Measuring total factor productivity for the United Kingdom

By Charlotta Groth, Maria Gutierrez-Domenech and Sylaja Srinivasan of the Bank's Structural Economic Analysis Division.(1)

A good understanding of productivity growth is important for understanding aggregate supply capacity, and so for the conduct of monetary policy. To understand the sources of supply capacity well, it is important to measure output and factor inputs correctly. This article summarises recent and ongoing research at the Bank of England on improved measures of factor inputs. This work explicitly accounts for changes in the quality of these inputs and for the flow of services available from them, as well as for the costs of adjusting the level and utilisation of the inputs over time. This research was presented at a workshop on 'measuring factor inputs' held at the Bank of England in December 2003.

Introduction

The aim of monetary policy is to keep inflation low and stable, in accordance with the target set by the Chancellor. A key influence on inflationary pressure is the balance between the demand for and the economy's capacity to supply goods and services. This capacity depends both on the quantities and qualities of the primary inputs into the production process—capital and labour—and on the efficiency with which they are combined. The latter concept is often referred to as total factor productivity (TFP). A good understanding of past and current productivity growth is thus important for understanding aggregate supply capacity, and so it is relevant for the conduct of monetary policy.

To understand the sources of supply capacity well, it is important to measure output and factor inputs, and therefore productivity, correctly. It is also crucial to recognise and adjust for the changing composition of the aggregate inputs, which may vary over time. This article discusses recent work at the Bank of England on improved measures of factor inputs, which accounts explicitly for changes in their quality and for the flow of services available from them, and for the costs of adjusting the level and utilisation of the inputs over time. These improved factor input estimates can then be used to obtain better measures of total factor productivity growth for the United Kingdom.

The Solow residual

The standard measure of total factor productivity growth is the Solow residual:[2] that part of output growth that cannot be accounted for by the growth of the primary factors of production, ie capital and labour.[3] The Solow residual (z) is calculated by subtracting the growth of the primary inputs (weighted by their respective shares in nominal output) from the growth of output:[4]

$$z = y - s_k k - s_l l \qquad \textbf{(1)}$$

where y is the growth rate of output, k is the growth rate of capital input, l is the growth rate of labour input and s_k and s_l are the shares of capital and labour in nominal output respectively.

Chart 1 shows a standard measure of the Solow residual for the United Kingdom.[5] The growth rate of TFP is calculated here using aggregate data, where the capital input is a capital stock measure and the labour input is total hours worked.[6] The growth rate appears to be

(1) We would like to thank John Fernald, Steve Nickell, Soledad Nuñez and Nick Oulton for valuable comments. We would also like to thank Pablo Burriel-Llombart and Jerry Jones for supplying us with the quality-adjusted labour series.

(2) Total factor productivity as defined in this article is also referred to in the literature as multi-factor productivity. (See the November 2003 *Bank of England Inflation Report* for some standard multi-factor productivity estimates for the United Kingdom.)

(3) For a 'biography' of total factor productivity, see Hulten (2001).

(4) This is a simplified version of the formula actually used in the empirical calculations. The full formula is given in the appendix in equation **(A.1)**.

(5) The sources for the data underlying the calculations are given in the appendix.

(6) A similar TFP growth measure, but using the number of people in employment as the labour input, is summarised in Table 3.A of the November 2003 *Bank of England Inflation Report*.

Chart 1
Growth of total factor productivity and GDP for the United Kingdom: 1980–2003

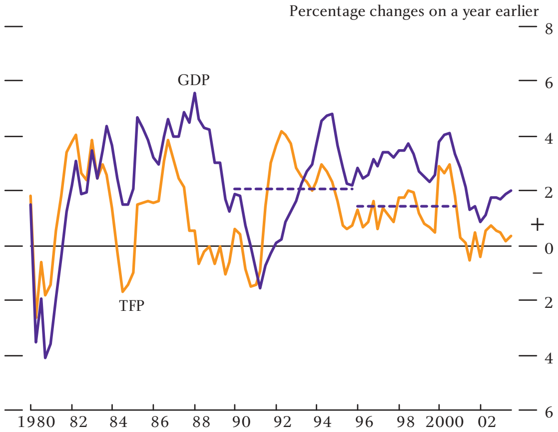

procyclical—it is positively correlated with GDP growth.(1) But over and above that, a slowing in the growth rate is noticeable in the second half of the 1990s (relative to the first half), in contrast to the United States, which experienced an increase in TFP growth in the late 1990s. Basu, Fernald, Oulton and Srinivasan (2003) discuss possible reasons for the differing productivity growth patterns in the United States and the United Kingdom.(2)

The Solow residual shown above provides us with just one estimate of total factor productivity growth in the United Kingdom. There are, however, a number of well-known measurement issues that need to be considered. First, capital and labour inputs need to be estimated correctly. For example, the capital measure should reflect the productive services available from the capital stock and needs to reflect factors such as the increased use of ICT capital; and the labour measure should reflect the changing composition and skills of the UK labour force. Second, because the movement of resources between industries also affects aggregate productivity, it is preferable to aggregate industry-level data rather than to use aggregated data directly.(3) Third, the basic Solow residual calculation in equation **(1)** assumes that the factors of production are flexible and fully employed. This may not be the case if there are costs involved in eg hiring and firing or in installing new machines and equipment (usually referred to collectively as adjustment costs). Also, if it is costly to adjust inputs, firms may respond to short-run fluctuations in demand by varying the rates at which their existing capital and labour are utilised. The remainder of this article summarises ongoing Bank of England research on each of these measurement issues and considers their impact on UK TFP growth.(4)

Measuring factor inputs

This section discusses measurement issues relating to the factor inputs used in the TFP calculations.

Capital services

The standard Solow residual is calculated as that part of output growth that cannot be accounted for by growth in capital and labour inputs. The measure of capital that is traditionally used is the stock of capital, which is a measure of economic wealth. As shown in the seminal work by Jorgenson and Griliches (1967), Jorgenson *et al* (1987) and Jorgenson and Stiroh (2000), what is in fact needed to measure productivity accurately is a measure of the flow of services that the capital stock generates. This issue was discussed in an earlier *Quarterly Bulletin* article (Oulton (2001)).

The main difference between a capital stock measure and a capital services measure is the way in which different assets are aggregated together. To create the aggregate stock of capital, different stocks of assets are weighted together by their asset (market) price weights.(5) In the capital services measure, on the other hand, different assets are weighted together by their rental price weights.(6) The rental price is the price that a user of the asset would have to pay to rent the asset for a period of time and, in a competitive market, it will reflect the value of the services which can be derived from the asset. The rental price is related to the price of the asset, but it also takes into account the opportunity cost of holding the asset, the cost of depreciation, and any capital gains or losses (including obsolescence) that are expected to be made by holding the asset over a period of time.

An important implication of using a services rather than a stock measure of capital input is that the services

(1) This is similar to the United States. See Burnside, Eichenbaum and Rebelo (1995) and Basu and Fernald (2000).
(2) They argue that unmeasured investments in intangible organisational capital—associated with the role of ICT as a 'general-purpose technology'—can explain the divergent US and UK productivity performance after 1995.
(3) See Stiroh (2002) and Bosworth and Triplett (2003) for an explanation of these effects.
(4) The focus of this article is on total factor productivity. Clearly, a corresponding labour productivity measure can be calculated.
(5) The asset price weight for each asset is calculated by multiplying the asset price by the asset stock and expressing it as a proportion of aggregate nominal wealth.
(6) The rental price weight for each asset is calculated by multiplying the Hall-Jorgenson user cost of capital for the asset by the asset stock and expressing it as a proportion of aggregate nominal profits.

measure will give more weight to assets for which the rental price is high in relation to the asset price. If the stocks of such assets are also growing more rapidly than those of other types of assets, the services measure of aggregate capital will grow more rapidly than the stock measure of aggregate capital. In recent years ICT assets have precisely had these characteristics: the growth rates of ICT assets have been high compared with those of non-ICT assets and their rental prices are also high in relation to their asset prices.(1) Altogether, this means that the flow of services from capital has recently been growing faster than the stock of capital.(2)

Chart 2 plots the growth rates of a services measure of capital (that accounts separately for ICT assets), against a stock measure of capital (based only on traditional ONS asset classifications: other buildings and structures, transport equipment, other machinery and equipment, intangible fixed assets) for the United Kingdom.(3) The growth of the capital services measure has been much higher than that of the capital stock measure over much of the past five years. This suggests that the Solow residual estimate in Chart 1 (which is based on a capital stock measure) may overestimate underlying total factor productivity growth over that period.

Chart 2
Growth of capital in the United Kingdom: 1980–2003

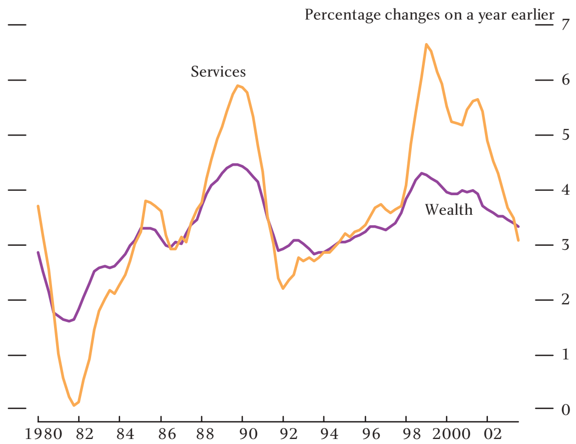

Quality adjustment of labour input(4)

In order to generate more accurate measures of TFP and aggregate supply, it is also necessary to derive a more accurate measure of aggregate labour input: one that takes into account the quality of labour and allows for changes in its composition over time.

The reason why it is important to adjust for labour quality is that a simple measure of labour input (total hours) disregards the fact that hours of work are not homogeneous: the output they can produce depends on the characteristics of individuals and of jobs. The standard measure of labour input does not capture potential changes in the quality of labour that are linked to changes in, for example, the educational composition of the workforce. For example, even if the amount of labour input (number of people or hours) remained fixed, a shift towards more skilled workers would increase supply capacity.

Determining the quality of labour inputs is not straightforward, since skills are difficult to measure directly. But if we assume that the labour market is competitive, 'quality' ought to be reflected in workers' wages since workers would be paid their marginal product. The disadvantage with this approach, however, is that wages might not be a good proxy for skills if there are significant imperfections in the labour market.

Deriving a better measure of labour inputs which reflects these factors requires dividing the working population into groups, according to characteristics linked to different levels of productivity (eg age, education and gender),(5) and weighting each group's total hours by its productive quality (ie by wages). In practice, the adjusted measure we use is an index (equation **(2)**), aggregating the growth rates of the number of hours of each group and weighting them by the group's contribution to total output:

$$\Delta \ln L_t = \sum_i \left(\frac{s_{i,t} + s_{i,t-1}}{2} \right) \ln \left(\frac{h_{i,t}}{h_{i,t-1}} \right) \qquad \textbf{(2)}$$

where $\Delta \ln L_t$ is the growth in the quality-adjusted labour input, $h_{i,t}$ is the number of hours of group i at time t, $s_{i,t}$ is the share in the wage bill of group i, and the weights in the index are given by the average shares in periods t and $t-1$.

(1) The reason for this is that ICT assets depreciate rapidly. The prices of most ICT assets have also been falling due to rapid technological change. This means that the rental price is high relative to the asset price, since the owner has to be compensated for both depreciation and capital losses.
(2) For details on the calculations of the stock and services measure of aggregate capital for the United Kingdom and the sensitivity of the calculations to various assumptions on the depreciation rate, and investment prices of individual assets see Oulton and Srinivasan (2003a).
(3) Chart 3.9 in the February 2004 *Bank of England Inflation Report* presents the same data for 1993–2003.
(4) This section is based on ongoing research undertaken at the Bank of England.
(5) The different groups are constructed by gender, age (16–24, 25–34, 35–44, 45–54, 55–64 (–59 for females)) and education (other qualifications, O level or equivalent, A level or equivalent, degree or equivalent).

This formulation assumes that firms behave competitively in the labour market, so that the contribution of each group of workers to total output is equal to its share of the wage bill: a group is given a higher weight if its members have a higher wage (higher marginal product reflecting higher quality) or work more hours or both. This implies that the quality-adjusted measure will increase by more than the unadjusted measure if the most productive groups of workers (as reflected in their relative wages) experience greater growth in the number of hours (holding the wage bill shares fixed) and/or if the groups with the highest wages experience an increase in their relative wages (holding growth in the number of hours fixed).

This approach parallels the capital services calculations, where each asset is weighted by its rental price weight: in the adjusted labour input measure, each type of labour is weighted by its share in the wage bill.

Chart 3 compares indices of unadjusted and adjusted measures of labour input where the adjusted labour input corrects for differences in age, education and gender. It is clear that the measure of labour input is biased downwards if there is no quality adjustment, especially from 1981 onwards.

Chart 3
Labour input: unadjusted and adjusted for quality

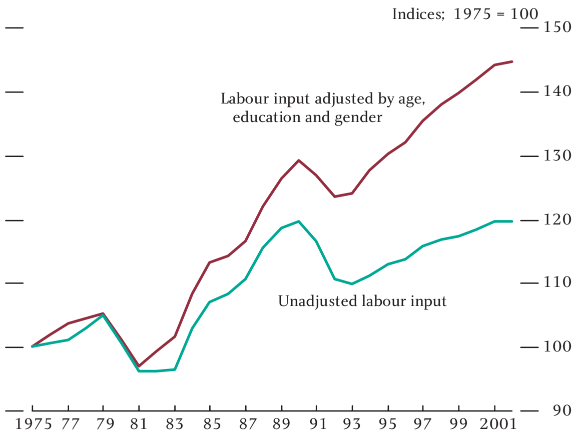

The difference between the two indices reflects important changes in labour composition (or quality of hours worked). In particular, changes in the educational composition of the workforce have contributed most to the increase in labour quality. This effect has been driven mainly by the fact that highly educated people have experienced the greatest rise in the number of hours worked over these two decades. Changes in the age distribution have had a small positive impact since young people, who are the least productive in terms of hourly wages, have accounted for a declining share of the workforce. Finally, changes in the gender distribution of the workforce have slightly reduced our measure of labour quality. The latter reflects the fact that more women have joined the workforce, but their wage bill has increased less, partly due to their relative preference for part-time jobs, which have tended to be less well paid per hour than equivalent full-time positions.

Because the adjusted measure of labour input shown in Chart 3 has risen faster than the unadjusted one, a large proportion of what would be considered as TFP growth using raw total hours (ie unadjusted labour input) can actually be attributed to labour input. That is, TFP growth is significantly lower once we allow for changes in labour quality.

There is another dimension of the data that also needs to be considered—namely, that of using disaggregated industry-level data to calculate aggregate productivity growth instead of using aggregate data directly. The following section discusses this issue.

Aggregate TFP growth calculated from industry data

The TFP growth rate shown in Chart 1 is calculated from aggregate data. An alternative aggregate TFP growth rate can be constructed by weighting industry-level TFP growth rates appropriately. As pointed out in Basu, Fernald and Shapiro (2001) and Bosworth and Triplett (2003) the two aggregate measures may not be identical if there are differing returns to scale across industries or heterogeneity across industries in the marginal products of identical factor inputs. It is thus preferable to calculate an aggregate TFP growth measure using industry data, since TFP growth calculated using aggregate data includes the above-mentioned scale and heterogeneity effects. The Bank of England industry data set was developed to address this and other issues. It contains data for 34 industries spanning the whole UK economy, for 1970 to 2000.[1]

Using this data set, the growth rate of aggregate TFP can be calculated by weighting industry-level TFP growth rates, which in turn are calculated using industry-specific gross output, capital services, labour and intermediate inputs measures.

(1) Oulton and Srinivasan (2003b), which is available on request, describe the Bank of England industry data set.

The Bank of England industry data set

This data set contains data for 34 industries spanning the whole UK economy, for 1970 to 2000. For each industry, there are data on gross output and inputs of capital services, labour and intermediates, in both nominal and real terms. Capital services cover four types of non-ICT assets (structures, plant and machinery, vehicles, and intangibles), and three types of ICT assets (computers, software, and telecommunications equipment). The real intermediate index is a weighted average of domestic purchases from all other industries and from imports. Labour services are measured as hours worked, both including and excluding labour quality adjustment, based on the work discussed above.

The data set is consistent with the official UK National Accounts (as given in the 2002 *Blue Book*, Office for National Statistics (2002)) in both real and nominal terms before the following adjustments were made. To derive series for real ICT investment (and thus ICT capital), US price indices were employed for computers and software, converted to sterling terms, to deflate investment in current prices. The main reason for this is that US price indices are believed to control better for quality, whereas the UK indices do not do so fully. Since technological progress is high for ICT goods, the quality rapidly improves, and US ICT price indices therefore fall at a faster rate than the official UK ones. Also, a large upward adjustment has been made to the official level of software investment.(1)

The approach to ICT has implications for the other variables in the data set. Changing the prices used for measuring real investment in computers and software means that the prices used to measure UK output of these products must also be adjusted. The upward adjustment to nominal software investment raises nominal GDP as measured from the expenditure side. To maintain consistency a corresponding adjustment is made to the income side of the accounts.

(1) This adjustment is discussed in Oulton (2002).

Chart 4
Growth of total factor productivity in the United Kingdom (using disaggregated data, capital services and quality-adjusted labour input): 1980–2003

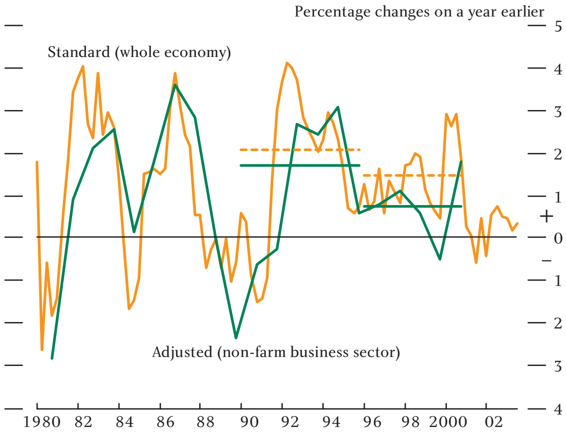

Source for adjusted (non-farm business sector): Basu, Fernald, Oulton and Srinivasan (2003).

Chart 4 presents an aggregate TFP growth estimate for the non-farm business sector in the United Kingdom. Since the aggregate (for the non-farm business sector) is calculated using a 'bottom-up' approach, the hard to measure government sector and agriculture are easy to exclude. Compared with the 'top-down' aggregate TFP growth measure in Chart 1, the non-farm measure shown in Chart 4 gives quite different point estimates for some years over the 20-year time period.(1) This indicates that there could be some heterogeneity of inputs across industries. However, the overall picture remains broadly similar. The growth rate is still procyclical and there is a slowdown in UK TFP growth in the 1990s, even after moving to a capital services measure, adjusting for labour quality and aggregating from industry-level data.

Adjustment costs and variable rates of utilisation(2)

So far, we have assumed that the factors of production can be adjusted costlessly in response to changes in economic conditions. The framework can, however, be extended to take into account costs of adjustment and variable rates of utilisation.

(1) The two lines in Chart 4 must be compared with caution: the standard measure is calculated using data consistent with the 2003 *Blue Book* (Office for National Statistics (2003)) whereas the adjusted measure (using data from the Bank of England industry data set) is calculated using data consistent with the 2002 *Blue Book* (Office for National Statistics (2002)).
(2) This section is based on ongoing research undertaken at the Bank of England.

Capital adjustment costs

The motivation for considering capital adjustment costs is that it may be costly for a firm to increase the amount of capital used for the production of output. One reason for this is that, when firms are investing in new capital, they may need to divert productive resources to installing the new capital rather than producing marketable output. This means that firms are essentially producing two types of products: the final product sold in the market, and the services used within the firm to install new capital. Marketable output may therefore be low during periods of high investment growth, and this would cause a downward bias in estimates of measured productivity growth.

Chart 5 shows the growth rates of business investment (measured in chained volume terms) and total factor productivity.(1) It suggests that there may be a relation between investment and productivity growth: productivity growth slowed during the late 1980s and during the second half of the 1990s, when investment grew rapidly.

Chart 5
Growth of business investment and total factor productivity in the United Kingdom: 1980–2003

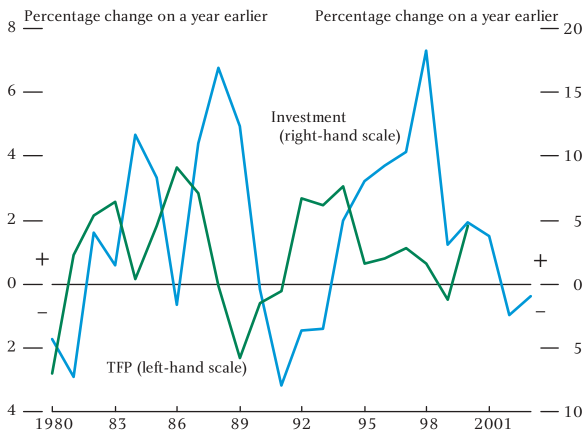

Source for TFP: Basu, Fernald, Oulton and Srinivasan (2003).

The measure of productivity growth can be extended to reflect these effects, by defining it as the fraction of output growth that cannot be accounted for by growth in the inputs, where output is defined as the joint product of observed market output and unobserved installation services. Let i be the growth rate of investment and let ϕ denote the (negative) elasticity of output with respect to investment. This measures the percentage change in marketable output that would occur following a percentage increase in investment. The Solow residual calculations can then be amended for adjustment costs—in equation **(1)**, the growth rate of total output (including services to install capital) now equals $y - \phi i$.

The effect on output of installing new capital is not directly observable. But we can estimate it indirectly, by relating the adjustment costs to observable variables. If a firm can adjust capital without incurring any costs, it will always make sure that its productive capital is at its long-run (or normal) level, at which the cost of using one extra unit of capital (given by the rental price of capital) equals the return to one more unit of capital in the production of output. When firms face adjustment costs, the optimal level of capital will still be one at which the cost of installing one more unit of capital equals capital's expected return. But the cost of installing capital now consists of both the rental price and an adjustment cost. And the marginal return to capital consists both of the return in the production of market output and of the contribution to lower adjustment costs in the future. So the optimal level of capital is determined by a dynamic condition, which links current capital to expected future levels of capital. This relation can be used to obtain an estimate of the marginal cost of adjusting capital, from which an estimate of the elasticity of output with respect to investment can be derived.

Ongoing work at the Bank of England focuses on estimating capital adjustment costs for the United Kingdom, for both non-ICT and ICT assets, using the Bank of England industry data set. The results suggest that capital adjustment costs are quantitatively important, and similar in magnitude to those obtained for the United States.(2) We find that, for every 1% increase in investment in aggregate capital, output falls by between 0.02% and 0.04%. If firms invest in traditional non-ICT capital, such as buildings and plant and machinery, output falls by more, while the opposite holds for investment in ICT capital.(3) The net impact on TFP growth, however, also depends on the growth rates of the different types of investment.

These results thus suggest that the standard measure of productivity growth underestimates actual productivity

(1) The TFP measure is the one shown in Chart 4, adjusted for capital services and labour quality, aggregated from industry data.
(2) See for example Shapiro (1986) and Basu, Fernald and Shapiro (2001).
(3) These estimates are based on average elasticities for the sample period (1979 to 2000).

growth in periods of high investment growth. In particular, the slowdown in UK total factor productivity growth in the late 1990s is less pronounced after taking into account capital adjustment costs, compared with the estimate of TFP growth that only adjusts for capital services and quality-adjusted labour services (as shown in Chart 4).

Variable rates of utilisation

If firms face adjustment costs in undertaking new investment and in hiring and firing workers, they may respond to short-run fluctuations in demand by adjusting the intensity with which labour and capital are used. For example, capital can be utilised more intensively by increasing the number of shifts, and labour can be used more intensively by increasing the effort of workers. The Solow residual would in this case overestimate productivity growth in periods when utilisation is growing rapidly, and *vice versa*. This would cause measured productivity to vary positively with the economic cycle, as Chart 1 suggests is in fact the case.

A measure of productivity growth that allows for these effects can be defined as the fraction of output growth that cannot be accounted for by growth in inputs or by growth in the utilisation of these inputs. Define s and e as the growth rates of the utilisation of capital and labour, respectively. Equation **(1)** can now be adjusted to take into account varying rates of utilisation by defining the growth of capital services as $k + s$ and the growth of labour services as $l + e$.(1)

It is not possible to observe the level of utilisation of capital and labour directly; the challenge is again to relate these unobserved variables to something that we can observe. An earlier *Quarterly Bulletin* article by Felices (2003)(2) discussed different approaches to measuring utilisation rates for labour inputs. Here we use an approach that derives links between observed variables and changes in the utilisation rates by using the optimality conditions faced by the firm.(3)

Consider a firm that would like to use more labour. The amount of labour can be thought of as a combination of the number of workers, the number of hours that each worker works, and the effort of each worker. If it is costly to hire more workers, the firm could alternatively consider increasing the number of hours worked, or worker effort. Since the alternative ways of increasing labour tend to come at a cost, it is optimal for the firm to consider all three margins at the same time. This means that the firm makes sure that the cost of a marginal increase in labour is the same irrespective of whether the firm hires more workers, increases the number of hours, or raises effort; when the number of hours is increasing, effort should therefore also be increasing. It should therefore be possible to use observed hours as a proxy for unobserved effort.

Similarly, the utilisation of capital is not observable. But to use capital more intensively, the firm has to use more labour, for example by increasing the number of hours or effort. Moreover, if capital wears out more quickly when utilisation is high, replacement investment should be high when capital utilisation is high. Also, when capital utilisation is rising, the use of intermediate inputs, such as energy inputs, should be increasing. Thus the growth of the number of hours, investment and intermediate inputs could be used as proxies for capital utilisation.

These relationships can be used to obtain an indirect estimate of utilisation. Ongoing work at the Bank of England focuses on this, by relating the growth rates of effort and capital utilisation to the growth rates of the number of hours, investment and intermediate inputs, again using the Bank of England industry data set. Because effort is unobservable, obtaining an appropriate proxy requires careful analysis of the data. For example, as discussed by Felices (2003), there has been a strong downward trend in the number of hours in the United Kingdom, driven by mainly structural factors. So hours worked appear to respond not only to cyclical factors, but also change for structural reasons, and taking this into account properly is important when measuring unobserved utilisation.

Initial results suggest that variations in utilisation of both capital and labour may be important and that, by adjusting for variable utilisation rates, the cyclical pattern in total factor productivity growth can be reduced. This is consistent with findings for the United States, as discussed in Basu, Fernald and Shapiro (2001).

(1) This is a simplified formula since we also need to correct the measure of productivity growth for costs of adjusting the capital stock and costs of changing the number of workers. For the exact formula, see Basu, Fernald and Shapiro (2001).
(2) An alternative approach to modelling and estimating utilisation rates for the United Kingdom is also discussed in Larsen, Neiss and Shortall (2002).
(3) This approach is discussed in Basu and Kimball (1997) and Basu, Fernald and Shapiro (2001).

Conclusions

The Solow residual is defined as that part of output growth that cannot be explained by the growth in the primary inputs. A standard estimate of total factor productivity growth for the United Kingdom appears to be procyclical and shows a lower growth rate in the late 1990s than in the first half of the decade.

There are, however, a number of well known issues related to the measurement of the factor inputs that we need to correct for. This article shows that these improvements in measurement could have a material impact on the estimates of total factor productivity growth. For example, using a services measure of capital input instead of a stock measure reduces estimated TFP growth for the United Kingdom in the late 1990s, since the services measure has grown faster than the stock measure. This difference is mainly due to the contribution of services from ICT capital. Using a quality-adjusted measure of labour input instead of an unadjusted measure also reduces TFP growth, since the quality-adjusted measure of labour input has been growing faster than the unadjusted one. This difference is mainly due to changes in the educational composition of the labour force. In contrast, correcting output growth to take into account costs of adjustment to changes in the level of capital input appears to increase TFP growth in periods of high investment growth, such as the late 1990s.

The net effect of these measurement improvements is complex and varies over time. While the overall picture before and after these corrections remains broadly similar, the point estimates are often different. It appears that, when all these improvements are made, the decline in the growth rate of aggregate total factor productivity in the late 1990s relative to the first half of that decade is reduced but not eliminated. In addition, if both capital and labour inputs are adjusted for differing degrees of utilisation over time, the correlation of total factor productivity growth with GDP growth is reduced.

This richer treatment of input measurement is also helpful in projecting future supply capacity. This is because it enables a higher proportion of capacity growth to be identified with measurable (and so forecastable) inputs rather than with the unidentified sources of growth represented by TFP. But even after taking into account this 'concealed increase in resource expansion' (Abramowitz (1956)), a significant part of output growth remains unexplained by the growth in inputs. Understanding this is the subject of future research.

Appendix

Sources and formula for the data in the charts

The formula used to calculate TFP growth is as follows:

$$z_t = y_t - 0.5^* (s_{k,t} + s_{k,t-1})k_t - 0.5^* (s_{l,t} + s_{l,t-1})l_t \qquad t = 1980, \ldots, 2000 \qquad \textbf{(A.1)}$$

TFP growth is calculated as the residual obtained from subtracting a Törnqvist index of the primary inputs (capital and labour) from the growth rate of output (value added).

When using industry-level data, the formula is modified so that the output measure is gross output, and an extra term $(0.5^*(s_{m,t} + s_{m,t-1})m_t)$ allowing for intermediate inputs (m) is subtracted from the right-hand side of equation **(A.1)**.

Chart 1: The variables used in the TFP calculations are defined as follows:

Output	GDP at factor cost: ONS code YBHH.
Capital	Wealth measure: Variant labelled ONS1 in Oulton and Srinivasan (2003a).
Labour	Total hours: ONS code YBUS.
Share of capital	1 – share of labour.
Share of labour	Assumed to be 0.7.

GDP: GDP at market prices: ONS code ABMM.

Chart 2: See Chart 3.9 of February 2004 *Bank of England Inflation Report*.

Chart 3: Bank of England estimates.

Chart 4: The growth rate of total factor productivity for the non-farm business sector is calculated by weighting industry-level TFP growth rates where the weights are the so-called 'Domar weights'—the share of each industry's gross output in aggregate value added. For each industry, the output measure is gross output, the capital measure is capital services, the labour input measure is total hours (adjusted by aggregate labour quality growth), intermediate inputs are taken into account and the share of each input (capital, labour, intermediate) is calculated as a proportion of nominal gross output.

The industry-level data are from the Bank of England industry data set and are described in Oulton and Srinivasan (2003b). The UK aggregate TFP measure (for the non-farm business sector) is summarised in Table 1 of and described more fully in Basu, Fernald, Oulton and Srinivasan (2003).

Chart 5: Chained volume measure of business investment: ONS code NPEL.

References

Abramowitz, M (1956), 'Resource and output trends in the United States since 1870', *American Economic Review*, Vol. 46, No. 2, pages 5–23.

Basu, S and Fernald, J G (2000), 'Why is productivity procyclical? Why do we care?', *NBER Working Paper no. 7940*.

Basu, S, Fernald, J G, Oulton, N and Srinivasan, S (2003), 'The case of missing productivity growth: or, does information technology explain why productivity accelerated in the United States but not in the United Kingdom', forthcoming in *NBER Macroeconomics Annual*, Vol. 18, 2003. (Also available as *NBER Working Paper no. 10010*.)

Basu, S, Fernald, J G and Shapiro, M (2001), 'Productivity growth in the 1990s: technology, utilization or adjustment', *Carnegie-Rochester Conference Series on Public Policy,* Vol. 55, pages 117–65.

Basu, S and Kimball, M S (1997), 'Cyclical productivity with unobserved input variation', *NBER Working Paper no. 5915*.

Bosworth, B P and Triplett, J E (2003), *Service sector productivity in the United States: Griliches' services volume revisited*, manuscript, The Brookings Institution, prepared for CRIW Conference in memory of Zvi Griliches. Available at www.brookings.edu/dybdocroot/views/papers/bosworth/20030919.pdf.

Burnside, C, Eichenbaum, M S and Rebelo, S T (1995), 'Capital utilization and returns to scale', in Bernanke, B S and Rotemberg, J J (eds), *NBER Macroeconomics Annual*, Cambridge MA, MIT Press, pages 67–110.

Felices, G (2003), 'Assessing the extent of labour hoarding', *Bank of England Quarterly Bulletin*, Summer, pages 198–206.

Hulten, C R (2001), 'Total factor productivity: a short biography', in Hulten, C R, Dean, E R and Harper, M J (eds), *New developments in productivity analysis*, Chicago, Chicago University Press.

Jorgenson, D W, Gollop, F M and Fraumeni, B M (1987), *Productivity and US economic growth*, Cambridge MA, Harvard University Press.

Jorgenson, D W and Griliches, Z (1967), 'The explanation of productivity change', *Review of Economic Studies*, Vol. 34, pages 249–83.

Jorgenson, D W and Stiroh, K J (2000), 'Raising the speed limit: US economic growth in the information age', *Brookings Papers on Economic Activity*, Vol. 1, pages 125–211.

Larsen, J, Neiss, K and Shortall, F (2002), 'Factor utilisation and productivity estimates for the United Kingdom', *Bank of England Working Paper no. 162*.

Office for National Statistics (2002), *United Kingdom National Accounts: The Blue Book 2002*, London: The Stationery Office.

Office for National Statistics (2003), *United Kingdom National Accounts: The Blue Book 2003*, London: The Stationery Office.

Oulton, N (2001), 'Measuring capital services in the United Kingdom', *Bank of England Quarterly Bulletin*, Autumn, pages 295–309.

Oulton, N (2002), 'ICT and productivity growth in the UK', *Oxford Review of Economic Policy*, Vol. 18, pages 363–79.

Oulton, N and Srinivasan, S (2003a), 'Capital stocks, capital services, and depreciation: an integrated framework', *Bank of England Working Paper no. 192*.

Oulton, N and Srinivasan, S (2003b), 'The Bank of England industry data set', Bank of England, *mimeo*. Available from authors on request. Also available at: www.nber.org/books/macro18/basu-et-al9-22-03_u.k.data-appendix.pdf.

Shapiro, M D (1986), 'The dynamic demand for capital and labour', *The Quarterly Journal of Economics*, Vol. 101, pages 513–42.

Stiroh, K (2002), 'Information technology and the U.S. productivity revival: what do the industry data say?', *American Economic Review*, Vol. 92, pages 1,559–76.