

Comments on papers by Eldridge and Sherwood and Gu and Ho

Jeremy Rudin*
Finance Canada

Both of the papers that I am commenting on are written by government employees, and in both cases the authors are very careful to say that their work represents their own opinions and not necessarily that of their employers. I want to go one step beyond that. My comments are based on discussions that I have had principally with my Finance colleagues, but other people as well, so it is quite possible that what I am about to say does not even reflect my own opinion, let alone that of my employer.

There is a lot to praise in both of these papers. I want to pick out three things in particular. First, both of these papers push very hard to get comparable data and comparable methods, and just as importantly, point out where they were not able to do that. Another thing I liked very much about the papers is that they both provide information on output per hour or, if you will, labour productivity. Even the Gu and Ho paper, which focuses on TFP estimates, includes labour productivity estimates as well. I am a big fan of labour productivity because it takes a minimal set of assumptions to estimate. So even when a paper is about TFP, I find it is very useful to have labour productivity estimates included for comparability. The other thing that I really liked about both of these papers is that both of them arrived on my desk in time for me to read them over two or three times before giving my comments. That is a courtesy you do not always get when you are a discussant and I think it should always be rewarded.

I am conscious of the fact that there are a number of people in this room who are experts in the measurement of productivity and that I am not one of them. So I am going to focus my remarks on the interpretation of the results and I will bring in some measurement issues to the extent that they are important for what I am going to talk about.

To avoid monopolizing the time that I have been allocated, I am going to focus on a single aspect of the interpretation of the results. Both of these papers can be read to say the following: roughly all of the difference in the growth of manufacturing productivity between Canada and the U.S. is due to the performance of two high-tech industries. They have slightly different names in the two papers; we will call them industrial machinery and electronic equipment.

I am not attributing this point of view to anyone in particular, I am simply saying that you could easily come to that conclusion from reading these two papers. You could get this “100% result” for Gu and Ho’s estimates of TFP growth over their full sample period.

* What follows is the transcript of my remarks at the conference, edited to make me appear to be more eloquent than I am. After giving my remarks, I learned that some of my comments were inaccurate or incomplete. I have therefore made some corrections and clarifications in the footnotes. I am grateful to Jeff Bernstein, Paul Darby, Erwin Diewert, Steven James, Alice Nakamura and Tim Sargent for discussions on these topics before and after the conference. Responsibility for the errors that remain is mine alone.

And you could get the same “100% result,” roughly speaking, from the Eldridge and Sherwood paper if you look at labour productivity in just the final period that they consider.

I want to argue that this is a fragile and misleading interpretation of these results. I have three reasons I want to put in front of you.

The first one is just arithmetic. When we decompose the gap in productivity growth into the contribution of individual industries, we have some industries that make a positive contribution to the gap, and some that make a negative contribution. Those that make a positive contribution are the industries where productivity is growing more rapidly in the U.S. than in Canada. Those that make a negative contribution are the industries where productivity growth is faster in Canada than in the U.S.

The result of the fact that some industries make a negative contribution is that, just as a matter of arithmetic, there are a lot of different ways to add this up. For example, in both of the papers we can throw out the two high-tech industries, and pick out all the other industries where productivity growth is faster in the U.S. than in Canada. What we will find is that a list of “low-tech” industries accounts for half of the gap over the periods in question. If that was all we knew about the papers we would think that the papers told a kind of a balanced story: half a low-tech story, and by implication, half a high-tech story.

Of course, it is also true that the two high-tech industries account for just about the entire gap. So we can explain 50% of the gap with a low-tech story and 100% of the gap with a high-tech story, and thus explain 150% of the gap.

If we had a finer industrial breakdown then we could probably find three or four different lists of industries, each of which explained 100% of the gap. This is all to say that any decomposition of this type is necessarily arbitrary.

My second argument is based on the fact that all productivity growth estimates begin with estimates of real output. And we know that measuring real output in the two high-tech industries is a big problem because of the rapidity of change in the industry. Eldridge and Sherwood point out in their paper that the impact of any problems that arise in measuring the real output of these two industries on the measure of productivity in manufacturing as a whole will be reduced to the extent that the output of the high-tech industries is intermediate input into the output of other industries.

That is an important point. But the industrial breakdown within the countries will be wrong if there are problems in the price deflators, and that will also affect the inter-country difference. The reason is that, in Canada, proportionately less of the high-tech intermediate inputs come from other domestic manufacturers.

This is a good example of a more general issue related to the comparability of the estimates for two countries in both papers. In deriving productivity estimates we often

have to make simplifying assumptions. When doing an inter-country comparison, it is probably a good practice, as the authors have tried to do in both of these papers, to use the same simplifying assumption in both countries. But a simplifying assumption introduces errors, by definition. And those errors may be systematically different in the two countries. So there is no guarantee that we will get truly comparable estimates for two countries just by using the same assumptions for both countries.

For example, if we use exactly the same deflators for the two high-tech industries in both countries but that deflator is wrong, we will move value added from one country to another and, therefore, we will not produce truly comparable estimates.

What is my third argument? Even if we correctly estimate productivity growth in the two high-tech industries, to be able to say how much those industries contributed to manufacturing as a whole we have to correctly estimate productivity growth in the sum of all of the other manufacturing industries. One major potential problem, and here I am picking up on a remark that Erwin Diewert made recently, is the measurement of service inputs into manufacturing. One reason that we like to focus on manufacturing productivity is that we think that the measurement problems there are not as bad as they are in services. But remember that the productivity estimates in these two papers are based on value added obtained by double deflating. That is, we have to subtract off the real value of service inputs into manufacturing in order to get manufacturing value added. So we cannot completely escape the problem of measuring service sector output just by looking at manufacturing.

If we are underestimating real service inputs in manufacturing then we are probably not making a big mistake in estimating value added for the entire economy, but we are moving value added erroneously from one sector to another. There is some circumstantial evidence that this could be important. In their paper, Gu and Ho report what I found to be a very surprising result: if we measure labour productivity in manufacturing using gross output, not value added, then over the time period that they use the growth rate of labour productivity is almost exactly the same in Canada and the U.S. It is only when we use the value-added basis for measuring labour productivity that we get the gap.

Now I agree that value added is a more meaningful basis for a labour productivity comparison than gross output. But the question remains, why the remarkable coincidence on a gross output basis? The implication is that intermediate inputs are growing more rapidly relative to gross output in Canada than in the U.S. Gu and Ho suggest the interpretation that outsourcing is becoming prevalent more rapidly in Canada than in the U.S. That is a plausible interpretation in that it matches the data. It does not, however, match the anecdotal evidence about outsourcing, which is that it is more important in the U.S. Nor does it fit with our Canadian predilection to think of ourselves as laggards in all sorts of areas.

One simple potential explanation is that underreporting of service inputs into manufacturing is more of a problem in the U.S. data than in the Canadian data. And one

circumstantial piece of evidence that suggests that this could be true is reported in the Eldridge and Sherwood paper, where they note that the last benchmark IO table available for Canada is 1995, while for the U.S. it is 1987.¹

Now it is only a short step from the interpretation of these papers as saying that all of the gap in productivity growth between Canadian and U.S. manufacturing is due to two high-tech industries, an interpretation that I am trying to dissuade you from using, to the contention that the best if not only way to close the gap is to get better at high-tech product innovation in Canada. This could be true. But I want to argue that, even if it is true, the data presented in these papers cannot be used to support that argument.

I have a shorter argument, and a longer one.

My shorter argument is this. In both papers, one-quarter of the gap can be attributed to the food processing industries. This is not high tech; it may not be glamorous. But if we have to come up with policy ideas for improving productivity in Canada while sitting in our offices, it may be the easiest industry to pick. I will bet that this will be the last time that the words “food processing industry” will be said at this conference.²

The longer argument is this. We do not want to try to make a story about product innovation story from labour productivity data for the reasons that we are familiar with. If we find a gap in labour productivity growth between two countries it could be because one country had faster capital growth, or a relative improvement in its capital or a relative improvement in its labour. There is no direct evidence about innovation in labour productivity data. Indeed, one of the motives for doing TFP estimates is to try to look inside labour productivity estimates to see what we can attribute them to.

But we immediately run a severe problem. To make TFP estimates we need a lot of auxiliary assumptions. And the auxiliary assumptions that are embedded in the Jorgensonian approach that Gu and Ho use, as sensible and as useful as they are, are not assumptions which fit well with a story about industry directed innovation.

These assumptions include that all of the factors are being employed to produce current output, and that they are all getting paid their marginal product. We also assume constant returns to scale, so that the factor payments exhaust revenue.³ We then use the factor shares as weights.

¹ After the session, Paul Darby pointed out that that the measured intermediate input intensity of manufacturing has been growing more rapidly in Canada than in the U.S. for some time.

² I lost the bet. Bart van Ark made a pointed reference to the food processing industry later in the conference.

³ This is too strong; we do not need all of these assumptions to *construct* TFP estimates. For example, among the many productivity indices that can be constructed using an index number, or axiomatic, approach is the Fisher productivity index. In practice, the values of the Fisher index will be very close to those of the translog productivity index used by Gu and Ho. Yet the Fisher index is derived without assuming constant

This is a really good approach, but it does not, as I said, fit well with a story about purposeful product, or even process, innovation. Product innovation is an uncomfortable fit because we are assuming, by estimating a different production function for each industry, that if an industry starts producing a different product it needs a different production function. We can try to get around that, to be sure, in the way we deflate the output, but it is clear that the results are going to hinge very much on how that deflation is done.

Leaving that aside, neither product nor process innovation fit comfortably in this framework because by assumption all our factors -- all our capital, all our labour -- are being employed to produce current output and they are all being paid their marginal product. The natural interpretation of TFP in this framework is that it is the exogenous shifting out of the production function, exogenous in the sense that it is not coming about because of the purposeful directed activity of the industry, all of which is going into producing current output. You could say that it is research and innovation, but only as an unintended consequence, something that is coming in from outside the industry, such as the accumulation of free knowledge.⁴

We know that the real world does not conform to this framework. We know that there is a lot of purposeful innovative activity going on. We know that constant returns do not hold in every industry, the electronic equipment industry might well be a good example, and so on. But what we do not know, well let's be fair, what *I* do not know, is how these departures from the framework will interact with other measurement issues and end up in the TFP numbers which are inexorably cranked out by the Jorgensonian approach. I do not think that there is any reason to presume that this mismatch between the framework and the real world necessarily ends up with a result that we should treat as an indicator of innovation.

On the contrary, I can give you two circumstantial reasons to think that these TFP numbers are not good indices of industry directed innovation. First, there are a lot of negative TFP growth numbers. We can find a few in the Gu and Ho paper, which looks at 2-digit industries. There is a very valuable companion paper by Wulong Gu and a couple of co-authors, Frank Lee and Jianmin Tang, which looks at 3-digit industries and shorter time periods, and it is full of negative estimates for TFP growth. That is not really

returns to scale. Thus it would have been more accurate to say that the *interpretation* of the results depends on the assumptions that are used to motivate the measure. On the multitude of possible TFP measures and their interrelationship, see Erwin Diewert and Denis Lawrence, "Measuring New Zealand's Productivity," New Zealand Treasury Working Paper 99/5, March 1999, Appendix A, available at <http://www.treasury.govt.nz/workingpapers/99-5.htm>. Alice Nakamura made this point during her remarks during the discussion period.

⁴ During the discussion period, Alice Nakamura pointed out that this interpretation of the Jorgensonian approach is that of Jorgenson. See, for example, Dale W. Jorgenson, *Productivity*, Vol. 1 (MIT Press, Cambridge, Massachusetts, 1995), p. xvii.

compatible with the view that TFP is an index of something that ought to go forward and not backwards, like innovation, technical progress, or the accumulation of knowledge.

Second, Gu and Ho find very strong TFP and labour productivity growth in the Canadian motor vehicle industry. I do not know where that comes from, but it does suggest, anecdotally at least, that high TFP growth is not necessarily associated with domestically-produced product innovations.

I am anxious to hear the comments from the floor so I am going to wrap up here. Just before I do that, I want to thank all the authors for very stimulating, very valuable papers. Thank you.