INTERVIEW

Chad Syverson

Editor's Note: *This is an abbreviated version of EF's conversation with Chad Syverson. For additional content, go to our website: www.richmondfed.org/publications*

Productivity growth drives economic growth, and for about the last 15 years, the United States and much of the world has experienced a significant productivity slowdown. The causes remain a puzzle to economists, and the predictions about when — or if — the United States will emerge from this slowdown vary widely.

Chad Syverson, an economist at the University of Chicago's Booth School of Business, has spent much of his career researching issues related to productivity at both the macro and micro levels. His research has shed light on why some firms are significantly more productive than others within the same industry, a long-standing question among economists working in the field of industrial organization. His work has also helped us better understand the process of learning by doing, why some firms have vertical ownership structures (and why those might not be very different from horizontal ownership structures), and the value of carefully done industry case studies. He recently has started researching the economics of artificial intelligence and what future developments in that area may mean for productivity growth.

Syverson joined the University of Chicago faculty in 2001, initially in the Department of Economics. In 2008, he moved to the university's Booth School of Business. He is currently an editor of the *RAND Journal of Economics* and was formerly an editor of the *Journal of Industrial Economics*. In addition to publishing prolifically in top professional journals, he is also the co-author of a microeconomics textbook with his colleagues Austan Goolsbee and Steven Levitt. Syverson earned undergraduate degrees in both economics and mechanical engineering and attributes his interest in productivity and firm dynamics to his engineering background.

Aaron Steelman interviewed Syverson in his office on the University of Chicago campus in June 2018.

PHOTOGRAPHY: JEFF SCIORTINO

EF: Some have argued that the productivity slowdown since the mid-2000s is due to mismeasurement issues — that some productivity growth hasn't been or isn't being captured. What does your work tell us about that?

Syverson: It tells us that the mismeasurement story, while plausible on its face, falls apart when examined. If productivity growth had actually been 1.5 percent greater than it has been measured since the mid-2000s, U.S. gross domestic product (GDP) would be conservatively $4 trillion higher than it is, or about $12,000 more per capita. So if you go with the mismeasurement story, that's the sort of number you're talking about and there are several reasons to believe you can't account for it.

First, the productivity slowdown has happened all over world. When you look at the 30 Organization for Economic Co-operation and Development countries we have data for, there's no relationship between the size of the measured slowdown and how important IT-related goods — which most people think are the primary source of mismeasurement — are to a country's economy.

Second, people have tried to measure the value of IT-related goods. The largest estimate is about $900 billion in the United States. That doesn't get you even a quarter of the way toward that $4 trillion.

Third, the value added of the IT-related sector has grown by about $750 billion, adjusting for inflation, since

the mid-2000s. The mismeasurement hypothesis says that there are $4 trillion missing on top of that. So the question is: Do we think we're only getting $1 out of every $6 of activity there? That's a lot of mismeasurement.

Finally, there's the difference between gross domestic income (GDI) and GDP. GDI has been higher than GDP on average since the slowdown started, which would suggest that there's income, about $1 trillion cumulatively, that is not showing up in expenditures. But the problem is that was also true before the slowdown started. GDI was higher than GDP from 1998 through 2004, a period of relatively high-productivity growth. Moreover, the growth in income is coming from capital income, not wage income. That doesn't comport with the story some people are trying to tell, which is that companies are making stuff, they're paying their workers to produce it, but then they're effectively giving it away for free instead of selling it. But we know that they're actually making profits. We might not pay directly for a lot of IT services every time we use them, but we are paying for them indirectly.

As sensible as the mismeasurement hypothesis might sound on its face, when you add up everything, it just doesn't pass the stricter test you would want it to survive.

EF: What might we learn from past examples of the diffusion process of general-purpose technologies, such as electricity, when considering future productivity trends?

Syverson: I think there are a couple of lessons. One is that it is not unusual at all to have an extended period — and by extended, I mean measured in decades — of slow productivity growth, even after a major technology has been commercialized and a lot of its potential has been recognized. You saw that with the internal combustion engine, electrification, and early computers. There was about a quarter-century of pretty slow productivity growth before you saw the first acceleration in productivity coming from those technologies.

The second part is that you don't necessarily have just one acceleration and then it's over. There were multiple accelerations from electrification separated by a decade. To me, that says that just because we've had one IT-related acceleration, that doesn't necessarily mean it's over. We can have a second wave. Technologies don't just have to come, give what they have to give, and then go away. You can get multiple waves.

Why that would happen is tied to some of the complementarity stories where the first set of gains is driven by direct replacement of the old technology with the new technology. The second wave comes when people recognize there are completely different ways of doing things that the new technology made possible. So it's not that you are simply swapping the old widget for a better one. You are actually doing completely different things now that you have the new technology. This is related to Paul David's widely cited work on how the electric motor didn't just directly replace the steam engine. It eventually led to a complete change in the way factories were designed once people realized you could put a little motor on every single machine. The work didn't have to be stacked on many floors around the single power source any more.

We are going out on a limb a little bit by saying this, but we think artificial intelligence checks the boxes for a general-purpose technology. And it seems that with some fairly modest applications of AI, the productivity slowdown goes away.

EF: Would you consider artificial intelligence (AI) a general-purpose technology? If so, how do you assess the view that the returns on investment in AI have been disappointing?

Syverson: It's way too early. There are two things creating this lag for AI. First, aggregate AI capital right now is essentially zero. This stuff is really just starting to be used in production. A lot of it is simply experimental at this point. Second, a lot of it has to do with complementarity. People have to figure out what sorts of things AI can augment, and we're not anywhere down that road yet.

Erik Brynjolfsson, Daniel Rock, and I are going out on a limb a little bit by saying this, but we think AI checks the boxes for a general-purpose technology. And it seems that with some fairly modest applications of AI, the productivity slowdown goes away. Two applications that we look at in our paper are autonomous vehicles and call centers.

About 3.5 million people in the United States make their living as motor vehicle operators. We think maybe 2 million of those could be replaced by autonomous vehicles. There are 122 million people in private employment now, so just a quick calculation says that's an additional boost of 1.7 percent in labor productivity. But that's not going to happen overnight. If it happens over a decade, that's 0.17 percent per year.

About 2 million people work in call centers. Plausibly, 60 percent of those jobs could be replaced by AI. So when you do the same kind of calculation, that's an additional 1 percent increase in labor productivity; spread out over a decade, it's 0.1 percent per year. So, from those two applications alone, that's about a quarter of a percent annual acceleration for a decade. So you only need maybe six to eight more applications of that size and the slowdown is gone.

EF: Many explanations have been offered about why we observe very large productivity differences among firms in the same industry. As the use of micro-productivity data has grown, do you think economists have been converging on a consensus?

Chad Syverson

➤ Present Position
Eli B. and Harriet B. Williams Professor of Economics, University of Chicago Booth School of Business (at Chicago Booth and, previously, the Department of Economics, University of Chicago, since 2001)

➤ Education
Ph.D. (2001), University of Maryland; B.A., Economics, and B.S., Mechanical Engineering (1996), University of North Dakota

➤ Selected Publications
"Challenges to Mismeasurement Explanations for the U.S. Productivity Slowdown," *Journal of Economic Perspectives*, 2017; "Healthcare Exceptionalism? Performance and Allocation in the U.S. Healthcare Sector," *American Economic Review*, 2016 (with Amitabh Chandra, Amy Finkelstein, and Adam Sacarny); "Vertical Integration and Input Flows," *American Economic Review*, 2014 (with Enghin Atalay and Ali Hortaçsu); "Toward an Understanding of Learning by Doing: Evidence from an Automobile Plant," *Journal of Political Economy*, 2013 (with Steven Levitt and John List); "Market Structure and Productivity: A Concrete Example," *Journal of Political Economy*, 2004; and numerous other papers

Syverson: An important fact is that the skewness of everything is increasing within industries. Size skewness, or concentration, is going up. Productivity skewness is going up. And earnings skewness is going up. To describe why our earnings are stretching out like this, why there is a bigger gap between the right tail and the median, I think you have to understand the phenomenon of increasing skewness in productivity and size. Is that technological? Is it policy? Is it a little bit of both? I don't think we really know the answer.

That said, I think it's less of a mystery now than it was when I started working on this many years ago back in graduate school. At that time, people would tell stories about maybe it's this, maybe it's that, maybe it's everything. There was a lot of speculation and not a lot of evidence. Since that time, I think the profession has been really good at systematically going after an answer.

The biggest change is the amount of work that has been done on management practice. There's still much more work to do, but increasing productivity dispersion seems related at least in part to management practices. Nick Bloom and John Van Reenen deserve a lot of credit for collecting systematic evidence on management practices in their World Management Survey program. The program has gathered information on tens of thousands of firms now. They and their co-authors have also been able to put supplemental management practice questions on the Census Bureau's annual survey of manufacturers.

So we have a lot more systematic data on that now, and there's no doubt productivity is correlated with certain kinds of management practices. People have also developed more causal evidence. There have actually been some randomized controlled trials where people intervened in management practices and saw productivity effects.

Is that all of the story? No, I don't think so. If I had to guess, it's probably 15 to 25 percent of the story. There's a lot more going on. I think part of it has to do with firm structure. I have done work on that.

I think we have gotten better at measuring quality differences in labor and a little bit better at measuring quality differences in capital, though I think capital mismeasurement is still the biggest issue with measuring productivity on the input side. A lot of work has also been done on the way we measure productivity on the demand side. We have learned about the importance of each side and what drives the fundamentals on both sides. That's going to help us get a more comprehensive answer of the causes of productivity dispersion within industries.

EF: Regarding management practices, it seems a little puzzling that lagging firms wouldn't have done more to replicate what more successful firms have done. You could imagine possible stories about why that may be the case, but it seems like an important question to answer.

Syverson: I agree, and there is some evidence we can look at from work done by Bloom and some colleagues. I'll call it the India experiments. They did a randomized controlled trial with textile producers in India. They provided management consulting practices to 28 plants — a small sample but still useful — and asked the management of every plant why they hadn't previously instituted some of the management practices that the consultants recommended. Basically, there were three classes of explanations. First, there was, I didn't know about them. The second was, I knew about them, but they're just not going to work here. The third was, they might work here, but I didn't have the time to put them into place. And then they tracked the plants over time and asked those who still had not adopted those practices why they hadn't. Obviously, plants are unlikely to still give the first answer, but you still had a lot giving answer two or three.

Now, maybe there's something special or unusual about the setting of that experiment. But I do think the fact that management is often just mistaken is a nontrivial factor. There is evidence coming out of this body of work that suggests companies don't know where they are in the distribution — they don't know whether they are well-managed or not. You can't fix yourself until you know you have a problem.

Also, I think even if you know you have a problem, a lot of firms can't simply say, well, we see this competing company over there has an inventory management tracking system that seems really useful, so we'll install it on our computers and our problems will be solved. That's not how it works. The firm that has adopted this practice has people trained in how to do it. It has changed its system, so that there's an interaction and a feedback loop between

what the system is recording and recommending and what you do. If you just say, OK, we're going to start collecting these data now and then do nothing else, you're not going to get the productivity benefits that the company with the complements is getting. I just think this stuff is way more complex than people might initially think.

An example I talk about in class a lot is when many mainline carriers in the United States tried to copy Southwest and created little carriers offering low-cost service. For instance, United had Ted and Delta had Song. They failed because they copied a few superficial elements of Southwest's operations, but there was a lot of underlying stuff that Southwest did differently that they didn't replicate. I think that presents a more general lesson: You need a lot of pieces working together to get the benefits, and a lot of companies can't manage to do that. It also typically requires you to continue doing what you have been doing while you are changing your capital and people to do things differently. That's hard.

EF: It is often argued that the health care sector is fundamentally different than other sectors of the economy — and that these differences might produce relatively less variation in productivity within the health care sector. What does your work suggest about the idea of health care "exceptionalism"?

Syverson: In general, we think companies that do a better job of meeting the needs of their consumers at a low price are going to gain market share, and those that don't, shrink and eventually go out of business. The null hypothesis seems to be that health care is so hopelessly messed up that there is virtually no responsiveness of demand to quality, however you would like to measure it. The claim is that people don't observe quality very well — and even if they do, they might not trade off quality and price like we think people do with consumer products, because there is often a third-party payer, so people don't care about price. Also, there is a lot of government intervention in the health care market, and governments can have priorities that aren't necessarily about moving market activity in an efficient direction.

Amitabh Chandra, Amy Finkelstein, Adam Sacarny, and I looked at whether demand responds to performance differences using Medicare data. We looked at a number of different ailments, including heart attacks, congestive heart failure, pneumonia, and hip and knee replacements. In every case, you see two patterns. One is that hospitals that are better at treating those ailments treat more patients with those ailments. Now, the causation can go either way with that. However, we also see that being good at treating an ailment today makes the hospital big tomorrow.

Second, responsiveness to quality is larger in instances where patients have more scope for choice. When you're admitted through the emergency department, there's still a positive correlation between performance and demand, but it's even stronger when you're not admitted through the emergency department — in other words, when you had a greater ability to choose. Half of the people on Medicare in our data do not go to the hospital nearest to where they live when they are having a heart attack. They go to one farther away, and systematically the one they go to is better at treating heart attacks than the one nearer to their house.

What we don't know is the mechanism that drives that response. We don't know whether the patients choose a hospital because they have previously heard something from their doctor, or the ambulance drivers are making the choice, or the patient's family tells the ambulance drivers where to go. Probably all of those things are important.

It's heartening that the market seems to be responsive to performance differences. But, in addition, these performance differences are correlated with productivity — not just outcomes but outcomes per unit input. The reallocation of demand across hospitals is making them more efficient overall. It turns out that's kind of by chance. Patients don't go to hospitals that get the same survival rate with fewer inputs. They're not going for productivity per se; they're going for performance. But performance is correlated with productivity.

All of this is not to say that the health care market is fine and we have nothing to worry about. It just says that the mechanisms here aren't fundamentally different than they are in other markets that we think "work better."

EF: What does your work tell us about why some firms benefit from common ownership of production chains, how those benefits can be measured, and how large those benefits might be?

Syverson: In a paper with Enghin Atalay and Ali Hortaçsu, we found that most vertical ownership structures are not about transferring the physical good along the production chain. Let's say you are a company that owns a tire factory and a car factory. When you look at instances analogous to that, most of the tires that these companies are making are not going to the parent company's own car factory. They are going to other car factories. In fact, when you look at the median pair, there's no transfer of goods at all. So the obvious question becomes: Why do we observe all this vertical ownership when it's not facilitating the movement of physical goods along a production chain? What we speculated, and then offered some evidence for, was that most of what's moving in these ownership links are not tangible products but intangible inputs, such as customer lists, production techniques, or management skills.

If that story is right, it suggests a reinterpretation of what vertical integration is usually about in a couple of ways. One, physical goods flow upstream to downstream, but it doesn't mean intangibles have to flow in the same direction. Management practices, for instance, could just as easily go from the downstream unit to the upstream unit.

The second thing is that vertical expansions may not

be as unique as we have thought. They may not be particularly different from horizontal expansions. Horizontal expansions tend to involve firms starting operations in a related market, either geographically or in terms of the goods produced. We're saying that also applies to vertical expansion. A firm's input supplier is a related business, and the distributor of its product is a related business. So why couldn't firms take their capital and say, well, we think we could provide the input or distribute the product just as well too? So, conceptually, it's the same thing as horizontal expansion. It's just going in a particular direction we call vertical because it's along a production chain. But it's not about the actual object that's moving down the chain.

We were able to look at this issue, by the way, because we had Commodity Flow Survey microdata, which were just amazing. It's a random sample of shipments from a random sample of establishments in the goods-producing and goods-conveying sectors of the U.S. economy. So, if you make a physical object and send it somewhere, you're in the scope of the survey. We get to see, shipment by shipment, what it is, how much it's worth, how much it weighs, and where it's going. And then we can combine that with the ownership information in the census to know which are internal and which are external.

EF: You have done a lot of work examining the concrete industry. Why concrete? And what can we learn about more general phenomena by looking at some pretty narrow industries?

Syverson: And not just concrete, but ready-mix concrete in particular. The reason is that it is a great laboratory for testing economic theory. It has a set of characteristics that not many industries have. One, it's geographically ubiquitous. Two, because of the transport costs and the perishability of the product, every one of these geographic markets is basically independent, and you can only ship this stuff so far. So every city is basically a different market. Three, almost all concrete is bought by the construction sector, but it's a small share of construction costs. What that means is that construction activity is basically an exogenous mover of concrete demand. Furthermore, there are a lot of firms in the concrete business, so even a modest-sized market is going to have multiple plants run by multiple companies. This means that it is like an economist having a laboratory full of petri dishes where you tweak each one and see what happens differently in response to different stimuli. On top of all that, the stuff is relatively easy to measure because it's physically homogeneous. It's not a differentiated product, so the prices are pretty comparable and the units are comparable. Just about everything you would want in an ideal, clean case study exists in this industry.

So that's why I have done so much work on concrete. What can we learn more generally? You hear jokes about people working in industrial organization (IO) looking at case studies and discussing the ketchup literature, or the yogurt literature, or in this case the ready-mix concrete literature. I have tried to be clear about what I think the broader lessons are from these case studies and what we can learn from them. One of the first studies I did on ready-mix concrete looked at whether variations in consumer scope for substitution show up in the equilibrium productivity distribution. In other words, is it indeed harder to be an inefficient producer in a market where customers can more easily find the more efficient producers? The answer is yes. I think that is a more general phenomenon; it's just one I can measure much better in that setting than in others. That said, I wrote a companion paper that does look across manufacturing industries and found similar things with different measures of substitutability to bolster the generalizability of the findings in the earlier paper.

Also, Ali and I looked at vertical integration between the cement and concrete industries. There is clearly an element of industry specificity to that work. But, on the other hand, those were sort of the poster-child industries for the market foreclosure literature. So if you thought that vertical mergers provide incentives for collusion and anticompetitive foreclosures, this is where you would see it. We looked, and we didn't find it. That might make you think differently about how likely you would find it in other industries too.

I understand the case-study method, why it's important and what advantages it has. I don't think people in IO should cede ground to those who question the value of individual case studies just because we haven't done case studies on the hundreds of other industries out there. We should use what we know from a case study, along with theory, to extend our understanding of economics as far as we can.

EF: You were given access to detailed production data from an auto assembly plant over the course of a year. What were those data able to tell you about the sources of learning by doing?

Syverson: Regarding the data, as a car is being made, there are things constantly being recorded in the factory's information system, either in an automated fashion or by workers manually inputting information. So Steve Levitt, John List, and I were able to see every step of the way whether the step went right or wrong. And then we looked at subsequent defect rates for every car that was made – about 190,000 over the course of a year.

Most of the empirical learning-by-doing literature has looked at unit costs, such as how many worker hours it took to make a unit, and then examined that over time and traced out the learning curve that way — how fast people adapted, for instance. Our more detailed data let us learn something about where the knowledge resided inside the organization and how it moved around.

There are a few facts that are important to understanding that in this setting. One is that a lot of learning happened early, as is pretty common. So, for example, defect rates fell 70 percent in the first two months of production. Now, as it happens, the factory only ran on one shift for the first two months of data we observed, and then starting in the eighth week, the second shift started. The second shift's training was to watch the first shift for one week. That was it. They weren't on the line itself. Once the second shift comes online, they are right at this new, lower defect level that the first shift achieved. So you immediately know that it's not just being on the line for a while that leads to improvements.

Two, there is a high correlation between defect rates for a particular operation across shifts. Operations don't go wrong with equal frequency. There is a right tail of processes that go wrong a lot of the time, and then there's a left tail where things never go wrong. That's true across shifts. So if some operation is problematic on the first shift, it's problematic on the second shift, even though the workers are different.

Three, we were able to see absenteeism every day at the factory and in which part of the production process the absent workers were placed. There is a positive relationship between absenteeism rates and defect rates along a set of operations on the line, but it's very weak.

So those three things suggest it's not the workers who are carrying the knowledge, which, again, is substantial. Defect rates over the course of the year came down 90 percent total.

What happened is the factory had a set of practices to take knowledge from the workers and as quickly as possible put it into the capital of the factory — either the physical capital, such as changing a faulty part on the line, or the organizational capital, such as workers conveying information to each other.

EF: Following the accounting scandals of the early 2000s, there were proposals to require companies to rotate auditing firms. You have looked at the possible effects of such a mandate. What did you find? Similarly, what is the potential impact if one of the Big Four firms were to fail, perhaps because of regulation or legal action?

Syverson: As you said, Joseph Gerakos and I looked at two things: mandated auditor rotation and what would happen if one of the Big Four were to fail. The two issues are related. A good way to start thinking about them is to ask whether companies choose auditors based on certain characteristics or do they just go with the lowest price. The answer is clear that the auditors are differentiated to the companies that hire them; companies are looking for the best match.

When you move around prices exogenously, you see the customer's willingness to substitute based on those changes in prices, and they're not nearly as willing to substitute one auditor for another as they would be if the auditors were not differentiated. So it's clear something is driving the value of the match-specific relationship. What does that mean? It means that if one of the Big Four were to fail, there would be losses suffered by the audited companies because you can't just swap one for the other and not lose that match-specific value. It also means if you mandate that they switch auditors after a certain number of years, you won't have that match-specific value anymore.

All that said, there is another side to the mandated switching policy. If you think too much coziness between firm and auditor can create the potential for corruption, there's value in eliminating that. We are not trying to measure that or saying that it's zero. We are simply saying that on the other side of the scale is a real cost.

EF: What do you think are some of the big open questions in IO and understanding firm dynamics?

Syverson: With IO, I would like people to pay greater attention to more general lessons we might be able to take from case studies. That could involve adding some comment in the paper and maybe writing a companion paper. I would also like people to avoid thinking that any empirical work that involves more than one industry is ipso facto flawed. I think there is a little too much stridency along that line — not across the board, but I would like to see people be more accepting of some broader approaches.

One really positive move I've seen in IO over the past 10 years is I think the field has moved toward answering more important questions. That's not to say the questions were unimportant before, but I think we're moving in a good direction. As I tell people at IO conferences, other fields are doing IO now. Look at macro and finance and development, just to name a few. They're trying to answer IO questions. And in part I worry that they're doing it because we haven't done enough. I think people working in IO can bring useful insights to the conversations people in other fields are having.

In terms of firm dynamics, I think we still have further to go to explain productivity dispersion, in particular what's creating this increase in skewness. I also think the micro aspects of the productivity slowdown are still a mystery. We have some understanding of these issues, but there's a lot we don't know.

EF: Do you think being an engineer might have affected your choice of research interests as an economist?

Syverson: There is no doubt. I got into productivity in grad school because of my engineering background. I was a mechanical engineer. I like looking at how systems work together to produce something and how those systems can be improved. Also, as an engineer, it's simply fun to go to factories and see how things are done. **EF**