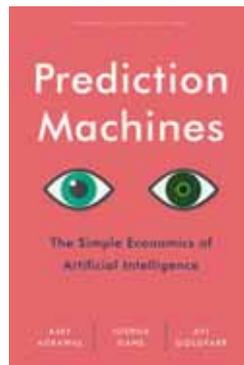


---

# VALUE OF JUDGEMENT

Artificial intelligence (AI) and automation make the seemingly impossible possible. They magically bring machines to life – driving cars, trading stocks, teaching children or providing healthcare. In *Prediction Machines*, Avi Goldfarb and two fellow economists debunk this magical myth and recast the rise of AI as nothing more than a drop in the cost of prediction. With this one masterful stroke their thinking transforms the debate, lifting the curtain on AI-as-magic and revealing how applying basic economic tools can transform how CEOs, CFOs and FDs view the AI revolution. By framing AI as merely cheap prediction, they reframe the debate. As this extract from *Prediction Machines* demonstrates, at the heart of this process is a requirement to unpack and understand the process of making decisions





*Prediction Machines* by Ajay Agrawal, Joshua Gans and Avi Goldfarb

We typically associate decision-making with big decisions. Should I buy this house? Should I attend this school? Should I marry this person? No doubt, these life-changing decisions, while rare, are important.

But we also make small decisions all the time. Should I keep sitting in this chair? Should I keep walking down this street? Should I keep paying this monthly bill? We handle many of our smaller decisions on autopilot, perhaps by accepting the default, choosing to focus all our attention on bigger decisions. However, deciding not to decide is still a decision.

Decision-making is at the core of most occupations. Schoolteachers decide how to educate their students, who have different personalities and learning styles. Managers decide who to recruit for their team and who to promote. Truck drivers decide how to respond to route closures and traffic accidents. Police officers decide how to handle suspicious individuals and potentially dangerous situations. Doctors decide what medicine to prescribe and when to administer tests.

Decisions like these usually occur under conditions of uncertainty. The teacher doesn't know for sure whether a particular child will learn better from one teaching approach or another. The manager doesn't know for sure whether a job applicant will perform well or not. The doctor doesn't know for sure whether it is necessary to administer a costly [medical] exam. Each of them must predict [the likely outcome].

But a prediction is not a decision. Making a decision requires applying

judgement to a prediction and then acting. Before recent advances in machine intelligence, this distinction was of academic interest because humans always performed prediction and judgement together. Now, advances in machine prediction mean we have to examine the anatomy of a decision.

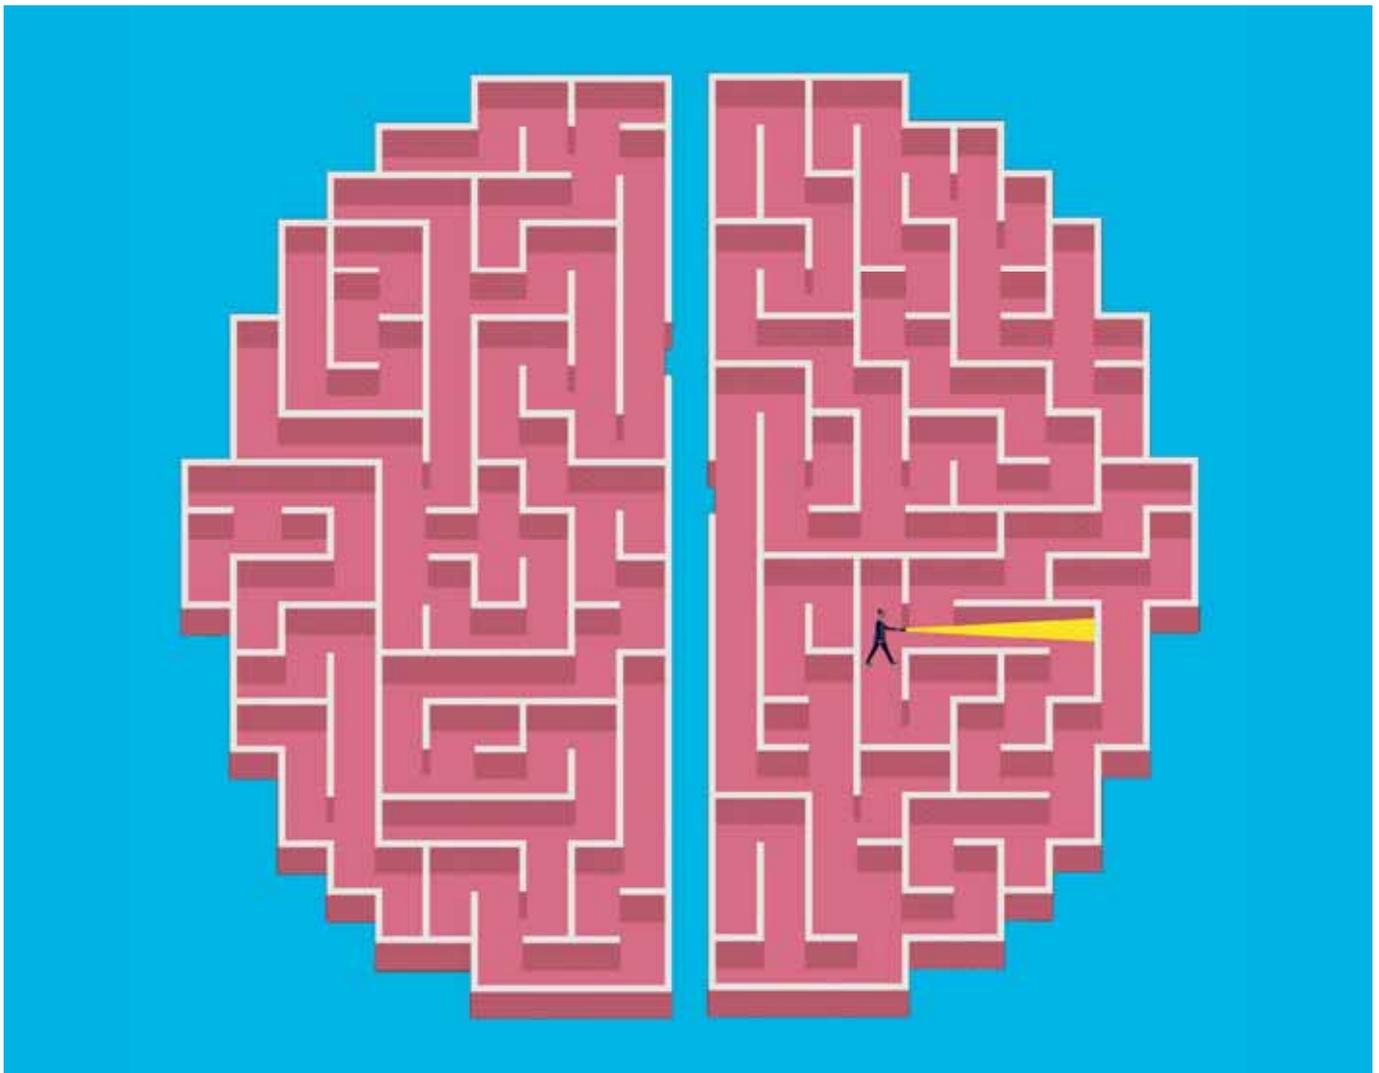
### **ANATOMY OF A DECISION**

Prediction machines will have their most immediate impact at the decision level. But decisions have six other key elements (see Figure 1). When someone (or something) makes a decision, they take input data from the world that enables a prediction. That prediction is possible because training occurred about relationships between different types of data and which data is most closely associated with a situation. Combining the prediction with judgement on what matters, the decision maker can choose an action. The action leads to an outcome, which has an associated reward or payoff. The outcome is a consequence of the decision. It is needed to provide a complete picture. The outcome may also provide feedback to help improve the next prediction.

By breaking up a decision into elements, we can think clearly about which parts of human activities will diminish in value and which will increase as a result of enhanced machine prediction. Most clearly, for prediction itself, a prediction machine is generally a better substitute for human prediction. As machine prediction increasingly replaces forecasts that humans make, the value of human prediction will decline. But a key point is that, while prediction is a key component of any decision, it is not the only component. The other elements of a decision - judgement, data, and action - remain, for now, firmly in the realm of humans. They are complements to prediction, meaning they increase in value relatively. For example, we may be more willing to exert effort by applying judgement to decisions where we previously had opted not to decide (for example, accepted the default) because prediction machines now offer better, faster and cheaper predictions. In that case, the demand for human judgement will increase.

### **LOSING THE KNOWLEDGE**

The Knowledge is a test London cabbies take to drive the city's celebrated black taxis. The test involves knowing the location of thousands of points and streets around the city and predicting the shortest



**By breaking up a decision into elements, we can think clearly about which parts of human activities will diminish in value and which will increase as a result of enhanced machine prediction**

or fastest route between two points at any time of day. The amount of information [required to be memorised] is staggering. To pass the test, potential cabbies need a near-perfect score. Passing the test takes, on average, three years, including time spent poring over maps but also riding around the city on mopeds memorising and visualising. Once they have achieved this, honoured green badge recipients are a font of knowledge.

A decade ago, London cab drivers' knowledge was a competitive advantage. No one could provide the same degree of service. People who would otherwise have walked would hop in a cab because the cab drivers knew the way. But today, a simple mobile GPS or satellite navigation system means all drivers have access to data and predictions that were once the cabbies' superpower. Today, these superpowers are available for free on most mobile phones. People don't get lost and they know the fastest route. Now the

phone is even better because it is updated in real time with traffic information. Cabbies who invested three years studying the Knowledge didn't know they would be competing with prediction machines. They took time to upload maps into their memory, test routes, and fill in the blanks with common sense. Now, navigation apps have access to the same map data and are able, through a combination of algorithms and predictive training, to find the best route whenever requested, using real-time data on traffic that the taxi driver cannot hope to know.

But the fate of cabbies rested not just on the ability for navigation apps to predict the Knowledge, but also on other elements to take the best path from A to B. First, the cabbies could control a motor vehicle. Second, they had sensors affixed to them - their eyes and ears most importantly - that fed contextual data to their brains to ensure they put the Knowledge to good use. But so did other people. No London

cabbie became worse at their job because of navigation apps. Instead, millions of other non-cabbies became better. The cabbies' Knowledge was no longer a scarce commodity, opening up cabbies to competition from ride-sharing platforms, such as Uber.

Other drivers with the Knowledge and predictions of the fastest routes on their phones meant they could provide an equivalent service. When high-quality machine prediction became cheap, human prediction declined in value, so the cabbies were worse off. The number of rides in London's black cabs fell. Others provided the same service. These others also had driving skills and human sensors, complementary assets that went up in value as prediction became cheap. One day self-driving cars might end up substituting for those skills and senses. The point is that understanding the impact of machine prediction requires an understanding of the various aspects of decisions.

**SHOULD YOU TAKE AN UMBRELLA?**

Until now, we've been a little imprecise about what judgement is. To explain it, we use decision trees. This device is especially useful for decisions under uncertainty, when you are not sure what will happen if you make a particular choice.

Let's consider a familiar choice you might face. Should you carry an umbrella on a walk? You might think that an umbrella is a thing you hold over your head to stay dry, and you'd be right. But an

**Judgement involves determining what we call the "reward function", the relative rewards and penalties associated with particular actions that produce particular outcomes**

umbrella is also a kind of insurance, in this case, against the possibility of rain. So, the following framework applies to any insurance-like decision to reduce risk. Clearly, if you knew it was not going to rain, you would leave the umbrella at home and if you knew it would rain, you would take it with you. At the root of the tree are two branches representing choices you could make: leave umbrella or take umbrella. Extending from these are branches representing what you are uncertain about: rain versus shine. Without a good weather forecast, you don't know. You might know that, at this time of the year, sun is three times more likely than rain. This would give you a three-quarters chance of sun and a

one-quarter chance of rain. This is your prediction. Finally, at the tips of the branches are consequences. If you don't take an umbrella and it rains, you get wet, and so on.

So, what decision should you make? This is where judgement comes in. Judgement is the process of determining the reward to a particular action in a particular environment. It is about working out the objective you're actually pursuing. Judgement involves determining what we call the "reward function", the relative rewards and penalties associated with particular actions that produce particular outcomes. Wet or dry? Burdened by carrying an umbrella or unburdened?

Let's assume that you prefer being dry without an umbrella (10 out of 10) more than being dry, but carrying an umbrella (eight out of 10) more than being wet (zero). This gives you enough to act. With the prediction of rain a quarter of the time and the judgement of the payoffs to being wet or carrying an umbrella, you can work out your average payoff from taking versus leaving the umbrella. Based on this, you are better off taking the umbrella (an average payoff of eight) than leaving it (an average payoff of 7.5).

If you really hate toting an umbrella (a six out of 10), your judgement about preferences can also be accommodated. In this case, the average payoff from leaving an umbrella at home is unchanged (at 7.5), while the payoff from taking one is now 6. Such umbrella haters will leave the umbrella at home.

This example is trivial. Of course people who hate umbrellas more than getting wet will leave them [at] home. But the decision tree is a useful tool for figuring out payoffs for non-trivial decisions, too, and that is at the heart of judgement. Here, the action is taking the umbrella, the prediction is rain or shine, the outcome is getting wet, and judgement is anticipating the happiness you will feel (payoff) from being wet or dry, with or without an umbrella. As prediction becomes better, faster, and cheaper, we'll use more of it to make more decisions, so we'll also need more human judgement and thus the value of human judgement will go up. ●



**Avi Goldfarb, chair, AI and healthcare, and professor of marketing, Rotman School of Management, University of Toronto**

**FIGURE 1: ANATOMY OF A TASK**

