

R. Preston McAfee

Just about every economist, of course, is excited about economics. And many economists are excited about technology. Few, however, have mashed those two interests together as thoroughly as Preston McAfee. Following a quarter-century career in academia at the California Institute of Technology, the University of Texas, and other universities, McAfee was among the first economists to move from academia to a major technology firm when he joined Yahoo in 2007 as chief economist. Many of the younger economists he recruited to Yahoo are now prominent in the technology sector. He moved to Google in 2012 as director of strategic technologies; in 2014, he joined Microsoft, where he served as chief economist until last year.

McAfee combined his leadership roles in the industry with continued research, including on the economics of pricing, auctions, antitrust, and digital advertising. He is also an inventor or co-inventor on 11 patents in such wide-ranging areas as search engine advertising, automatically organizing collections of digital photographs, and adding user-defined gestures to mobile devices. While McAfee was still a professor in the 1990s, he and two Stanford University economists, Paul Milgrom and Robert Wilson, designed the first Federal Communications Commission auctions of spectrum.

Among his current activities, McAfee advises the FCC on repurposing satellite communications spectrum and advises early-stage companies. The latter include Telescent, a network switching company; Prysm Group, a blockchain governance company; Merlin, an online employment market; CG, a digital security company in stealth mode; OpenX, a digital advertising exchange; and the Luohan Academy, a not-for-profit research institute created by Alibaba. He also serves on the visiting committee of the MIT Institute for Data, Systems, and Society and on the boards of the Pardee RAND Graduate School and the Mathematical Sciences Research Institute.

McAfee served as editor of *Economic Inquiry* for six years and co-editor of the *American Economic Review* for nine years and is a founding co-editor of the economics and computer science journal *ACM [Association for Computing Machinery] Transactions on Economics and Computation*.

He is also a confirmed iconoclast. In the pages of the *Journal of Economic Literature*, he opined that “the most important reason for China’s success” was that “China ignored the advice of Harvard economists.”

David A. Price interviewed McAfee in Washington, D.C., in November 2018.



EF: How did you become interested in economics?

McAfee: When I was a high school student, I read *The Worldly Philosophers* by Robert Heilbroner. It’s a highly readable history of economic thought. I didn’t know anything about economics — I didn’t even know who Adam Smith was — and I found it fascinating. I was pretty familiar with the science of atoms and electrons and planets and stars, but the idea of a science of people was not something I had encountered.

EF: You were one of the first academic economists to move to a major technology company when you joined Yahoo as chief economist. You’ve since spent more than a decade as an economist at major technology companies. What has changed in the way that economic research is used in these firms?

McAfee: The major change is the relevance of microeconomics — the study of individual markets.

Economists have had a big role in companies doing macroeconomics for forever, worrying about inflation, GDP, and how those broad aggregates influenced demand for the firm’s products. Microeconomists bring a very different skill set and answer very different questions.

That’s a major change in roles. Amazon, for instance, has more than 150 microeconomists. A really big thing there, and at Microsoft and at Google, is the problem of causality.

Microeconomists have been studying how to get at causality — what caused something as opposed to what’s just correlated with it — for 40 or 50 years, and we have the best toolset.

Let me give an example: Like most computer firms, Microsoft runs sales on its Surface computers during back-to-school and the December holidays, which are also the periods when demand is highest. As a result, it is challenging to disentangle the effects of the price change from the seasonal change since the two are so closely correlated. My team at Microsoft developed and continues to use a technology to do exactly that and it works well. This technology is called “double ML,” double machine learning, meaning it uses machine learning not once but twice.

This technique was originally created by some academic economists. Of course, as with everything that’s created by academic economists, including me, when you go to apply it, it doesn’t quite work. It almost works, but it doesn’t quite work, so you have to change it to suit the circumstances.

What we do is first we build a model of ourselves, of how we set our prices. So our first model is going to not predict demand; it’s just going to predict what decision-makers were doing in the past. It incorporates everything we know: prices of competing products, news stories, and lots of other data. That’s the first ML. We’re not predicting what demand or sales will look like, we’re just modeling how we behaved in the past. Then we look at deviations between what happened in the market and what the model says we would have done. For instance, if it predicted we would charge \$1,110, but we actually charged \$1,000, that \$110 difference is an experiment. Those instances are like controlled experiments, and we use them in the second process of machine learning to predict the actual demand. In practice, this has worked astoundingly well.

The pace at which other companies like Amazon have been expanding their microeconomics teams suggests that they’re also answering questions that the companies weren’t getting answered in any other way. So what’s snowballing at the moment is the acceptance of the perspective of economists. When I joined Yahoo, that was still fairly fragile.

EF: In both your academic work and in your published work as a corporate economist, you’ve done a lot of research on market design, including auction design. And of course, you collaborated on the design of the FCC wireless spectrum auctions. What are some of the main things you’ve learned about designing markets?

McAfee: First, let’s talk about just what market design is. It’s a set of techniques for improving the functioning of markets. Specifically, it uses game theory, economic theory, experimental research, behavioral economics, and

From an economic perspective, the frequently encountered goal of recreating a market, entrepreneurial or otherwise, inside a firm involves a misunderstanding of the reason for a firm to exist. If a market can work inside a firm, there shouldn’t be a firm in the first place!

psychology, all of those disciplines, to make markets work better.

In politics, you have people who don’t want to use markets, and then you have people who say just let the market do it — as if that didn’t have any choices attached to it. But in fact, often how you make a market work

determines whether it works well or poorly. Setting the rules of the game to make markets more efficient is what market design is all about. Thus, whether to hold an auction, whether to sell or lease, who bears responsibility for problems, and what information is communicated to whom are all questions answered by market design. At least four Nobel Prizes have gone for developments in this area.

One thing we learned is to design for mistakes by participants. People will make mistakes, and to encourage participation and efficient outcomes, it is desirable that those mistakes not be catastrophic.

Moreover, there is a trade-off between the potential efficiency of a market and the generation of mistakes. Give people the ability to express complex demands, for example, and the potential efficiency rises, because people can express exactly what they want. But the number of mistakes will rise as well, and the actual performance can decline. I often find myself supporting a simpler design for this reason; I push back on complexity unless that complexity buys a lot of efficiency.

When we designed the PCS [personal communications services] auctions, the spectrum auctions, we were aware that if you made them complicated, people weren’t likely to function that well. We had empirical evidence of that.

Take a situation where you have seven properties up for auction. One regime is that I bid independently on each of the properties, and if I am the winning bidder on all seven, I get the seven. Another is to allow the bidder to submit a contingent bid — to say I only want all seven. That’s called package bidding or combinatorial bidding. We were aware that in practice those don’t work so well, because it winds up taking a long time to figure out who should win what.

But there is some potential loss from not having a package. Because if, let’s say, I’m selling shoes, most people don’t have much use for a single shoe. So you would not want to sell the shoes individually, even though there are a few people who want only the left shoe or the right shoe. And in fact, I am a person who would like to get different sizes in a left shoe and a right shoe. So there’s this trade-off between simplicity, which makes it easier for most, and expressiveness. There is value in that simplicity not only in terms of getting to an answer more quickly, but also in helping bidders avoid mistakes.

Another example is a second-price auction, where you don’t pay what you bid; if you’re the highest bidder, you pay the second-highest bid, as opposed to paying your

own bid. It has a certain resilience to it. There was a guy who actually submitted a bid that was 1,000 times higher than he intended. Just added three zeroes by accident. But in that auction, if you're paying not your bid but the next highest bid, it takes two to make the mistake in order for that to actually cause him to go broke. He wouldn't have gone broke under the second-price auction, whereas he would under the first-price auction. In that specific instance, we had put in a withdrawal rule that allowed him, at some penalty but not a ruinous penalty, to withdraw.

EF: Much of the economic research that has been publicly discussed by technology companies has focused on outward-facing decisions such as pricing and, as we discussed, market design. Are tech companies also using research to structure the incentives of their employees, and is there more they can be doing?

McAfee: I've hired a lot of people over the years, more than 50 anyway, probably more than 60. And among those have been several people, some quite distinguished economists, who decided that the first thing they wanted to do was get involved in compensation.

Your leverage regarding compensation is greatest in the sales force. If you've got a salaried engineer, let's say, there's not as much you can do.

But in sales, the financial incentives are large and strong. I try to prevent economists on my teams from ever messing with sales force compensation, because there's no quicker way to be fired. The sales force is very persuasive. That's their job; they're supposed to be persuasive.

There was a case where we had an executive vice president come to us and say, "We really want to run some experiments and learn about the sales force." As I said, I did my best to keep my team out of such matters, but when management comes to me and asks for help, I feel I have to oblige. Not only that, I had people chomping at the bit wanting to get involved. We designed some incentives and then what happened next was fully predictable, which is that the EVP got fired. Fortunately, my team was safe because it hadn't come from them.

My teams have worked with HR on other issues. There's always some ongoing work with HR. It can be on

R. Preston McAfee

► Selected Positions

Chief Economist and Vice President, Microsoft, 2014-2018; Director, Google Strategic Technologies, 2012-2014; Chief Economist, Vice President, and Research Fellow, Yahoo Research, 2007-2012; J. Stanley Johnson Professor of Business, Economics, and Management, California Institute of Technology, 2004-2009; chaired professor, University of Texas at Austin, 1990-2003

► Education

Ph.D. (1980), Purdue University; B.A. (1976), University of Florida

► Selected Publications

"Machine Learning in an Auction Environment," *Journal of Machine Learning Research*, 2016 (with Patrick Hummel); "Evaluating the 'Big Deal' Journal Bundles," *Proceedings of the National Academy of Sciences*, 2014 (with Ted Bergstrom, Paul Courant, and Michael Williams); "The Cost of Annoying Ads," *Proceedings of the 22nd International Conference on the World Wide Web*, 2013 (with Dan Goldstein and Sidharth Suri); "Capacity Choice Counters the Coase Conjecture," *Review of Economic Studies*, 2008 (with Thomas Wiseman); "Signaling Character in Electoral Competition," *American Economic Review*, 2007 (with Navin Kartik); "Coarse Matching," *Econometrica*, 2002; "Bidding Rings," *American Economic Review*, 1992 (with John McMillan); "Correlated Information and Mechanism Design," *Econometrica*, 1992 (with Philip Reny)

promotion, recruiting, collaborating — anything but compensation.

EF: Based on the literature and on your own experiences at Google and Microsoft especially, what is the role of entrepreneurship within large tech companies and has it been evolving?

McAfee: Most tech companies have been extremely skewed toward trying to encourage entrepreneurship, as well as giving a lot of lip service to tolerating failure, so as to recreate entrepreneurial activity inside the firm. The "maximize entrepreneurship" approach works pretty well for certain kinds of projects, in particular the kind where a small team can build a functioning product. But there are other products where it is a terrible idea — do you really want to fly in an airplane where each piece was designed and built by separate entrepreneurial teams aiming to maximize their own success?

Indeed, the economic theory of the firm suggests that firms arise when markets don't work well. We know markets work well when complementarities are weak and tend to fail when complementarities are strong. The term "complementarity" is economics jargon for synergy. As a result, the economic theory of the firm suggests that when complementarities are strong, we should see firms arise to internalize these complemen-

tarities and use nonmarket control — dictators, hierarchies, committees, and so on — to direct activities. Thus, from an economic perspective, the frequently encountered goal of recreating a market, entrepreneurial or otherwise, inside a firm involves a misunderstanding of the reason for a firm to exist. If a market can work inside a firm, there shouldn't be a firm in the first place!

Four firms — General Motors, Standard Oil, DuPont, maybe Sears — developed the multidivisional firm. These were firms where pieces of the firm operated as separate firms. And they were doing that just because they had gotten to the stage where they were too large for any one person to operate. It's unsurprising that Silicon Valley's version of the multidivisional firm is to say we're going to run a venture capital firm inside.

I'm generally a voice, not all that successful a voice, against this trend. And the reason is, first, Silicon Valley's

venture capital is an extremely finely tuned machine. It works extraordinarily well. And if you think about business strategy 101, one of the first rules is that if you've got a competitive market doing something, buy it from them, don't do it yourself.

There are a few exceptions. You might want to do it yourself if the market won't produce the quality you need. Also, we've had actually a long-running challenge where American companies like Cisco will subcontract to Chinese manufacturers that eventually go into business against them — so you might not want buy it where you're going to create future competitors.

But otherwise, in general, no. Venture capital does a great job, and it's a competitive market. So the idea of trying to replicate venture capital inside the company is usually misguided.

EF: How do you expect the exploitation of big data and machine learning to affect market structure and competition?

McAfee: AI is going to create lots of opportunities for firms in every industry. By AI, I mean machine learning, usually machine learning that has access to large volumes of data, which enables it to be very clever.

We're going to see changes everywhere: from L'Oréal giving teenagers advice about what makeup works best for them to airplane design to logistics, everywhere you look within the economy.

Take agriculture. With AI, you can start spot-treating farms for insect infestation if you can detect insect infestations, rather than what we do today, which is spread the treatment broadly. With that ability to finely target, you may be able to reduce pesticides to 1 percent of what you're currently using, yet still make them more effective than they are today and have them not deteriorate so rapidly in terms of the bugs evolving around them.

If you look back at the history of big firms, what you see is that when there are these big innovations — electricity and the automobile are good examples — these innovations fundamentally change the way things are done. So what we see and will continue to see is that companies in the face of AI technology have to change their way of doing things. We expect to see a lot of entry into these spaces from firms that have mastered an adjacent technology and can use AI to push themselves into a business. Meanwhile, the existing firms of course are going to fight back, and in some cases they'll push into other areas. This will likely be very disruptive. You'll also get the creation of completely new markets.

Some of those markets are likely to be ones in which a single firm becomes dominant. Digital commerce was an example of this; there was a period when there were lots of companies in digital commerce, but Amazon has clearly stepped out as the leader.

We will also see a lot of mergers and acquisitions. If you

look at the history of merger waves, they tend to follow disruptive technologies. Indeed, all of them followed extensive technological change except the 1980s merger wave, which came about from deregulation. Such merger waves arise as firms struggle to change their business model, due to the changing environment the technological change brought about, and purchase new capabilities via merger. I expect to see a large merger wave from AI, lasting a decade or more, that could change competition in many or even most sectors.

The provision of AI technology is itself quite competitive. Google, Microsoft, Amazon, and IBM offer general AI technologies that, while somewhat differentiated, are competitive with each other, and a plethora of small firms offer more specialized technologies. When electricity disrupted industry, typically there was only one local provider. When business machines disrupted industry, there was one dominant vendor, IBM. But with AI, there are three or four strong vendors. That is positive both for advancing the technology and for maintaining competition. Competition among AI vendors will limit the antitrust problems in other verticals as they adapt to AI. Indeed, the shortage today is in humans: ML experts to implement and operate AI and data scientists to clean the data, prepare pipelines, and structure the output.

EF: What are the implications of machine learning, if any, for regulators?

McAfee: It is likely to get a lot harder to say why a firm made a particular decision when that decision was driven by machine learning. As companies come more and more to be run by what amount to black box mechanisms, the government needs more capability to deconstruct what those black box mechanisms are doing. Are they illegally colluding? Are they engaging in predatory pricing? Are they committing illegal discrimination and redlining?

So the government's going to have to develop the capability to take some of those black box mechanisms and simulate them. This, by the way, is a nontrivial thing. It's not like a flight recorder; it's distributed among potentially thousands of machines, it could be hundreds of interacting algorithms, and there might be hidden places where thumbs can be put on the scale.

I think another interesting issue now is that price-fixing historically has been the making of an agreement. In fact, what's specifically illegal is the agreement. You don't have to actually succeed in rigging the prices, you just have to agree to rig the prices.

The courts have recognized that a wink and a nod is an agreement. That is, we can agree without writing out a contract. So what's the wink and a nod equivalent for machines? I think this is going somewhat into uncharted territory.

EF: Is part of the difficulty that's emerging the result of machine learning in particular? As opposed to a

company making decisions based on an algorithm that's in code or using an econometric model?

McAfee: Yes. If you're using a deep neural net, which is a way of simulating how brains might work, it's really hard to say what the factor was, and actually you're seeing a bunch of interesting examples of this.

Deep neural nets are what have gotten people excited about artificial intelligence now. AI is a field that came and went repeatedly. People were excited in 1980. They get excited and then it never delivers. But this time was different, and what was different was the deep neural net and its capabilities.

Let me give the example of classifying photos. With deep neural nets, both Google and Microsoft can classify photos better than humans. The way we measure this is that we first have humans classify the photos — this the Golden Gate Bridge, that's a dog running in a field. We have humans do it and then we have machines do it. Then we show a human the photo and the two answers, and we ask which one is better. And the machines win. That is, the human picks the machine's interpretation over the human interpretation.

So they use a deep neural net, which is a kind of statistical process that's just wildly complicated because it has multiple layers — 150, 170, 200 of these layers that each have numerical weights attached, so there may be thousands of parameters in each layer and hundreds of layers. It's a wildly complicated system. It doesn't look like a regression where I can say, "Oh yeah, the coefficient on income in a loan is 0.2."

EF: What should antitrust policy be doing more generally, if anything, to respond to the dominance of some online firms in terms of market share?

McAfee: I disagree with those who find the antitrust laws inadequate. With few exceptions, I find our laws adequate for preventing monopolistic mergers, sanctioning anti-competitive behavior, and potentially offering the powerful ability to break up a firm that abuses its dominance.

I do sometimes question the application of the laws. There have been many tech acquisitions where the target might have grown into a serious competitor for the acquirer. Facebook, Instagram, and WhatsApp all offer competing services. Perhaps more of a recognition that tech firms in adjacent markets grow into challengers is warranted, though even the merger guidelines recognize the potential for entry.

We can address monopoly power, even when legally acquired, with regulation. I realize this is incredibly unpopular at the moment, but regulation is a pendulum that swings back and forth. When electricity generation, with its sizeable scale economies, was subject to monopolization, we responded both by regulating private provision and by creating municipal utilities. We should

do the same with Internet provision and for exactly the same reasons.

Of course, a lot of the discussion today is focused on FAANG — Facebook, Apple, Amazon, Netflix, and Google. I see the issues somewhat differently. First, let's be clear about what Facebook and Google monopolize: digital advertising. The accurate phrase is "exercise market power," rather than monopolize, but life is short. Both companies give away their consumer product; the product they sell is advertising. While digital advertising is probably a market for antitrust purposes, it is not in the top 10 social issues we face and possibly not in the top thousand. Indeed, insofar as advertising is bad for consumers, monopolization, by increasing the price of advertising, does a social good.

Amazon is in several businesses. In retail, Walmart's revenue is still twice Amazon's. In cloud services, Amazon invented the market and faces stiff competition from Microsoft and Google and some competition from others. In streaming video, they face competition from Netflix, Hulu, and the verticals like Disney and CBS. Moreover, there is a lot of great content being created; I conclude that Netflix's and Amazon's entry into content creation has been fantastic for the consumer. Who would have thought that tech geeks could actually teach Hollywood, with a century of experience, a thing or two?

That leaves Apple, and the two places where I think we have a serious tech antitrust problem. We have become dependent on our phones, and Apple does a lot of things to lock in its users. The iMessage program and FaceTime are designed to force people into the Apple ecosystem. Also, Apple's app store is wielded strategically to lock in users (apps aren't portable), to prevent competition with Apple services, and to prevent apps that would facilitate a move to Android. My concern is that phones, on which we are incredibly dependent, are dominated by two firms that don't compete very strongly. While Android is clearly much more open than Apple, and has competing handset suppliers, consumers face switching costs that render them effectively monopolized.

So there are issues as to how the antitrust laws should be applied, but by and large, the framework of antitrust is fine. We shouldn't want competition for competition's sake; we want competition because it delivers innovation and good and cheap products. That's how the antitrust laws have been interpreted, and so I'm happy with that.

Going back to Facebook and Google, the reason people are worried is along the lines that our ability to communicate with Grandma is through only this one company. That's what we're worried about. It's not actually an antitrust issue, though. The same with fake news: We want companies to be more responsible, but I don't think the antitrust laws are a solution to that. That's a place where we should, as a society, look at what regulations are appropriate.

A good way to arrive at what those regulations should look like is by doing experiments. The fact that Europe

and California have adopted forms of data protection is a good idea. It's good for us to see some experiments.

The second place I'm worried about significant monopolization is Internet service. In many places, broadband service is effectively monopolized. For instance, I have only one company that can deliver what anyone would reasonably describe as broadband to my house. The FCC says I have two, but one of these companies does not actually come to my street.

I'm worried about that because I think broadband is a utility. You can't be an informed voter, you can't shop online, and you probably can't get through high school without decent Internet service today. So that's become a utility in the same way that electricity was in the 1950s. Our response to electricity was we either did municipal electricity or we did regulation of private provision. Either one of those works. That's what we need to do for broadband.

EF: The notion of regulation or public provision makes sense from your perspective in the broadband market. Does it also make sense in the provision of, let's say, social media?

McAfee: I'd be pretty leery about government provision of social media. Partly because it's a scale play — you need to run a pretty large network. With electricity and with broadband, you can actually run a municipal-level service and you can have local control and you can meet the needs of the local community, but that doesn't really work for a phone system or a social media system. So I would tend to look more toward regulation for that reason, to make sure it serves the national interest.

EF: What was the most surprising part of your transition from being an academic economist to being an economist in a high-tech corporate setting?

McAfee: There's a school of thought that government is inefficient because it can be, while firms, subject to markets, are forced to be efficient. The thing that shocked me the most was how inefficient large firms can be. Sure, there is government waste, but it is commensurate with size and clarity of mission. In one sense, I already knew that large firms could be inefficient — the failure of Kodak and Blockbuster are examples — but it is another thing to live through it.

I have a much deeper appreciation that slow optimization is a better model of human behavior than full optimization, and indeed, I've often used evolutionary models rather than optimization models in my work. People do respond to incentives, and they respond faster to stronger incentives, but along the way there are lots of mistakes and bad choices and hysteresis.

EF: What are the best and worst things about working in a place like Microsoft or Google?

McAfee: The thing I liked best was access to real problems. As a professor, I would dream up problems and solve them. I tried to pick problems whose solutions were likely to be valuable, and I had reasonable success at doing that. But it is another thing entirely when a multibillion-dollar business is measurably improved by a change your research suggested.

Indeed, one way of framing the answer is that, 300 years ago, scientists wrote each other letters of their findings, and these letters came to be reprinted in volumes for others to see. Eventually, these volumes become journals, and universities start to hire people who wrote lots of these letters. At that point, the writing of letters, as opposed to the making of discoveries, becomes a way of advancing in a scientific career, and you start to see "literature-driven" contributions, which are often uninteresting or not useful or both. As a corporate economist, in contrast, I and my team would typically be handed an existing problem, and if we made substantial progress in resolving it, we would write something up for a journal. In that way, I felt much more grounded in reality and actual success rather than academic success.

The worst aspect was firing people. Universities fire a lot of assistant professors, but the process is structured so that committees make decisions and there is no individual responsibility. Firing people is awful, even when it turns out they needed the change and are ultimately better off for it.

EF: Who have been your main influences?

McAfee: I learned to be a modern economist from John McMillan, my long-term co-author and author of *Reinventing the Bazaar*, which I think is the best book on market design. John made ideas operational and was a fabulous expositor. I now spend a full third of my research time on exposition — ideas will never persuade if not articulated well.

Paul Milgrom's perspective on economic theory — his relentless focus on high-value insights, his often uncanny ability to simplify and get at the root cause, and his mastery of statistics underlying economic analysis and its role in economics — continues to be a crucial influence. I would be happy to produce even 1 percent of his theoretical insights.

And I learned a great deal from my boss at Yahoo, who I followed to Google, Prabhakar Raghavan. Prabhakar now leads advertising engineering at Google. Let me describe an outstanding thing he taught me. A manager's job is to make his or her team successful. Full stop. It isn't even to get a job done, though the team's success may require getting some job done. By defining your job as making the team succeed, you focus on what is blocking the team and how to remove those blocks. You acknowledge and advertise the team's contributions within the company. You are no longer the leader but the cheerleader. Upper management loves managers whose teams are successful, and I was well-rewarded for the success of my teams. **EF**