

Does *Prediction Machines* Predict Our AI Future? A Review[†]

AQ1

LAURENCE KOTLIKOFF*

AQ2

Abstract to come^{???} (JEL C45, C55, E26, M10, M20, O30)

Three of the world’s top artificial intelligence (AI) experts—Ajay Agrawal, Joshua Gans, and Avi Goldfarb (henceforth, AGG), all business professors at the University of Toronto’s Rotman School of Management—have produced a big-think business book on the second (after climate change) scariest economic issue of our day: the AI revolution.

AGG view current AI as almost exclusively improving prediction, hence their title. When they consider driverless trucks, they see prediction—a vast array of sensors sending voluminous instructions to the truck to step on the gas, turn on the lights, hit the brakes, honk the horn, use the signals,

et cetera—all based on intelligent guesses about the optimal coordinated use of these driving instruments. The instructions reflect deep machine learning based on vast amounts of training data, reinforced (updated) by real-time data.

When they consider online advertising, AGG think prediction. They think about Facebook’s and Google’s massive surveillance of our web browsing and online purchasing, which leads to real-time, pop-up ads predicted to entice us to buy what we want, if not necessarily what we need.

When AGG consider drone warfare, they think about all the data and analysis used to predict the enemy’s position, civilian casualties, and the best targeting method. And when they consider robots replacing factory workers, they see software telling hardware to do X, Y, and Z based on what the software

* Boston University.

[†] Go to <https://doi.org/10.1257/jel.20191519> to visit the article page and view author disclosure statement(s).

predicts will produce the best product as quickly, cheaply, and inhumanely (as in not requiring humans) as possible.

AI is hard to wrap one's brain around. It's refreshing, then, to have AGG drive a stake in the ground and claim that "The current generation of AI provides tools for prediction and little else" (p. 133). Unfortunately, the authors don't provide a clear definition of AI, but for good reason. It's not easy to define. Saying, as AGG do, that the current generation of X does Y when you haven't clearly defined X is dicey. It seems like the authors are inviting readers to share their position that, "We can't say what AI is, but we know it when we see it."

Since AGG don't precisely define AI, let me take a stab. *AI constitutes tools for calculation, the storage of prior calculations, the retrieval of stored calculations, including the calculations of predictions, and the use of these calculations to guide and assist machines and people in performing a virtually unlimited number of tasks.*

If you accept my AI definition, it's tough to agree with AGG's forceful statement that today's AI is all about prediction. Today's AI embeds yesterday's AI. Hence, it seems impossible to even define today's AI. Moreover, we have new AI being developed all over the planet on a continual basis. How can the authors know that virtually all of the new AI—the new calculations, the new calculation capacities, and the new means of accessing calculations—involves either the use or production of predictions?

To make matters concrete, consider getting driving directions and calculating distances from point A to point B. There are a finite number of routes and yesterday's global positioning system (GPS) technology could and did identify and store them all. In addition, it could and did compute or retrieve prior calculations of the shortest route between the points and provided directions for that route. This is AI with no resort to prediction or production of predictions.

Today's GPS will predict the best route in the presence of uncertain congestion. That's an improvement, but it's still just an added calculation to the set of existing calculations. Moreover, the code for predicting the shortest route surely accesses the code that came before. That is, the new prediction-based AQ4 calculations surely build on the old non-prediction based calculations.

The new GPS calculations—advising which route to take—are important, but seem less fundamental than the old calculations. Indeed, as the authors confess, predictions can undermine themselves if they induce behavior, for example, everyone taking the "shortest" route, which is contrary to what the prediction assumed. In the extreme, predictions leading to recommended behavior that change the predictions can make the predictions worthless. If, for example, everyone driving from point A to point B is using Waze, the equilibrium will be to be indifferent between a set of paths from A to B. In this case, the old technology is all one needs because we can randomly choose among that set of routes.

The authors are focusing, as one would expect in a business book, on the next big thing. But their focus on today's hot new AI applications may obscure new AI that is not prediction based, but which could change the economy far more fundamentally. Take checkout machines at CVS, which are getting better all the time. Initially, they were side shows. Then they became the main means of checking out, but with plenty of human minders in their distinct CVS uniforms. Now the human minders are disappearing. Soon the checkout machines will also be history, as scanners will record our CVS purchases, not probabilistically, but precisely as we exit the store. In the process, they will electronically charge our accounts.

Another example is transacting in major cities in China, which is almost exclusively done today by cell phone. There is no

prediction involved, but this change, which will spread worldwide, spells the end of the underground economy. Street vendors in China can't survive without the scanning technology they need to process transactions. This requires them to obtain a checking account from a bank, which can, and surely does, automatically deduct taxes on a transaction by transaction basis and also report the street vendors' annual receipts to the tax authorities.

This is the boring, mundane side of AI that the authors poo poo. This may reflect their excitement with the AI incubation lab they run in Toronto, whose firms are surely selected to focus on prediction. But for a country like Russia, which needs a 34 percent payroll tax to fund its pension system—a rate that's terribly high because the underground economy is so large—moving to digital transactions will mean a massive increase in payroll tax revenues and permit a major decline in the payroll tax rate. This will also provide the government extra income tax and value-added tax (VAT) revenue, which it can use to provide badly needed public goods. I view this particular means by which old AI can transform the developing world as far more important to the world economy than, say, driverless vehicles.

AQ5

The above said, the authors appear as close to the AI ground as anyone. So, they may very well be right. These days, new work in AI may really be almost all about prediction. Thus, having done my reviewer duty to pick discernable nits, let me consider what the book tells us about prediction machines.

The authors' major, if commonsense, point is that the ever-expanding capacity to use big data and machine learning to make predictions will reduce the prices of products that rely on predictions and undermine commercial enterprises that rely on traditional ways of forming predictions.

AGG point to the focus of big tech companies on prediction. Tesla's vehicles are

amazing prediction machines that, the authors tell us, even use on-board radar to keep us from hitting obstacles that aren't visible until it's too late. Amazon and other companies are considering sending us products we haven't ordered, but that we are highly likely to want—all predicted based on our spending patterns and those of others like us. Google is shoving ads in our faces predicated on the content of our emails, the nature of our searches, and our conversations with Siri. These ads are predicted to grab our attention and empty our wallets. New, improved prediction machines are diagnosing our illnesses. They are helping sports teams choose which athletes to hire to maximize the chances of a winning season. They are also helping political groups, domestic and foreign, decide where to place ads and how to design the ads to influence our decisions to vote, and for whom and what to vote. And they are helping us find love by assisting dating apps to improve the chances we find a successful match. Weather forecasting, while far from perfect, is now dramatically better thanks to prediction machines. So too is detection of the fraudulent use of our credit cards and the decision about a would-be borrower's creditworthiness. Perhaps the most important prediction machine raised by AGG is real-time monitoring of your heart using a smart watch that checks your pulse. *Prediction Machines* is full of mind bending examples of these and other new technologies that are changing, or will change, economic and social life as we know it.

However, as AGG make clear, prediction machines are going to creatively destroy vast numbers of existing products and jobs. The black cabbies in London are an interesting example raised by the authors. Their *knowledge*—memorization of every street and route between points A and B in London—was obsolesced by GPS. But before GPS was married to prediction of where traffic was heaviest, the black cabbies retained

something special—the experience of traffic patterns and congestion and knowledge of shortcuts to avoid traffic. Uber drivers, equipped with Waze, can now compete with the cabbies and do so at a cheaper price, since they don't need to spend years learning *the knowledge*. If Uber drivers aren't already better in navigating the streets of London, they surely will be as Waze acquires ever more data to reinforce its stored knowledge. As a result, black cabbies will be priced out of the market and, almost surely, lose their middle-class lifestyles.

Another profession headed for the dustbin is radiology. Prediction machines are making radiologists redundant. Why? Because they are rapidly generating far more reliable predictions of the presence of tumors and other forms of morbidity than humans peering at shadowy films using their retinas. Indeed, in a few decades, the job title of radiologist may be as quaint as switchboard operator. Yes, a few radiologists may be retained to make sure the prediction machines aren't missing something that wasn't part of their training (what the authors call unknown knows), but no one can now safely choose radiology as a profession.

Yet a third example of worker replacement by prediction machines is translators. Surely lip-reading sensors embedded in our glasses coupled with AirPods in our ears will soon let us speak with anyone on the planet regardless of their language. The translation machines will be prediction machines because they will need to predict what words are being spoken based not only on the movement of the speaker's lips, but the topic, location, and circumstance of the conversation.

The trillion-dollar question, which the book raises in its middle/gut, is the impact of the old and new AI on workers and the economy. Unfortunately, the answer AGG deliver is a combination of "We don't know." and "What, me, worry?"

Here are AGG's answers to five big AI-anst questions.

- "Will there still be jobs? *Yes.*"
- "Will (AI) generate more inequality? *Perhaps.*"
- "Will a few large companies control everything? *It depends.*"
- "Will countries ... forfeit our privacy and security to (achieve) a competitive advantage?" *Some will.*"
- "Will the world end?" *You still have ... time to (read) this book.*"

The What, me, worry? attitude reminds me of a recent tour I took of a 150 year-old distillery in Scotland. The guide proudly told our group that five computer operators now do the work of 500 workers. I heard no audible groans, but also no cheers. The guide showed no sign of emotion, but surely worried at a deep level that her position would soon be filled by one of those perpetually smiling information robots now roaming our malls.

AGG do point to the decline in labor's share of national income and the actual and potential loss of jobs. But they also argue that human judgment and intervention will always be a necessary factor. One example they provide is having an adult travel along in a self-driving school bus to prevent the children from fighting. Another is needing a human to ride shotgun on a convoy of trucks, which might be stopped by highwaymen trying to pilfer the trucks. If the trucks are programmed not to run anyone over, it will be easy enough to stop the trucks. But this will assist the thieves, hence the need for Wyatt Earps. AGG also claim that as long as people can hide their true feelings from the prediction machines, the prediction machines will be unable to know everything we know, leaving us some advantage. And they tell us that humans will always be needed to police machines.

AGG seem to ignore the potential for TV monitors and other devices placed inside the bus to deter bad behavior or drones to ride

shotgun over the trucks. Their suggestions that humans hide from the machines but also police them are also rich. It's clear that AGG needed to work overtime to imagine new things for obsolesced humans to do.

More telling are AGG's few and far-between mentions of the fact that AI, while it may not eliminate jobs (in the end, everyone can become a cleric), is already eliminating good jobs—jobs that pay decent wages. Median real weekly earnings are only 7 percent higher than they were 40 years ago. Even with the economy's decent growth, real wages have grown, since 2015, at only one-quarter of 1 percent per year.¹ There also appears to be strong evidence of a secular increase in wage inequality.² The precise role of AI in wage stagnation and wage inequality isn't clear, but when massive numbers of good jobs go “poof” from one day to the next, it's surely a big part of the story.

AQ7 My guess is that down deep, AGG are as scared about AI as anyone on the planet precisely because they see AI's full potential to automate virtually every job on the planet. Their discussion of Rio Tinto's automation of its massive mining operations in Australia is illustrative of the special knowledge the authors bring to the table. They point out that the company now has over 70 self-driving trucks that operate autonomously, that is, they don't need even a human to control them remotely. The company's stated goal is to have not a single human involved in the digging, loading, or shipping of iron ore. In other words, Rio Tinto is seeking 100 percent automation of its mining operations.

AQ8 ¹U.S. Bureau of Labor Statistics. 2015–19. “Employed Full Time: Median Usual Weekly Real Earnings: Wage and Salary Workers: 16 Years and Over—LES1252881600Q.” FRED—Federal Reserve Bank of St. Louis. <https://fred.stlouisfed.org/series/LES1252881600Q> (accessed MONTH DAY, YEAR).

²Gould, Elise. 2019. “State of Working America Wages 2018.” *Economic Policy Institute*, February 20. <https://www.epi.org/publication/state-of-american-wages-2018/>.

AGG claim that prediction machines are overcoming the final roadblock to Rio Tinto achieving this goal.

Although AGG reference actual and potential microeconomic impacts of prediction machines, they ignore macroeconomic impacts. What happens when real wages stagnate? There is less saving, less investment, larger current account deficits, and lower growth. There are also more fights over the pie both within and across generations. This is the postwar history of the US macroeconomy in a nutshell.

Another way to put the macroeconomic problem is to ask who will buy the products that the new/old AI produces if firms achieve their stated objective: producing without workers. This is a case of supply dynamically destroying, rather than creating, its own demand. This reversal of Say's law is illustrated in Benzell et al. (2019), which shows that simple dynamic models of AI can produce long-run immiseration.³ In their model, AI improves the economy's technology, but reduces wages, saving, and investment. In the long run, there is so little capital that output and living standards can be substantially lower than they would have been without the AI.

The authors aren't macroeconomists and their book is targeted primarily to the business community. So, one can understand them not delving too deeply into where AI is taking us. But one issue that deserved more space is the use of prediction machines to exercise state control over human behavior.

When I was a 16-year-old exchange student in Germany, my German guest father told me what he felt was among the very worst things about the Nazis. It was, he said, their ability to prevent parents from telling their children the truth about who the Nazis were and why they were evil. Doing so could

have led their children to repeat the lesson in school, leading to their parents' imprisonment or worse. Fifty-two years later, that conversation haunts me as I read stories emanating from China about "social" tracking of Chinese individuals.

Prediction Machines may seem like a fascinating book about the next big economic thing. But it's actually a terrifying book when you realize that governments are already using the new AI to predict who you are, what you think, how to manipulate your thinking, and to decide whether you should be denied employment or incarcerated because of your preferences and political views.

I close by answering the question: Do I recommend this book? Definitely. It's a must read for business leaders who need to know where AI is heading and how best to

harness the new technology. It's a must read for economists; it forces us to think more deeply about the essence of AI and its connection to prediction. And it's a must read for the public who need to know the enormous dangers the old/new AI poses to our own and our children's economic futures and freedoms.

REFERENCES

- Agrawal, Ajoy, Joshua Gans, and Avi Goldfarb. 2018. *Prediction Machines: The Simple Economics of Artificial Intelligence*. Boston: Harvard Business Review Press.
- Benzell, Seth G., Laurence J. Kotlikoff, Guillermo LaGarda, and Jeffrey D. Sachs. 2019. "Robots Are Us: Some Economics of Human Replacement." <https://kotlikoff.net/wp-content/uploads/2019/04/Robots-Are-Us-2019.pdf>.

PLEASE ANSWER ALL AUTHOR QUERIES (numbered with “AQ” in the margin of the page). Please disregard all Editor Queries (numbered with “EQ” in the margins). They are reminders for the editorial staff.

AQ#	Question	Response
1.	Please verify that we have included the correct affiliation in the * footnote. Also, did you wish to add any acknowledgments to this footnote?	
2.	Since review essays are longer and more detailed than regular book reviews, we publish them more like an article. As such, please provide an abstract of around 100 words or fewer.	
3.	It is JEL style to define most initialisms and acronyms at their first appearance in a paper, followed by the shortened form in parentheses, if it is used again. (Subsequent uses then just refer to the shortened form.) Did we choose the correct definition for GPS?	
4.	We follow the Chicago Manual’s suggestion of reserving i.e., e.g., and similar Latin abbreviations for footnotes and parenthetical statements. Thus, we have changed these in regular text to “that is” and “for instance” or “for example” as appropriate.	
5.	Did we choose the correct definition for VAT?	
6.	Are these questions and answers quoted from the book, or are the extrapolations? If quoted, please provide the page number(s). If they are not quoted, we will remove the quotation marks.	

PLEASE ANSWER ALL AUTHOR QUERIES (numbered with “AQ” in the margin of the page). Please disregard all Editor Queries (numbered with “EQ” in the margins). They are reminders for the editorial staff.

AQ#	Question	Response
7.	In the sentence beginning with “My guess is...,” may we delete “on the planet” at the end of the sentence, since the same phrase appears earlier in the same sentence? The new sentence would read: “My guess is that down deep, AGG are as scared about AI as anyone on the planet, precisely because they see AI’s full potential to automate virtually every job.”	
8.	In footnote 1, please provide an access date for this data.	
9.	There was a marker for footnote 3 here, but no text in the actual footnote. Did you intend to include a footnote here and, if so, what should it say?	