Artifical Intelligence In The Boardroom

by Ajay Agrawal, Joshua Gans and Avi Goldfarb

The pace of major, market-shifting technology change has increased over the past century. The latest wave, artificial intelligence (AI), brings many previous shifts together and takes them in a whole new direction. AI, along with related trends like machine learning and big data, will demand top-level strategic discussion in the boardroom to tap its dawning business opportunities—and to avoid its potential dangers.

In January 2007, when Steve Jobs paced the stage and introduced the world to the iPhone, not a single observer reacted with, "Well, it's curtains for the taxi industry." Fast forward to 2018 and that appears to be precisely the case. Over the last decade, smartphones evolved from being simply a smarter phone to an indispensable platform for tools that are disrupting or fundamentally altering all manner of industries.

Even Andy Grove, who famously quipped that "only the paranoid survive," would have to have been pretty paranoid to foresee how far and wide the smartphone would reach into some very traditional industries.

Boards realize how hard it is to forecast when, where, and how the most disruptive business changes will take place, but also the importance of staying ahead of the curve.

Recent developments in artificial intelligence and machine learning have convinced us that this innovation is on par with the great, transformative technologies of the past: electricity, cars, plastics, the microchip, the Internet, and the smartphone. From economic history, we know how these general-purpose technologies diffuse and transform. We also realize how hard it is to forecast when, where, and how the most disruptive changes will take place. At the same time, we have learned what to look for, how to be ahead of the curve, and when a new technology is likely to transition from something interesting to something transformative.

What does this mean for AI in the boardroom? Keeping up with the now daily barrage of AI-related news would overwhelm most boards. How should board chairs, who set the agenda and influence the allocation of time and attention of members, choose what AI-related topics to focus on? Thankfully, there are several methods for sorting through the noise.

AI-related business topics, as opposed to technical topics, can be grouped into three general categories.

☐ *This first category is opportunities for operational efficiency created through the deployment of AI tools.* This category is the largest in terms of commercial activity, although it requires the least amount of attention from the board.

A tidal wave of new AI tools that promise enhanced efficiency take advantage of advances in data collection, processing power, and machine learning algorithms. These tools do things like automate document processing, control robots, and respond to customer service queries. They often enable one person to do the same work that previously required many.

Despite the value of these tools, this category does not require any more attention from the board than other productivity-enhancing topics. The board must ensure that systems are in place to make certain the company is operating efficiently, employing whatever tools are necessary to achieve that.

☐ *The second category for AI concerns risks.* Corporate boards must ensure that systems are in

Ajay Agrawal, Joshua Gans *and* **Avi Goldfarb** *are professors in the Rotman School of Management at the University of Toronto. This article is excerpted from their forthcoming book* Prediction Machines: The Simple Economics of Artificial Intelligence *(Harvard Business Review Press).* *[www.rotman.utoronto.ca]*

place to assure that their organizations are protected from the downside risks associated with deploying AI.

□ *The third category concerns industry disruption.* This can represent either an opportunity or a risk. Although companies usually employ AI tools to better execute their strategy, sometimes an AI tool can change the economics of an industry so fundamentally that it leads to redefining the strategy itself. We opened with an example of a disruptive technology—the iPhone. Significant changes to strategy are of direct importance to corporate boards.

The most common question board members ask us is: "How will AI affect our strategy?" We use a thought experiment to answer that question. Most people are familiar with shopping at Amazon. As with most online retailers, you visit its website, shop for items, place them in your cart, pay for them, and then Amazon ships them to you. Right now, Amazon's business model is shopping-then-shipping.

During the shopping process, Amazon's AI offers suggestions of items that it predicts you will want. The AI does a reasonable job. However, it is far from perfect. In our case, the AI accurately predicts what we want to buy about five percent of the time. We actually purchase about one of every twenty items it recommends. Considering the millions of items on offer, that is not bad.

Now, imagine that the Amazon AI collects more information about us and uses that data to improve its predictions, an improvement akin to turning up the volume knob on a speaker dial. Rather than volume, though, it is turning up the AI's prediction accuracy.

At some point, as Amazon turns the knob, the AI's prediction accuracy crosses a threshold, changing the company's optimal business model. The prediction becomes sufficiently accurate that it is more profitable to ship you the goods that it predicts you will want rather than wait for you to order them.

You then would not need to go to other retailers, and the fact that the item is there may well nudge you to buy more. Amazon gains a higher share of wallet. Clearly, this is great for Amazon, but it is also great for you. Amazon ships before you shop, which, if all goes well, saves you the task of shopping entirely. Cranking up the prediction dial changes Amazon's business model from shopping-then-shipping to shipping-then-shopping.

If this is a better business model, then why is Amazon not doing it already? Because if implemented today, the cost of collecting and handling returned items would outweigh the increase in revenue from a greater share of wallet. Today, we would return 95 percent of the items it ships to us. That is annoying for us and costly for Amazon. The prediction is not yet good enough to adopt the new model.

Better AI predictions will attract more shoppers, more shoppers generate more data to train the AI, more data will lead to better predictions, and so on, creating a virtual cycle.

We can imagine a scenario where Amazon adopts the new strategy even before the prediction accuracy is good enough to make it profitable because the company anticipates that at some point it will be. By launching sooner, Amazon's AI will get more data sooner and improve faster. Amazon realizes that the sooner it starts, the harder it will be for competitors to catch up.

Better predictions will attract more shoppers, more shoppers will generate more data to train the AI, more data will lead to better predictions, and so on, creating a virtual cycle. Adopting too early could be costly—but adopting too late could be fatal.

Our point is not that Amazon will or should do this (although skeptical readers may be surprised to learn that Amazon already filed for and was granted a U.S. patent for "anticipatory shipping"). Instead, the salient insight is that turning the prediction dial has a significant impact on strategy.

In this example, it shifts Amazon's business model from shopping-then-shipping to shipping-then-shopping, generates the incentive to vertically integrate into operating a service for collecting the increased number of product returns (perhaps a fleet of trucks), and accelerates the timing of investment due to the economics of a first-mover advantage. All this is due simply to turning up the prediction accuracy on AI.

AI rises to the level of a strategic rather than operational decision, which is important because strategic decisions are directly relevant to corporate boards.

☐ ***How AI can change business strategy.*** The Amazon thought experiment neatly illustrates three ingredients that together cause AI adoption to rise to the level of being a strategic rather than operational decision. This is important because strategic level decisions are directly relevant to corporate boards, whereas operational decisions are not.

First, a strategic dilemma or trade-off must exist. In the Amazon case, the quandary is that shipping-then-shopping may generate more sales, but simultaneously produce more goods consumers want to return. When the cost of returning items is too high, then the ROI (return on investment) is lower than for the traditional approach. This explains why, in the absence of some technological change, Amazon maintains their traditional business model, just like almost every other retailer.

Second, the problem can be resolved through the reduction of uncertainty. In the Amazon case, it is about consumer demand. If you can accurately forecast what an individual will purchase, especially if delivered to their doorstep, then you reduce the likelihood of returns and increase sales. Uncertainty reduction hits both the benefit and cost side of the dilemma.

This type of demand management is not new. It is one of the reasons that physical stores exist. Physical stores cannot forecast individual customer demand, but they can forecast the likely demand from a group of customers. By pooling together the customers who visit a location, physical stores hedge demand uncertainty among individuals. Moving to a shipping-then-shopping model based on individual dwellings requires better predictions (precisely in the wheelhouse of AI) about individual customer demand, which can overcome the competitive advantage physical stores have.

Third, companies require an AI that can reduce uncertainty enough to change the balance in the strategic dilemma. In the Amazon case, a very accurate model of customer demand may make the shipping-then-shopping business model worth doing. Here, the benefits of increased sales outweigh the costs of returns.

AI allows consumer targeting that can be highly effective, but could also prove discriminatory and defamatory.

☐ ***What about AI risk?*** Latanya Sweeney was the chief technology officer for the U.S. Federal Trade Commission, and is now a professor at Harvard University. She was surprised when a colleague Googled her name to find one of her papers and discovered ads suggesting she had been arrested. Sweeney clicked on the ad, paid a fee, and learned what she already knew—she had never been arrested.

Intrigued, she entered the name of her colleague Adam Tanner, and the same company's ad appeared but without the suggestion of arrest. After more searching, she developed the hypothesis that maybe black-sounding names were triggering the arrest ad. Sweeney then tested this more systematically and found that if you Googled a black-associated name like Lakisha or Trevon, you were 25 percent more likely to get an ad suggesting an arrest record than if you searched for a name like Jill or Joshua.

Such biases are potentially damaging. Searchers might be looking for information to see if someone is suitable for a job. If they find ads with titles like "Latanya Sweeney, Arrested?" the searchers might have doubts. It is both discriminatory and defamatory.

Why was this happening? Google provides software that allows advertisers to test and target particular keywords. Advertisers might have entered racially associated names to place ads alongside, although Google denied that. Another possibility is that the pattern emerged as a result of Google's algorithms, which promote ads that have a higher "quality score" (meaning they are more likely to be clicked).

AIs likely played a role there. For instance, if potential employers searching for names were more likely to click on an arrest ad when associated with a black-sounding name than other names, then the quality score associated with placing those ads with such keywords might rise. Google is not intending to be discriminatory, but its algorithms might amplify prejudices that already exist in society. Such profiling exemplifies one risk of implementing AI.

☐ ***Liability.*** Discrimination might emerge in even subtler ways. Economists Anja Lambrecht and Catherine Tucker, in a 2017 study, showed that Facebook ads could lead to gender discrimination. They placed ads promoting jobs in science, technology, engineering, and math (STEM) fields on the social network, and found Facebook was less likely to show the ad to women. This was not because women were less likely to click on the ad or because they might be in countries with discriminatory labor markets.

On the contrary, the workings of the ad market discriminated. Because younger women are valuable as a demographic on Facebook, showing ads to them is more expensive. So, when you place an ad on Facebook, the algorithms naturally place ads where their *return* per placement is highest. If men and women are equally likely to click on STEM job ads, then it is better to place ads where they are cheaper—with men.

Harvard Business School professor, economist, and lawyer Ben Edelman explained to us why this issue could be serious for both employers and Facebook. While many tend to think of discrimination as arising from disparate treatment—setting different standards for men and women—the ad-placement differences might result in what lawyers call “disparate impact.” A gender-neutral procedure turns out to affect some employees who might have reason to fear discrimination (a “protected class” to lawyers) differently from others.

A person or an organization can be liable for discrimination, even if it is accidental. A court found that the New York City Fire Department discriminated against black and Hispanic applicants becoming firefighters with an entrance exam that included several questions emphasizing reading comprehension. The case was eventually settled for about $99 million. Blacks’ and Hispanics’ lower performance on the exam meant that the department was liable, even if the discrimination was unintentional.

One challenge with AI is that unintentional discrimination can happen without anyone in the organization noticing. Predictions generated by AI appear to be created from a black box.

So, while you may think you are placing a neutral ad on Facebook, disparate impact might be emerging regardless. As an employer, you could be liable. Of course, you don’t want to engage in discrimination, even implicitly. One solution for Facebook is to offer tools for advertisers to prevent discrimination.

A challenge with AI is that such unintentional discrimination can happen without anyone in the organization noticing. Predictions generated by AI appear to be created from a black box. It is not feasible to look at the algorithm underlying the prediction and identify what causes what. To figure out if AI is discriminating, you have to look at the output. Do men get different results than women? Do Hispanics get different results than others? What about the elderly or the disabled? Do these different results limit their opportunities?

To prevent liability issues, if you discover unintentional discrimination in the output of your AI, you need to fix it. You need to figure out why your AI generated discriminatory predictions. Yet if AI is a black box, then how can you do this?

Some in the computer science community call this “AI neuroscience.” A key tool is to hypothesize what might drive the differences, provide the AI with different input data that tests the hypothesis, and then compare the resulting predictions. Lambrecht and Tucker did this when they discovered that women saw fewer STEM ads because it was less expensive to show the ad to men.

The point is that the black box of AI is not an excuse to ignore potential discrimination or a way to avoid using AI in situations where discrimination might matter. Plenty of evidence shows that humans discriminate even more than machines.

Algorithmic discrimination can easily emerge at the operational level but end up having strategic and broader consequences. Strategy involves directing those in your organization to weigh factors that might not otherwise be obvious. This becomes particularly salient with systematic risks, like algorithmic discrimination, that may have a negative impact on your business.

Showing the STEM ads to men and not women bolstered short-term performance (in that the ads the men saw cost less) but created risks due to the resulting discrimination. The consequences of increasing risks may not become apparent until too late. Thus, a special task for boards is to ensure systems are in place to anticipate these risks and procedures are in place to manage them.

☐ ***Three steps to get started.*** When we work with corporate boards or CEOs and their leadership teams on getting started with AI, we follow three general steps.

▫ *Review the organization's work flows that turn inputs into outputs.* Identify the steps in each work flow where the deployment of an AI could add value by making predictions better, faster, and cheaper than are currently possible, thereby reducing uncertainty and enhancing decision-making.

▫ *Estimate the ROI from either building or buying AI to perform each of the identified prediction tasks.* Then, rank order the AIs from highest to lowest ROI, creating an "AI Heat Map" for the organization. Start investing in the AIs at the top of the list and then work down.

▫ *"Science fiction" each of the AIs.* Project the consequences of turning the dial to determine which, if any, could transform the organization's strategy as the prediction accuracy improves. If any of these scenario planning exercises suggest potential industry disruption, then plan accordingly. ■