Economic Commentaries

SVERIGES RIKSBANK

NO. 1, 2009

Understanding Productivity Growth

Björn Andersson, Susanto Basu, Mikael Carlsson and John Fernald

Björn Andersson is at the Applied Research Division of Sveriges Riksbank, Susanto Basu is at Boston College and is affiliated with the National Bureau of Economic Research, Mikael Carlsson is at the Research Division of Sveriges Riksbank, and John Fernald is at the Federal Reserve Bank of San Francisco.

At the end of 2007, the Riksbank arranged a research workshop on productivity growth in specific firms and industries. This commentary sums up the discussions and conclusions from the workshop and provides outlines of the papers that were presented. One of the conclusions of the workshop was that productivity data for firms or industries may be very useful for forecasting new waves of productivity growth as well as the persistence of productivity trends.

Productivity growth due to technological innovations makes it possible for firms to produce more goods and services with the same amount of input factors such as capital and labor. Higher productivity growth therefore allows an economy to grow faster without increased cost pressure and, ultimately, inflationary pressure. As a consequence, monetary policy can be more expansionary if productivity is expected to increase rapidly than if productivity growth will be low. The rate of productivity growth is therefore an important factor when central banks decide on the appropriate stance of monetary policy, and every central bank is intensely interested in understanding and predicting productivity growth.

On November 30th and December 1st, 2007, the Riksbank arranged a workshop where researchers from universities and central banks together with experts from statistical agencies discussed current research on productivity growth. The main focus of the workshop was on total factor productivity (TFP) at the disaggregate level (that is, in specific firms or industries) and its connection to information and communication technology (ICT) and human, organizational and other forms of intangible capital.[1] The workshop ended with a panel discussion on what was viewed as the most important challenges for measuring and learning about productivity growth. This commentary sums up the discussions and conclusions from the workshop and provides outlines of the papers that were presented.

The conclusions of the workshop had both promising and cautionary elements. In a variety of ways, many of the papers presented made the cautionary point that disaggregated data are unlikely to be more useful than aggregate macroeconomic data in revealing the current state of productivity growth in a country. First, disaggregated data usually become available later than do economy-wide statistics. Second, disaggregated data sets usually do not cover the production of the entire economy, so extrapolating from disaggregated to economy-wide statistics typically requires some assumptions about aggregation. Thus, contemporaneous analysis that is immediately useful for monetary policy will probably continue to be based on aggregate productivity data.

The conference suggested, however, that disaggregated and cross-country data may be very useful for forecasting new waves of productivity growth and for predicting their persistence. For example, the United States experienced a large increase in productivity growth rates over the decade from 1995 to 2005, relative to its experience over the previous 20 years. So did Sweden. Since the developed countries appear to share the same long-run productivity trends, it would be surprising if major

[1] Total factor productivity is the change in output (production) not caused by changes in inputs (capital, labor, energy, materials and service inputs). It reflects the joint effects of factors such as new technologies or improvements in the organization of production.

European countries did not experience the same productivity boom at some point in time. Firm-level and cross-country analysis of the sort discussed at the conference help to identify the countries that are likely to experience U.S.-style productivity growth rates in the near future.

Disaggregated analysis can also help policy-makers forecast the persistence of productivity trends, which is useful for monetary and fiscal policy. If sectoral analysis shows that many industries are being affected by a single technological innovation, it is more likely that the innovation is a "General-Purpose Technology", such as the invention of electricity, automobiles and computers. As shown by economic historians, such innovations have transformational effects that can be slow in coming, but often last for generations once they arrive. A combination of statistical investigation and case studies at the level of firms and industries can be very helpful in identifying a new General-Purpose Technology in its early years, and in alerting policy makers to the promises and the transitional costs of such wide-spread transformations. Unsurprisingly, the workshop participants agreed that we have much yet to learn about the sources and effects of productivity growth. But they also agreed that we have learned a great deal, especially by using firm-level data sets that have become available in recent years. While a single paper or finding rarely has a major impact on the conduct of policy, over the longer term the introduction of new ideas, buttressed by robust statistical findings, has often profoundly changed the ways in which policy makers approach their tasks. The field of productivity analysis is no exception.

Papers presented at the workshop

Sector Specific Technical Change

Susanto Basu, John Fernald, Jonas Fisher and Miles Kimball

John Fernald (from the Federal Reserve Bank of San Francisco) presented a unified framework linking the recent (and more traditional) literature on growth accounting with the recent literature on the importance of variations in the economy's ability to produce different final goods.[3] Economic theory tells us that the composition of technological change, in terms of which final goods production is affected, matters for the dynamics of the economy's response to technology change. In fact, in a model with two sectors, one sector producing investment goods and the other consumption goods, all of the interesting business cycle movements stem from variations in the productivity of the investment-goods producing sector.

When measuring technological change in the production of investment goods versus consumption goods, the recent macroeconomic literature has used a "top-down" approach in which the relative price of consumption to investment goods is used as a proxy.[4] In this paper the authors instead propose a "bottom-up" methodology to measure final-use sector technologies by aggregating industry estimates of technological change. The main empirical results suggest that measures based on relative prices may overestimate the relative increase in investment-specific technical changes. Moreover, as predicted by theory, the response of the labor input (hours worked) seems to be mainly driven by productivity changes in the investment-goods producing sectors as opposed to the consumption-goods producing sectors. However, in the short run labor input falls when productivity increases. This response is not in line with the simple neoclassical two-sector model. All in all, this paper provides an important step forward in thinking about technological progress on the sectoral production level, the link to final-use sectors and macroeconomic outcomes.

2. The objective of growth accounting is to decompose the increase of production in a company, industry or country into contributions from different factors of production and from productivity.
3. In special cases, the relative-price measures are appropriate; but in general, they are not.

Regulation, Competition and Productivity Convergence

Paul Conway, Donato de Rosa, Giuseppe Nicoletti and Faye Steiner

Paul Conway (from the OECD) turned the attention to the role of differences in anti-competitive product market regulation in explaining cross-country differences in productivity. In particular, this study investigates the link between regulations and the international diffusion of productivity changes using a statistical model that relates productivity growth to improvements in the global technology frontier and the speed of the process in which countries or sectors "catch up" with countries or sectors with the best technology.

The results suggest that regulations do slow down the adjustment process towards best practice. This, in turn, gives scope for remaining cross-country differences in regulations to explain the recent observed divergence of productivity in OECD countries. That is, since regulated economies adjust more slowly than deregulated ones, an improvement in the best practice, like the emergence of new General-Purpose Technologies in the 1990s, would cause productivity differences across more or less regulated countries. In fact, the estimates of the model suggest that the productivity gains from further reform may be considerable, especially in countries that operate at some distance from the technology leader. The authors also look directly at two channels with which regulations may affect international technology diffusion: the adoption of ICT and the location decisions of multi-national enterprises. In both instances the effect of regulation is found to be negative.

Long-term Productivity Growth in Canada and the United States 1961 to 2006

John Baldwin and Wulong Gu

Wulong Gu presented work with John Baldwin (both from Statistics Canada) on long-term productivity growth in Canada. Baldwin and Gu describe the new Canadian industry-level productivity accounts and identify broad trends over the period from 1961 to 2006. The authors also decompose output and labor productivity growth and draw a comparison to the United States. Over the period, labor productivity in the two countries grew at about the same rate, but Canadian growth exceeded that of the United States up to the early 1980s. Since then U.S. labor productivity growth has exceeded that of Canada. The gap has widened particularly after 2000. Moreover, the sources of labor productivity growth in the two countries differ. Investment and improvements in skills are more important sources in Canada, whereas TFP growth was larger in the U.S. Overall, the lower relative labor productivity growth in Canada in recent years seems to be a result of smaller increases in capital intensity and slower TFP growth.

Americans Do I.T. Better: US Multinationals and the Productivity Miracle

by Nick Bloom, Raffaella Sadun and John Van Reenen

John Van Reenen (from the London School of Economics) presented a paper studying the reason for the recent productivity divergence between Europe and the U.S. Two broad classes of explanations have been raised in the literature. The first line of reasoning is that the U.S. productivity miracle is due to some natural advantage for firms of being located in the U.S., for example the geographic, business or demographic environment (such as more space or younger workers). The second idea

is that U.S. firms have better organizational or management practices. This paper provides evidence of the latter explanation by looking at the relative productivity performance of U.S. multinationals relative to non-U.S. multinationals operating in the U.K.

The idea is that since a U.S. multinational located in the U.K. lacks any natural advantage of operating in the U.S., one should not see any differences between these firms and other firms operating in the U.K. if the first explanation were correct. However, the results suggest that U.S. multinationals do have higher productivity than non-U.S. multinational firms. Moreover, this is mainly due to the higher productivity of their ICT capital. Also, the results indicate that firms taken over by U.S. multinationals increase their productivity, whereas firms with the same characteristics taken over by non-U.S. multinationals do not.

In order to strengthen the case that it is actually different management and organizational structure of U.S. firms that can explain these results, the authors presented evidence drawn from a new cross-country survey on management and organization practices matched with data at the firm level. The preliminary results suggest that U.S. firms are differently organized than non-U.S. firms both at home and abroad, and that the higher U.S. ICT productivity can be explained by these differences.

Intangible Assets in France and Germany

Janet Hao, Vlad Manole and Bart van Ark

In his talk, Vlad Manole (from The Conference Board) argued that although many statistical agencies have started to incorporate intangible investment in the official statistics, several important areas are still left out, such as R&D, innovations to organizational structure, and branding and marketing. Lacking this information it is impossible to evaluate the role of this type of investments for productivity measures and growth.

This paper extends the body of international evidence by measuring intangible assets for France and Germany in three broad categories: Computerized information, innovative property and economic competencies. While remembering that measuring intangibles is often complicated, the results indicate that intangible investment do correspond to a sizable part of economic activity. The authors find that in 2004, Germany spent about 7.5 per cent of GDP on intangibles and France about 8.5 per cent. With these results in mind, the road ahead is to refine the measurement of intangibles and try to uncover the subtleties of the interaction between, on one side, tangibles and intangibles and, on the other, innovation and productivity growth.

The Value of Risk: Measuring the Services of U.S. Commercial Banks

Susanto Basu, Robert Inklaar and Christina Wang

Susanto Basu (from Boston College) reported on a research program focused on measuring output of the banking sector (in both nominal and real terms). Measuring the financial services provided by commercial banks is difficult because many of the services do not have explicit prices, but are implicitly priced. Basu argued that in order to arrive at the conceptually correct value of these services, current National Accounts practice needs to be modified to account for the risky nature of most bank loans. When accounting for the risk premium (using a conservative estimate) the authors show that nominal bank output in the U.S. National Income and Product Accounts

was overestimated on average by 21 per cent between 1987 and 2003. Moreover, since the risk premium is time varying it affects the measurement of bank output growth as well. For real banking output, Basu pointed out that an activities count based on various categories of transactions (as used by the Bureau of Labor Statistics) is always a correct measure. However, instead of the current practice of using employment shares as weights for aggregation, nominal output shares should be used to estimate total real bank output in accordance with economic theory.

Endogenous Skill Bias in Technology Adoption: City-Level Evidence from the IT Revolution

Paul Beaudry, Mark Doms and Ethan Lewis

Mark Doms (from the Federal Reserve Bank of San Francisco) discussed the two-way interaction between the adoption of ICT and labor market conditions – in particular, the supply of relatively skilled labor. One motivation is the cross-country studies that find large differences in ICT adoption across countries; understanding those differences is challenging, however, because so many potential explanatory factors differ across countries. But Doms noted that if you look across cities within a country – thereby holding many of the background factors fixed – you also see large differences in ICT adoption.

The authors show that the education of the workforce is the factor that is most strongly related to personal computer (PC) intensity; they also provide a model of endogenous technology adoption to explain these results. In the model, when the new technology (in particular, PCs) comes along, cities with relatively abundant (and, thus, relatively cheap) skilled labor adopt PCs more aggressively. As a result, returns to skill increase most in those cities. These (and other) theoretical predictions for adoption rates and relative wages are supported by the cross-city data. The paper thus builds a bridge between two different strands of the literature: one explaining how the relative supply of skilled labor affects technology adoption, and another studying how technological change influences the relative demand for skilled labor.

Cross-Border Flows of People, Technology Diffusion and Aggregate Productivity

Thomas Barnebeck Andersen and Carl-Johan Dalgaard

Carl-Johan Dalgaard (from the University of Copenhagen) highlighted the vast differences in TFP across countries, which, in turn, provides a proximate explanation for the vast differences in GDP per capita across countries. An important part of the differences in TFP presumably reflects true technological differences, raising the question of why countries differ in their ability to adopt new innovations. Andersen and Dalgaard argue that, empirically, cross-border flows of people are centrally important in explaining flows of ideas and knowledge. The idea is that societies that are more exposed to foreign influences, as measured by inflows and outflows of travelers, can better obtain technologies and organizational strategies from abroad. In cross-country data, Andersen and Dalgaard find that, even after taking account of the fact that travel intensity may depend on a country's TFP, an increase in travel intensity substantially raises the level of TFP in a country. In particular, the country with the highest travel intensity has a level of TFP that is more than twice the level in the least integrated country. Indeed, the authors find that when travel intensity is taken into account, other variables that measures trade – which are commonly found to be important in explaining cross-country differences in TFP – become insignificant. This raises questions about a large body of existing work.

Cross Country Differences in Productivity: The Role of Allocative Efficiency

Eric Bartelsman, John Haltiwanger and Stefano Scarpetta

Eric Bartelsman (from VU University Amsterdam) focused on productivity dynamics at the firm level. He began by summarizing empirical evidence from a number of countries that shows that there is wide and persistent dispersion in productivity levels across firms also within very narrowly defined industries. Moreover, more productive businesses have a larger market share, and over time, outputs and inputs within sectors shift from less productive to more productive businesses. In order to explain these results, the authors seek to model and quantify the frictions that allow reallocations to affect productivity. In particular, they consider a model where frictions and distortions can explain why different firms operate with different productivity levels.

The authors also explore whether calibrated versions of their model can generate outcomes that match the patterns of resource allocation and productivity observed across countries and sectors and over time. They find the results are subtle. Simple measures of allocative efficiency capture important aspects of the model economy. However, the simple measures do not capture other implications of the distortions. For example, some firms never enter the market because of the distortions, and so never show up in any of the statistics. The theory and evidence help advance our understanding of what underlies the black box of aggregate productivity. In addition, having a well-founded model is necessary for analyzing the impact of different policy measures in response to the productivity gaps we see across countries.

The Effects of Organizational Change on Firm Productivity

Christina Håkanson

Christina Håkanson (from the Institute for International Economic Studies at Stockholm University) focused on what happens within firms when they invest in ICT. A considerable literature suggests that when firms invest in ICT, they also invest in complementary organizational changes. Håkanson provides empirical support for this view, using a unique Swedish panel dataset that combines firm-level survey indicators of organizational changes with data on TFP and ICT investments. She looks at the effects on productivity before and after the organizational changes take place. Her preliminary results suggest that productivity growth depends in part on the interaction of ICT investments with organizational change. The results are consistent with the view that, to benefit from ICT, firms need to change how they organize their operations and tasks.

Summary of panel discussion

The workshop was rounded up by a panel discussion on the important challenges for measuring and learning about productivity growth. In the opening statement, Deputy Governor Svante Öberg (from Sveriges Riksbank) pointed to the importance of measuring and forecasting productivity growth precisely for the conduct of monetary policy. Öberg also discussed the considerable difficulties involved and the value of learning more about the structural forces at play in order to increase the precision of productivity measurements and forecasts. Wulong Gu (from Statistics Canada) then continued with presenting the work being done in the Productivity Program of Statistics Canada. The major issues ahead for this project are first to improve the

measurement in the service sector and secondly to gauge the role of intangible and infrastructure capital in productivity growth.

Joaquim Oliveira Martins (from the OECD) discussed the data work that is being done at the OECD, for example, maintaining the STAN and EU-KLEMS databases.[4] He also discussed the need to coordinate efforts and the creation of an OECD micro-data centre. Oliveira Martins continued by presenting the work being done at the OECD on the link between ICT and innovations as well as on measurement issues related to the health and education sector. Hans-Olof Hagén (from Statistics Sweden) rounded off the panel session by discussing the major data imperfections related to input and output measurements (which are crucial for measuring productivity) as well as the studies being carried out at Statistics Sweden on both aggregate and disaggregate data to learn more about the drivers of productivity growth.

References

Baldwin, John and Wulong Gu (2007), "Long-term productivity growth in Canada and the United States 1961 to 2006", Canadian Productivity Review No. 013.

Barnebeck Andersen, Thomas and Carl-Johan Dalgaard (2007), "Cross-border flows of people, technology diffusion and aggregate productivity", unpublished manuscript, University of Copenhagen.

Bartelsman, Eric, John Haltiwanger and Stefano Scarpetta (2007), "Cross country differences in productivity: The role of allocative efficiency", unpublished manuscript, VU University Amsterdam.

Basu, Susanto, John Fernald, Jonas Fisher and Miles Kimball (2007), "Sector specific technical change", unpublished manuscript, Federal Reserve Bank of San Francisco.

Basu, Susanto, Robert Inklaar and Christina Wang (2006), "The value of risk: Measuring the services of U.S. commercial banks", unpublished manuscript, Boston College.

Beaudry, Paul, Mark Doms and Ethan Lewis (2006), "Endogenous skill bias in technology adoption: City-level evidence from the IT revolution", unpublished manuscript, Federal Reserve Bank of San Francisco.

Bloom, Nick, Raffaella Sadun and John Van Reenen (2007), "Americans do I.T. better: US multinationals and the productivity miracle", Discussion Paper No 788, Centre for Economic Performance, London School of Economics and Political Science.

Conway, Paul, Donato de Rosa, Giuseppe Nicoletti and Faye Steiner (2006), "Regulation, competition and productivity convergence", Working Paper No. 509, OECD Economics Department.

Håkanson, Christina (2007), "The effects of organizational change on firm productivity", unpublished manuscript, Stockholm University.

Hao Janet, Vlad Manole and Bart van Ark (2007), "Intangible assets in France and Germany", unpublished manuscript, The Conference Board.

4. OECD's STAN (STructural ANalysis) Database provides output, input and trade data by economic activity, which enables detailed analysis of industrial performance across countries. EU-KLEMS is a project funded by the European Commission. It is a database on productivity by industry for EU member states with a breakdown into contributions from capital (K), labor (L), energy (E), materials (M) and service (S) inputs.