Volume Title: New Developments in Productivity Analysis

Volume Author/Editor: Charles R. Hulten, Edwin R. Dean and Michael J. Harper, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-36062-8

Volume URL: http://www.nber.org/books/hult01-1

Publication Date: January 2001

Chapter Title: After â€œTechnical Progress and the Aggregate Production Functionâ€•

Chapter Author: Robert Solow

Chapter URL: http://www.nber.org/chapters/c10126

Chapter pages in book: (p. 173 - 178)

5

After "Technical Progress and the Aggregate Production Function"

Robert M. Solow

My never-failing source of guidance on occasions like this is a telephone call that Paul Samuelson once made. He had agreed to be the lunch speaker at a meeting of some business group. Talking to the organizer, Paul asked: "Should I talk for 30 minutes?" The reply was, "Thirty minutes would be optimal. Twenty minutes would be better."

That 1957 paper of mine is now a little over forty years old, about the same age as my youngest child. Like my children, it has aged well, and has produced grandpapers. It is very nice of the organizers of this conference to take note of it and invite me to comment on developments since then, and perhaps still to come.

It goes without saying that a lot of effort, by many hands, has gone into whittling away at the residual. Progress is being made any time some part of the residual can be imputed to some measurable input, or to some adjustment of measured output. This process began very quickly with the work of Ed Denison, and it continues today, even at this conference. I hardly need to mention the important advances that have been made in this way, but they include the introduction of human capital as an explicit input (or the adjustment of hours worked to account for the varying mix of skills); the attempt to define and measure a stock of technological knowledge that could be entered in a production function along with stocks of tangible and human capital; the effort to get a better measure of the flow of capital services as distinct from the mere existence of a stock (in which connection I should mention the almost single-handed campaign of Murray Foss to get a grip on the length of the work-week or work-year

Robert M. Solow is Institute Professor of Economics, Emeritus, at the Massachusetts Institute of Technology.

of capital goods); and the need to include the depletion of natural resources either as an input to production or as a deduction from output. I am undoubtedly forgetting some lines of research that lead in this general direction.

Sometimes this kind of work has to focus on measurement difficulties as the main obstacle to progress. This is no trivial matter. Recent papers by Ruth Judson and by Peter Klenow and Andres Rodriguez-Clare have shown that alternative ways of measuring human capital can make a big-time difference in the plausible interpretation of economic growth, so it is really important to come to some scientific agreement on the best way to deal with human capital as an input (and as an output, don't forget) and then to implement it. When comparable attention comes to be paid to natural—including environmental—resources, similar subtleties may arise.

Sometimes, however, this process of whittling away at the residual runs up against obstacles that are not simply measurement problems, but are conceptual problems or modeling problems. I want to come to one or two of those in a few minutes, but first I just want to mention a couple of conceptual issues that are too complicated to discuss in a lunch talk. For instance, there is the question of how to deal with the quality of capital goods. I suppose that if you play fast and loose with this notion, you could get rid of most of total factor productivity in a hurry. I *think* that is what Zvi Griliches and Dale Jorgenson once proposed to do. That was too bad; there was a lot of important stuff in that paper, and some of it got hidden behind that one minor and misguided side issue. To go very far down that line you have to think you can identify the source of improved quality of capital goods. And to do that you have to have greater faith than I can muster that whoever gets to appropriate the benefits of improved quality is also the source of improved quality. I will come back to this quality-of-capital question in a moment, in a more restricted context where it can be dealt with clearly.

That leads directly to a broader conceptual issue that I want to mention but don't want to discuss seriously. In accord with the spirit of the time in economics, there is a tendency to embed the estimation of technical change and the aggregate production function in a general equilibrium model, usually a very special general equilibrium model. Not only does this fit in with the spirit of the time, but it also may provide relatively easy answers to otherwise difficult specification problems. I find I am suspicious of this way of proceeding.

You might say that I am partially responsible for starting that habit forty years ago by cheerfully using observed factor shares as estimates of elasticities. Yes; and it made me uncomfortable then. As the process goes further, my discomfort increases. The problem is not that I am hostile to general equilibrium theory, or even to competitive general equilibrium the-

ory in its place. I am not. The generic problem is that any kind of empirical application then becomes a test of a very broad combined hypothesis, including the hypothesis that the whole economy behaves as if it were doing something very special. You never know how to interpret the results, whether they are good or bad.

I think I will spend the rest of my time talking about a couple of other conceptual issues that arise in the continuing project of analyzing productivity growth into its measurable components and the residual. The first is—to quote the title of a very relevant paper by Dale Jorgenson—the embodiment question.

In the year following my 1957 paper, I wrote another paper called "Investment and Technical Progress." It was done for a conference at Stanford and took two years to get into print, but I think it was actually written in 1958. It seemed to me then that the 1957 model might have grossly understated the importance of old-fashioned capital investment as a vehicle for bringing new technology into productive operation: No amount of clever jet-engine technology could affect productivity unless airlines bought jet aircrafts—that sort of thing. The new paper produced a clean model in which *all* new technology had to be embodied in new gross investment before it could have any influence on production or productivity.

I liked the idea, but it went nowhere. Nobody ever suggested that it wasn't entirely plausible. How could they? It *is* plausible, common sense even: If you don't like the jet-engine example, how about numerically controlled machine tools? (You see the connection to the quality-of-capital problem.) The problem was not plausibility; it was that embodiment seemed to cut no empirical ice at all, and if you couldn't find the embodiment effect leaving a significant trace in data, then it wasn't really so interesting.

There are simple reasons why it might not show up even if it were really there. For instance, a steady state with embodiment looks like a steady state without embodiment; embodiment works through changes in the average age of capital equipment, and in or near a steady state the average age of equipment will be just about constant. So, unless the rate of investment fluctuates substantially, you would not see the embodiment effect at work. Or, to look in a different direction, if a substantial part of new technology can be retrofitted to old capital at a small expense, then the embodiment effect can be unimportant, or it can be swamped by other influences. (I have been told that a lot of chemical engineering technology is actually designed to allow retrofitting.) Anyway, decades went by and embodiment languished. Sad. (Ed Wolff thought he saw it once, and more recently Larry Lau had a sighting; it occurred to me that if it could be connected to the Clinton family, Kenneth Starr would either find it or invent it, or at least leak it. The general impression remains: no dice.)

But now all of a sudden the whole idea has revived, under the influence

of an entirely new approach, as you will all learn when you hear the Jovanovic-Greenwood paper this afternoon, which builds on earlier work by Chuck Hulten, Jeremy Greenwood, Zvi Hercovitz, Per Krusell, and no doubt others. I don't want to spoil their story. Think of what I am about to say as a preview of coming attractions.

Imagine a two-sector economy. In one sector capital goods are produced by labor alone. In the other sector already accumulated capital goods (or their services) combine with the remaining labor to produce consumer goods. I want to tell two different stories about this economy. In one story, technological progress steadily reduces the labor input required to produce a unit of capital goods in the investment sector. The capital goods don't change at all. They enter the production function for consumer goods in the same old way. They become cheaper relative to consumer goods as time goes by, for obvious reasons.

The second story is more complicated. Suppose that the investment sector continues to produce "machines" with the same labor requirement as at the beginning. You could perhaps hear that I put machines in quotation marks, because I want to suppose that newer machines are more productive than old ones when employed in the manufacture of consumer goods. In other words, there is capital-embodied technological progress. When I say that it takes the same old amount of labor to produce a machine, I must be measuring the output of machines crudely—say, by counting the number of legs and dividing by four. It will come out all right at the end. The output of consumer goods is calculated by imagining an efficient allocation of labor to machines of different vintages, and then adding up the output over surviving vintages. (This is the kind of model I produced in 1958.)

Now the interesting thing is that these two stories are essentially indistinguishable from the macroeconomic point of view, by which I mean you can't tell them apart if you look only at aggregates and price indexes. You can see why that might be so: In story number two newer machines are more productive than older machines. In story number one that is not literally true; but if you think of machines as congealed labor, newer investment-sector labor is more productive (of consumer goods) than is older investment-sector labor. Under any reasonable pricing rule, the relative price of machines will be falling in story number one because the labor cost of producing a machine is perpetually falling, and there are constant returns in the production of machines. In story number two, the labor cost of producing a machine (in quotation marks) is constant, but the machines get more productive. If you asked for the labor cost of producing a machine of constant quality, it would be falling. And it would be falling at the rate of embodied technical progress. (Of course, if you look closely you can see the difference between a jet and a propeller.)

The new approach talks of investment-specific technological progress,

because that term is meant to cover both stories. In both stories, it is true that society can only take advantage of technical progress to the extent that it invests. That is what the whole idea is about, after all. And it then becomes fairly natural to *measure* the rate of investment-specific technical progress by looking at the rate at which the relative price of quality-adjusted capital goods falls. According to this scheme the rate of investment-specific technical change has averaged close to 4 percent a year during the postwar period. (Needless to say, this does not translate one-for-one into output growth—not nearly.) What is even more interesting is that there is no sign of a post-1973 slowdown in the investment-specific component of technical change. But I will leave all this to the paper by Greenwood and Jovanovic, which you can look forward to.

What I do want to call to your attention is that this body of research suggests and confirms the importance of two general points I have been trying to make. The first is the way conceptual issues and measurement issues get intertwined. In the case of embodiment, an awful lot hangs on our ability to estimate the price trend for quality-corrected, or productivity-corrected, capital goods. Bob Gordon's heroic campaign to do just that was an indispensable foundation for this new line of research. It has to be carried forward and extended, as the Bureau of Labor Statistics is now doing.

The second point is more worrisome: It is the way that empirical research gets tied up with complex hypotheses about pricing and other aspects of market behavior. The idea of a dichotomy between measuring and modeling is breaking down. This is not a remark about productivity research but about economic research, and maybe about research in general. Progress in measuring and understanding investment-specific or capital-embodied technical change will then be tied up with different stories about the way the economy functions. I said that this is worrying, but it's not such a big deal. I only want to urge that research not get tied to any one particular picture of the way the economy functions. It would be a good idea to try out some alternative theories of factor pricing, for example, so we can have an educated idea about how sensitive the measurement outcome is to additional assumptions about the market environment from which observations are presumed to arise.

I am coming to the end of the time I allowed myself. The last thing I want to do is to throw out another pet idea and hope that it can play a role in future research on productivity. Bits of experience and conversation have suggested to me that it may be a mistake to think of R&D as the only ultimate source of growth in total factor productivity. I don't doubt that it is the largest ultimate source. But there seems to be a lot of productivity improvement that originates in people and processes that are not usually connected with R&D. Some of it comes from the shop floor, from the ideas of experienced and observant production workers. This should

probably be connected with Arrow's "learning by doing" or with the Japanese slogan about "continuous improvement." There is another part that seems to originate in management practices—in design, in the choice of product mixes, even in marketing. Notice that this is not just straightforward enhancement of productive efficiency. All this talk about value creation that one hears from business consultants may be more than a buzzword; it may even be important. We need to understand much more about how those kinds of values get reflected in measured real output, and whether they can be usefully analyzed by our methods.

This is an inexhaustible subject; but patience is a scarce resource. Thank you for yours.