1 Introduction

DAVID G. MAYES

It seems only a few years ago that the focus of concern over the subject of productivity lay on the slowdown which affected much of OECD after the first oil shock in 1973. Worries were expressed that the slowdown might be enduring and there was extensive debate over the cause (see, for example, Wenban-Smith, 1981). The 1980s have seen a turnround in this performance with productivity accelerating again. In the case of the UK the extent of the acceleration has been so great that it has been called the Thatcher miracle (Muellbauer, 1986; Oulton, 1987), reflecting the major policy change which took place during the period. To some extent this acceleration is merely a catching up of the progress foregone during the 1970s and represents a return to the longer-run rates of productivity growth observed in the 1950s and 1960s. However, hypotheses abound to explain behaviour in the 1980s. In this book we focus on one group of such analyses, those which draw on detailed data on individual businesses across the whole of industry.

The book is the culmination of nearly two decades of research into productivity using data drawn from the Annual Censuses of Production in the UK and other major OECD countries. The research has involved international collaboration throughout and the current book, although it emphasises the experience of the UK, contains work on the United States, Canada, Japan, Australia, Belgium, Norway, Sweden and Mexico. The genesis of the work has varied across the participating countries.In the case of the UK it began with work at the National Economic Development Office in 1976, undertaken as part of the Industrial Strategy being developed by the government of the day. NEDO was concerned to identify industries where there was potential for productivity improvement, which could be addressed by its working groups. In particular it wanted to look for

industries where there was a long 'tail' of less productive firms. In those cases it was suggested that it would be possible to assist firms in attaining the best practice for that industry, by a combination of information and analysis of the ways in which such transformations could be achieved.

The UK analysis typified that in the other countries in that it required access to the confidential data obtained from the Annual Censuses of Production. This was obtained with the assistance of what was then the Business Statistics Office (now the Central Statistical Office) in Newport. Information relating to individual businesses cannot be released because of the terms of the Statistics of Trade Act which guarantee confidentiality so the results were presented as parameters of the distribution of productivity in each of over 150 industries within manufacturing. From the early days, the work benefited from collaboration with Professor Richard Caves at Harvard, who was instrumental in bringing together many of the foreign collaborators.

The current joint work has been made possible by research funding in the UK by the Economic and Social Research Council who provided a grant to Mayes and Lansbury, by a grant from the Fulbright Commission to set up a symposium to discuss the results at the National Institute of Economic and Social Research in London in February, 1994, and by support from the British Academy for the travel costs of the Australian and Japanese participants. This book is derived from the papers discussed at that symposium.

What is productivity?

This book is largely about labour productivity – output per unit of labour input. Just looking at labour productivity presents an easy opportunity for distortion as output will vary according to the amount of capital available. A firm with more up to date equipment will be able to produce more output per person. In the same way a firm that works two or three shifts will be able to produce more per unit of capital than one which works a single shift. However, it is labour productivity which is the main policy variable which governments concentrate on as it is this which determines incomes per head. Nevertheless the work of O'Mahony in Chapter 8 and Oulton (1994) considers total factor productivity as well, exploring the extent to which the results for labour alone can be rather misleading.

For many comparative purposes it is therefore necessary to try to eliminate the various factors which might cause labour productivity to vary other than the activities of labour itself in the production process. Since output

involves the transformation of inputs using labour and capital, the simplest element required for a helpful comparison is to ensure that inputs, whether of materials or capital in the form of buildings, equipment and machinery are fully compatible. This similarity also needs to extend to the product and to the techniques used to produce it.

This is a very tall order as with outside commodities a large part of the ability to compete comes from differentiating the product. It is thus difficult to make highly detailed comparisons across firms. Indeed one might feel that the only way to make such comparisons was to pick different establishments from within the same organisation as in the comparisons of the production of photographic paper and water heaters in Chapter 9 by Chris Harris and of bank branches by Henry Tulkens and Amador Malnero in Chapter 10. Here, however, the problem is different – if we look at comparisons within a single firm many of the sources of variation which are most important are omitted. Management styles and structures will tend to be the same. Training programmes will tend to be common to all plants within the group. It is just these differences which are most important in between-company studies as these variables are subject to choice by the firm.

It is for this reason that Mason and van Ark in Chapter 6, while considering just a small number of firms in detail, pick for their comparison, not plants that are closely matched (along the lines explored in Mayes, 1983, for example) but firms that are each typical of the industry in the country from which they come. Thus in the case of the comparison of biscuit-making which is considered in detail in the chapter, the type of biscuit commonly produced in the UK is different from that commonly produced in Germany or the Netherlands. In the UK the typical biscuit is a relatively straightforward product, at most involving a chocolate coating. In the continental case a much more complex product is involved, often entailing multi-stage baking. It is thus not surprising that the latter should be more labour intensive.

Mason and van Ark therefore approach the comparison differently by looking at common and identifiable activities in the production process, such as maintenance, reprogramming machinery, supervision and so on. Clearly it may still be difficult to get exact comparisons as more complex machines are often more difficult to maintain or operate, particularly in the engineering firms they compare. (This is by no means a universal finding as newer machines attempt to have more user friendly programming structures. They also incorporate more self-checking and adjustment, requiring less operator intervention. As a result the craft skills involved in production may actually be reduced, despite the ability to produce more

sophisticated products. Indeed the greatest skills may come from the successful operation of traditional wood fired ovens in baking.)

There is a further and more pertinent reason for focusing on labour productivity which stems from the accumulated result of the programme of research undertaken at the National Institute of Economic and Social Research in London (clearly summarised in Prais *et al.*, 1970). Although in some instances it is possible to explain the UK's lower productivity levels compared with its competitors in Europe and North America in terms of material inputs, the age of machinery and the nature of the products and lengths of production runs employed, there is still a pervasive shortfall in productivity. Although it is difficult to pin down the exact causes of this shortfall (for this reason this type of discrepancy has been labelled as X-inefficiency by Leibenstein, 1976) most of the factors which have been identified are alterable and hence this is of great policy significance for those who wish to see productivity and competitiveness improved, with its consequent benefits for increasing real incomes and reducing unemployment.

These factors include the skills of the labour force and the nature of its training, production organisation, human resource development programmes and the whole range of management operations. Mason and van Ark show that not only is the stock of skills in the UK low compared with the Netherlands and Germany, particularly at the intermediate level, but that this lack of skills has a clear impact on efficiency and the ability to undertake various tasks. In particular it reduces the flexibility of the labour force. This results in interruptions to the smooth flow of work, the need to employ larger numbers of staff in order to get the necessary specialisms (while more qualified employees could do all of the tasks themselves) and a lower ability to innovate and bring through new ideas and processes into production. In the past the lack of flexibility in the labour force was in part due to deliberate rigidities agreed with trade unions. (While such agreements may originally have been a means of ensuring the quality of those to be employed, many had become outdated.)

Determinants of productivity

The level of productivity which prevails is largely the result of a combination of choices made by firms and the efforts of those that work in them. The primary determinant of productivity lies in product and process innovation including the skills of the labour force as well as the technology embodied in machinery. In many instances it is possible to map out what is achievable with a given range of inputs using a particular technology.

While not so rigorous a concept as we observe in the operation of specific processes, such as electricity generation, it is nevertheless possible to map out what is feasible. With the important role played by human intervention in the production process, particularly in service industries the scope for departure from the feasible standard is very considerable.

The chapters by Melanie Lansbury and David Mayes and by Henry Tulkens and Amador Malnero illustrate two approaches to this and Chris Harris describes some others. The Tulkens and Malnero case takes a large number of bank branches and maps out the relationship between their inputs and outputs. Since bank branches have both multiple inputs and multiple outputs it is necessary to express both inputs and outputs as indices after weighting the components. This then enables us to graph actual performance and identify those establishments with the highest productivity (figure 1.1). It is thus possible to see what can be achieved on the basis of the best which has been achieved. Lansbury and Mayes employ a rather different technique by suggesting that it is possible to estimate a production function for each industry they consider. This suggests that a specific and identifiable relationship exists linking inputs to outputs, which is common across firms in the industry that use the same technology.

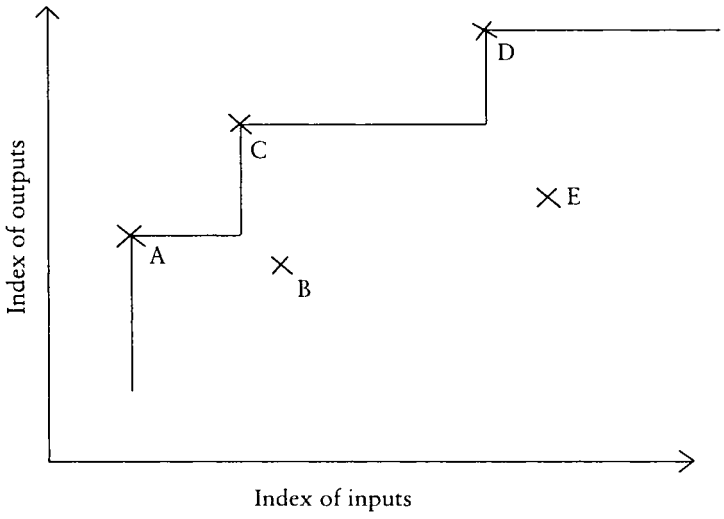

Figure 1.1 *Frontier of actual performance*

The continuous line shows the frontier of performance actually achieved. Points A, C and D lie on the frontier. C produces more output with less input than B and clearly out-performs it, similarly for D over E. No attempt is made to link A and C or C and D and suggest that intermediate points on that line should theoretically be attainable; only actual recorded performance is admitted.

Both these approaches have drawbacks. The Tulkens and Malnero method requires quite a large data set of establishments with similar sizes if the 'frontier' of best attainable productivity is not to be too ragged. The Lansbury and Mayes 'frontier production function' on the other hand imposes strong assumptions on the behaviour of the industry. So while it may produce a productivity frontier for a wide range of inputs, that frontier may be a very inexact representation of the data.

Technical, allocative and X-inefficiency

Actual productivity will depart from that theoretically achievable for a number of reasons but Farrell (1957) has distinguished neatly between suboptimality because the firm has chosen an inappropriate combination of labour and capital and inefficient use of the combination which has been chosen. Thus in the second case, even given the shortfall in productivity due to a poor combination of inputs it would be possible to have higher productivity if they had been used more efficiently. He labels the two sources of inefficiency 'allocative' and 'technical' inefficiency respectively. Lansbury and Mayes concern themselves primarily with the second area but Walfridson and Hjalmarsson distinguish both sources. Indeed they take the analysis further in two respects, first by allowing for the fact that the scale of operation will also have an effect on productivity and second by placing the whole analysis in a dynamic framework. Tybout and Westbrook (1994) also provides a three-way split, separating out scale effects from changes in share and a residual which itself comprises both the shift in the frontier and movements with respect to it.

Farrell's decomposition is not the only way of looking at departures from optimality. Leibenstein (1976) attempted a different definition of the extent of waste stemming from inefficient production by coining the term 'X-inefficiency' – that element in the system which is left when all the factors relating to the nature of outputs, inputs and the production technology have been taken into account. In Chapter 15 Torii explores a means of relating allocative, technical and X-inefficiency. These attempts to disentangle the different sources of discrepancies in productivity and more importantly their contribution to changes in measured productivity lie at the heart of the

analysis in this book. As explained in the next section, although it may be possible to separate out these factors in theory, it is much more difficult in practice and we normally either have to isolate only some factors or make quite strong assumptions in order to perform the decomposition.

Economies of scale

The consequences of scale are straightforward and well known, enabling a spreading of fixed costs over a larger amount of output. While typically this is applied to capital, it applies just as well to labour. In larger establishments the workforce can specialise and they can run larger batches without having to reset equipment. Such gains also occur across time as the workforce can learn how to operate new production techniques with increasing skill and hence can increase their productivity through learning. Such learning curves tend to become flatter the more experience develops but refinements in the process or the next innovation can make them important on a continuing basis.

This ability to spread costs across outputs also applies to spreading costs across related activities. We have already noted that multi-skilling enables employees to do a number of jobs and hence avoid idle time. Most firms also produce a range of products. In these circumstances production of one product benefits from production of the other. This may involve joint use of equipment or joint use of staff. Lessons learnt in one production process can be applied in another. A business operating without these interlinkages will tend to have lower productivity although it may be possible to obtain them externally, say by marketing agreements or by buying in components required in small numbers from others who by supplying a large range of firms can themselves reap the economies of scale. All these external economies will lead to variation in productivity among firms if they are not equally available to all of them.

Dynamic adjustment

While productivity, even with a fully efficient workforce, will vary with the nature of their inputs, their combination and their scale, these outcomes do not occur within a static framework. The firm is striving to achieve an optimum position not just in a changing world but in a world where its past actions govern the scope for actions in the present and future. Thus it can be a slow process to change fixed capital, to implement new techniques or even products. Although there may be fairly substantial scope for altering the size of the labour force it too has restrictions on the rate it can change, both through training and through learning by experience.

Walfridson and Hjalmarsson do not incorporate a fully dynamic framework but they go a long towards doing so.

Perhaps the simplest example of this adjustment problem is revealed by the behaviour of productivity over the economic cycle. When the economy is growing at its fastest, productivity also tends to grow rapidly as firms strive to meet that demand. Firms are also likely to be expanding capacity rapidly during the upturn, providing newer machinery that will help raise productivity even further. However, as growth slows those that have most recently expanded capacity find that they cannot utilise it fully. Because of the costs of hiring and firing, firms will only adjust their labour force downwards relatively slowly in the hope that the fluctuation will be short-lived. As a consequence their productivity will fall. When the costs rise too far, labour has to be cut back markedly, restoring productivity levels.

The behaviour of productivity over the cycle is one of the preoccupations of this book. During the 1980s in the UK productivity rose rapidly both in the sharp downturn at the beginning of the period and during the stronger growth of the second half of the decade. It was only with the recession at the end of the period that productivity growth began to slip back. It is this consistency in prolonged productivity growth which contributes to the Thatcher miracle. Part of the argument is that the severity of the 1979–82 recession was so great that harsh decisions were made to restructure companies, in many cases making up for failures to act over an extended period. What Lansbury and Mayes show in Chapter 2 is that there are important differences in the structure of the components of productivity growth between the periods of output decline and growth. In the recession important contributions to productivity growth in the economy as a whole came from the closure of the worst performing plants rather than from an increase in the share of growing plants with above average productivity, which characterised the recovery. There is also a considerable switch in the industries within manufacturing which were making the most important contributions to growth.

Decomposing productivity growth by category of business

The most common disaggregation of businesses is clearly by industry as this is the form in which the statistics are available. Similarly decompositions are often possible by size. Beyond that it is usually very difficult to explain the contribution of particular types of firm. Chapters 11 and 12 are an exception to this. The importance of different groups of establishments to the overall change in productivity is explored in more detail by Martin Neil Baily, Eric Bartelsman and James Haltiwanger for the USA

and John Baldwin for Canada. In those two instances, because they have full databases in which they can trace the performance of individual plants across time, they can see how plants of different types contribute to productivity growth. In the case of Canada, for example, Baldwin notes that it is the increasing importance of small firms that has helped reduce aggregate productivity growth as small firms tend to have below average productivity. The authors are also able to distinguish performance by ownership – both according to whether the business is foreign or domestic owned or whether it is independent or part of a larger group. (Takeovers and mergers themselves can alter productivity performance.)

Measuring productivity

As we have noted, it is one thing to be able to say that we can in theoretical terms decompose productivity into a series of component determinants. It is quite another to suggest that we can distinguish these influences in practice. To summarise, we need to be able to measure inputs: labour, capital and materials (in the widest sense of the word); outputs, whether of goods or services; and the technologies employed in a way that we can make comparisons across firms, industries, countries and over time – a short list but a formidable problem in almost all its aspects.

Starting with inputs, labour presents a straightforward series of problems. A measure of the amount of labour available would probably be best measured in person-hours rather than just number of people employed as the latter poses several difficulties:

- employees may work different numbers of hours in the day. A one hour reduction in the working week might appear to be a fall in productivity
- some employees will be full-time and others part-time. Even if treating them as single group can be avoided it will usually only be possible to multiply the part-time employees by a standard factor to compute their full-time equivalent. If there were no variations among the firms to be studied in how they use part-time employees then there would be no problem, but this is highly unlikely.

However, hours themselves are not without problems. It has been observed that the productivity of standard and overtime hours differs (Leslie and Wise, 1980). It appears that because overtime hours are better paid than standard hours there is a drop off in productivity in the last few standard hours to ensure that overtime is activated. There is thus no clear homogeneity in either hours or people. We are all aware that our productivity drops as we become tired from working long hours.

People themselves are not homogeneous, nor is their role in the production process. We might therefore wish to distinguish between tasks and skills as does Harris in Chapter 9 and Mason and van Ark in Chapter 6. In general, however, it has to be said the data available to us relating to the whole range of plants across industry do not permit the use of refined labour input variables.

Capital itself is a heterogeneous item and we would normally wish to distinguish between buildings, works, land, equipment and vehicles and would probably wish to separate out computer and IT equipment from the rest, as is the case in the national balance sheets. The problem in this case, however, lies with valuation. Capital depreciates. We may know what it cost to buy – historic cost – and in some cases we will know its value written down for depreciation purposes or even then revalued to replacement cost. However, this will tend to come from detailed studies and only be available for a limited number of enterprises. Aggregate official figures for industries in the UK draw on the figures for capital *investment* available each year and build up the capital *stock* figures from then, making assumptions about life lengths of assets and the distribution of scrapping round the average life (this is described in detail in Mayes and Young, 1994). The accuracy of the results is sensitive to the assumptions and it has been widely argued that the published UK capital stock figures may be erroneous by as much as a third (see the summary in *Ibid.*). However, more recent work at the National Institute in a project for the Central Statistical Office suggests that the extent of overestimation may be relatively small.

When it comes to individual businesses no such data exist on a coherent basis. Although the same investment figures have been recorded each year, the data are not held in the form of a business database where individual businesses can be identified in successive years. (The National Institute is now undertaking a joint project with the Central Statistical Office, with the assistance of the Economic and Social Research Council, which will establish just such a database. When this is complete it will then be possible to undertake longitudinal studies at the level of the individual establishment for the UK of the same sort that are undertaken in Chapters 11 to 14 for Canada, the USA, Sweden and Norway.) Unfortunately it is also difficult to make use of company accounts in generating the appropriate variables as a single company may include several different businesses or establishments spread over a variety of industries. Following Mayes and Young (1994) the Central Statistical Office is investigating whether more appropriate information can be obtained by direct estimation but even if implemented it will be several years before it becomes available.