

More on Measured Productivity and the Labor Share

15 Aug 2016 [2 Comments](#)

Last week, [Matt Yglesias](#) dug up an old [post](#) of mine about labor's share of output and the measurement of productivity. The short version was that the decline in labor's share may be part of the explanation for why measured productivity growth has slowed down in the same period of time. The original post was a bit math-heavy, and several people on Twitter were looking for a more intuitive explanation of what was going on. I was at the beach (sorry, *down the shore*) last week, so this response is coming in a little late, but here's my attempt at explaining how this works. No math.

Let them measure cake

Productivity is just output divided by input(s). "Labor productivity" is output divided by the number of workers (or hours worked). "Total factor productivity" (TFP) or "multifactor productivity" is output divided by an index that measures a combination of inputs, like labor and capital. But it is always output divided by some measure of inputs.

The issue we have is in measuring inputs. For TFP, we divide by an index of inputs. How do you form that index? If it doesn't conform exactly to how inputs are used in production, then our measure of TFP is going to be influenced by the amount of the inputs we use.

The analogy I like here is that GDP is like cake, and inputs are like ingredients. It's easy to measure "total ingredient

productivity” (TIP) by dividing the volume of cake by an index of the various ingredients included. Let’s say that index of ingredients is very crude, just the total weight of ingredients used. Add the weight of flour to the weight of the eggs and so on.

If I add another cup of flour - which should be around 120 grams - then my index of inputs goes up by 120 grams. My cake volume will probably increase very slightly, and it will end up being kind of crispy and dry. My TIP falls due to a change in the use of one of the inputs. There was no technological change (e.g. no new step in the recipe), and yet TIP changed.

This happens because my index of ingredient inputs does not treat flour the same way that the actual recipe (the production function) treats flour. My index treats flour as a perfect substitute for other ingredients, and the cake production function uses flour as something close to a perfect complement to other ingredients. Because of this mismatch, my measure of TIP is mechanically influenced by the amount of ingredients I use.

It is this kind of mismatch that makes our measure of TFP mechanically influenced by the amount of labor and capital we use. The production function for GDP uses inputs in a way slightly different from how we incorporate them the index of inputs that we divide by to get TFP. Labor’s share of revenue matters here because we often use it to guide us in how to incorporate labor (and implicitly capital) into our index of inputs.

If labor’s share of revenues is low (affecting our index of inputs), but labor’s role in production is high (affecting GDP), then measured TFP is affected as well. The low labor revenue share means our input index is *understates* labor’s role in production, and hence some of the growth in labor inputs gets picked up as measured TFP growth.

Because of how we typically construct the input index, if we understate labor’s role, then we *overstate* capital’s role. This means we falsely attribute some productivity growth to

capital accumulation, and hence measured TFP growth is lower because of our overstatement of capital's role.

As the labor share falls, this exacerbates the issue. Because capital tends to accumulate faster than labor (except maybe [not any more?](#)), the drag on measured TFP growth from overstating capital's role is bigger than the boost to measured TFP growth from understating labor's role. So as labor's share of revenue falls, so does measured TFP growth.

In the original post I went through all this stuff about markups and how those are related to labor's share. That isn't necessary to make the direct point that if we misstate labor's (and capital's) importance in our input index, then inputs matter for measured TFP.

More technical, but still no math

If you're happy with the above explanation, feel free to stop there. This section just gives you a little better idea of what I mean by "misstating" labor and capital's role in production, and how that influences measured TFP.

Start with labor productivity, GDP divided by labor. I think most readers would be fine with the idea that both technological improvements *and* increases in capital raise labor productivity. The first is using a better "recipe", and the second is increasing the other ingredients besides labor.

Intuitively, it seems obvious that capital matters for labor productivity, because we have this notion that they are complements to some extent. But the mechanical reason it matters is that GDP - the numerator in labor productivity - has some positive elasticity with respect to capital. At the same time, the denominator in labor productivity - labor - has a zero elasticity with respect to capital. Hence our measure of labor productivity responds positively to changes in the amount of capital, as the numerator rises while the denominator doesn't change.

Technology improvements also increase GDP holding

constant inputs, so labor productivity also rises due to technological improvements. This would be like coming up with a better recipe for cake (whisk the egg whites separately?) that used the same inputs as the old recipe.

When we measure labor productivity, we are thus measuring some combination of capital accumulation and technological improvement. If capital accumulation slowed down (which it looks like it did recently), so would measured growth in labor productivity, and vice versa. Labor productivity isn't a pure measure of technological improvement. And this is because the elasticity of output with respect to capital in the numerator is different than the elasticity of the input, labor, with respect to capital.

Okay, now let's take on total factor productivity (TFP). Recall that all this involves is dividing GDP by an index of inputs, rather than just labor. When we form that index of inputs, we have to specify the elasticity of the index with respect to each input.

If we pick exactly the right elasticity for capital in our input index, then in response to an increase in capital, our index in the denominator will rise by exactly the same percent as GDP rises in the numerator. The effects will cancel, and there will be no effect of the increase in capital on our measured TFP. This same logic works for labor as well.

Given the exact right elasticities in our index of inputs, TFP will *only* respond to changes in technology. It wouldn't tell us exactly what the new technology involved (whisk the egg whites separately?), but it would tell us that this new technology raised GDP holding all inputs constant.

Before we even talk about whether we get these elasticities right, what happens if we get them *wrong*? Then the elasticity of GDP with respect to an input (in the numerator) doesn't match the elasticity of our index with respect to the input (in the denominator), and our measure of TFP is not neutral with respect to changes in that input.

If we understate an elasticity in our index, then any increase in the associated input will increase measured

TFP. This is the exact analogue of the situation I described with capital growth and labor productivity. There, we understated the elasticity with respect to capital in our denominator (i.e. we set it to zero), and so more capital meant higher measured labor productivity.

Similarly, if we overstate an elasticity in our index, than an increase in the associated input will *decrease* measured TFP. If our index says that the elasticity for labor is 0.8, but the elasticity of GDP with respect to labor is actually 0.6, then measured TFP will *fall* (with an elasticity of -0.2) as we add workers.

Note that getting the elasticities in our index wrong means that measured TFP deviates from true technological growth. But it doesn't mean that technological growth itself is affected.

So why might we be getting the elasticities in our index wrong? For several reasons. The elasticities embedded in the numerator, GDP, are the actual elasticities of output with respect to each input. The elasticities are like a number that captures the nature of the chemical reactions between flour, baking powder, sugar, milk, eggs, and the other inputs. In the cake example we might be able to figure these out by doing a bunch of chemistry experiments. In the economy, we'd have to know the exact engineering specifications for every business in the country.

Since we don't know the exact elasticities, we try to infer them from what we can observe. This is like trying to infer the chemical reactions involved in baking from observing how different cakes come out of the oven. Based on some minimal economic theorizing - namely that firms try to minimize their costs and face competitive factor markets - we think that the elasticity for an input should be equal to its share of total costs. If total wages are 60% of total costs, then this would imply the elasticity was 0.6, for example.

There are a whole lot of reasons this might be getting things wrong. If factor markets are not competitive, then the

share of costs paid to an input would not necessarily equal its elasticity. If firms are not cost minimizing, then we'd have the same problem. We might substitute an input's share of *revenues* in place of its share of *costs*, and this would only be correct if we also assumed firms were competitive and earned zero profits. We might assume that capital's share of revenues was just one minus labor's share, which again only works if firms are competitive and earn zero profits.

There is almost no way that we are getting the elasticities in our index of inputs exactly right. Even the BLS, who tries really hard to get good estimates of labor's share of costs, isn't getting them right. And because we are not getting the elasticities in our index right, our measures of TFP growth are influenced by growth in capital and labor.

What I did in the original post was to look at a particular way in which we are not getting the elasticities right. The crucial part of that was assuming that capital's elasticity in our input index was equal to one minus labor's elasticity. Now we're back to the situation I described in the prior section, where if labor's share in revenues is low, then we are probably understating labor's role in our input index, and overstating capital's role in our input index.

Because capital tends to grow faster than labor, the effect of overstating capital's role (which lowers TFP growth) outweighs the effect of understating labor's role (which raises TFP growth). Hence TFP growth is lowered by some amount, and as labor's share falls, this gets worse. TFP growth will be smaller than technological growth, and the gap will grow as labor's share falls. Slower TFP growth doesn't necessarily imply that technological growth is slow. It may be that our measure of TFP growth isn't just picking up technology, but also changes in inputs.

But we probably care anyway

From everything I described, it sounds like measured TFP growth is just wrong, so why should we care about it? If we could generate an index of inputs that used the right

elasticities, then our measure of TFP would be exactly equal to a measure of technology. And that's what we care about, right?

Not necessarily. Susanto Basu and John Fernald theorized in a [paper](#) that we should actually care about measured TFP even if it is not exactly a measure of technology. In their setting, we have information on both labor's share of revenues and capital's share of revenues (so we are not just assuming it is one minus labor's share). If we use these revenue shares as our estimates of the elasticities in our input index, then they will be less than the elasticities in the production function for GDP, and given this understatement, *both* capital and labor growth will add to measured TFP growth.

This sounds wrong again, just in a different way. But what Basu and Fernald point out is that while not a good measure of technological growth, measured TFP growth would be a good measure of *welfare* growth. Their theory is that what matters for welfare is both the quantity of consumption goods we provide (which comes from generating GDP) but also the cost of providing the inputs (spending time working, or saving output rather than consuming it).

The revenue shares of labor and capital may not accurately reflect their role in production, but they do accurately reflect the cost of providing those inputs. A low labor share means lower wages, for example, which means that the benefit of providing labor is low and we get less labor provided. Similarly for capital. If the revenue share is low then the benefit of providing capital is low, and hence we get less capital accumulation.

However, we do get welfare from the things that labor and capital help produce. And what Basu and Fernald point out is that the value of those things is higher than the cost of providing the inputs when the revenue shares are low. Which means that any additional provision of labor or capital actually adds to welfare, even if it doesn't change technology. Measured TFP picks up this benefit of

providing additional inputs on welfare, and hence is more important than just technological growth.

What do I mean that the value of output is higher than the cost of providing inputs? Think of when we have increasing returns to scale. If we doubled our input provision, we'd get more than twice as much output, and hence we get more value in output than we paid in costs of input. In my setting here, with IRS the elasticities in the production function add up to more than one, while the revenue shares used in the input index can only ever add up to one. And hence measured TFP captures both technological improvement as well as some adjustment for increasing supplies of inputs.

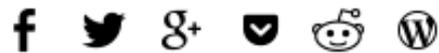
Alternatively, we might not have increasing returns, but we may have firms with market power. In that case, firms are charging a markup over marginal cost, meaning the price you pay for output is higher than the marginal cost of producing it. Again, if we added inputs, the value of output produced would be higher than the cost of providing it. In terms of accounting, the elasticities in the production function may add up to one (constant returns), but with markups some of the output is paid out as profits/rents, so the revenue shares of labor and capital have to add up to less than one. Again, we understate the elasticities in our input index, and measured TFP rises with additional inputs of either capital or labor.

Regardless of the source, measured TFP growth is a better metric for welfare growth than just technological growth. With either IRS or market power, we *want* to use a measure of TFP that is “wrong” in order to capture welfare gains.

In this case, a falling labor share could also lower measured TFP growth (and hence welfare growth), but it is a little fuzzier whether this follows for sure. It depends on how markups respond, and whether the lower labor share is offset by increased profit shares or increased capital shares.

The - very long in coming - conclusion is that measured TFP growth is not “wrong”. There is information in measured TFP growth even if it does not precisely measure technology growth. And measured TFP growth can be influenced by the labor share. When that falls, it can pull down measured TFP growth.

How big is that effect? That’s a different question, and something that I’m playing around with at getting a better handle on. Stay tuned.



Did Real Manufacturing
Output Grow Faster than We
Thought?



Is there evidence of balanced
growth?



