

Commentary: Causes of Changing Earnings Inequality

Kevin Murphy

I'll try to say a few things that Robert Lawrence didn't cover and talk about this paper generally. And again, I'll try to go through the paper in the way in which it was presented.

The paper really begins with a description of the facts. And I think the paper does an excellent job of talking about what's happened in the United States where we've seen an increased premium for education, an increased premium for skilled workers, and greater inequality within those groups—what we call within-group inequality. We have also seen greater instability in earnings, (that is, greater variability in earnings for the same person over time). These (education, experience, within-group, and instability) are four dimensions on which inequality has tended up. There is one dimension—and the paper makes a lot of reference to it—on which inequality has tended to go down: inequality between men and women. The male-female differential has shrunk over time. This fact plays a key role in Professor Snower's analysis. In fact, it plays a pretty critical role in the analysis, but I'll get back to that a little bit later.

One of the things, though, that I want to make sure I do cover—and it is covered in the paper, but it is also something that people often miss—is that the growth in inequality in the United States is a pervasive

phenomenon. Sometimes you hear about it in terms of the super-rich. Sometimes you hear about it in terms of the underclass, but neither of these is the real story. The story is a widening of the income differentials across the spectrum. The 75th percentile did better than the 50th percentile, which did better than the 25th percentile. And so it's a spreading out of the entire distribution. And I think that's what set people on the road of thinking about big phenomena. They are looking for things that can explain why, simultaneously, we saw a spreading out of wages among men, among women, among college graduates, among high school graduates—between college and high school graduates, between the upper end and the middle, between the middle and the bottom—across the board. And the kind of story that people have told is the model that is laid out in Professor Snower's paper. He talks about the supply and demand model. The easy way to think about it, I think, is that you have a growing supply of skilled labor. Simultaneously, you have growing demand for skilled labor. When demand grows faster than supply, the premium for skills goes up. When supply grows faster than demand, the premium for skills goes down. What you don't want to focus on is whether supply is going up, or supply is going down; is demand going up or is demand going down? The truth is, through all periods, not just in the recent periods, the supply of skilled labor has been going up *and* the demand for skilled labor has been going up. That's true in all countries. That's true in all parts of the economy, in all time periods, in the United States.

Whether it is a period when supply is growing faster or slower than demand is going to determine what's going on with the skilled premium. And so we're left with this question: Where does that demand growth come from? It's clear that we can measure the supply growth pretty directly. But where does the demand growth come from? And I think Professor Snower's right to say there have been three types of hypotheses that people have focused on. They've focused on deindustrialization. Think of that as the shift, say from manufacturing to services. They also have focused on globalization—increased trade, particularly with less-skilled countries, countries with different factor endowments from the United States, a factor that one might expect to push factor prices in a particular way. As emphasized in

Robert Lawrence's work, trade is going to have an effect. In simple models, trade effects have to work through prices. But it's also important to remember that if prices have an effect, they're going to work back through quantities within a domestic economy. That is, you have to be in an equilibrium. And, therefore, both prices and quantities adjust in any such equilibrium. But anyway, that's a globalization story.

And finally, we get to the technical change story. And what is technical change? Technical change, at least the way I was taught, was the ability to get more output from the same inputs. And that brings me to the notion of organizational change brought up in the paper. The paper talks about things in these terms: we've reorganized production; we've changed production from a division of labor to a kind of integration of tasks among individuals. I take that to say, you have the same set of individuals and you've changed them in a way that gives you more output—so to me, that is the purest form of technical change. And it's odd that we don't often think that way, because often we say, well, what was technical change? Well, that was the introduction of the tractor; that was the introduction of something else. But that's really a mixture of technical change and capital equipment change. And really, organizational change to me—and from an economic perspective—is technical change in its purest form. Now, you might say this is an issue of semantics. But I do think there are a couple of important implications that have come out of that. And first of all, let's think about what that means.

This is where I think the insights of the paper are really important. Think about what identifying technical change with a phenomenon like organizational change means for measurement. When people talked about demand changes—changes in the output market—the major focus of deindustrialization, they thought that this would be measured in terms of, for example, a shift from manufacturing to services. You'd see reduced demand for manufactured products, increased demand for services. But if the change in product demand, as emphasized in the paper, was an increased demand for variety or timeliness or customization, that would be a product demand shift just like a shift from manufacturing into services, but it would be

going on within manufacturing, going on within services. We would demand more specialized, more timely, more customized manufactured goods and more timely, more customized services, as well. So when we use our traditional measures, we would almost certainly miss it. And I think that's an important lesson of this paper. It says, if this is the kind of demand change that we've seen, then our traditional industrial breakdowns aren't going to capture it very well.

The same is true for technological change. Technological change, as Professor Snower said, is a residual. But when you try to get behind the residual, how do you do it? Well, one of the common ways is to look at occupational mix and say we ought to see the occupational mix changing to reflect the new technology. So we look for growth in those kinds of occupations that need educated labor and falling relative employment in those kinds of occupations that need unskilled labor.

Now think about an example that he talks about in the paper. He gives the example, what if we delayered management? What if we got rid of a lot of middle managers? You might say to yourself, well, middle managers tend to be pretty educated. So in the traditional view, you might say that this is something that ought to reduce the demand for skilled labor. But, as he correctly points out, if what you're really doing is moving those tasks—those coordination, synthesis, and analysis tasks—out of the middle management area and into the more production-oriented areas, moving them into those integrated jobs, you are raising rather than lowering the demand for skilled labor. So I think that there is a difference between this organizational change view and our traditional view of technical change. Not in the theoretical construct, because I think it fits squarely in what we thought of as technical change. But more in terms of the implications it has for how we have to *measure* the impact of that technological change. I think that's the real lesson brought out in the paper.

It's also important to realize that past technological revolutions have brought on their own share of organizational changes. When we introduced the factory system, enormous organizational changes took place. So, it's not really unique to the more recent period. One

way to summarize organizational change would be an increase in the demand for workers who can analyze, synthesize, and respond to a variety of situations. And I think he's right to say that that's something closely related to the demand for education. You go back to the work of T.W. Schultz. He really emphasized this organizational attribute of education—that education taught you how to deal with a variety of situations. Think about the dissertation work of Finis Welch. He did the same thing. He showed that educated farmers did well precisely because they made good organizational decisions. If that's where we're going now, then it's clear why that would feed back in terms of demand for education.

So I think the model that Professor Snower lays out, and the framework that he lays out, actually fits very well within the supply and demand framework. In fact, often when I've heard people discussing technical change, this is what they've had in mind. But, because of the same measurement issues described above, some of the analysis ran into a type of inconsistency. They would start their paper discussing these issues and then they'd turn to the traditional measurements, because that's all they had when they actually went to look at the data.

Now there are a couple of criticisms that come up in the paper, and I want to address each of those. First of all, it is important to discuss the interpretation of the narrowing of the male-female gap. The male-female gap is narrowing for a variety of reasons, many of which are outside the scope of what we're talking about here. Many have to do with changes in the relationship of women to the labor market, in terms of increased education, increased labor force attachment, changes in the types of jobs that females are engaged in, changes in the occupational mix for women, and the like. This is probably the most important factor and has little to do with technical change, organizational or otherwise. But I think that there is at least some relationship to technical change. If you look at who works with whom, what people are working together, you'll see that women, particularly women with a high school degree and above, have disproportionately worked in the same areas of the corporate or industrial environment as have educated men. So that when you see an increase

in a demand for education, I don't think it's surprising that the demand for workers who work with those same individuals also goes up. So the fact that the information technology revolution would affect women, given where they worked, is not so unusual. The one group of women I think you would expect not to have benefited very much would have been the very low-skilled women, women with less than a high school degree. And I think the evidence helps support that idea.

Now let's go on and talk a little about earnings variability. This is, I think, one of the most misunderstood areas when people talk about inequality. Remember what most of earnings variability is: the way people get big changes in their income is by changing their jobs. They either get a much better job or they lose a job and move to another one. In a world in which the skilled premium is larger, you get a bigger premium for being in a highly skilled situation than you would for a low-skilled job. It's not surprising. In fact, it would be almost impossible to believe that in a world where inequality was higher at a point in time, you wouldn't see more earnings variability. That is, the same move of losing my job and moving to a temporary low-skilled job is going to have greater consequences in today's economy when the wage gap between those two jobs is larger, than it would have before. And the evidence, I think, is right on with that. It says that if you looked at percentile mobility, in terms of variability, you wouldn't see an increase. I'm just as likely to lose my job at the 75th percentile and go to the 25th percentile, as I was before, maybe even a little less. But the point is, the consequence of that same percentile move, in terms of the actual wage change, is much greater than it was before. So I don't think the growth in instability, in that sense, is at all a mystery. In fact, I would be very surprised if one did not see that going up at the same time that you saw an increase in inequality in general. Thus, I do not think that the growth in earnings instability is a challenge to the supply and demand view in any way.

The other criticism, I think, comes when we talk about unemployment and the growth of unemployment among the unskilled. There are a couple of issues here. One is just to think about measurement. When the unemployment rate for college graduates goes from 1½

percent to 2½ percent and the unemployment rate for high school graduates goes from 5 percent to 8 percent, did it go up more for college graduates than high school graduates (which you would get if you look at ratios)? I don't think so. I think it's an issue of measurement. And also, I think the facts are that for European countries, unemployment rates do go up for high-skilled workers. That doesn't mean they went up as much for high-skilled workers as they did for low-skilled workers. I think that's part of the factual side. Certainly, in the United States, there was tremendous concentration of the growth of unemployment among the less skilled.

Now, I would like to take up the point that recent technological changes have redefined the notion of skill. First, let me say that I think this is true of all technological change. When we use terms like high-skill or low-skill, it's always a very shorthand notation, a temporary notion of what represents being "skilled" at a particular time. Maybe even skill is a bad term. I've found when I go out and talk to more popular audiences about this concept that, when you talk about skilled labor, people tend to think about skilled craftsmen and the like. They don't think about educated workers who have cognitive and analytical skills as the quintessential skilled worker. As I said before, technical change has been going on for a long time. And it's been accompanied by organizational change as well. That's nothing new. That gets us back to the supply side. Again, I can't say enough about how the supply side is really important. It appears that in the United States in the 1970s we saw a suppression of the growth in the college premium precisely because of the fact that the fraction of college workers grew so rapidly. Comparing Canada and the United States over the more recent period, we see the same thing. Supply grew rapidly in Canada—grew but much less so in the United States. And as a result, we saw a widening of college-high-school premiums in the United States and not so much in Canada.

Finally, I just want to put one last spin on things, which is to talk about inequality from a different perspective. There is always this notion to think about the inequality "problem." And I agree that inequality creates a lot of problems. But if I were to tell you that the return to human capital investment would be 60 percent higher today

than it was a decade and a half ago, you'd be hard pressed to say that, in itself, was a problem. So I think there are certainly elements of the growth in inequality that can be viewed in a positive light, particularly if people are able to respond to the incentives generated by rising inequality—such as taking advantage of the tremendous increase in the return for going to college.

Finally, and I almost forgot this, I want to agree with Robert Lawrence in the sense that this very redefinition of skill is inherent in technological change. That we're really changing the kinds of outputs and the way we're producing them has an impact, because that relates closely to the notion of whether we're overestimating or underestimating the growth in the demand for skill and technological change from a measurement perspective. It may be that in the current period we really are mismeasuring the economic growth, if, in fact, this type of organizational change is going on. This is an important contribution of the paper.

It also causes us to think about long-term growth in general. And this, I think, is an important point. We all know how to think about technological progress and the accumulation of physical and human capital as part of the general way in which economies progress. But you want to think about skill-biased technical change as an equally important part of that process. It is the skill bias of technical change that allows us to continue to accumulate higher and higher levels of human capital without having diminishing rewards to that investment. And, if anything, the rewards have continued to rise because skill-biased technical change (including organizational change) has actually outpaced growth in the supply of skilled labor. That means we can continue to invest in human capital and reap significant rewards on those investments.