

Goodfriend Memorial Lecture: Productivity and Structural Change

By [Richard Rogerson](#)


Economic Brief

September 2025, No. 25-34

Introduced in 2023, the Goodfriend Memorial Lecture series honors the legacy of Marvin Goodfriend, long-time Richmond Fed economist, research director and senior policy advisor. The lecture was delivered as part of the Richmond Fed's [Collaboration of Research Economists \(CORE\) Week](#) model, which brings together Richmond Fed economists and visiting economists from a range of disciplines for seminars, conferences, networking and collaboration.

On Aug. 5, 2025, Richard Rogerson delivered the Marvin Goodfriend Memorial Lecture titled "Productivity and Structural Change." Rogerson is the Charles and Marie Robertson Professor of Public and International Affairs at Princeton University.

The following is a lightly edited transcript of the lecture.



I'm thrilled to be here, so thanks to the organizers. As was pointed out this morning, I'm going to talk about macro development type stuff. Marvin Goodfriend, of course, was famous most for his contributions to monetary theory but, as was pointed out this morning, wrote two excellent papers on macro development. So, it's definitely connected to some of his work.

So, in terms of background, two salient features about the development process. First, that development is basically about productivity growth. And when I say productivity here, I mean I'm going to talk about labor productivity. So, I'm not doing the "Is it TFP or human capital?" whatever, just labor productivity. So, I think that's not controversial.

And the second is development is associated with a process of structural change. The poorest countries are doing almost all agriculture. The richest countries, almost nobody in agriculture, a whole bunch of activity in services.

So, those are two facts. And I want to talk about this movement out of agriculture. And so development necessitates that these countries move workers out of agriculture. And I'll say, it's not that you can't be rich doing agriculture. But the fact is, if you look at rich countries, people working in the agricultural sector, a lot of them do very well for themselves. There's just no way to do that at scale, and there's not enough land for everybody to be rich in agriculture. So, if you want to be rich for a typical country, you're going to have to move people out of agriculture. And there's kind of two questions I want to address, and I'm not necessarily going to give definitive answers, but these are two questions. And I was talking to some people over break; there's some strong priors about this.

So, one is: What forces shape this movement out of agriculture? What's responsible for that? And the second is: At the early stages of development, you've got to move people out of agriculture, (so) does it matter where they go to in particular? Does it matter if they go to manufacturing or services?

And I see two longstanding ideas. The first one, which I say, kind of goes back to the Lewis view of development that things outside of agriculture are kind of critical to getting activity out of agriculture. And the second longstanding idea is somehow it's critical that when you move people out of agriculture, you move them to manufacturing.

I have something here. I bet nobody in this room even knows what it is: Kaldor's First Law of Growth. That's what he asserted, as the first law of growth, is that growth is directly related to how much you move activity into manufacturing. So, I'm going to talk about both of those.

Three objectives for what I'm going to talk about today. One of the things I want to do is basically use this as, I want to let you know about two recent datasets, which I think are truly phenomenal for people who want to think about macro development from this kind of structural change perspective. It's kind of like having the equivalent of the [Penn World Tables](#), but now at a sectoral level. And if you think of macro development, where would the research on this be without the Penn World Tables that allow us to make these comparisons across countries? That this is what that data, and then I'm going to use that data to talk about these two questions: 1) Is there something special about manufacturing? and 2) What drives the reallocation out of agriculture?

So, that's where I'm headed. So, let me talk about the two new and important datasets. So, both of these are maintained by the GGDC, and one of them is called the [Economic Transformation Database](#). The other is called the [Productivity Level Database](#). I'll tell you a little bit about those.

So, this Economic Transformation Database is a panel dataset of countries, 51 countries, data from 1990 to 2018. You can see the geographic distribution. One of the key points here is that there's a whole bunch of relatively poor countries. And what this dataset gives you is data at a 12-sector breakdown on unemployment, nominal value added, and real value added in 2005 domestic prices for these 12 sectors.

And most importantly, I mean, data collection is tough. It's especially tough in developing countries. But they've made great efforts to try and harmonize this across countries and in particular especially on the employment side. One of the big challenges is, in developing countries, there's lots of informal own employment and kind of tracking that turns out to be critical, and they make special effort to do all of that. So, it's not to say that any of this is perfect, but it's kind of trying to do all the things right to make these things comparable.

Here's just to show you kind of the distribution of countries in terms of where they are in terms of levels. This is in 1990 relative to the U.S. And you can see almost all the countries in this sample are less than 0.5 of the U.S. But there's a whole bunch that are, you know, less than 0.2 of the U.S., more than half the countries. So, if you're interested in kind of that early stage of development, there's just a lot of representation here. And then in addition to that, there's actually also a lot of heterogeneity in terms of development experiences in terms of how fast these countries grow over these 30 years. So, I show you something about the distribution of growth. Right. So that's all just kind of good to have that there in terms of thinking about issues. So that dataset has these 51 countries.

I said I was going to talk about these two issues. That said here, so I'm going to draw on kind of two to three different projects that I have with different co-authors. So, some of the work is with Akos Valentinyi, Berthold Herrendorf. Some of it also includes Martina Kirchberger. And then one of the projects involves Carolina Piazza, who's a graduate student at Princeton. So I have a lot of we's here. And the we's can be different groups of people on different slides.

So, for some purposes, having all these poor countries is great, but you also want to have some of the rich countries in there. And so you kind of want to add those in. That's not that big of a deal because the data that's in the ETD, getting that for rich countries is kind of readily available from various places.

But we're going to add that in to give us a sample of 71 countries, which then includes 90 percent of the world's population, 90 percent of world GDP. So the ETD, economic transformation, is great for lots of things, but it has one big drawback. And the one big drawback is it gives you all this country-level stuff, but it doesn't actually give you anything that allows you to compare across countries at the sector level.

And the thing that's missing. The whole big contribution of the Penn World Tables was to give you the price indices that you can use to compare outputs across countries. So, we need that. And so, the ETD does not have that. So, that's a limitation. And that's where this other data set comes in: the Productivity Level Database of 2023. There was an earlier version of this — I forget which year — which had a much smaller representation of countries. So, the PLD 2023 gives you sectoral PPPs, for the same 12 sectors as what I said before, for three benchmark years: 2005, 2011, 2017. And it actually has that for 68 of the 71 countries that I had referred to. So, the three it's missing are three African countries you can see there. So, I'm going to call that the expanded Economic Transformation Database for these 68 countries. So that's the data.

Now, as is well known, even at the aggregate level, there's benchmark years. And then you need to fill things in and outside of the benchmark years. And that the intuitive idea is outside of a benchmark year. You look at country level growth rates to kind of fill in beyond that. The issue is if you've got 2005 and 2011, you could fill in 2007 using 2005 going forward, 2011 going backwards. So, we do that in a weighted way so that you get the same answer in 2000. You don't just pick one. For before 2005, you just extrapolate back from 2005, etc. So, using that you get comparable sectoral productivity the whole time from 1990 to 2018. Okay. That's all the countries. We don't have to go through the list.

So, I want to talk about those two issues. So, the first one is: Is manufacturing special? So, I say there's a long history for that idea. If you read policy statements, you kind of see this all the time. That kind of it's somehow critical for the poorest countries in the world to expand manufacturing, that that's the route to becoming a rich country.

I think one of the origins of this idea is going back to the Industrial Revolution. The Industrial Revolution was, you look at England at the time, it was largely agricultural. Where did all the technological innovation come from? It was all about manufacturing stuff, textiles, energy, etc. That's kind of what happened. I think there's this idea that somehow copying what happened in the Industrial Revolution — since that's what took England from a poor country into an advanced economy — that somehow that's the path. But as I say, I don't think that's a very persuasive or compelling idea, because at the time, that was the

frontier stuff that was happening. Nowadays, a developing economy, there's a whole bunch of frontier stuff. It's not clear you should follow the path from hundreds of years ago. So, I think there's kind of a lack of compelling arguments for this idea.

But one idea that has been out there in the literature is that there is something special about manufacturing, and there's, in fact, something special about productivity and manufacturing. And it was a paper by Rodrik that I'll talk a lot about from 2014 QJE titled "Unconditional Convergence in Manufacturing," where he argued that manufacturing exhibited unconditional convergence. And this there was this kind of intuitive idea. If you've got a lot of people in manufacturing, and manufacturing displays unconditional convergence, having a lot of people in manufacturing is how you catch up to the rich countries, because you've got unconditional convergence going on there. So, that's an idea that's out there that would be coherent. And so I want to talk about that idea.

The first thing to note — and there's many papers that have made this point — I think the argument that it's good to get out of agriculture, that's a pretty easy argument to provide compelling evidence on. And so here's a picture. So, this is looking, I forget which, I think this is the 1990 cross-section from the data. This is looking at agricultural productivity gaps versus aggregate productivity gaps and then the same thing for the nonagricultural sector. In this picture, if you're below the 45-degree line, that means the gap in agriculture is much bigger than the gap in at the aggregate level. And if you look at nonagriculture, you see a lot of the points — especially at the low end here — are above the 45-degree line. So, the gaps in nonagriculture are less than the aggregate gap. So, this has been made by many people in the literature. So, this is the paper by Restuccia and co-authors in the JME.

And if you know, Caselli's chapter in the Handbook of Economic Growth, he kind of stressed that point. And so this just tells you in a mechanical sense, if you move a bunch of people in a poor country from agriculture into nonagriculture and just assume that when you're doing that all these workers have the average productivity, you will catch up in terms of Penn World Table type calculations.

So, poor countries are both very bad at agriculture and they have a lot of people in agriculture. And so that's a kind of a double whammy in terms of aggregate differences. So, there's kind of no doubt from a statistical perspective that moving people out of agriculture seems like a good thing. So, the question is: Does it matter where you move them in nonagriculture? So that's what we want to talk about.

Audience Question: Richard, I guess this is always — I'm just thinking that highly mechanized agriculture doesn't need a lot of people. And at the same time, it's very productive. And so it seems like that's, like the poor country it's bad that agriculture has a lot of people in agriculture, but it's sort of like, not really bad. It's just like a more mechanized agriculture gives you less people.

This is just, so I think you're asking about why they have so many people in agriculture. You're telling me if they mechanized, they wouldn't need so many people in agriculture. This picture is neither against, just think of this as a pure statistical statement. I mean, I'm going to move people. I'm just statistically moving some people from agriculture to nonagriculture assuming the productivities are unaffected by that. So I mean, if I move people out of agriculture, there will be less food in that counterfactual.

So, this is not about the driving forces behind right now. This is a purely statistical decomposition. But we'll come back to that later on.

So, here's a quote from Bhagwati, which I kind of like. This is his response to Kaldor's first law. You can see him making fun of Kaldor. He basically was oblivious to what was going on in the world by virtue of living in his Oxbridge bubble. But again, it's kind of this point that in today's world, what would somebody look at to say that, oh, all the technological advances happen in manufacturing. So, unless you're in manufacturing, you couldn't possibly be at the frontier. So, just kind of not there.

So, let me talk a little bit about this Rodrik paper, "Unconditional Convergence in Manufacturing." So, he wrote this before any of these datasets I told you about. And I said, you kind of need these datasets to answer this question.

So, what did Rodrik do? Given limitations at the time, he did something quite creative. He went to the UNIDO dataset — some would call INDSTAT — which has data on manufacturing, value added manufacturing, employment, all the value added in U.S. dollars. He basically assumed a law of one price holds to make these things comparable across countries. And so then used that and found substantial evidence for substantial amounts of unconditional convergence. I note the paper has over a thousand citations. And I'll say this is one of the worst reasons for writing a paper is that you read something and it kind of annoyed you and so you felt you had to write a paper, as opposed to being issue driven. But you see so many cites, if you kind of think about macro development, this paper is cited all the time as kind of proving, supporting that it's so important to put people into manufacturing. It kind of grates on you after a while.

So, we had to write this paper just to kind of get to the bottom of this. And you can see the conclusion, which is exactly this kind of Kaldor thing: "Aggregate convergence fails due to the small share of manufacturing employment in low income countries and the slow pace of industrialization." This idea, if only they had more people in manufacturing or had moved more people into manufacturing, everything would have been better at an aggregate level. And so that's kind of the punch line from his paper and what gets quoted a lot.

So, let me tell you a little bit about this UNIDO data. So, the UNIDO dataset is not harmonized. It's not, you know, a single entity trying to put stuff together across countries. It's basically data that draws on surveys at individual country levels. These surveys are not

harmonized across countries at a point in time and not really harmonized actually across time within a country. And also there's no attempt to cover the entire manufacturing sector. They actually typically survey either large or registered manufacturing establishments. And kind of one of the first things you see in terms of what does this mean in terms of what the coverage is — and I'll show you a table here — this is looking at the, as a subsample where you've got countries that are both in UNIDO, which has its own limitations in terms of which countries they're in which year, and the ETD that I told you about earlier. So, 30 countries that are there for the 1990 to 2018 period.

And I'm taking the ratio of employment as measured by UNIDO, which they acknowledged is not the universe but its subsample, versus, what's in the EETD for manufacturing, which at least is intended to count the universe of manufacturing employment. And these are the ratios. And what you can see is, out of the 30, 14 of them have, well 13 of them, almost half of them, have less than 50 percent coverage. So, you kind of know from the outset there's something highly specialized about the UNIDO data. So, we'll come back to that.

Now, the title of Rodrik's paper, "Unconditional Convergence in Manufacturing," which sounds rather bold and strong. There is some small print in the paper and, in particular, he does at one point say because of these issues of data coverage, the results should be viewed as applying to the organized formal parts of manufacturing. So, I want to think about that statement. Many issues there.

So, the first thing is if this is across countries, some comparison of formal organized manufacturing sectors across countries, you would say there must be a systematic definition for what it means to be formal and organized. But in fact that would require harmonization for even thinking about that. And so that's just not there. So, that's kind of an issue. As I said, a typical criterion is establishment size. But the establishment size threshold is different across countries and varies even across time within a country. And we know that firm dynamics are very different across the size distribution, which means it's quite likely that these things are highly influenced by selection issues.

But there's actually even a deeper question here. And I think this is research that was done after Rodrik wrote his paper and it had come out. And that is that there's maybe this old idea that there's a formal sector. And in the formal sector, you have formal firms that hire formal workers. And then there's the informal sector, which is informal firms and/or self-employed. And then the informal firms hire informal workers.

It turns out that's actually not the way the world works. I reference Ulysea there. He studied Brazil in particular, but in his *Handbook* chapter, he talks about that actually in these poor countries with lots of informal workers, firms that are formal firms actually hire all kinds of informal labor. And this creates a huge problem for productivity, because on these surveys, why is it that a formal firm is hiring informal workers? It's because there's some regulations that they find it's advantageous for them to avoid by virtue of hiring informal workers. So, they're interviewed, surveyed about how many people you have

working here. You have to expect it's highly unlikely they're going to report the informal workers they have working there, which means that the measure of employment is inaccurate at these firms. And it's inaccurate in a way that if they were reporting less employment than they have, they're going to report total sales, this kind of total sales, but then they're going to give you a biased measure of their employment, ignoring the informal workers. This is going to elevate productivity. And so if you're interested in productivity per se this could be hugely problematic. So, you should be very concerned about this.

The last point I make is I don't think there's any doubt that in any poor economy, there are some kind of modern productive firms and then a whole bunch of not so modern, not very productive firms. But this is not unique to manufacturing. This is true across the economy. This is true in agriculture. There's a whole bunch of subsistence farming. And then there's a whole bunch of corporate farms. This is true in manufacturing. It's also true in services, you've got supermarkets in large cities, and then you've got all kinds of people with carts and selling vegetables in the market and at stands, etc.

So, even if it's the case that there's a part of manufacturing where the productivity dynamics are distinctive, it doesn't mean there's anything special about manufacturing, because it's probably also true if we looked at modern agriculture. Modern agriculture looks pretty good, and modern services or formal services looks pretty good. So, none of it would speak to that.

So, what I want to do next in terms of this issue is I'm going to just look to see, using the EETD data, which is the comprehensive data for the manufacturing sector. Now with PPPs, what does it say about unconditional convergence in manufacturing? And then I want to talk a little bit about: Should we even take at face value that Rodrik's results mean what he wants them to say?

And then I will also show you that there is, in fact, from a productivity perspective, something special about manufacturing, but it's actually something special in a negative way. And then I'll talk about a couple other things.

So, here's just a picture which already tells you where things are headed. This is a picture of either the 90/10 percentile gap or the cross-sectional standard deviation of log productivity for this EETD dataset from '90 to 2018. And the dark line is aggregate GDP, and the dashed line is what's going on in manufacturing.

For those of you who haven't seen some recent papers, so obviously a whole bunch of papers talking about the lack of convergence at the aggregate level, but actually the post-1990 period has actually seen a, I'll say, a very modest amount of unconditional convergence at the aggregate level. And I think people make a big deal because used to be nothing and now there's something. But what they're finding is actually incredibly small. In

the paper, we have a number. But for a country that starts at 10 percent of the frontier over this 28-year period, the unconditional convergence would take them from 10 percent to something like 16 percent.

So, there's something there, but it's actually very, very small. So, people are excited about it because there's some of that going on. But it's a very small effect. But the key point here is you can see manufacturing is actually going in the other direction here. So, already kind of makes you a little bit doubtful.

We're going to run the standard regression that people run to look for unconditional convergence. And this is, L_{jt} here means log of labor productivity. So, change in labor productivity on the level plus some time dummies. You can either, you know, like or not like those, but this is kind of standard in the literature, and these are the numbers you get.

And I'm going to do this for the aggregate. I'm going to do it for agriculture. I'm going to do it for market services. I could do it for this dataset. You can do it for all 12 sectors. As we all know, even there's a bunch of service sectors where we think that output is not well measured to begin with. So, you could say it's just not that interesting to compare those. So, that's why I'm just going to focus on what I'm calling market services. So, I've taken out, oh question.

Question from audience: What is conditional convergence conditional on?

Unconditional convergence is basically just saying is everybody converging in some sense to the same thing. Conditional convergence is every country has their own level. And so every country is. Yeah. So, I'm not going to — I have the conditional convergence there, but I'm not going to talk about any of that today.

Question from audience: Can I also add, is there any argument we should expect for manufacturing if there will be unconditional convergence? Because manufacturing is easy to copy or scalable? So, what makes manufacturing really special?

Well, I think if you read Rodrik's paper, having found unconditional convergence, he said, oh, this makes perfect sense that we found this, for the reasons you said: that manufacturing is like easily copyable, you know, but that's him giving the rationale for the result after he found it.

So hard to, I mean, there's actually a book called *The Power of Productivity*. Lewis is the author. I think he used to work for McKinsey. It's a great book to read if you're interested in productivity. He has a chapter, and he makes this point, which I thought was so salient. He said many people think that somehow if there's a frontier technology out there, it's actually so easy to reproduce. And he said it's just not, if I told you, go build a state-of-the-art auto assembly plant, it's just not that easy. Even if you can find lots of information to actually make it happen in a in a randomly chosen country, it's just not that easy a thing to

do. So Rodrik relied on that and said, this is very intuitive that we see unconditional convergence in manufacturing because we all know it's so easy to copy these things and transport them around the world.

Question from the audience: Is it the same type of manufacturing across these countries? Because I can sort of see it. If you started digging deeper, you have different kinds of manufacturing in different countries, and then it's not so clear.

So, it's the aggregate of manufacturing. And so this data — the data that I'm using here — does not allow you to go deeper. So, Rodrik, the UNIDO data does allow you to go deeper. And Rodrik actually did stuff at the two-digit level in addition to the aggregate. So that's a nice thing to do if you can do it. So and certainly it's of interest to do that. But this data doesn't allow you to do that.

Okay. So, that's the setting. And then what are the key results? In this regression, if you have unconditional convergence, it means that the higher your initial level, the less you grow, the further you are away, the more you grow. That's a negative coefficient. And you can see you get a negative coefficient for the aggregate. But again, it's very modest. You get negative, modest coefficient for agriculture. Same thing for market services. Manufacturing is the one that's actually not statistically significant. So, in this dataset, manufacturing is special. It's the one sector in which, over this time period, there's no evidence for unconditional convergence.

Now, why do we get a different result than Rodrik? So, I mentioned before, Rodrik doesn't actually have PPPs. What he does in the background is say, let's assume law of one price, and then if we have everything in U.S. dollars, we can actually compare those. We don't have to worry about different price indices. You can actually check that with this dataset to see. I guess if you take the PPPs as the truth, you can see how well that law of one price holds. Or more generally, if you're nervous, don't trust the PPPs, you can just see how different these things are. It turns out they're actually not that different. I'm not going to show you the pictures. You can statistically reject that those two things are equivalent. But in fact, the difference between them is not very significant economically. So, you get a small bias from that, but not much.

Second issue is there's different time periods, different samples of countries. And we all know there's whatever's going on with convergence. Some countries might be converging, and it may happen in some periods and not others. So, it could be there's no inconsistency. He studied '65 to 2005 with a rotating sample of countries based on who was available at which time period. What I'm going to do to address that is look at the period '95 to 2005 where there's overlap and there's a set of countries that are in both datasets. So, I can get 48 countries that are both in the UNIDO data that Rodrik used and in our EETD. And for those 48 countries, I can actually run this for the period '95 to 2005. And what you see here

is when you do it with the EETD, you actually get a positive coefficient for that subsample, statistically significant at the 5 percent level. So, actually a little bit of divergence, which you kind of saw in one of the pictures I showed you.

When you use the UNIDO data, you actually get a fairly substantial amount of unconditional convergence, which is there's really just something about the data that's very different here. So, it's not just about different samples and different groups of countries, even when you have the same sample time period, same countries, the two datasets, they give very different pictures of what's going on.

And one of the things I want to comment on now is why you should — I think, you know, one of my main messages is I think people should not use UNIDO data to think about productivity. And one of the reasons has to do with this coverage ratio I told you about. This is actually showing you changes in the coverage ratio across these countries. So, the y-axis is the coverage ratio in 1990. This is in 2018. If you think that this is all just informality, then we kind of know informality tends to decrease with development. And so you'd kind of expect there to be some number of points above the 45-degree line. There's also a bunch of countries below the 45-degree line, some of them a lot below the 45-degree line.

So, there's a lot of variation in coverage rates over time. Here's just a couple countries to show you kind of weird things happening. You can see Mexico has a huge jump up at some point, kind of strongly suggestive that something just happened with the way they're surveying in terms of who's being included. You see other countries with massive drops over time. And then here's a picture that says in the data, it turns out in the UNIDO data, there's actually a relationship between the change in the coverage ratio and the recorded productivity growth. And when the coverage ratio is going down, productivity growth is actually high. So, that should make you very nervous that there's something going on with the surveys. And as a result of that, productivity has been distorted. And that's creating quite an issue. And what we actually do is for that UNIDO time period, if you actually remove the countries that have had the largest drops in their coverage ratio, you actually lose the unconditional convergence result. So again, just on the surface that kind of looks like it's quite possible that all of that is kind of spurious stuff driven by kind of measurement issues inside that dataset.

Question from audience: So, Richard, this part obviously the, you know, the data has all these problems. But when you show that the convergence in the productivity sector is nonsignificant. But I was thinking that it's sort of like one of the things you could show is like the variance across countries in productivity. Because I'm thinking suppose productivity is already uniform in manufacturing, then I just have to move people there to catch up.

I think I'm going to show you this picture. So, it could be that you want to move people into manufacturing because everybody in manufacturing is actually very close to the frontier. And actually, here's a picture where I've now split. I showed you the picture before for

nonagriculture, now we split nonagriculture. I'm showing you manufacturing and market services separately, but they're basically the same. So, it's not that there's something wrong with moving people into manufacturing. I'm not saying don't move people into manufacturing for that reason. But there's nothing in the levels, and in fact, in terms of the, to the extent there's anything to the unconditional convergence, you're actually better off moving them into services because market services actually had a little bit of unconditional convergence in the subsequent 28 years. But you're right that it could have been that there's no one conditional convergence, but you should still go there because this is actually such a great place to be. But that's just also not true.

I want to switch topics to the last thing, but I will say that's just showing you that manufacturing is the one sector without unconditional convergence, or one of the few without.

I will say two other things about — so one of the nice things about the EETD is it has a whole bunch of poor countries, and it has a whole bunch of heterogeneity in development experiences post 1990. One thing you could do is ask: Does it jump out from the data in any sense that countries that either moved a lot of people into manufacturing or had a lot of people in manufacturing did much better? It turns out there's lots of heterogeneity in terms of how many people they had in manufacturing, how many people they moved into manufacturing. And it turns out you get kind of nothing from those. So, this is just pure scatterplot. But it's just if you're thinking that this is like so obvious and it's just like staring you in the face and so you have to accept it, it's not staring you in the face.

But one thing which is there in the data is manufacturing productivity growth is very highly correlated with aggregate productivity growth. All the sectors have that property. And here's the cross correlation. The one thing which is kind of striking here is how the nonagriculture sectors kind of bunch relative to agriculture. So, you see very little correlation between improvements in agriculture and improvements in nonagriculture. But inside nonagriculture, you see very high correlation. Manufacturing has some of the highest correlations. Now whether that means being productive at manufacturing has lots of spillovers or just getting good at things then applies to lots of activities. That's an unknown question. And I don't know the answer to that at this point, but productivity in manufacturing is certainly correlated with good things.

So, the last five minutes I'm going to talk about the second question, this issue of movement of labor out of agriculture and what drives that. This relates to some of the stuff Kristina talked about. And I'm going to kind of have two kind of crossover points with what she did, one of which I'm going to maybe be in contradiction to what she said, another where you can view what she's doing as deeper than the superficial stuff I'm going to show here.

In terms of modeling structural change, kind of the name of the game, because people started with the one sector growth model and they thought, this is like a pretty good framework for thinking about macro development in terms of some of the broad facts. The so-called Kalder facts. But then you want to add structural change on to that. You kind of need to build that into a multisector model. And the name of the game has been to think about a world where productivity, sectoral productivity is evolving exogenously. And now we need to find a structure preferences so that we can get something that kind of looks like balanced growth at the aggregate level at the same time that we get the patterns of structural change. And so the search is for some kind of stable preferences that will provide that fit to the data. And there's a whole bunch of papers in that literature: Kongsamut et al., Ngai and Pissarides, Boppart, Comin, etc., all following this approach. Kind of different versions of preferences to think about that. So, one of the things when people kind of validate what they're doing, sometimes they just pick one country. So, it's sometimes people just study the U.S. and say, "Here's my model. I fit the U.S., so, that's a good model." Sometimes they fit averages. They say, "I can describe the average evolution so that's success." Sometimes they look at a whole bunch of countries, but when they look at a whole bunch of countries, they're actually in the background adding fixed effects at the country level. So, they're not really saying my one structure fits all the countries. They're saying kind of subject to some country fixed, and the reason they're implicitly doing country fixed effects is they didn't have the productivity comparisons to make things comparable across countries.

And so, here's a picture I find kind of interesting. A little bit off topic in some ways. This is looking at manufacturing value-added shares for five countries. But I want to focus on these four. You've got China, Brazil, Mexico and Thailand, all at kind of similar ranges of GDP per worker. And these are the trajectories, the dashed lines are kind of after you add in net trade in value added. And what I want to note is you kind of see all those slopes kind of look the same, but the levels are dramatically different. And so lots of people who said, "My model fits many countries," they're kind of saying, "I get the slopes right, but they're kind of silent on those levels." Well, it turns out this dataset where we now have the productivity levels that are comparable countries allows you to actually try and go deeper than that. I'm going to do that today, but not with manufacturing. I'm going to do it with agriculture.

So, I have a paper with Gollin and Parente a long time ago, which was a very simple model to try and paint the picture that thinking about the movement out of agriculture could be useful for thinking about aggregate differences. And the essence of the model captures this intuition I mentioned before. If you're really bad at agriculture or poor countries are both really bad at agriculture and have a lot of workers there, why do they have a lot of workers in something they're so bad at? For a moment, think about a closed economy. If you require a certain amount of food, then the worse you are at producing food, the more people you need there to produce the food you had.

So, we wrote down a model for that is the whole story. We said people must have a certain amount of food. They don't want any food above and beyond that. They just need that amount. And so, in a model where you actually only have labor, then there's a clear 1-to-1 relationship between productivity and agriculture, share of workers in agriculture. So, you get -1 relationship in logs. And so we kind of did that a little bit to make a stark point.

There's a paper in *Economics Letters* by Ungor who actually said, let me actually just go to the data with this for a bunch of countries. And data is the solid line, and the dotted is the model. That's for those countries. Here's these countries. You can see here Latin America, Peru obviously doesn't fit very well, but the others fit pretty well. Here's your rich countries and those all fit well. So you kind of said this is actually not such a crazy idea. But one of the features of this model is that agriculture employment is completely determined by agricultural productivity. And that kind of fits not so badly. But Ungor is also subject to this critique that in the background it's a little bit like he's got a country-specific subsistence term because his exercise basically said, "I'm going to assume each country in the first year of data, I'm going to choose parameters such that I get a perfect fit." This is like choosing country-specific subsistence levels, but if you don't have comparable productivity levels, you can't kind of do the right thing. So, let's actually do that. So, I'm going to do that.

So, the first thing I'll show you is these are the three cross-sections. And I'm already over time, so I'll speed up. In the cross-section, if you just regress agricultural employment here on agricultural employment measured in agricultural labor productivity in PPP, you get an R^2 of 0.81 and something incredibly stable in terms of that relationship.

If you add nonagricultural product productivity, all you do is you get the R^2 up from 0.81 to 0.84. So, as I say in the data, it looks like agricultural productivity gets you a long ways, and adding nonagricultural productivity does not get you far at all. So good news. Agricultural productivity, highly predictive; nonagricultural productivity adds very little. Relationship is stable across time. But in the simple model you should get a -1 coefficient. I got -6. That's bad news. What's the good news? I'm going to show you a slight generalization that actually works pretty well. The slight generalization is just go to this Stone-Geary form. This is the equation. And I know I'm out of time.

So, I want to show you a picture here. The dots in this picture are the data, log of agricultural productivity. The yellow line is if you ran a regression that first goes through the origin. And it actually does not so bad. The black line is the regression fit I was showing you before. And the red crosses are that Stone-Geary form. And the point is the red line maps close does a good fit to the black regression line. So, this is a simple model that actually fits those cross-section results. And then this is if you take that regression from one of the cross-sections and then just run it for the entire time series.

This is showing you the time-series fits of the data to that model across three regions. So, if the model was doing a really good job, everything would be on the 45-degree line. Each color is a country. What you can see is Latin America really works well. This is Asia, and

there's kind of two countries out of the 21 that don't fit very well, but the rest are pretty good. And Africa, you've got like three or four that don't fit that well. But basically this goes pretty far to fitting. And so I'm going to stop there because I'm over time and keeping you from dinner and stuff. So anyways, thank you.

Richard Rogerson is a professor of public and international affairs at Princeton University.

To cite this Economic Brief, please use the following format: Rogerson, Richard. (September 2025) "Goodfriend Memorial Lecture: Productivity and Structural Change." Federal Reserve Bank of Richmond Economic Brief, No. 25-34.

This article may be photocopied or reprinted in its entirety. Please credit the author, source, and the Federal Reserve Bank of Richmond and include the italicized statement below.

Views expressed in this article are those of the author and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System.



Topics

[Economic Growth](#)

[Marvin Goodfriend's Legacy](#)

[Production and Investment](#)

Subscribe to Economic Brief

Receive a notification when Economic Brief is posted online.

By submitting this form you agree to the [Bank's Terms & Conditions and Privacy Notice](#).

Email Address

Subscribe

@ Contact Us

RC Balaban

© 1997-2025 Federal Reserve Bank of Richmond