driverless cars is likely fragile. Much of it is premised on an assumption that returns are around the corner. That assumption is shaky. For one thing, autonomous cars can't yet overcome human incompetence. In March 2018, for example, an Uber self-driving car was involved in a fatal crash in Tempe, Arizona. The car failed to prevent the all-too-human collision of a jaywalker and a "safety" driver watching *The Voice* on her phone.³⁸ The crash threw Uber's self-driving car program into disarray.³⁹ In the automotive industry and elsewhere, progress with AI will likely be slowed by a widespread urge to blame machines for failing to cope with human misbehavior. Self-driving cars will likely have to bypass the "as good as humans" stage, and proceed straight to the "much better than humans" stage, before they're widely permitted. That is a big hurdle.

At any rate, if crashes or security breaches or unforeseen complications spread the belief that fully autonomous cars are more than a decade away, it will become hard for companies to justify continuing to spend hundreds of millions of dollars a quarter developing them.⁴⁰ Money could dry up quickly, pushing the driverless-car revolution out even further.

And we mustn't forget the remaining technical hurdles. Laser sensors are still expensive. Filtering relevant from irrelevant sensor data is still difficult. The process of spotting and avoiding things like rushing ambulances is still unreliable. Driverless cars integrate an immense amount of hardware and software, and much work remains to ensure that such complex

³⁸ Heather Somerville and David Shepardson, *Uber Car's 'Safety' Driver Streamed TV Show Before Fatal Crash: Police*, Reuters, <https://www.reuters.com/article/us-uber-selfdriving-crash/uber-cars-safety-driver-streamed-tv-show-before-fatal-crash-police-idUSKBN1JI0LB> (June 22, 2018); see also Alison Griswold, *Waymo's Self-Driving Car Crashed Because its Human Driver Fell Asleep at The Wheel*, Quartz, <https://qz.com/1410928/waymos-self-driving-car-crashed-because-its-human-driver-fell-asleep/> (Oct. 2, 2018) ("The safety driver unwittingly turned off the car's self-driving software by touching the gas pedal. He failed to assume control of the steering wheel, and the Pacifica crashed into the highway median.").

³⁹ Mike Isaac, et al., *Uber's Vision of Self-Driving Cars Begins to Blur*, N.Y. Times, <https://www.nytimes.com/2018/08/19/technology/uber-self-driving-cars.html> (Aug. 20, 2018) ("The crash in March . . . altered everything. Since then, Uber has steadily narrowed the scope of its autonomous vehicle operations.").

⁴⁰ *Id.* ("[Uber] needs to persuade investors that it can eventually create a sustainable profitable business. The self-driving efforts, which have been losing \$100 million to \$200 million a quarter, do little to help that case.").

mechanisms are durable. “Technology developers are coming to appreciate that the last 1 percent is harder than the first 99 percent.”⁴¹

It is not yet safe to assume that driverless cars are coming soon.⁴²

Internet of Things. “Internet of things” is shorthand for the growing network of common physical devices, such as home appliances, that contain “smart” features. An owner of Samsung’s Family Hub refrigerator, for example, can view the inside of the fridge remotely.

The hottest internet-of-things products are virtual assistants such as Amazon Echo and Google Home. These smart speakers can answer questions, play music, and control other smart home appliances.

Smart home devices are helpful and convenient, but they are not exactly passing the Turing Test. They struggle to understand forms of context that humans grasp unthinkingly. In consequence they are still “as likely to botch your request as they are to fulfill it.”⁴³ “They’ll misconstrue a question, stress the wrong syllable, offer a bizarre answer, apologize for not yet knowing some highly knowable fact.”⁴⁴

The first home robots have hit the market, but these, too, are rudimentary. “A truly great home robot is still a long way off. It will need better sensors and processing, plus a far more precise way to pick up and manipulate objects.”⁴⁵

Translation. Google first provided a language-translation program in 2006. In 2016 it switched to a system that uses deep neural networks. The

⁴¹ Marshall, note 14 above.

⁴² *Id.*; Wolmar, note 34 above (“Of the 20 or so exhibitors I spoke to [at an automobile AI convention], not a single one believed autonomous cars would be on our roads within a decade.”).

⁴³ Judith Shulevitz, *Alexa, Should We Trust You?*, The Atlantic, <https://www.theatlantic.com/magazine/archive/2018/11/alexa-how-will-you-change-us/570844/> (Nov. 2018.)

⁴⁴ *Id.*

⁴⁵ David Pierce, *The Robot Revolution is Coming, And Dang is it Cute*, Wall St. J., <https://www.wsj.com/articles/the-robot-revolution-is-coming-and-dang-is-it-cute-1540731600> (Oct. 28, 2018).

new version “demonstrated overnight improvements roughly equal to the total gains the old [version] had accrued over its entire lifetime.”⁴⁶

But Google Translate has no memory, no imagination, no life-experience, no world-awareness. It lacks, in short, all the knowledge of context that humans take for granted. Although it can decode words, it can’t grasp concepts. True “translation” remains elusive. As Douglas Hofstadter exhaustively illustrates, Google Translate remains a cross between Mrs. Malaprop and a gibberish machine.⁴⁷ The program is constantly tripped up by its ignorance of information extraneous to the cold text it’s presented.

This problem is not likely to be solved merely with more data and more processing power. “Having ever more ‘big data’ won’t bring you any closer to understanding,” Hofstadter explains, “since understanding involves *having ideas*, and lack of ideas is the root of all the problems for machine translation today.”⁴⁸

AI Apocalypse Revisited. When discussing the risk that AI will run amok, pessimists sometimes invoke a hypothetical about a “paperclip maximizer” that, ruthlessly following its prime directive to make paperclips, converts the earth into a giant paperclip mill. Musk discusses a strawberry-picking AI, allowing him wryly to ask whether we’ll bumble into creating “strawberry fields forever.”⁴⁹ But as Steven Pinker observes, these scenarios “depend on the premises that (1) humans are so gifted that they can design an omnipresent AI, yet so moronic that they would give it control of the universe without testing how it works, and (2) the AI would be so brilliant that it could figure out how to transmute elements and rewire brains, yet so imbecilic that it would wreak havoc based on elementary blunders of misunderstanding.”⁵⁰

⁴⁶ Gideon Lewis-Kraus, *The Great A.I. Awakening*, The N.Y. Times Mag., <https://www.nytimes.com/2016/12/14/magazine/the-great-ai-awakening.html> (Dec. 14, 2016).

⁴⁷ Douglas Hofstadter, *The Shallowness of Google Translate*, The Atlantic, <https://www.theatlantic.com/technology/archive/2018/01/the-shallowness-of-google-translate/551570/> (Jan. 30, 2018).

⁴⁸ *Id.*

⁴⁹ Dowd, note 20 above.

⁵⁰ Steven Pinker, *Enlightenment Now: The Case for Reason, Science, Humanism, and Progress*, p. 300 (2018); cf. Dowd, note 20 above (“Some sniff that [Elon] Musk is not truly part of the whiteboard culture and that his scary scenarios miss the fact that we are living in a world where it’s hard to get your printer to work.”).

And we have seen bold claims about AI before. “Herbert Simon, one of the pioneers of AI, forecast in 1965 that computers would be able to do any work a human was capable of within 20 years.”⁵¹

Dramatic predictions make headlines. Then, when they fail, they’re promptly forgotten. (Remember the Y2K bug?)

We should be wary of any forecast that “identif[ies] a dominant contemporary trend” and assumes that it will “continue unchecked.”⁵² True, “it is in the nature of prophecies of doom that all but the last are falsified”; but it is also the case that “problems beget solutions” and that “harmful trends” are often “self-correcting.”⁵³

The reality, moreover, is that AI still requires much handholding from humans; and even with that help, it still struggles to understand basic aspects of circumstance. AI systems are *idiots savants*, able to learn only narrow tasks. It’s quite possible, in fact, that current forms of deep learning will never learn anything *other* than narrow tasks.⁵⁴

“Many AI researchers,” Pinker reports, “are annoyed by the latest round of hype (the perennial bane of AI) which has misled observers into thinking that Artificial General Intelligence is just around the corner.”⁵⁵ Each AI system still has “little ability to leap to problems it was not set up to solve. . . . A game-playing program is flummoxed by the slightest change in the scoring rules.”⁵⁶ AI remains constrained and brittle.

⁵¹ Waters, note 13 above.

⁵² Richard A. Posner, *The End is Near*, *The New Republic* (Sept. 12, 2003).

⁵³ *Id.*

⁵⁴ Waters, note 13 above (“The more complex the situation, the less clear the link between an action and its result. Reinforcement learning only works when one gets an instant reward for doing the right thing.”); Clive Thompson, *How to Teach Artificial Intelligence Some Common Sense*, *Wired*, <https://www.wired.com/story/how-to-teach-artificial-intelligence-common-sense/> (Nov. 13, 2018) (“But some heretics argue that deep learning is hitting a wall. They say that, on its own, it’ll never produce generalized intelligence, because truly humanlike intelligence isn’t just pattern recognition.”).

⁵⁵ Pinker, note 50 above, at 298.

⁵⁶ *Id.*; see also Phil Wainewright, *Why Humans Will Always be Smarter Than Artificial Intelligence*, *diginomica*, <https://diginomica.com/2018/02/15/why-humans-will-always-be-smarter-than-artificial->

B. Present Benefits of AI.

Although the long-term trajectory of AI is unknown (and unknowable), the present benefit of AI is obvious—and it is immense. For example:

- AI is indispensable for combatting viruses, spam, and malware. “With AI-powered security systems, it is now possible to flag viruses before they are officially announced or even given a name.”⁵⁷
- AI is a key tool for blocking fake news. “Humans can’t move quickly enough to identify and act on misinformation before it goes viral on a platform the scale of Facebook’s, with billions of posts produced per day.”⁵⁸ “The vast majority of Facebook’s efforts against fake news are powered by artificial intelligence, not humans.”⁵⁹
- AI is lowering information costs. Anyone who has used YouTube regularly over the years has witnessed its ability to recommend relevant videos—for example, a music video similar to the music video being played—improve dramatically. AI is also a crucial tool for flagging videos that violate YouTube’s terms of use.⁶⁰
- Smart home devices are saving energy. A smart thermostat, for example, can be turned off remotely. It can

intelligence/ (Feb. 15, 2018) (“We overestimate [AI’s] achievements and underestimate our own performance because we rarely stop to think how much we already know. All of the context we bring to interpreting any situation is something we take for granted.”).

⁵⁷ Amir Husain, *The Sentient Machine: The Coming Age of Artificial Intelligence*, pp. 69-86 (2017).

⁵⁸ Georgia Wells and Lukas I. Alpert, *In Facebook’s Effort to Fight Fake News, Human Fact-Checkers Struggle to Keep Up*, Wall St. J., <https://www.wsj.com/articles/in-facebooks-effort-to-fight-fake-news-human-fact-checkers-play-a-supporting-role-1539856800> (Oct. 18, 2018).

⁵⁹ *Id.*

⁶⁰ David Meyer, *AI is Now YouTube’s Biggest Weapon Against The Spread of Offensive Videos*, Fortune, <http://fortune.com/2018/04/24/youtube-machine-learning-content-removal/> (Apr. 24, 2018).

also learn its owner’s schedule and adjust the temperature accordingly.⁶¹

- Generative-design software is introducing fresh perspectives into fields such as engineering. Fed millions of data points about the stresses a racecar faces as it speeds around a track, for instance, such software will design a highly asymmetric chassis. Free of the constraints sometimes imposed by experience, such software can “come up with novel solutions that never would have occurred to us.”⁶²

These are concrete, non-scary benefits that AI is already providing. For the foreseeable future, AI will remain nothing more than a useful tool.⁶³

II. Guidelines For Regulating (And Not Regulating) AI.

“We don’t know” is the proper response to many questions about the future of AI.

There are grounds, however, for privileging hope over fear. “At the beginning of every cycle of innovation,” writes Mark Zuckerberg, “there’s a temptation to focus on the risks that come with a new technology instead of the benefits it will bring.”⁶⁴ But, he continues, “whoever cares about saving lives should be optimistic about the difference AI can make. If we slow down progress in deference to . . . unfounded concerns, we stand in the way of real gains.”⁶⁵

⁶¹ CNET, *Nest Learning Thermostat (3rd Gen) Review*, <https://www.cnet.com/reviews/nest-learning-thermostat-third-generation-review/> (May 11, 2018); Nest, *Nest Learning Thermostat*, <https://nest.com/thermostats/nest-learning-thermostat/overview/>.

⁶² McAfee & Brynjolfsson, note 8 above, at 113-16.

⁶³ Cf. Stanford University, One Hundred Year Study on Artificial Intelligence, *Artificial Intelligence And Life in 2030*, p. 8, https://ai100.stanford.edu/sites/default/files/ai_100_report_0831fnl.pdf (Sept. 2016) (“AI will likely replace tasks rather than jobs in the near term. . . . AI will also lower the cost of many goods and services, effectively making everyone better off.”).

⁶⁴ Mark Zuckerberg, *Ensure That Artificial Intelligence Helps Rather Than Hurts Us*, Wired, <https://www.wired.com/2016/10/obama-six-tech-challenges/> (Oct. 18, 2016).

⁶⁵ *Id.*

The last thing the FTC (or any other agency) should do is try to steer the progress of a nascent but promising new technology. As Zuckerberg observes, “We didn’t rush to put rules in place about how airplanes should work before we figured out how they’d fly in the first place.”⁶⁶

This comment does not contain comprehensive “solutions.” It’s not just that they’re not needed; it’s that they can’t be crafted. Experts are terrible at predicting the path even of established industries.⁶⁷ Crafting a set of rules to ensure that AI turns out “right” is a fool’s errand.

Rather than try to “solve” anything, the FTC should simply proceed with some basic principles—in truth, restraints—in mind.

Be Predictable. Innovative new technology does not require innovative new tests of the boundaries of the FTC’s power. Quite the contrary. To ensure that innovation has room to run, the FTC should stick to narrow enforcement actions that punish true bad actors.

The FTC’s recent enforcement action against LabMD, Inc., illustrates what *not* to do.⁶⁸ In violation of LabMD’s computer policy, a LabMD employee downloaded a peer-to-peer file-sharing application onto a company computer. The employee then designated her “My Documents” folder as available for sharing on the peer-to-peer network. That folder contained a 1,718-page document with 9,300 LabMD customers’ birth dates, social-security numbers, and lab-test codes. A data-security company discovered the exposure and alerted LabMD, at which point the file-sharing application was removed.

The FTC filed an administrative complaint. The FTC ultimately concluded that the exposure of the file to a peer-to-peer network had caused intangible privacy harm. The exposure amounted, in the FTC’s view, to an “unfair practice” under FTC Act § 5. The FTC ordered LabMD to fix this

⁶⁶ *Id.*

⁶⁷ See Philip E. Tetlock and Dan Gardner, *Superforecasting: The Art And Science of Prediction* (2016); Louis Menand, *Everybody’s an Expert: Putting Predictions to The Test*, *The New Yorker*, <https://www.newyorker.com/magazine/2005/12/05/everybodys-an-expert> (Dec. 5, 2005).

⁶⁸ *LabMD, Inc. v. FTC*, 894 F.3d 1221 (11th Cir. 2018).

“unfair practice” by implementing a “reasonable” data-security program. The FTC’s order was to remain in effect until at least 2036.

Vacating the order, the Eleventh Circuit faulted the FTC for using the document’s exposure “as an entry point to broadly allege that LabMD’s data-security operations [we]re deficient as a whole.”⁶⁹ The disruption of this camel-nose-under-the-tent approach was exacerbated, the court concluded, by the FTC’s demand that LabMD “meet an indeterminable standard of reasonableness” rather than simply abstain from specific acts or practices.⁷⁰ The FTC improperly anointed itself a fine-imposing micro-manager of LabMD’s data-security program.

The FTC’s action against LabMD was at once aggressive and vague. Companies hoping to avoid LabMD’s fate were offered no guidance about what data-security measures to implement. They learned only two things. One: they should each have a “reasonable” but “comprehensive” policy. Two: if the FTC decides that a policy is lacking, the consequences are dire. It is impossible, armed only with this information, to set a predictable or manageable budget for data security. To the question, “How much should my company spend on data security?” the FTC in effect answered, “More.”

This approach—especially when it’s used to punish *intangible* harms—is highly problematic. Companies will overspend on data security. Companies will hesitate to implement AI-driven security measures that the FTC can, with the benefit of hindsight, fault as newfangled. And, finally, with such uncertain (but certain to be high) penalties for data breaches, companies will avoid testing creative, ingenious, socially beneficial ways of using data.

When dealing with applications of AI, the FTC should seek to remedy only concrete harms. Prosecuting intangible harms (if there is such a thing) will lead to over-enforcement of the FTC Act and under-encouragement of invention.

And the FTC should focus (or let others focus) on the real bad guys. Hackers who use bots to engage in ad fraud—artificially inflating the number

⁶⁹ *Id.* at 1230.

⁷⁰ *Id.* at 1236.

of clicks an advertisement receives—are using AI malevolently.⁷¹ The authorities should focus on this type of AI misuse, not on legitimate AI innovation. (True, hackers are hard to find, whereas companies operate in plain sight. That’s not a good reason to pick on companies.)

A company rolling out a new use of AI should know that if the AI is causing no *actual* harm, and if no fraud or theft or vandalism is occurring, the company need not expect to hear from the FTC. A simple, predictable enforcement regime will enable AI innovation to flourish. Conversely, “inappropriate regulatory activity”—such as haphazard FTC lawsuits—“would be a tragic mistake.”⁷²

Adhere to FTC Act § 5(n). The FTC should not treat its power to prosecute “unfair” practices as a roving commission to “do good.” Fortunately, the FTC Act itself provides useful guidance on how to avoid this kind of mission creep. Section 5(n) (15 U.S.C. § 45(n)) provides:

The Commission shall have no authority under this section . . . to declare . . . that [an] act or practice is unfair unless the act or practice [1] causes or is likely to cause substantial injury to consumers which is [2] not reasonably avoidable by consumers themselves and [3] not outweighed by countervailing benefits to consumers.

In other words, the FTC may prosecute as “unfair” only unavoidable and unjustifiable sources of substantial injury.

The FTC should stringently adhere to § 5(n) when it encounters AI. Until AI develops further, we will not know the full extent of its benefits, the true extent of its harms, or the scope of humans’ ability to adapt to it. No discrete application of AI should be curtailed or banned as “unfair” based on rank speculation.

⁷¹ Lucy Handley, *US And UK Join Up to Tackle Ad Fraud, a \$50 Billion Problem*, CNBC, <https://www.cnbc.com/2018/10/23/us-and-uk-join-up-to-tackle-ad-fraud-a-50-billion-problem.html> (Oct. 23, 2018).

⁷² Stanford University, *One Hundred Year Study*, note 63 above, at 10 (“Faced with the profound changes that AI technologies can produce, pressure for ‘more’ and ‘tougher’ regulation is probably inevitable. . . . [But] poorly informed regulation that stifles innovation, or relocates it to other jurisdictions, would be counterproductive.”).

Take advertising—a core area of FTC concern.⁷³ The retailer and marketing pioneer John Wanamaker supposedly said, “Half the money I spend on advertising is wasted; the trouble is I don’t know which half.” The internet helped chip away at this problem by providing a platform for easy A/B testing. The marketer creates two online advertisements, ad “A” and ad “B,” identical in all respects but one. She then sees which ad draws more clicks. She then repeats this process over and over, tweaking this and then that, gradually optimizing the attractiveness of her advertisement.

AI is revolutionizing this process. With AI a computer can generate its own A/B tweaks, run many A/B-tests at once, and run many A/B variables at once (making an A/B test more of an “A-to-Z test”).⁷⁴ Above all, AI can optimize ad variables and ad placements based on consumers’ shopping and internet browsing habits.

Do better advertisements benefit consumers? Arguably they do. It is easy to forget what a nuisance it was, in the early days of the internet, to be followed around by pop-up ads for irrelevant products. Better-placed, more relevant advertisements improve a website’s user experience. They also lower information costs, enabling consumers to find useful goods and services more efficiently. Showing ads for car insurance to a car shopper can save her some time. Even better, *not* showing ads for car insurance to someone who neither owns nor wants a car avoids a deadweight social loss.

Better advertising is part of a larger movement toward frictionless transactions. Amazon “sends us ads for products we probably want when we’re already in a buying mood, allows us to click on those products, and, without even making us reenter our credit-card number or address, ships those products to our front door.”⁷⁵

⁷³ Federal Trade Commission, *Advertising And Marketing on The Internet: Rules of The Road*, <https://www.ftc.gov/tips-advice/business-center/guidance/advertising-marketing-internet-rules-road> (Dec. 2000).

⁷⁴ Husain, note 57 above, at 147-49; Emily Alford, *How AI Could Make A/B Testing a Thing of The Past*, ClickZ, <https://www.clickz.com/how-ai-could-make-a-b-testing-a-thing-of-the-past/216302/> (Aug. 10, 2018).

⁷⁵ Alana Semuels, *The Amazon Selling Machine*, The Atlantic, <https://www.theatlantic.com/technology/archive/2018/10/amazon-selling-machine/574045/> (Oct. 25, 2018).

Hand-wringing about unbridled consumerism often accompanies advances in retail, and the advance wrought by AI in advertising is no exception.⁷⁶ Such concerns stand, however, on an unfortunate underlying assumption—that people can’t manage their affairs on their own and need to be protected from “too much” convenience.

There is no administrable legal standard the FTC can use to determine whether “too many clicks, too much time spent, and too much money spent” on a website such as Amazon.com are “bad for our collective financial, psychological, and physical health.”⁷⁷ And even if one existed, the adoption of such a standard would transform the FTC into a Committee of Virtue, charged with defining the point at which people have spent “too much” money or wasted “too much” time on what they like.

Those who worry about Amazon assume in effect that it generates false consciousness. Amazon provides not advertisements that add value, but advertisements that “*feel* like” they add value.⁷⁸ But the FTC is neither equipped nor authorized to speculate in this fashion about the “true” desires of the masses. (We say “masses,” of course, because it is almost always those *other* people whom the would-be central planners claim need “protection” from low transaction costs.)

So far as we know, people are not living in back alleys, huddled for warmth around barrel fires, because they could not stop shopping on Amazon.com. This to be sure is not what Amazon’s detractors say is going to happen. But unless it really does happen, third-party critiques of consumers’ online shopping habits will retain a distinct whiff of high-brow disdain for the behavior of “common” people. Disdain for the citizenry’s free choices is not a ground for FTC intervention.

Does AI-driven, non-deceptive advertising—or any other AI-driven innovation—cause “substantial injury” that is neither “reasonably avoidable

⁷⁶ See, e.g., Alex Salkever, *Amazon Has a Massive New Division—And No One’s Paying Attention to It*, *Fortune*, <http://fortune.com/2018/10/08/amazon-advertising-business-platform/> (Oct. 8, 2018) (“As Amazon’s artificial intelligence algorithms divine more and more of our inner desires, the system will grow more and more adept at pushing our buttons . . . Inducing hyper-consumption [in this way] could do significant harm to Americans who . . . remain mired in credit card debt.”).

⁷⁷ *Id.*

⁷⁸ *Id.* (emphasis added).

by consumers themselves” nor “outweighed by countervailing benefits”? When confronted with such a question, the FTC should say, “Cough up the evidence.” Concrete evidence—not philosophical ruminations. In particular, the evidence that harms outweigh benefits—and are likely to continue to do so—should be substantial. The side seeking FTC intervention should have to overcome a presumption that, in point of fact, we don’t know how the challenged technology, and its effect on society, will evolve.

Understand Incentives. Technology companies have nothing to gain from suffering data breaches or exposing embarrassing information about their users. And although they admittedly have an incentive to use algorithms to sell merchandise and optimize advertising, they also have an incentive (1) to do so in a fashion users like and (2) to avoid gaining a reputation as tricksters or Svengalis.

Government agencies do not suffer for mistakenly restricting innovation consumers want in the way that technology companies suffer for introducing innovation that consumers do not want.⁷⁹ The parties with the most to lose from failing to strike the right balance between innovation, transparency, and privacy—the balance, that is, that *consumers*, not politicians or pundits or theoreticians, want—are the technology companies themselves.

There is more. If a company tries charging a fee for an ad-free product, but the experiment flops with consumers, the company can quickly revert to providing a free product with data-optimized advertisements. But if the government declares (1) that a consumer may simply opt-out of data sharing yet (2) that a company may not treat consumers who opt-out differently than consumers who don’t, the result is likely to be fewer free products on the internet—whether consumers like it or not.⁸⁰ It is, in short, much easier to change a business practice than to amend a law or regulation. The FTC should be careful not to issue regulations that lock sub-optimal rules in place.

⁷⁹ Frank H. Easterbrook, *The Inevitability of Law And Economics*, Legal Education Review, <http://www.austlii.edu.au/au/journals/LegEdRev/1989/2.html> (1989).

⁸⁰ See Corbin K. Barthold, *Problems With The California Consumer Privacy Act*, Daily Journal, <https://www.callawyer.com/articles/348062-problems-with-the-california-consumer-privacy-act> (June 26, 2018).

The major technology companies are working hard to strike the right compromises.⁸¹ It might be said that the companies are simply reacting to new laws passed by the European Union and California. To some degree this is true, but these governments are also a bit like rain dancers jumping in front of the storm clouds. There was, for example, zero chance that, absent government intervention, Facebook would have taken *no* action in response to the Cambridge Analytica scandal.

A word of caution. Although technology companies should be afforded some room to respond to incentives and self-regulate, the focus should be on what such companies do rather than on what they say. Like all companies, technology companies have an incentive to try to raise rivals' costs. A technology company that relies less on data than its rivals, for example, is incentivized to push for strict government regulation of data. This is an odious form of rent-seeking.

Respect Revealed Preferences. “Americans, if one may judge from their behavior (or ‘revealed preference,’ in the language of economics), are not greatly worried about advertisers’ obtaining personal information about them. Otherwise they would not buy books from Amazon.com or use Gmail.”⁸²

Many people will immediately, almost instinctively, think this statement—a statement written by the prominent retired federal judge Richard Posner—is somehow wrong. But why?

Is it that people are being bamboozled? Such a conclusion assumes selective rationality: the companies know what they’re doing but the customers don’t. Our systems of economics and government rely on a contrary assumption—that people are, on the whole, rational and autonomous. And research bears this assumption out. People use free services such as Google and Facebook because they would rather purchase content with data and advertisement-viewing than with money.⁸³

⁸¹ See Google, *Data Transparency*, <https://safety.google/privacy/data/>; Facebook, *Data Policy*, <https://www.facebook.com/policy.php>.

⁸² Richard A. Posner, *Just Friends*, *The New Republic*, <https://newrepublic.com/article/76815/just-friends> (July 21, 2010).

⁸³ McKinsey & Company, *Consumers Driving The Digital Uptake*, pp. 5-6, 14-16, 19-23 https://archive.iab.com/www.iab.net/media/file/White-Paper-Consumers-driving-the-digital-uptake_FINAL.pdf (Sep. 2010).

Activists have every right to raise a hue and cry about the importance of data and the peril of data collection. But consumers need not listen and agree. As the Niskanen Center notes in its comment, consumers are free to conclude that the “data exhaust” they produce as they surf the web has no inherent value.⁸⁴ There is nothing wrong with seeing free email, free search engines, and free social media as the product of a miraculous Pareto-optimal alchemy. Consumers get to trade something they don’t value for something they do.

If a search engine’s advertisements are irrelevant or intrusive, a consumer is free to use a different service. And if a social-media platform fails to block noxious content, a consumer is free to go elsewhere. She is also free, by the way, to stay right where she is. In the debate over the impact of “fake news,” many Americans are sometimes depicted as unreflective rustics.⁸⁵ But it’s unlikely that internet users as a group are irredeemably ignorant. Surveys suggest, for example, that the furor over “fake news” has taught consumers to be cautious about news published by untraditional sources.⁸⁶

Is the issue with Posner’s statement not that consumers need protection, but simply that regulators could by intervening make internet products *better*? But why would the FTC know better than companies and consumers what the best AI or data or privacy policy should look like? Would consumers be better off if regulatory limits on Amazon’s use of data and AI made its products more expensive?⁸⁷ Given people’s predilection for shopping on Amazon.com, the answer is, “Not likely.” But at the very least the answer

⁸⁴ Ryan Hagemann, Submission No. FTC-2018-0056-D-0019, p. 15, https://www.ftc.gov/system/files/documents/public_comments/2018/08/ftc-2018-0056-d-0019-155108.pdf (Aug. 20, 2018).

⁸⁵ See, e.g., Evan Osnos, *Reading The Mueller Indictment: A Russian-American Fraud*, The New Yorker, <https://www.newyorker.com/news/news-desk/reading-the-mueller-indictment-a-russian-american-fraud> (Feb. 16, 2018) (“At the heart of the Russian fraud is an essential, embarrassing insight into American life: large numbers of Americans are ill-equipped to assess the credibility of the things they read.”).

⁸⁶ BusinessWire, *‘Fake News’ Reinforces Trust in Mainstream News Brands: Hits Reputation of Social Media Sources*, <https://www.businesswire.com/news/home/20171030006313/en/Fake-News-Reinforces-Trust-Mainstream-News-Brands> (Oct. 30, 2017).

⁸⁷ See Dirk Auer, *The Amazon Investigation And Europe’s “Big Tech” Crusade*, Truth on the Market, <https://truthonthemarket.com/2018/10/21/the-amazon-investigation-and-europes-big-tech-crusade/> (Oct. 21, 2018).

is, “We don’t know.” And “if you don’t know what is best, let people make their own arrangements.”⁸⁸ The FTC should be careful not to damage or destroy transactions in the name of “improving” them.⁸⁹

But isn’t there a problem with contracts of adhesion? There is, perhaps, when a company can hide what it does with consumers’ data. Measures requiring data transparency have something to be said for them. Requiring companies to disclose what data they collect, and what they do with it, lowers transaction costs by clarifying the implicit price of the companies’ products.

A case for data transparency, however, is not a case for widespread regulation of what can and cannot be done with data. Contractual terms of data use should be treated like anything else that affects a product’s price. As Frank Easterbrook explains:

Just as no one would think of saying that the buyer of a computer with a 500 MB disk really is “entitled” to a 750 MB disk, or a faster disk, on the ground that disk size and speed is a “contract of adhesion,” so it is foolish to complain about contract terms. These all are mediated by price. “Better” terms (as buyers see things) support higher prices, and sellers have as much reason to offer the terms consumers prefer (that is, the terms that consumers find cost-justified) as to offer any other ingredient of their products. It is essential to enforce these terms if markets are to work.⁹⁰

Consumers may not know exactly what terms and prices they want—but they know better than the government does. When it regulates in a way that forecloses options—that, for example, turns a pay-with-data product into a pay-with-money product—the government just creates contracts of adhesion of its own.

⁸⁸ Frank H. Easterbrook, *Cyberspace And The Law of The Horse*, 1996 U. Chi. Legal F. 207, 210 (1996).

⁸⁹ Thomas Sowell, *Intellectuals and Society*, p. 102 (2009) (“To say, as John Dewey did, that there must be ‘social control of economic forces’ sounds good in a vague sort of way, until that is translated into specifics as the holders of political power forbidding voluntary transactions among the citizenry.”).

⁹⁰ Easterbrook, note 88 above, at 214-15.

AI is improving internet security⁹¹ and lowering transaction costs.⁹² It is generating consumer surplus. And it is continuing to evolve. Consumers appear happy with all of this. The FTC should not throw a wrench in things. Rather than try to “match an imperfect legal system to an evolving world we understand poorly,” we should “permit the participants in this evolving world to make their own decisions.”⁹³

Neither private-sector incentives nor consumers’ revealed preferences justify an absolute laissez-faire approach to AI regulation. Hackers should be punished. Online scams should be eliminated. But the FTC should resist nebulous calls to “do something” about AI.

Take Stock of History. It’s normal to worry about the future. But being faced with new things is not itself a new thing. It’s instructive to look at how fears over new things tend to pan out.

Videogames, rock and roll, telephones, and railroads—even the waltz—have each had their turn as a bugbear of busybodies.⁹⁴ Regardless, however, of whether the menace is a change in culture, a change in economics, or a change in technology, the end result is usually the same. Fear arises; driven by “thought leaders” and the media, it crescendos; as dystopia fails to materialize, it fades; and, eventually, it’s forgotten altogether. The fusspots move on to decrying the next big change.

When the first Kodak portable cameras appeared, concern spread about the “Kodak fiend”—a hypothetical person sneakily snapping photos of people in a state of “hilariousness” and then sharing them.⁹⁵ “What with Kodak

⁹¹ Husain, note 57 above, at 69-86.

⁹² Hagemann, note 84 above, at 4 (AI is ensuring that “individuals receive better-targeted, and perhaps fewer, ads while continuing to enjoy zero-priced digital services”).

⁹³ Easterbrook, note 88 above, at 215-16.

⁹⁴ See Tom Standage, *The Culture War*, Wired, <https://www.wired.com/2006/04/war/> (Apr. 1, 2006); Nick Gillespie, *That Time Al And Tipper Gore Teamed Up to Brand Prince a Public Menace*, Reason, <https://reason.com/blog/2016/04/21/that-time-al-and-tipper-gore-teamed-up-t> (Apr. 21, 2016); Joseph Hayes, *The Victorian Belief That a Train Ride Could Cause Instant Insanity*, Atlas Obscura, <https://www.atlasobscura.com/articles/railway-madness-victorian-trains> (May 12, 2017).

⁹⁵ Daniel Castro and Alan McQuinn, *The Privacy Panic Cycle: A Guide to Public Fears About New Technologies*, Information Technology & Innovation Foundation, pp. 10-12, <http://www2.itif.org/2015-privacy-panic.pdf> (Sep. 2015).

fiends and phonographs and electric search lights,” declared an 1890 newspaper article, “modern inventive genius is certainly doing its level best to lay us all bare to the gaze of our fellow men.”⁹⁶

Similar fears arose when Google started gathering images for its Street View system in 2007. The technology had changed, but the worry that people would lose “anonymity as they move about their daily lives” remained the same.⁹⁷ But the clamor over Street View didn’t last. The product is helpful, and, once again, the concerns were more smoke than fire.

Granted, privacy activists’ concerns are usually not pure hysteria, and their activism can smooth some rough edges. In response to privacy concerns, for instance, Google took to blurring people’s faces in Street View.⁹⁸ Note, however, how well such a move illustrates the value of one of our earlier guidelines: understand incentives. Had the government imposed onerous restrictions on Street View, the program might be a minimally useful shell of its current self. Instead Google was allowed to adjust to privacy concerns on its own, and, over time, Street View flourished in its full form and became accepted as the handy product it is.

When presented with a new form of AI, the FTC should meticulously separate potentially valid concerns from mere fears of novelty. If the FTC rushes to act in response to a fleeting panic, two unfortunate things will happen: (1) in the short term, innovation will be hindered; and (2) in long term, the FTC will look foolish.

The FTC should remember, in short, that fears about new technologies “generally dissipate over time without the need for policymakers to intervene.”⁹⁹

Favor American Interests. The Chinese and Russian governments do not obsess over whether algorithms are “transparent” or “inclusive” or “fair.”

⁹⁶ *Id.* at 11-12.

⁹⁷ *Id.* at 26.

⁹⁸ Stephen Shankland, *Google Begins Blurring Faces in Street View*, CNET, <https://www.cnet.com/news/google-begins-blurring-faces-in-street-view/> (May 13, 2008).

⁹⁹ Castro and McQuinn, note 95 above, at 10.

They have no compunction about developing and using AI, as both a tool and a weapon, to the fullest extent possible.¹⁰⁰ “The one who becomes the leader in this sphere,” says Vladimir Putin, “will be the ruler of the world.”¹⁰¹

As Nicholas Thompson and Ian Bremmer explain, the stakes in the race to develop the best AI are high:

A country that strategically and smartly implements AI technologies throughout its workforce will likely grow faster, even as it deals with the disruptions that AI is likely to cause. Its cities will run more efficiently, as driverless cars and smart infrastructure cut congestion. Its largest businesses will have the best maps of consumer behavior. Its people will live longer, as AI revolutionizes the diagnosis and treatment of disease. And its military will project more power, as autonomous weapons replace soldiers on the battlefield and pilots in the skies, and as cybertroops wage digital warfare.¹⁰²

China enjoys “two fundamental advantages over the US in building a robust AI infrastructure.”¹⁰³ As an authoritarian government, it has (1) easier access to its citizens’ data and (2) the ability to direct the activity of its tech companies.¹⁰⁴ These are not advantages we should seek to emulate. We have our own advantages, such as a dispersed network of innovation and a free-market economy that encourages entrepreneurship. But because China gains from centralization, we can’t afford to throw away the gains of decentralization. We can’t afford to implement burdensome restrictions and regulations that China would never dream of adopting.

¹⁰⁰ See Christopher Mims, *Inside The New Industrial Revolution*, Wall St. J., https://www.wsj.com/articles/inside-the-new-industrial-revolution-1542040187?mod=hp_lead_pos8 (Nov. 12, 2018) (“[China] is leveraging Big Data to create a form of surveillance capitalism—in which economic value is created and controlled through the harvesting of data about every aspect of people’s daily activities—the likes of which the world has never seen.”).

¹⁰¹ Nicholas Thompson and Ian Bremmer, *The AI Cold War That Threatens Us All*, Wired, <https://www.wired.com/story/ai-cold-war-china-could-doom-us-all/> (Oct. 23, 2018).

¹⁰² *Id.*

¹⁰³ *Id.*

¹⁰⁴ *Id.*; Shelly Banjo, *China Protectionism Creates Tech Billionaires Who Protect Xi*, Bloomberg, https://www.bloomberg.com/amp/news/articles/2018-03-06/how-china-protectionism-creates-tech-billionaires-who-protect-xi?__twitter_impression=true (Mar. 6, 2018).

If the Chinese government supports AI innovation while the American government punishes it, we will have no grounds for complaint when China becomes the world's dominant military and economic force.

* * *

Elon Musk might well be right that, in the long run, humans will give way to digital beings. In *the long run*. Love that outcome or hate it, it's not something to lose sleep over.

We are nowhere close to achieving general AI—AI that can understand context; AI that can plan; AI that can adjust, as humans so deftly do, to an ever-shifting world. The best AI systems are still just (very good) pattern-finders. They categorize pictures; they create search results; they reduce heating bills; they run ride-sharing platforms; they recommend other songs you might like. They help.

Let them keep helping.



Corbin K. Barthold
Washington Legal Foundation
2009 Massachusetts Ave., NW
Washington, D.C. 20036
(202) 588-0302