What Have We Learned About Productivity in the Last Two Decades?

A Review Article on *New Developments in Productivity Analysis*

Andrew Sharpe
Centre for the Study of Living Standards

In 2001, the Conference on Research in Income and Wealth (CRIW) published an important research volume on productivity issues entitled *New Developments in Productivity Analysis*, edited by Charles R. Hulten, Edwin R. Dean, and Michael J. Harper.[1] Many leading researchers in the productivity field contributed papers or comments to the volume. It had been over two decades since the last CRIW volume on new productivity developments was published (Kendrick and Vaccarra, 1980).

The appearance of this volume represents an opportunity to assess the state of knowledge in productivity research over the last two decades, as represented by the papers in the volume, many of which constitute work on the frontier of productivity research. A key question is in which areas of the productivity field has our understanding increased and in which areas has progress been slow. The objective of this review article is to identify the major advances in productivity research, and to indicate the areas in which progress has been more limited.

The volume contains 15 papers, 10 comments on papers, and a short introduction by the editors, and runs to 631 pages. This review article is divided into eight sections. Part one discusses the history of the concept of total factor productivity (TFP), often referred to as the residual and the theoretical backbone of productivity analysis. Part two looks at the advances that have taken place in the crucial area of data, including the evolution of the BLS productivity program, the development of micro-data bases, data gaps, and dynamic factor demand models. The third section discusses the revival of the vintage capital or embodiment approach to economic growth. Part four examines the relationship between productivity and the business cycle. Part five looks at international productivity comparisons and productivity convergence. Part six summarizes the lessons from the McKinsey service sector productivity studies. Part seven outlines the integration of natural resources and the environment into the TFP framework. Part eight looks at what we do not know about productivity.

The History of the Residual

In the volume's lead paper, co-editor Charles Hulten provides an excellent overview of total factor productivity or the residual, which he considers the "workhorse of empirical growth analysis." He traces the history of the concept from the development of the theoretical link developed between the production function and index number approach in Solow's seminal 1957 paper to its current use as one of the most monitored statistics produced by the U.S. government. While noting that total factor productivity is commonly identified with technical change, he argues that this is not an appropriate interpretation.[2] Rather the residual is "a measure of our ignorance," encompassing the effects of technical and organizational innovation, but also measurement error, omitted variables, aggregation bias, and model misspecification.

Hulten discusses the strengths and weaknesses of both the index number or non-parametric and econometric or parametric approaches to TFP estimation, concluding that it is pointless to debate whether the benefits of one approach outweigh the costs because there is no reason to view the two approaches as competitors. Indeed, the indexes of output and input used in the two approaches are the same. Hulten recommends a joint approach, with the residual disaggregated into terms for increasing returns to scale, costs of adjusting factor inputs, technical innovation, unclassified trend productivity, and measurement error. Hulten does recognize, however, that the tide has begun to turn against the aggregative non-parametric approach for a number of reasons, including the shift in interest from the aggregate or industry level to the firm or plant level, and the shift from the competitive model of industrial organization to non-competitive models.

Hulten admits that the TFP framework did not furnish a consensus explanation of the post-1973 productivity slowdown, but points out that it is unclear whether this shortcoming reflects problems inherent in the character of the residual or rather problems inherent in the data to which the TFP framework is applied.

An excellent example of the uncertainty associated with TFP estimates is provided in the article, but not discussed. For the total U.S. economy, TFP is estimated to have increased at a 0.25 per cent average annual rate from 1949 to 1969, accounting for only 9 per cent of real GDP per capita growth of 3.05 per cent. In contrast, over the 1948-73 period in the private business sector, TFP is estimated to have advanced at a 2.1 per cent average annual rate, accounting for 64 per cent of growth in real GDP per capita of 3.3 per cent. While the periods and coverage of the economy differ somewhat, the large discrepancy in the magnitude of the TFP growth rates and the relative importance of the contribution to output growth is very surprising, and disconcerting. It undoubtedly reflects differences in the methodologies used to construct the estimates. The lesson may be that while TFP remains an essential tool for productivity analysis, the range of values that can be calculated for TFP estimates is very large. The range of estimates for the less sophisticated concept of labour productivity is much smaller.

Data, Data, Data

A key theme found in almost all papers in the volume is the importance of data for productivity analysis. This of course is not a new concern, but recognition of the complexity of data issues and their crucial role for our understanding of productivity trends appears to be growing. Many papers stress that econometric results based on faulty data have little meaning. This section of the review discusses the BLS productivity program, the development of firm micro-data sets data gaps, and dynamic factor demand models.

BLS Productivity Program

The Bureau of Labor Statistics (BLS) is responsible for the development of official productivity statistics in the United States and is a world leader in productivity measurement and research. A very useful overview of the BLS productivity program is provided in the paper by Edwin R. Dean and Michael J. Harper, both of whom have been or currently are associated with the program. One of the most important innovations of the program has been the development of official multifactor productivity (MFP) estimates for the U.S. private business sector, first published in 1983. The BLS was the first statistical agency in the world to produce such data. MFP estimates have required the development by the BLS of appropriate estimates of capital services and labour composition. The authors review the main elements of this work, which has largely taken a production theory approach.

The authors provide a very revealing comparison between what is called BLS "production theory" measures of outputs and inputs and measures based on more traditional measurement techniques. The former produce much smaller estimates of multifactor productivity growth than the latter: 0.0 per cent versus 0.8 per cent per year in the 1979-90 period and 0.3 per cent versus 1.8 per cent in the 1990-94 period (more recent estimates are not available). Output growth is smaller in the "production theory" approach and labour and capital input growth are greater. The authors do not discuss the reasons for these discrepancies nor make the case for the superiority of the production theory estimates, if such a case can be made. The reader comes away with the impression that once one moves beyond basic labour productivity[3] to multifactor productivity, productivity growth estimates are not particularly robust, being extremely sensitive to the choice of methodology.

The BLS productivity program has done a great deal of work at the industry level and now publishes labour productivity measures for around 500 industries. However, because of data problems, productivity estimates are not produced for all industries in the business sector or total economy. This incomplete coverage, particularly in service industries, means that BLS industry productivity measures are not particularly useful for comprehensive analysis of productivity trends by industry and calculations of sectoral or industry contributions to aggregate productivity growth. For this purpose, the estimates of output and labour input produced by the Bureau of Economic Analysis (BEA) are much more useful, even though the reliability of some of BEA's estimates of real value added and gross output may be problematic because of measurement problems. In my view, from the perspective of productivity research, it is better to work with industry productivity estimates that are comprehensive, even if part of the data is recognized to be of poor quality, than to be limited to incomplete data sets of higher quality.

One development that has been a boon to productivity researchers has been the posting of BLS productivity data and BEA output and input data on their respective websites (www.bls.gov and www.bea.gov) for free and user-friendly download. Both agencies are to be congratulated for this initiative, which is particularly appreciated by researchers in countries such as Canada where the statistical agencies charge large sums of money for data access.

Availability of Micro-data Bases

The development of micro-data bases has been a great benefit to productivity research in the past two decades. The Center for Economic Studies (CES) at the U.S. Bureau of the Census has been responsible for producing much of the establishment-level and firm-level data. The paper in the volume by Lucia Foster, John Haltiwanger, and C.J. Krizan, all associated with CES, discusses the lessons from microeconomic evidence for aggregate productivity growth. A key theme that emerges from their discussion is the importance of the large-scale reallocation of outputs and inputs that is continuously taking place across individual producers, both between and within sectors. The pace of reallocation varies over time both secularly and cyclically and across sectors. Because of the large differentials in the levels and growth rates of productivity across establishments within the same sector, these reallocations can have a major effect on aggregate productivity growth. For example, it was found that the primary source of productivity growth between 1987 and 1992 for the automobile repair shop industry was the exit of very low-productivity plants.

Micro-data sets have been developed from a variety of sources, including government mandated reporting. The paper by Denny Ellerman, Thomas M. Stoker, and Ernst Berndt uses information collected by the Mine Safety and Health Administration to construct a plant-level data base on productivity in the American coal industry. Through the analysis of this data base, they show how trends in coal prices have been the driving force behind developments in labour productivity. The hike in coal prices in the 1972-78 period related to the OPEC price shock led to the opening of mines that were not only smaller, but also geologically inferior, lowering aggregate labour productivity. The fall in coal prices in the 1980s saw the closing of many of these low-productivity operations and the revival of productivity growth in the sector. The authors are careful to stress that the decline in labour productivity in the 1970s did not represent technical regress as it was price-induced. Indeed, technological progress in coal mining continued to advance during this period.

Data Gaps

Despite the progress of the past two decades, particularly in the development of MFP estimates and the micro-data bases, productivity researchers continue to encounter many data gaps. In his paper, Erwin Diewert reviews a number of ideas on productivity measurement which he believes should be embraced by statistical agencies. He points put a number of deficiencies or gaps in our current statistical system. Perhaps the most glaring is the lack of data on interindustry leasing of capital and flows of business services because of the lack of surveys in this area. He also makes a case for the inclusion of inventories and land as factors of production in productivity studies and the development of estimates of knowledge capital. The latter is defined to include stocks of patents, R&D expenditures, and employer education and training.

Finally, he assesses the relative merits of a central statistical agency model based on his dealings with Statistics Canada versus those of a decentralized model based on his experience with U.S. statistical agencies. He concludes the case is much stronger for the former and in fact recommends the establishment of a central statistical agency (Statistics USA) that would bring under one roof the statistical units of U.S. government departments including the Department of Labor and Department of Commerce. While such a recommendation may make sense from the point of view of the users of the U.S. statistical system, Diewert does not address the political obstacles that such a reorganization might face.

Dynamic Factor Demand Models

The index number approach to productivity measurement only provides meaningful estimates of technical change when restrictive conditions are met. These conditions include constant returns to scale, competitive markets, full utilization of all inputs, and instantaneous adjustment of all inputs to their desired demand levels. In their paper, M. Ishaq Nadari and Ingmar R. Prusha make the case that dynamic factor demand models, which are flexible and can incorporate a wide variety of different assumptions, can provide great insight into the nature of productivity growth. They point out that such models generate a very rich set of critical information about the structure of production, sources of productivity growth, and the impact of technical change and policy instruments on output, input demand, and productivity growth.

As the discussant of the paper, Dale Jorgenson expresses a certain degree of skepticism about the dynamic factor demand approach and its attempt to measure the effect of departures from the standard assumptions. In particular, he argues that there have been no empirical examples of the successful implementation of complex specifications. The authors dispute this view, pointing to the decomposition of the conventional measure of TFP in the U.S. electrical machinery industry provided in their paper.

The Revival of the Vintage Capital or Embodiment Approach

Robert Solow was a luncheon speaker at the conference where the papers in this volume were first presented. His address is included in the volume and, as usual, provides much insight. The central theme of Solow's address is that he thought his famous 1957 paper on technical change and the aggregate production function grossly understated the importance of capital investment as a vehicle for bringing new technology into productive operation.[4] In fact, Solow notes that he published a paper three years later (Solow, 1960) that developed a model in which all new technology had to be embodied in new gross investment before it could affect productivity. Although based on common sense, the idea went nowhere. This was because embodiment did not show up in the data since in a relatively stable investment environment, the average age of the capital stock is constant.

The paper by Jeremy Greenwood and Boyan Jovanovic is an example of the revival of interest in vintage capital models to which Solow refers. In contrast to the standard assumption of growth models that all capital is the same, the paper starts from the premise that advances in technology tend to be embodied in the latest vintages of capital, making new capital more productive than old capital, even when the old capital was new. In other words, there can be no technical progress without investment. They argue that a vintage capital model does much better than the dominant growth-accounting framework in explaining four features of U.S. postwar growth, namely the post-1973 productivity slowdown, the falling price of capital goods relative to consumption goods, the large productivity gap between best-practice and average plants, and the recent rise in wage inequality.

The innovation of the paper relative to the standard vintage capital model is the introduction of a technology-specific learning curve on the part of users of capital goods and lags in the diffusion of new technologies. The authors then assume that the vintage-specific efficiency of investment accelerated in the 1970s with the advent of information technologies, that these new technologies have steep learning curves, and that diffusion takes time. The productivity slowdown is thus explained as a period of above-normal unmeasured investment in human capital specific to the information technologies, with increased wage inequality a result of the rise in skills premia again associated with IT. Different vintages of capital within a sector can account for large intra-sectoral productivity differences. One policy implication of this analysis is the importance of fostering investment, including policies that reduce the costs of acquiring new equipment.

The vintage approach is much closer to reality than the standard approach and more work needs to be done in this area. But as the discussant Barry Bosworth noted, it is hard to believe that the lag between the purchase of the machines and the time required to learn to use them is so long that an acceleration of technological change has the perverse effect of reducing TFP growth.

Productivity and the Business Cycle

An issue that has bedevilled productivity researchers for years is the relationship between productivity growth and the business cycle. The conventional wisdom is that productivity is procyclical. The key questions are why and how do we best decompose the cyclical and structural components of productivity growth. The paper in the volume by Susanto Basu and John Fernald directly addresses these concerns.

The authors note that economists have traditionally not been particularly interested in the cyclical behaviour of productivity as they believe that cyclical fluctuations wash out over the cycle. They consequently have focused attention on long-run average productivity growth for the analysis of growth and welfare. However, the authors point out that more recently productivity fluctuations have taken centre stage in modeling output fluctuations and are now viewed as an essential part of the cycle. Procyclicality is considered closely related to the propagation mechanisms underlying business cycles.

The authors propose four explanations for the procyclical nature of productivity. The first explanation is that procyclical productivity may reflect procyclical technology, as in the aggregate technology shocks of real business cycle models. Second, imperfect competition and increasing returns may lead productivity to rise whenever inputs rise. Third, utilization of inputs may vary over the cycle. Fourth, reallocation of resources across uses with different marginal products may contribute to cyclicality.

The authors then attempt to identify the importance of the four explanations, using a methodology[5] that builds on earlier work by Robert Solow and Robert Hall. The major findings are: that variable utilization and cyclical reallocations appear to explain much of the cyclicality of aggregate productivity; that technological improvements reduce input use, a result inconsistent with real business cycle models, but consistent with sticky-price models; and that reallocations are welfare relevant, not biases.

A key issue that the paper does not resolve is how to determine the most appropriate methodology to separate trend and cycle components of productivity growth. I was reminded of the importance of this issue and the lack of professional consensus on it by the debate over the sources of the labour productivity acceleration in the second half of the 1990s in the United States. The Council of Economic Advisors (2001:27) concluded that almost none of the acceleration in productivity after 1995 was cyclical. This result was based on an econometric model in which hours worked adjust gradually to changes in output. This model found that by 1995 strong demand had already pushed productivity about 2 percentage points above where it would have been otherwise so that in the 1995-2000 period the cyclical component edged up only slightly. In contrast, Robert Gordon (2000) found that about one half of the acceleration in labour productivity growth after 1995 was due to the strength of the business cycle. A key priority for productivity researchers must be to reduce the range of estimates for the impact of the business cycle on productivity.

International Comparisons of Productivity Levels and Productivity Convergence

Interest in the measurement of productivity levels across countries has increased in recent years, in part due to the debate on international convergence of productivity levels. Indeed, productivity levels in some countries, particularly in East Asia, have converged toward those of the productivity leader, while in other countries, they have not. Insight into the forces behind this convergence process is essential for an understanding of the overall development process. Two papers in the volume address this issue. Nazrul Islam discusses different approaches to international comparisons of total factor productivity and Dale W. Jorgenson and Eric Yip present international comparisons of patterns of economic growth over the 1960-95 period.

Islam notes that international differences in TFP have traditionally been studied with a time series growth accounting approach, but because of data constraints this approach has been limited to developed countries, even though the concept of convergence in principle applies to all countries. He notes that two new approaches have been developed for international TFP comparisons, namely the cross-section growth accounting approach and a panel regression approach, and provides a detailed exposition these approaches. He makes the obvious point that TFP differences are largely based on technology differences, despite the fact that economic theory has assumed that technology is identical across countries and that differences in capital intensity account for productivity differences. Again, the separation of the impact of technology and capital accumulation on productivity is untenable.

Islam reviews the debate on TFP growth in East Asia economies. The accumulation view, started by Alwyn Young, argues that TFP growth in Singapore and Hong Kong has been limited and that economic growth has been largely driven by factor accumulation, not technological progress. More recently, a revisionist view has developed, finding high TFP levels in these countries, which implies that past TFP growth has been strong. Indeed, Islam estimates that Hong Kong has the world's highest TFP level (54 per cent greater than that of the United States) and Singapore the fifth highest.

The motivation for the Jorgenson-Yip paper was the disappearance of TFP growth in the first half of the 1990s in G7 countries, a development that has been reversed in the second half of the 1990s with the revival of productivity growth, at least in the United States. The authors present estimates of output per capita, input per capita, and productivity for G-7 countries for the 1960-95 period based on a methodology that adjusts capital input for changes in quality through changes in asset types and ownership sectors, and changes in labour input though quality or composition changes in the demographic characteristics of the workforce including sex, educational attainment, and employment status.

Jorgenson and Yip find that Canada in 1970 was the first country to overtake the United States in terms of TFP levels, although only temporarily, as by the mid-1980s the United States had regained the lead. The productivity leader in 1995 was France, with a TFP level about 5 per cent above that of the United States. France's advantage was largely based on its low labour input per capita, about one half that of the United States due to levels of hours worked per capita and labour quality about three quarters that of the United States. Indeed, France had by far the lowest hours worked per capita and labour quality in the G-7.

I find it rather ironic that the G-7 country with the poorest quality labour has the highest TFP level, with the former causing the latter (less input increases TFP, ceteris paribus). From this perspective, a strategy to increase TFP would be to attempt to produce the same output with poorer quality labour. Yet I doubt this strategy would find many takers. From the point of view of societal welfare, what we want to maximize is labour productivity, not TFP as it is the former that determines the material conditions of society.

McKinsey Service Sector Productivity Studies

With the service sector accounting for well over two thirds of output and experiencing below-average productivity growth, and with productivity measurement difficult in many service industries, the issue of service sector productivity has received increased attention in recent years.[6] The paper by Martin N. Baily and Eric Zitzewitz provides lessons for measurement from the service sector productivity comparisons that the McKinsey Global Institute has conducted. Instead of testing economic theory, the McKinsey studies try to provide the best estimate of both the magnitude of the productivity differences among industries in different countries and the explanations for those differences at the production process level through plant visits and expert interviews.

The authors summarize their case studies of productivity trends and determinants in retail banking, telecom, retailing, and airlines, providing fascinating insights into international productivity dynamics in these sectors. They point out, for example, that the hub-and-spoke system in the U.S. airline industry has actually had a negative effect on measured labour and capital productivity as it leads to extremely peak-driven operations, with unfavourable implications for labour and capital utilization.

They conclude that from the point of view of international productivity comparisons, physical output measures work surprisingly well. They also find that current business systems would be impossible without information technologies (IT). Indeed, they stress that a substantial increment to productivity is associated with IT use even though it represents only a small component of total capital, providing evidence that the contribution of IT to productivity growth can be greatly underestimated by assuming that IT's income share is proportional to its contribution. Finally, they consistently find that intense price competition in an industry is the best way to stimulate higher productivity.

The McKinsey approach to productivity analysis with its real world orientation has shed much light on productivity trends in the sectors they have studied. The lessons from this approach are extremely useful and relevant both for researchers and for policy-makers and should be integrated into the traditional production function approach to productivity analysis.

The Integration of Natural Resources and the Environment into the TFP Framework.

The depletion of natural resources and the degradation (or improvement) of the environment have traditionally not been integrated into the total factor productivity framework. In my view, the most innovative and ambitious paper in the volume is by Frank M. Gollop and Gregory P. Swinand. They attempt to construct a total resource productivity (TRP) framework that can account for changing environmental quality. They argue that just as Solow's 1957 seminal article represented a transfer of intellectual dominance from labour productivity to total factor productivity, a comparable shift in thinking is needed today to broaden the concept of productivity to include non-market resources.

They note that TRP requires choosing between competing production and welfare based paradigms as externalities and market failures cannot be ignored. They define TRP as the net growth in social output within the welfare function, a household-based production approach that they argue is wholly consistent with the evolution of productivity measurement over the past 40 years.

The authors develop their model theoretically and then apply it to the U.S. farm sector. Pollution is defined as total pesticide doses. Estimates of the marginal abatement cost of improving ground water quality by one dose as well as the marginal social value of a unit of clean water are developed for the TRP calculation. The results are intriguing. They find that growth in pollution produced slower TRP growth in the 1972-79 period relative to TFP growth while a decline in pollution between 1979 and 1995 meant that TRP growth exceeded TFP growth.

The authors recognize that a lack of consensus on the marginal social value of a unit of a cleaner environment means that statistical agencies cannot currently be expected to produce TRP estimate based on a particular value for the marginal social value. But they feel one option would be for agencies to produce TRP estimates based on a range of values, at least until a consensus emerges on the most appropriate value.

A second paper in the volume by Eldon Ball, Rolf Fare, Shawna Grosskopf, and Richard Nehring also addresses the issue of the integration of undesirable outputs into a production model and applies the model to the U.S. agricultural sector. They use activity analysis to construct a nonparametric representation of technology that allows the identification of the production frontier and deviations from and shifts in the frontier. Their approach requires no information on input and output prices and shares, a useful property where undesirable outputs such as pollution are unmarketed. They construct an index that allows for the contraction of undesirable outputs and the expansion of "goods" or desirable outputs. Firms are credited for the reduction in undesirable outputs as well as for increases in goods outputs. The authors, not surprisingly, find that measured productivity differs when undesirable output is accounted for. In their preferred model based on the Malmquist-Luenberger index, they report higher productivity growth for U.S. states with declining trends in water contamination resulting from the use of pesticides and chemical fertilizers.

What We Do Not Know About Productivity

In the final paper in the volume, Zvi Griliches, who was a leader in the productivity research field and director of the NBER productivity and technical change program, addresses the issue of what we need to know about the sources of productivity growth. This article was one of his last before he passed away in 1999.

Griliches argues that the current productivity accounting framework is incomplete as a number of productivity-enhancing activities use resources and improve human capital, but are not included in national output. Two variables that are missing are health investments and specific training. The current manner in which R&D is integrated into the framework can also be improved upon.

Griliches stresses that longer term productivity growth comes from the discovery of new resources, new ways of doing things, and the exploitation of investment opportunities that such discoveries create. But new knowledge does not arrive in a steady stream. Rather it is created in a poorly understood process with often clustered outcomes which create new opportunities for investment and allow the economy to approach a new equilibrium growth rate. The long-term productivity growth rate comes from a collection of such traverses, lurching from one equilibrium to another. Griliches recommends that more attention be given to the study of how the economy makes such traverses.

Recognizing that the post-1973 productivity slowdown has never been adequately explained, Griliches suggests that insight might be gained by turning our searchlights from recent data to data from the 1950s and 1960s. Maybe we were overestimating productivity growth at that time. What did the BLS know then that it has forgotten?

In addition to the three above points, Griliches identifies a number of measures that could be taken to improve productivity statistics. These include: better measures of construction output; a census of capital equipment; more data on actual hours worked by people and machines; improved information on the length of life of the capital stock and economic depreciation patterns; and better measurement of the production and distribution of information and new knowledge.

As noted earlier, a number of productivity knowledge gaps and challenges have been identified earlier in the papers. They include: the reduction of the range of estimates for the impact of the business cycle on productivity; the forging of a compelling link between the micro and macro levels, which requires overcoming the barrier that micro-level research is generally based on imperfect competition while macro-level research assumes perfect competition; and the development of a more convincing vintage capital model of productivity growth.

One word that does not appear in the volume is "new economy." This may reflect the fact that the CRIW conference for which the papers in the volume were originally prepared took place in March 1998. The U.S. economy has changed considerably since then and these changes have had a major effect of the questions productivity researchers ask. The key development has been the acceleration of productivity growth after 1995, a phenomenon only recognized in late 1998 and 1999. This development has obviously raised many questions, including the sources of the acceleration, its sustainability in the United States and its transferability to other countries. None of the papers in this volume directly address these issues. But the large number of widely available studies recently published on these topics means that this omission is not a critical weakness of the volume.

Conclusion

The papers in this volume present strong evidence that progress is being made in productivity measurement and analysis. In my view, the best examples of this progress are the paper by Frank Gollop and Gregory Swinard on the development of total resource productivity to account for changing environmental quality and the paper by Martin N. Baily and Eric Zitzewitz providing insight into international productivity differences in service industries through case studies and expert interviews. Other examples of what we have learned include a much greater understanding of the dynamics of productivity growth at the firm-level as demonstrated in the paper by Lucia Foster, John Haltimanger and C. J. Krizan, thanks to the development of firm-level micro-data bases, and the development of the methodologies for international TFP comparisons, as shown in the paper by Nazrul Islam as well as the paper by Dale Jorgenson and Eric Yip.

Like past NBER productivity studies, this volume represents an important contribution to our knowledge base on productivity. The papers are in general of high quality and some of them are outstanding. The editors are to be congratulated for putting together such a first rate volume. I do hope, however, that we will not have to wait another 21 years for the next NBER volume on new developments in productivity analysis. A 5-7 year cycle for such a "new developments" conference and volume would much better serve the productivity research community.

Notes

1 The Volume was Number 63 in the Studies in Income and Wealth series of the National Bureau of Economic Research and published by the University of Chicago Press ($80 US hard cover).

2 Richard G. Lipsey has also strongly emphasized this point. See Lipsey and Carlaw (2000).

3 Of course, even basic labour productivity estimates require real output measures, whose estimation involve a number of complex methodological issues (index number issues, quality adjustment methods, etc.). But at least official constant price GDP estimates provide a benchmark for real output. Most productivity researchers do not adjust these estimates in their work.

4 Solow also observes that he is increasingly uncomfortable with estimating technical change and the aggregate production function in a general equilibrium model and consequently with the use of observed factor shares as estimates of output elasticities.

5 The methodology allows for imperfect competition as well as variations in the work week of capital and labour effort. The production function residual is taken as a measure of sectoral technology shocks. Increasing returns to scale and markups of price over marginal cost are assumed. It is recognized that the marginal product of an input may differ across uses because of variations in market power, with the result that aggregate productivity growth depends in part on which sectors change inputs.

6 See, for example, the volumes edited by Griliches (1992) and Diewert, Nakamura, and Sharpe (1999).

References

Council of Economic Advisors (2001) *Economic Report of the President*, January (Washington, D.C.: U.S. Government Printing Office).

Diewert, Erwin, Alice Nakamura, and Andrew Sharpe, eds. (1999) *Service Sector Productivity and the Productivity Paradox*, *Special Issue of the Canadian Journal of Economics*, Volume 32, No. 2, April.

Gordon, Robert J. (2000) "Does the New Economy Measure up to the Great Inventions of the Past?" *Journal of Economic Perspectives*, 14 (4), Fall, pp. 49-74.

Griliches, Zvi, ed. (1992) *Output Measurement in the Service Sectors*, Studies in Income and Wealth Volume 56, National Bureau of Economic Research (Chicago: University of Chicago Press).

Kendrick, John W. and Beatrice Vaccara, eds. (1980) *New Developments in Productivity Measurement and Analysis*, Studies in Income and Wealth Volume 44, National Bureau of Economic Research (Chicago: University of Chicago Press).

Lipsey, Richard G, and Kenneth Carlaw (2000) "What Does Total Factor Productivity Measure?" *International Productivity Monitor*, Number One, Fall, pp. 31-40.

Solow, Robert (1957) "Technical Change and the Aggregate Production Function," *Review of Economics and Statistics* 39 (3) pp. 312-20.

Solow, Robert (1960) "Investment and Technical Progress," in Kenneth Arrow, Samuel Karlin, and Patrick Suppes (eds.) *Mathematical Models in the Social Sciences*, pp. 89-104 (Stanford, California: Stanford University Press).