

Joshua Gans

On managing pandemics, allocating vaccines, and low-cost prediction with AI

Economist Joshua Gans spent the past quarter century researching issues that range from digital currencies to the economics of scientific publishing, from antitrust policy to entrepreneurship, from net neutrality to artificial intelligence. Last spring, he became one of many millions who found themselves stuck in lockdown and thinking about the coronavirus. He found an outlet for his energies in researching and writing about policy responses to the crisis. The resulting book, *The Pandemic Information Gap: The Brutal Economics of COVID-19*, will be published by MIT Press in November. In a departure from usual publishing practice, reflecting the urgency of the topic, an early version of the book was released online in April under the title *Economics in the Age of COVID-19*.

A native of Australia, Gans came to the United States in 1990 to pursue his Ph.D. at Stanford University. Today, he is a professor at the University of Toronto's Rotman School of Management, where he teaches entrepreneurial strategy and the economics of artificial intelligence. Gans is also chief economist of the Creative Destruction Lab, a program for advanced technology startup companies. The organization, founded at the Rotman School and with branches at other universities, provides mentoring and networking opportunities to selected companies in technology areas that include artificial intelligence, blockchain, energy, and space.

In addition to his book on the SARS-CoV-2 pandemic, Gans is the author or co-author of, among other books, *Innovation + Equality* (MIT Press, 2019), *Prediction Machines* (Harvard Business Review Press, 2018), and *The Disruption Dilemma* (MIT Press, 2016).

David A. Price interviewed Gans via videoconference in June 2020.



EF: How did you become interested in economics?

Gans: I was interested in science fiction in high school. I read a novel by Isaac Asimov called *Foundation*. I saw what was going on in that book and in economics as sort of similar and quite interesting.

Foundation has a premise that a hero character invents a science called psychohistory. In psychohistory, you can't predict individuals, but you can predict large movements in society and social forces on a galactic scale, because you know, why not? (*Laughs.*)

The book got me interested in the possibility of being able to predict with social science in the same way that physicists were able to predict movements of planetary bodies and so on. Economics turned out to be nothing like that, but that's another matter.

I didn't think of economics as a profession until much later, but that's when I started getting interested in studying it.

MAKING SENSE OF THE CORONAVIRUS

EF: What led you to write your new book on the economics of the coronavirus? Had you done research in this area before?

Gans: What led me to write it is I didn't know what else to do. Back in March, I was stuck at home, so I decided to write a book.

On any scale of normal scholarly credentials, I didn't have any background for this book. I had done some health economics and studied some of the other topics in this book, like innovation. But beyond that, no. The main reason I decided to do it was that I figured at this time everybody who was a real expert was going to be busy. (*Laughs.*)

My idea was to explain what's going on from the eyes of an economist. The challenge was that of course things were moving very quickly. From conception to publication was a couple of days over a month, which is kind of ridiculous. MIT Press had a lot to do in that time, also. They had to have it peer reviewed because they won't just publish anything. They had to have it copy edited. They opted to do a whole lot of things in parallel that they normally do sequentially.

Another move that was unusual was that when the book went out for peer review, MIT Press also posted the draft online. Everybody could see it and comment on it. Those comments turned out to be quite valuable. With those comments and some further thinking and research, I've now written a version of the book that's twice the size, which will come out in November.

EF: Did you change your mind about anything since writing that first draft?

Gans: Yes. What's reflected in the book that's coming out is that I now see these pandemics as manageable things. Policymakers have to react right away and stay the course, but pandemics can be managed. If I had to guess how history is going to judge this period, the judgment is going to be that this shouldn't have been a two- to three-year calamity, it should have been a three-month calamity.

The need for testing aggressively at the beginning had to be appreciated. You aggressively isolate people you find who are infected, you trace who they had contact with, and you aim for quick, complete suppression. The countries that had had experience with pandemics — Hong Kong, South Korea, Taiwan, most of Africa — got it right away. They knew what the problems would be if they didn't do anything about it. So experience with viruses was definitely a factor. But Canada had that and didn't quite get its act together quickly enough. Some provinces were better than others. Quebec was way too slow and has had the worst problem. Australia and New Zealand lucked out because of their distance, which gave them time to understand what to do.

But once the virus breaks out, then you've got a problem. Then you've got to do the complete lockdown. And we're seeing places that did a complete lockdown — like they did in Italy, France, and Spain — squash it all the way

“I think this pandemic has surely disrupted everything in terms of the development of AI for normal business practices. **Because we don't know what normal is anymore.** The problem with having everything rely on a statistical model is that if you have a major structural break, those models break too.”

down. Locking down is terribly painful; that's why you don't want to go through it in the first place. But you may have to. So there's a separate factor, which is resolve — how far are you willing to go to push the spread down.

EF: Looking at this set of choices that you've outlined, where has the United States been and where do you think it should be or should've been?

Gans: Early in the crisis, people in the United States and Canada were not talking about the virus as something we needed to suppress completely. The discussion was mainly, “We're going to push down the curve, and then we'll wait for a vaccine.” But the evidence both historically and now with this virus is that, as I said, you can achieve suppression in months if you act quickly. You have to keep working at it because if you don't have a vaccine, the disease can crop up again, but it's manageable.

In the United States, different states are using different policies. Most states appear to be following the doctrine of pushing down the curve and waiting for a vaccine. But there are some states that have opted to do nothing. That doesn't mean you get everybody riding around and getting ill, because people exercise their own judgment, but it means you get these outbreaks and ups and downs as a result. And it's not just states in the United States; Sweden and Brazil also did that. For me, it's an odd thing to be doing.

RATIONING A VACCINE

EF: When a vaccine is ready, presumably there won't be enough right away for everyone who wants it. If that happens, what's the best way to allocate it?

Gans: This is a huge issue that's coming. The CDC already has a list of how to allocate flu vaccines based on how essential you are and how at risk you are.

The essential part of course makes sense. Everybody we decided was essential in March should be considered essential and get the vaccine first. But on the at-risk side, we get into really interesting issues. Normally, it would be pregnant women and young children who would get the vaccine first. It doesn't look like that's necessarily the at-risk population this time around.

But does that mean you want to give it to the most at risk — the elderly — up front? That's not as clear either, because the elderly aren't running around in public and getting exposed.

Who else would you want to give it to? You'd want to give it to people who are in close quarters. Prisoners would be obvious choices on moral and practical grounds.

Then there's the debate about whether to use market forces — willingness and ability to pay — versus something else, like a lottery. My guess is, officially, it'll be a lottery. I'd rather have a lottery but allow people to sell their dose to somebody else who's further down, who got a worse ticket. At least that would be aboveboard and clear. And if you're someone who's poor who can stay at home when the vaccine is in short supply, you can benefit from staying at home instead of getting a vaccine.

Whatever the right policy, the issues should be discussed and understood. Another reason I would like to have the discussion about rationing is that I would like governments to see how bad rationing is going to be — because one of the best ways to get rid of a rationing problem is to have no scarcity.

There are also the international issues: Which country gets the vaccine, what are their intellectual property rights, what are their manufacturing capabilities? Not everyone is going to build all their own plants. What's going to happen?

Normally, what would happen is all the countries of the world would be getting together and deciding on that allocation right now. There are some things going on there, but it seems that the United States, Russia, China, and India aren't participating in that discussion. So that doesn't look like it's going to end well.

EF: When you look at future treatments, do the same issues play out in the same way?

Gans: The issue of treatments is a little bit easier because you don't need enough for everybody. You just need enough to treat the sick. And fortunately, at any given time, there aren't that many people sick. Unless, of course, the virus goes out of control and there are a lot of people sick, with intensive care units filling up — that's going to create scarcity on the treatment side. That was the whole discussion back in March: Let's not let that happen. Let's keep the infection rate low so we can treat everybody.

As it turned out, overrunning of hospitals was avoided by the skin of our teeth. If we had waited another week, it would've happened.

AI AND THE COST OF PREDICTION

EF: Let's turn to your work on artificial intelligence. You've argued that AI will reduce the cost of prediction in much the same way that the web reduced

Joshua Gans

► Present Positions

Professor of Strategic Management and Jeffrey S. Skoll Chair of Technical Innovation and Entrepreneurship, Rotman School of Management, University of Toronto; Chief Economist, Creative Destruction Lab, University of Toronto

► Selected Additional Affiliations

Research Associate, National Bureau of Economic Research; Research Affiliate, Center for Digital Business, Sloan School of Management, Massachusetts Institute of Technology

► Selected Past Positions

Professor of Management, Melbourne Business School, University of Melbourne (2000–2011)

► Education

Ph.D. (1995), Stanford University; B.Econ (1989), University of Queensland

the cost of communication and search. How will it do that, and why is it important?

Gans: Artificial intelligence is a term that gets bandied around to mean all sorts of things. We have a pop culture version; we have technical versions.

At the University of Toronto, we have a startup program I'm involved in called the Creative Destruction Lab. The program doesn't make financial investments; we connect the accepted companies with investors and advisers. We were seeing novel kinds of software applications coming up in 2013 and 2014. People were saying the technology was "artificial intelligence," but it wasn't clear to us for a while what they meant. It turned out that it actually was much more familiar than that. It was ultimately just an advance in statistics.

But it was a big advance, an advance that took advantage of the computational power and large datasets we now have. It was about being able to take a bunch of data and use it for the purposes of prediction.

Some tasks may be obviously based on prediction, like forecasting demand. But a lot of tasks that don't seem like prediction problems can be framed as prediction problems, such as a computer being able to look at a photo and tell you what's in it. You aren't actually requiring the computer to know if a photo has a frog in it. You're asking the computer: What's your best prediction of what a human would call what's in it?

That best guess is based on the computer having seen a million photos that people have labeled as containing a frog and another million photos that they haven't. That's enough for machine-learning algorithms to work out whether a new photo has a frog in it or not.

It turned out a lot of tasks that had been thought of as hard to implement on computers — image recognition, natural language processing, predictions about human behavior — were within the range of machine learning and became really cheap.

One of the companies we met with, called Atomwise, was using artificial intelligence to predict whether a particular protein was more likely to bind with other molecules for the purposes of developing drugs. That is the sort of innovation that could really speed up the drug discovery process. And when that company came through, no one had heard of these artificial intelligence tools. They ended up getting frustrated and went to Silicon Valley, where they raised a whole lot of money, and they are now hugely successful. But we learned from that that maybe we should find out more.

Anyplace where you want to use prediction, it's going to become a lot cheaper, which means you'll use more prediction and you'll find more applications for it.

I think this pandemic has surely disrupted everything in terms of the development of AI for normal business practices. Because we don't know what normal is anymore. The problem with having everything rely on a statistical model is that if you have a major structural break, those models break too. If you were using one to forecast demand, it's bloody useless now.

EF: Regarding the public's awareness of AI, is AI still ahead of where people think it is?

Gans: No, I think we're on the other side of the hype cycle now. There are AI uses coming out all the time. It's getting nice and boring.

But there are exceptions. For instance, we have facial recognition engines that can identify people, most of the population now, which is the scarier end of this kind of technology. We're getting a bit of that.

EF: You've written that although data normally have decreasing returns to scale, with AI they may have increasing returns to scale. Why is that?

Gans: Normally, it's decreasing returns to scale. Get a bit more data, it doesn't help you predict much.

The situation in which data can have increasing returns to scale — economies of scale — is when you can get data on a wider variety of things, including some things that are very rare. For instance, Google, because of its reach, gets a lot of queries that no one's ever asked there before — queries that Microsoft doesn't get. So if Google is using AI, it can train off those more remote results. And so to that extent, there's an increasing return to scale.

EF: What do you think AI will mean for concentration of markets?

Gans: When a development in productivity like AI comes along, invariably people say, "It's going to reinforce existing power." But if it's really a big change, it doesn't tend to do that. Why? Because it's reducing the cost of something. And no one has a monopoly over the hardware, the software, or even really the data to generate AI products at the moment.

So I think it's not going to reinforce existing power. But if it follows the normal patterns, there will be a big company, probably not one of the current ones, that will eventually come out of this as the market leader and we

"A lot of tasks that don't seem like prediction problems can be framed as prediction problems, such as a computer being able to look at a photo and tell you what's in it. **You aren't actually requiring the computer to know if a photo has a frog in it.** You're asking the computer: What's your best prediction of what a human would call what's in it?"

will say, "Ah, that's where the monopolist was." If I could predict which company it will be, I would invest in them, but I can't. What I can predict is that that will happen, because it's always what happens.

AI AND WHITE-COLLAR JOBS

EF: Information technology in general is sometimes said

to be skill-biased, which is a shorthand way of saying it favors educated workers. Is that equally true of AI? What will AI mean for white-collar jobs?

Gans: No one knows yet. You can come up with stories either way.

The way I look at it is that AI is prediction and prediction is a component of decision-making — but it's not the only component of decision-making. In many applications, you still need people with the judgment to evaluate what the trade-offs are of what they're looking to do. Does that come from people who have the highest education? Possibly, but it's not a given.

Part of the AI trend is taking very narrowly specified tasks and automating them. For example, some call centers are suited to that. Then there are other activities that we normally think of as requiring extensive education, such as reading legal documents. Where it may have taken you hundreds of hours to analyze a set of documents without AI, now it will take you, say, two hours. That makes whoever is doing that two hours of work immensely productive, so that's good for them. But the open question will be, are there really enough legal documents to be reviewed to keep everybody occupied who was previously occupied with them?

Historically, we end up with more legal document reviews to do. Or those people have found something else to do. So I'm on the optimistic side that we'll have enough time such that we won't see mass unemployment or anything like that as a result of AI. But I find it hard to predict who is safe.

EF: One reads about efforts in China to establish a leadership position in AI. Do you have any view about who's going to dominate in this field?

Gans: It's always hard to think about issues of national dominance. I find them uncomfortable and not that useful. The only issue that's interesting here is that if China has an advantage, it has an advantage because it can collect data so easily. Here, we haven't been comfortable giving up that level of data to some organization or a government.

I think what will happen is there will be some areas — facial recognition, general surveillance, and things like

that — that China will be better at because they will do more of it.

What the United States is doing and what the defense departments are doing, we don't know. Where that spills over, we don't know. I don't think the Chinese are going to get as good as the United States at targeting ads. (*Laughs.*)

NEGOTIATING WITH CHILDREN

EF: Another area that's been of interest to you is economics in parenthood. In your book *Parentonomics*, you said that parents are in a weak negotiating position vis-a-vis their children when it comes to messy rooms. Why?

Gans: That's because you care about the mess in the room and the children do not. It is much easier to negotiate an outcome where you can find things that people care about equally: You care about X as much as I care about Y. So to negotiate with a child to clean up a messy room, you have to be able to find in that negotiation bundle something that the child cares as much about.

Now, in the time since I wrote the book, I've found the most useful thing that I have that the child cares a lot about is the access to the Wi-Fi. I have a button that I can press to cut my children off from the internet. Suffice it to say, that's all I need. I may encounter resistance; I might encounter a child saying, "Fine! Shut off the internet, I don't need it!" But a few hours later, I'm getting a clean room.

So there's new technology that has changed the balance. The iPad and other such devices are a parent's dream. They are reducing the cost of punishment.

EF: You have experienced higher education in three countries — Australia, Canada, and the United States — as a student, a professor, or both.

Gans: Right.

EF: What do you think are their main relative strengths and weaknesses?

Gans: Well, Canada and Australia are the same in the sense it's mostly public universities. So it's not as expensive. But then again, there's the perennial issue of somebody proposing to cut the budget and everybody panics.

The places I've experienced in the United States are not representative. My experience has been in the elite institutions. And it's a bit of a mystery as to how the whole thing works. Why is it that so many resources are devoted to a relatively small number of people? These institutions tend to be smaller, they get the smaller classrooms, the professors have less teaching, higher salaries, etc., etc.

So you sit there and ask yourself, why is that persisting? I can see what everybody's learning, and it's not that much different between the elite places and the other places. Yet you have people willing to pay many times more. You get the sense that there is a sorting going on and that people were paying to be members of a better club. Whereas in Toronto and Melbourne, the universities are huge. Sixty, 70,000 people. That's not so exclusive a club to be a member of.

I don't know the value of the club membership, but you asked about what the differences are. Those are the differences, I think.

EF: What are you working on now?

Gans: I'm finishing up a textbook — a longstanding textbook on entrepreneurship. I'm just about to pack off that updated version of the pandemic book to MIT Press. Then I'm not quite sure what I'm doing next. Probably whatever it was I was doing before the virus. I can barely remember. **EF**