

Prediction Machines



The Simple Economics of
Artificial Intelligence

AJAY
AGRAWAL

JOSHUA
GANS

AVI
GOLDFARB

Harvard Business Review Press
Boston, Massachusetts

2018

Contents

Cover

Title Page

Copyright

Dedication

Acknowledgments

1. Introduction: Machine Intelligence
2. Cheap Changes Everything

Part One: Prediction

3. Prediction Machine Magic
4. Why It's Called Intelligence
5. Data Is the New Oil
6. The New Division of Labor

Part Two: Decision Making

7. Unpacking Decisions
8. The Value of Judgment
9. Predicting Judgment
10. Taming Complexity
11. Fully Automated Decision Making

Part Three: Tools

12. Deconstructing Work Flows
13. Decomposing Decisions

14. Job Redesign

Part Four: Strategy

15. AI in the C-Suite

16. When AI Transforms Your Business

17. Your Learning Strategy

18. Managing AI Risk

Part Five: Society

19. Beyond Business

Notes

Index

About the Authors

2

Cheap Changes Everything

Everyone has had or will soon have an *AI moment*. We are accustomed to a media saturated with stories of new technologies that will change our lives. While some of us are technophiles and celebrate the possibilities of the future, and others are technophobes who mourn the passing of the good ole days, almost all of us are so used to the constant drumbeat of technology news that we numbly recite that the only thing immune to change is change itself. Until we have our AI moment. Then we realize that this technology is different.

Some computer scientists experienced their AI moment in 2012 when a student team from the University of Toronto delivered such an impressive win in the visual object recognition competition ImageNet that the following year all top finalists used the then-novel “deep learning” approach to compete. Object recognition is more than just a game; it enables machines to “see.”

Some technology CEOs experienced their AI moment when they read the headline in January 2014 that Google had just paid more than \$600 million to acquire UK-based DeepMind, even though the startup had generated negligible revenue relative to the purchase price but had demonstrated that its AI had learned—on its own, without being programmed—to play certain Atari video games with superhuman performance.

Some regular citizens experienced their AI moment later that year when renowned physicist Stephen Hawking emphatically explained, “[E]verything that civilisation has to offer is a product of human intelligence ... [S]uccess in creating AI would be the biggest event in human history.”¹

Others experienced their AI moment the first time they took their hands

off the wheel of a speeding Tesla, navigating traffic using Autopilot AI.

The Chinese government experienced its AI moment when it witnessed DeepMind's AI, AlphaGo, beating Lee Se-dol, a South Korean master of the board game Go, and then later that year beating the world's top-ranked player, Ke Jie of China. The *New York Times* described this game as China's "Sputnik moment."² Just as massive American investment in science followed the Soviet Union's launch of Sputnik, China responded to this event with a national strategy to dominate the AI world by 2030 and a financial commitment to make that claim plausible.

Our own AI moment came in 2012 when a trickle and then a surge in the number of early-stage AI companies employing state-of-the-art machine-learning techniques applied to the CDL. The applications spanned industries—drug discovery, customer service, manufacturing, quality assurance, retail, medical devices. The technology was both powerful and general purpose, creating significant value across a wide range of applications. We set to work understanding what it meant in economics terms. We knew that AI would be subject to the same economics as any other technology.

The technology itself is, simply put, amazing. Early on, famed venture capitalist Steve Jurvetson quipped: "Just about any product that you experience in the next five years that seems like magic will almost certainly be built by these algorithms."³ Jurvetson's characterization of AI as "magical" resonated with the popular narrative of AI in films like *2001: A Space Odyssey*, *Star Wars*, *Blade Runner*, and more recently *Her*, *Transcendence*, and *Ex Machina*. We understand and sympathize with Jurvetson's characterization of AI applications as magical. As economists, our job is to take seemingly magical ideas and make them simple, clear, and practical.

Cutting through the Hype

Economists view the world differently than most people. We see everything through a framework governed by forces such as supply and demand, production and consumption, prices and costs. Although economists often disagree with each other, we do so in the context of a common framework. We argue about assumptions and interpretations but not about fundamental concepts, like the roles of scarcity and competition in setting prices. This approach to viewing the world gives us a unique

vantage point. On the negative side, our viewpoint is dry and doesn't make us popular at dinner parties. On the positive side, it provides a useful clarity for informing business decisions.

Let's start with the basics—prices. When the price of something falls, we use more of it. That's simple economics and is happening right now with AI. AI is everywhere—packed into your phone's apps, optimizing your electricity grids, and replacing your stock portfolio managers. Soon it may be driving you around or flying packages to your house.

If economists are good at one thing, it is cutting through hype. Where others see transformational new innovation, we see a simple fall in price. But it is more than that. To understand how AI will affect your organization, you need to know precisely what price has changed and how that price change will cascade throughout the broader economy. Only then can you build a plan of action. Economic history has taught us that the impact of major innovations is often felt in the most unexpected places.

Consider the story of the commercial internet in 1995. While most of us were watching *Seinfeld*, Microsoft released Windows 95, its first multitasking operating system. That same year, the US government removed the final restrictions to carrying commercial traffic on the internet, and Netscape—the browser's inventor—celebrated the first major initial public offering (IPO) of the commercial internet. This marked an inflection point when the internet transitioned from a technological curiosity to a commercial tidal wave that would wash over the economy.

Netscape's IPO valued the company at more than \$3 billion, even though it had not generated any significant profit. Venture capital investors valued startups in the millions even if they were, and this was a new term, "pre-revenue." Freshly minted MBA graduates turned down lucrative traditional jobs to prospect on the web. As the effects of the internet began to spread across industries and up and down the value chain, technology advocates stopped referring to the internet as a new technology and began referring to it as the "New Economy." The term caught on. The internet transcended the technology and permeated human activity at a fundamental level. Politicians, corporate executives, investors, entrepreneurs, and major news organizations started using the term. Everyone began referring to the New Economy.

Everyone, that is, *except economists*. We did not see a new economy or a new economics. To economists, this looked like the regular old economy. To be sure, some important changes had occurred. Goods and services could be distributed digitally. Communication was easy. And you could

find information with the click of a search button. But you could do all of these things before. What had changed was that you could now do them cheaply. The rise of the internet was a drop in the cost of distribution, communication, and search. Reframing a technological advance as a shift from expensive to cheap or from scarce to abundant is invaluable for thinking about how it will affect your business. For instance, if you recall the first time you used Google, you may remember being mesmerized by the seemingly magical ability to access information. From the economist perspective, Google made search cheap. When search became cheap, companies that made money selling search through other means (e.g., the Yellow Pages, travel agents, classifieds) found themselves in a competitive crisis. At the same time, companies that relied on people finding them (for example, self-publishing authors, sellers of obscure collectibles, homegrown moviemakers) prospered.

This change in the relative costs of certain activities radically influenced some companies' business models and even transformed some industries. However, economic laws did not change. We could still understand everything in terms of supply and demand and could set strategy, inform policy, and anticipate the future using off-the-shelf economic principles.

Cheap Means Everywhere

When the price of something fundamental drops drastically, the whole world can change. Consider light. Chances are you are reading this book under some kind of artificial light. Moreover, you probably never thought about whether using artificial light for reading was worth it. Light is so cheap that you use it with abandon. But, as the economist William Nordhaus meticulously explored, in the early 1800s it would have cost you four hundred times what you are paying now for the same amount of light.⁴ At that price, you would notice the cost and would think twice before using artificial light to read this book. The subsequent drop in the price of light lit up the world. Not only did it turn night into day, but it allowed us to live and work in big buildings that natural light could not penetrate. Virtually nothing we have today would be possible had the cost of artificial light not collapsed to almost nothing.

Technological change makes things cheap that were once expensive. The cost of light fell so much that it changed our behavior from thinking about whether we should use it to not thinking for even a second before flipping

on a light switch. Such significant price drops create opportunities to do things we've never done; it can make the impossible possible. So, economists are unsurprisingly obsessed with the implications of massive price drops in foundational inputs like light.

Some of the impacts from producing cheaper light were easy to imagine, and others less so. What might be affected when a new technology makes something cheap is not always precisely obvious, whether the technology is artificial light, steam power, the automobile, or computing.

Tim Bresnahan, a Stanford economist and one of our mentors, pointed out that computers do arithmetic and nothing more. The advent and commercialization of computers made arithmetic cheap.⁵ When arithmetic became cheap, not only did we use more of it for traditional applications of arithmetic, but we also used the newly cheap arithmetic for applications that were not traditionally associated with arithmetic, like music.

Heralded as the first computer programmer, Ada Lovelace was the first to see this potential. Working under very expensive light in the early 1800s, she wrote the earliest recorded program to compute a series of numbers (called Bernoulli numbers) on a still-theoretical computer that Charles Babbage designed. Babbage was also an economist, and as we will see in this book, that was not the only time economics and computer science intersected. Lovelace understood that arithmetic could, to use modern startup lingo, “scale” and enable so much more. She realized that applications of computers were not limited to mathematical operations: “Supposing, for instance, that the fundamental relations of pitched sounds in the science of harmony and of musical composition were susceptible of such expression and adaptations, the engine might compose elaborate and scientific pieces of music of any degree of complexity.”⁶ No computer had been invented yet, but Lovelace saw that an arithmetic machine could store and replay music—a form that defined art and humanity.

That is precisely what happened. When, a century and a half later, the cost of arithmetic fell low enough, there were thousands of applications for arithmetic that most had never dreamed of. Arithmetic was such an important input into so many things that, when it became cheap, just as light had before, it changed the world. Reducing something to pure cost terms has a way of cutting through hype, although it does not help make the latest and greatest technology seem exciting. You'd never have seen Steve Jobs announce “a new adding machine,” even though that is all he ever did. By reducing the cost of something important, Jobs's new adding machines were transformative.

That brings us to AI. AI will be economically significant precisely because it will make something important much cheaper. Right now, you may be thinking about intellect, reasoning, or thought itself. You might be imagining robots all over or non-corporeal beings, such as the friendly computers in *Star Trek*, letting you avoid the need to think. Lovelace had the same thought, but quickly dismissed it. At least insofar as a computer was concerned, she wrote, it “had no pretensions to originate anything. It can do whatever we know how to order it to perform. It can follow analysis; but it has no power of anticipating any analytical relations or truths.”⁷

Despite all the hype and the baggage that comes with the notion of AI, what Alan Turing later called “Lady Lovelace’s Objection” still stands. Computers still cannot think, so thought isn’t about to become cheap. However, what will be cheap is something so prevalent that, like arithmetic, you are probably not even aware of how ubiquitous it is and how much a drop in its price could affect our lives and economy.

What will new AI technologies make so cheap? *Prediction*. Therefore, as economics tells us, not only are we going to start using a lot more prediction, but we are going to see it emerge in surprising new places.

Cheap Creates Value

Prediction is the process of filling in missing information. Prediction takes information you have, often called “data,” and uses it to generate information you don’t have. Much discussion about AI emphasizes the variety of prediction techniques using increasingly obscure names and labels: classification, clustering, regression, decision trees, Bayesian estimation, neural networks, topological data analysis, deep learning, reinforcement learning, deep reinforcement learning, capsule networks, and so on. The techniques are important for technologists interested in implementing AI for a particular prediction problem.

In this book, we spare you the details of the mathematics behind the methods. We emphasize that each of these methods is about prediction: using information you have to generate information you don’t have. We focus on helping you identify situations in which prediction will be valuable, and then on how to benefit as much as possible from that prediction.

Cheaper prediction will mean more predictions. This is simple economics: when the cost of something falls, we do more of it. For example, as the computer industry began to take off in the 1960s and the cost of arithmetic began to fall rapidly, we used more arithmetic in applications where it was already an input, such as at the US Census Bureau, the US Department of Defense, and NASA (recently depicted in the film *Hidden Figures*). We later began to use the newly cheap arithmetic on problems that *weren't* traditionally arithmetic problems, such as photography. Whereas we once solved photography with chemistry, when arithmetic became cheap enough, we transitioned to an arithmetic-based solution: digital cameras. A digital image is just a string of zeros and ones that can be reassembled into a viewable image using arithmetic.

The same goes for prediction. Prediction is being used for traditional tasks, like inventory management and demand forecasting. More significantly, because it is becoming cheaper it is being used for problems that were not traditionally prediction problems. Kathryn Howe, of Integrate.ai, calls the ability to see a problem and reframe it as a prediction problem “AI Insight,” and, today, engineers all over the world are acquiring it. For example, we are transforming transportation into a prediction problem. Autonomous vehicles have existed in controlled environments for over two decades. They were limited, however, to places with detailed floor plans such as factories and warehouses. The floor plans meant engineers could design their robots to maneuver with basic “if-then” logical intelligence: if a person walks in front of the vehicle, then stop. If the shelf is empty, then move to the next one. However, no one could use those vehicles on a regular city street. Too many things could happen—too many “ifs” to possibly code.

Autonomous vehicles could not function outside a highly predictable, controlled environment—until engineers reframed navigation as a prediction problem. Instead of telling the machine what to do in each circumstance, engineers recognized they could instead focus on a single prediction problem: “*What would a human do?*” Now, companies are investing billions of dollars in training machines to drive autonomously in uncontrolled environments, even on city streets and highways.

Imagine an AI sitting in the car with a human driver. The human drives for millions of miles, receiving data about the environment through their eyes and ears, processing that data with their human brain, and then acting in response to the incoming data: drive straight or turn, brake or accelerate. Engineers give the AI its own eyes and ears by outfitting the car with sensors (e.g., cameras, radar, lasers). So, the AI observes the incoming

data as the human drives and simultaneously observes the human's actions. When particular environmental data comes in, does the human turn right, brake, or accelerate? The more the AI observes the human, the better it becomes at predicting the specific action the driver will take, given the incoming environmental data. The AI learns to drive by predicting what a human driver would do given specific road conditions.

Critically, when an input such as prediction becomes cheap, this can enhance the value of other things. Economists call these “complements.” Just as a drop in the cost of coffee increases the value of sugar and cream, for autonomous vehicles, a drop in the cost of prediction increases the value of sensors to capture data on the vehicle's surroundings. For example, in 2017, Intel paid more than \$15 billion for the Israeli startup Mobileye, primarily for its data-collection technology that allows vehicles to effectively see objects (stop signs, people, etc.) and markings (lanes, roads).

When prediction is cheap, there will be more prediction and more complements to prediction. These two simple economic forces drive the new opportunities that prediction machines create. At low levels, a prediction machine can relieve humans of predictive tasks and so save on costs. As the machine cranks up, prediction can change and improve decision-making quality. But at some point, a prediction machine may become so accurate and reliable that it changes how an organization does things. Some AIs will affect the economics of a business so dramatically that they will no longer be used to simply enhance productivity in executing against the strategy; they will change the strategy itself.

From Cheap to Strategy

The single most common question corporate executives ask us is: “How will AI affect our business 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 to buy. 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 5 percent of the time. We actually purchase about

one of every twenty items it recommends. Considering the millions of items on offer, that's not bad!

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. But rather than volume, it's turning up the AI's prediction accuracy.

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

With that, you won't 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.

Of course, shoppers would not want to deal with the hassle of returning all the items they don't want. So, Amazon would invest in infrastructure for the product returns, perhaps a fleet of delivery-style trucks that do pickups once a week, conveniently collecting items that customers don't want.⁸

If this is a better business model, then why hasn't Amazon done 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. For example, today we would return 95 percent of the items it ships to us. That is annoying for us and costly for Amazon. The prediction isn't good enough for Amazon to adopt the new model.

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 profitable. 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 virtuous cycle. Adopting too early could be costly, but adopting too late could be fatal.⁹

Our point is not that Amazon will or should do this, although skeptical

readers may be surprised to learn that Amazon obtained a US patent for “anticipatory shipping” in 2013.¹⁰ 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 product returns (including a fleet of trucks), and accelerates the timing of investment. All this is due simply to turning up the dial on the prediction machine.

What does this mean for strategy? First, you must invest in gathering intelligence on how fast and how far the dial on the prediction machines will turn for your sector and applications. Second, you must invest in developing a thesis about the strategic options created from turning the dial.

To get started on this “science fictioning” exercise, close your eyes, imagine putting your fingers on the dial of your prediction machine, and, in the immortal words of Spinal Tap, turn it to eleven.

The Plan for the Book

You need to build foundations before the strategic implications of prediction machines for your organization become apparent. That is precisely how we structured this book, building a pyramid from the ground up.

We lay the foundation in part one and explain how machine learning makes *prediction* better. We move to why these new advances are different from the statistics you learned in school or that your analysts might already conduct. We then consider a key complement to prediction, data, especially the types of data required to make good predictions, and how to know whether you have it. Finally, we delve into when prediction machines perform better than humans and when people and machines might work together for even better predictive accuracy.

In part two, we describe the role of prediction as an input into *decision making* and explain the importance of another component that the AI community has so far neglected: judgment. Prediction facilitates decisions by reducing uncertainty, while judgment assigns value. In economists’ parlance, judgment is the skill used to determine a payoff, utility, reward, or profit. The most significant implication of prediction machines is that

they increase the value of judgment.

Practical matters are the focus of part three. AI *tools* make prediction machines useful and are implementations of prediction machines designed to perform a specific task. We outline three steps that will help you figure out when building (or buying) an AI tool will generate the highest return on investment. Sometimes such tools slot neatly into an existing work flow; at other times, they motivate redesigning the work flow. Along the way, we introduce an important aid for specifying the key features of an AI tool: the AI canvas.

We turn to *strategy* in part four. As we describe in our Amazon thought experiment, some AIs will have such a profound effect on the economics of a task that they will transform a business or industry. That's when AI becomes the cornerstone of an organization's strategy. AIs that have an impact on strategy shift the attention on AI from product managers and operations engineers to the C-suite. Sometimes, it's hard to tell in advance when a tool will have such a powerful effect. For example, few people predicted, when they tried it for the first time, that the Google search tool would transform the media industry and become the basis of one of the most valuable companies on earth.

In addition to upside opportunities, AI poses systemic risks that may hit your business unless you take preemptory actions. Popular discussion seems to focus on the risks AI poses to humanity, but people pay much less attention to the dangers AI poses to organizations. For instance, some prediction machines trained on human-generated data have already "learned" treacherous biases and stereotypes.

We end the book in part five by applying our economists' tool kit to questions that affect *society* more broadly, examining five of the most common AI debates:

1. Will there still be jobs? *Yes.*
2. Will this generate more inequality? *Perhaps.*
3. Will a few large companies control everything? *It depends.*
4. Will countries engage in race-to-the-bottom policy making and forfeit our privacy and security to give their domestic companies a competitive advantage? *Some will.*
5. Will the world end? *You still have plenty of time to derive value from this book.*

KEY POINTS

- Economics offers clear insights regarding the business implications of cheaper prediction. Prediction machines will be used for traditional prediction tasks (inventory and demand forecasting) and new problems (like navigation and translation). The drop in the cost of prediction will impact the value of other things, increasing the value of complements (data, judgment, and action) and diminishing the value of substitutes (human prediction).
- Organizations can exploit prediction machines by adopting AI tools to assist with executing their current strategy. When those tools become powerful, they may motivate changing the strategy itself. For instance, if Amazon can predict what shoppers want, then they may move from a shop-then-ship model to a ship-then-shop model—bringing goods to homes before they are ordered. Such a shift will transform the organization.
- As a result of the new strategies that organizations pursue to take advantage of AI, we will be faced with a new set of trade-offs related to how AI will impact society. Our choices will depend on our needs and preferences, and will almost surely be different across different countries and cultures. We structured this book in five sections to reflect each layer of impact from AI, building from the foundations of prediction all the way up to the trade-offs for society: (1) Prediction, (2) Decision making, (3) Tools, (4) Strategy, and (5) Society.