Assessing the extent of labour hoarding

By Guillermo Felices of the Bank's Structural Economic Analysis Division.

The strength of employment during the recent slowdown is sometimes taken as evidence of labour hoarding. But the extent of such hoarding is difficult to measure. This article reviews different definitions of labour hoarding and a variety of ways of measuring it using aggregate data. Most of these measures indicate that labour has been underutilised during the recent slowdown, implying that firms have indeed hoarded labour to some extent. However, the magnitude of the reduction in utilisation differs across these measures. The evidence also suggests that the recent decrease in utilisation has been limited compared with previous episodes in which labour utilisation was significantly below trend.

Introduction

Firms are said to hoard labour when they choose not to adjust their employment of labour in line with short-run fluctuations in demand for their product and, instead, allow their utilisation of labour to vary over the cycle. Such behaviour would be sensible if firms face costs in adjusting the size of their workforce. It would result in employment being less volatile than output, and this in turn could explain why labour productivity, in the United States and the United Kingdom, tends to vary positively with the economic cycle.(1)

There are several reasons why the study of labour hoarding is important. Because the utilisation of labour falls when labour is hoarded, accounting for changes in labour utilisation can help provide a more accurate measurement of changes in labour input. This in turn leads to better estimates of the inputs that account for output growth, and hence to more accurate measures of total factor productivity.(2) Hoarding of labour might also affect wage pressures: if firms can increase their labour input during an upturn without recruiting extra staff, then wage pressures might be more muted.(3) Measurement of the variation of labour hoarding can therefore be an important part of assessing the state of the labour market, and so a useful input into decisions on monetary policy.

The possibility of labour being hoarded by firms during the recent slowdown in GDP growth is indicated by the continuing strength of employment growth. Over the period 2000 Q1 to 2002 Q1, total employment (measured by the Labour Force Survey) increased at an average rate of 0.9% a year and the employment rate remained close to 74.4%, its average over this period, whereas annual GDP growth fell from 3.4% in 2000 Q1 to 0.7% in 2002 Q1. The steady growth of employment contrasts with the falling and erratic evolution of average hours worked during the recent slowdown (see Chart 1).(4)

Chart 1
Employment, average hours and GDP growth

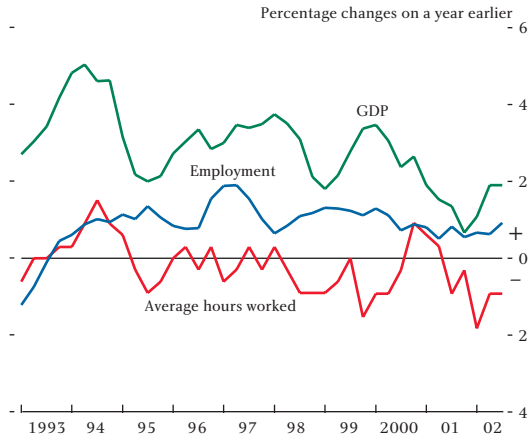

(1) Basu and Kimball (1997) and Basu and Fernald (2000), among others, have studied this correlation extensively. Basu and Fernald (2000) point out that other explanations for the procyclical behaviour of labour productivity are technological progress, imperfect competition, increasing returns and resource reallocation.
(2) Total factor productivity is calculated by subtracting the weighted growth of factor inputs (capital and labour) from output growth and it is often used as a measure of the rate of technological change in the economy.
(3) Darby, Hart and Vecchi (2001) argue that this 'intensive' measure of unemployment, together with the 'extensive' unemployment rate, add to our understanding of the wage/unemployment relationship.
(4) The data used in this article do not include 2003 Q1 because some of the underlying series were not available at the time of publication. This does not create a major problem as the analysis focuses on the behaviour of labour hoarding during the latest slowdown in GDP growth, which approximately covered the years 2000 and 2001.

Chart 2 illustrates the strong procyclicality of output and labour productivity growth. The procyclical behaviour of labour productivity per person reflects the fact that employment adjusts less than output over the cycle. This simple indicator is a commonly used measure of labour utilisation, as it is assumed that labour intensity increases with labour productivity. The recent slowdown would then point to lower labour utilisation. Indeed, year-on-year growth in labour productivity calculated from employment (using Labour Force Survey data) fell from 2.2% in 2000 Q1 to 0.1% in 2002 Q1. The behaviour of labour productivity in hours, however, was rather different. In particular, labour productivity in hours has been higher than its per-person counterpart for most of the period between 2000 and 2002, and has remained higher since. This is the result of the persistent downward trend in average hours worked evident in Chart 1.

Chart 2
Labour productivity and GDP growth

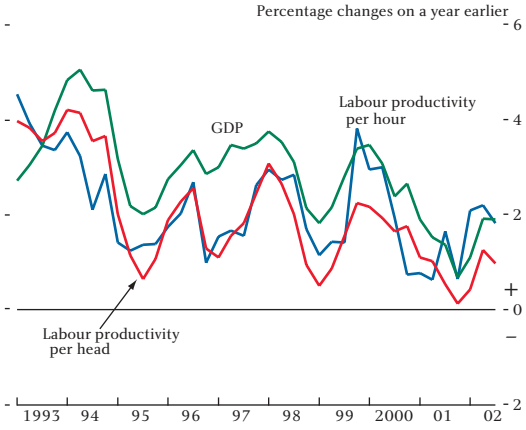

At first sight, this pattern is consistent with firms 'hoarding' labour to some extent. Furthermore, in the face of a slowdown, firms may face an incentive to reduce average hours, especially if they decide to hang on to scarce skilled labour. But average hours worked have been falling for many years, and since 1996 at an average annual rate of 0.4% (see Chart 1). So they do not appear to be responding only to cyclical factors, and may have been trending downwards for structural reasons. In recent years, these reasons include the effect of the Working Time Directive(1) and, possibly, workers' preferences for fewer and more flexible hours, or for part-time work.

The rest of this article examines the concept of labour hoarding in more detail and looks at several ways of measuring it. The next section compares two definitions of labour hoarding. This is followed by a discussion of the method used to measure 'normal' labour utilisation, and an assessment of the extent of labour hoarding in recent years for a variety of measures. The final section presents the main conclusions.

Labour hoarding in heads or hours?

Labour input can be measured either in terms of the number of people employed, or the total number of hours worked. So one can think of both heads and hours-based measures of labour hoarding. According to Hamermesh (1993), labour hoarding can be defined as 'a less than proportionate decrease in worker hours(2) in response to a negative demand shock.' However, he goes on to argue that 'we can be fairly sure that labour hoarding in response to negative shocks is the result of slower adjustment of heads employment than of hours'. Horning (1994) interprets labour hoarding as the 'retention during recessions of workers not needed for current production', as an optimal response to the costs firms face to hire and fire labour. Becker (1975) mentions the incentive to retain workers with specific training in response to temporary negative demand shocks, as an example of these firing costs.(3)

Labour hoarding is a reflection of the intensity with which labour input is used when the amount of labour is costly to adjust. So the key point of interest is how labour utilisation evolves over the cycle. Labour utilisation can differ depending on whether we assume labour is being hoarded in heads or hours. Labour input can be characterised as the product of N, the number of people employed; h, average hours worked per head; and e, the level of effort with which total hours worked (Nh) are applied. If one defines labour input as total hours worked, then the labour utilisation rate can be thought of as the (average) effort rate, e, applied by the workforce over those hours. However, if one defines labour input in terms of the number of people the firm employs, then the utilisation rate of that labour will be the product of the average number of hours worked and the effort rate that applies in those hours, he. Although both definitions should be related,

(1) The Working Time Directive, which came into effect in October 1998, introduced a 48-hour limit on the number of hours an individual can be asked to work in a week.
(2) Worker hours refer to total labour input or, in other words, total hours worked.
(3) A review of studies that seek to measure labour hoarding using micro data includes Fay and Medoff (1985), Oi (1962), Becker (1975), Parsons (1972) and Mincer (1962). Most studies at the macro level use labour hoarding to explain part of the cyclical variation of productivity. The most relevant are Summers (1986), Basu and Kimball (1997), Imbs (1999), Basu and Fernald (2000), Basu, Fernald and Shapiro (2001), and Larsen, Neiss and Shortall (2002).

they may differ if there are costs of adjusting average hours and if the behaviour of average hours changes significantly over time. Indeed, the latter definition will probably trend downwards over time given the long-term decline in average hours worked.

Measuring labour utilisation empirically is further complicated by the fact that labour effort is not observable. By contrast, average hours worked and the number of people in employment are observable and commonly used labour market statistics. The problem then boils down to proxying labour effort with observed variables in order to assess labour utilisation over the cycle.

How can normal utilisation be measured?

Since there are costs to adjusting input quantities (capital and labour), it is the input utilisation rate, not the input quantities, that adjusts to economic shocks in the short run. Given costs of changing heads employment, firms will in the first instance alter the intensity with which the labour input is used. So variations in labour utilisation away from its long-run or normal level can be interpreted as variations in the extent of hoarding: the more intensely the labour input is utilised, the less hoarding one should expect there to be.

Chart 3 illustrates a simple way of assessing whether labour utilisation is above or below its long-run level, taking labour productivity as the proxy for labour utilisation.(1) In the chart, labour productivity, measured as output per head, varies around an upward-sloping linear trend that could reflect its long-run equilibrium. If this were the case, then labour productivity above this line would reflect labour utilisation above its long-run trend, or conversely, labour hoarding below trend. Chart 4 shows the difference between labour productivity and its fitted linear trend. Data points above the zero line represent labour intensity above trend (hoarding below trend). Similarly, underutilisation of the workforce (higher hoarding) occurs when the data lie below zero.

A critical issue with this approach, of course, is whether a straight trend line is a good measure of the 'true' level of productivity consistent with long-run utilisation of the workforce.(2) Despite its simplicity, a linear trend is not necessarily the best representation of the long-run behaviour of labour productivity. A time-varying trend appears to be a more suitable way of tracking the long-run changes in labour productivity over time. A Hodrick-Prescott filter provides estimates of such time-varying trends.(3) Chart 3 shows the trend calculated using this procedure. A 'filtered' series is then obtained by subtracting the non-linear trend from the productivity series. Chart 4 shows the cyclical behaviour of the 'filtered' labour utilisation series. The resulting series measures the percentage deviation of the logarithm of the series from the estimated trend. Needless to say, the implications for labour utilisation over the cycle can be quite different from those derived on the basis of a linear trend.

Chart 3
Labour productivity per person: actual and trends

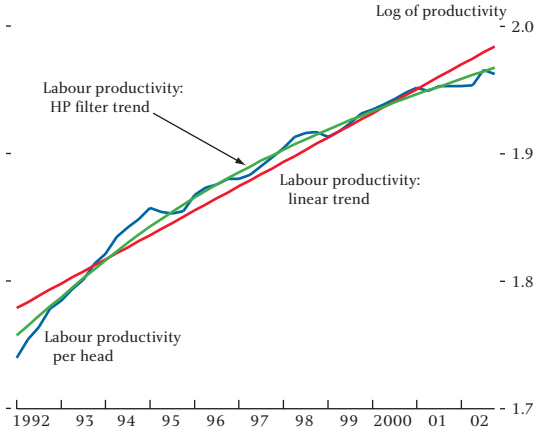

Chart 4
Labour utilisation (hoarding) based on labour productivity in heads

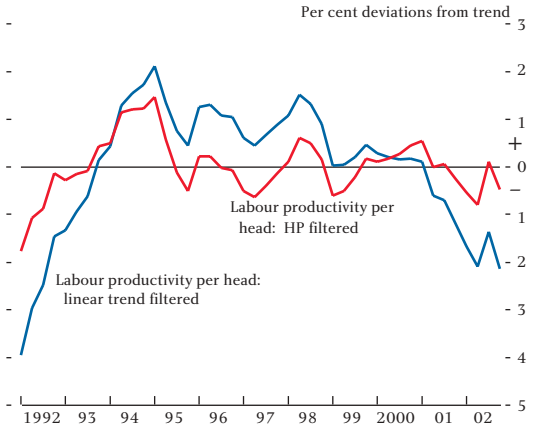

(1) The assumption behind this measure is that labour is utilised more intensely when labour productivity is higher.
(2) A variant of this method has previously been used by Darby, Hart and Vecchi (2001), Fair (1985), and Fay and Medoff (1985). These authors interpolate straight lines between the peaks in the sample.
(3) One can choose the smoothing parameter of the Hodrick-Prescott filter to affect the curvature of the trend. We use a smoothing parameter of 1600 for quarterly data, following standard practice. One should bear in mind that this filter has a greater margin of error at the beginning and end of the sample period, therefore the filtered series for the 1992–2002 period would leave us uncertain about labour utilisation movements around the end of the early-1990s' recession.

The main drawback associated with the labour productivity per person measure of labour utilisation is that it is affected by the downward trend in average hours worked. Analysing labour productivity *per hour* avoids this criticism, but other potential problems affecting both measures remain. In particular, labour productivity measures neglect the effect of other factors of production that may affect labour productivity for reasons other than cyclical movements in labour utilisation. Cyclical changes in total factor productivity, the capital stock and its utilisation rate, as well as the returns (increasing or decreasing) associated with the different inputs might affect labour productivity in heads or hours, for reasons unrelated to changes in labour intensity.

Productivity measures might also be affected by the fact that aggregate output data include only regular production of marketable output and do not consistently include other necessary work which supports production, such as painting the factory or repairing the machinery. This type of work can be particularly important during downturns. Measured labour productivity could therefore fall due to output mismeasurement and not necessarily due to firms hoarding labour.

An alternative to these measures is average hours worked. If firms face relatively high costs of altering the size of their workforce, then one should expect effort and average hours to move together as they are relatively cheaper to adjust. Given that effort is not observable, then (detrended) average hours should represent a good proxy for effort.

These three measures are *ad hoc* in the sense that they are based on commonly used series that proxy labour utilisation. Two additional measures are also studied here. These are based on optimising models that can be used to derive effort series. All of these measures assume that workers, and not hours, are hoarded, except for the labour productivity per hour measure, and the one based on the consumption to output ratio.

As in the previous example, these series are detrended using a Hodrick-Prescott filter to calculate the percentage deviations around the long-run trend representing 'normal' utilisation. These deviations are then standardised to express them as a fraction of the maximum absolute deviation in the sample period. This facilitates comparisons across measures. All of these series are compared with the logarithm of GDP, which is filtered and detrended in the same way as the utilisation series. The exact functional forms of the utilisation series are presented in the appendix.

Was labour hoarded in the latest downturn?

This section compares our five measures of labour utilisation in order to assess the extent of labour hoarding in recent years. Chart 5 presents these measures using quarterly data for two different time periods. The charts on the right-hand side compare the measures from 1992 to 2002, while the charts on the left use data from 1970 to 2002 (or from 1984 to 2002, depending on data availability). The main reason for this split is the lack of quarterly hours data prior to 1992.

Annual hours data are available from 1984 to 1991; therefore, we interpolate them to extend the sample period as much as possible.(1) The pre-1992 data provide a longer time span and a lower margin of error in the filtering process. Nevertheless, their reliability could be affected because they have not yet been officially adjusted using the results of the 2001 Population Census. This in turn could lead to some inaccuracy when comparing the extent of labour hoarding pre and post-1992.

Ad-hoc empirical measures of utilisation

Utilisation measure based on labour productivity in heads

This measure assumes that labour is utilised more intensely when labour productivity in heads is higher. The first chart on the right-hand side of Chart 5 shows that the utilisation of the workforce has been below trend for most of 2001 and 2002. The chart on the left confirms this result, despite the minor discrepancy observed in 2002 due to the fact that it uses a different underlying employment series.(2) This chart also shows that the extent of labour hoarding in recent years could be small compared with previous episodes. As explained earlier, the main drawback of this measure is that movements in average hours due to structural reasons

(1) We interpolate the annual observations from 1984 to 1992 using a series for hours worked in manufacturing. Total hours worked are the product of average hours and the number of people in employment. The latter is not available from the Labour Force Survey (LFS) prior to 1992. Therefore, we interpolate the annual observations from the LFS using the quarterly pattern of the series for Workforce Jobs.
(2) The 1970–2002 series uses the Workforce Jobs (index) series, while the 1992–2002 series uses LFS data. The latter series is not available on a quarterly basis before 1992.

Chart 5
Measures of labour utilisation (standardised percentage deviations from trend)

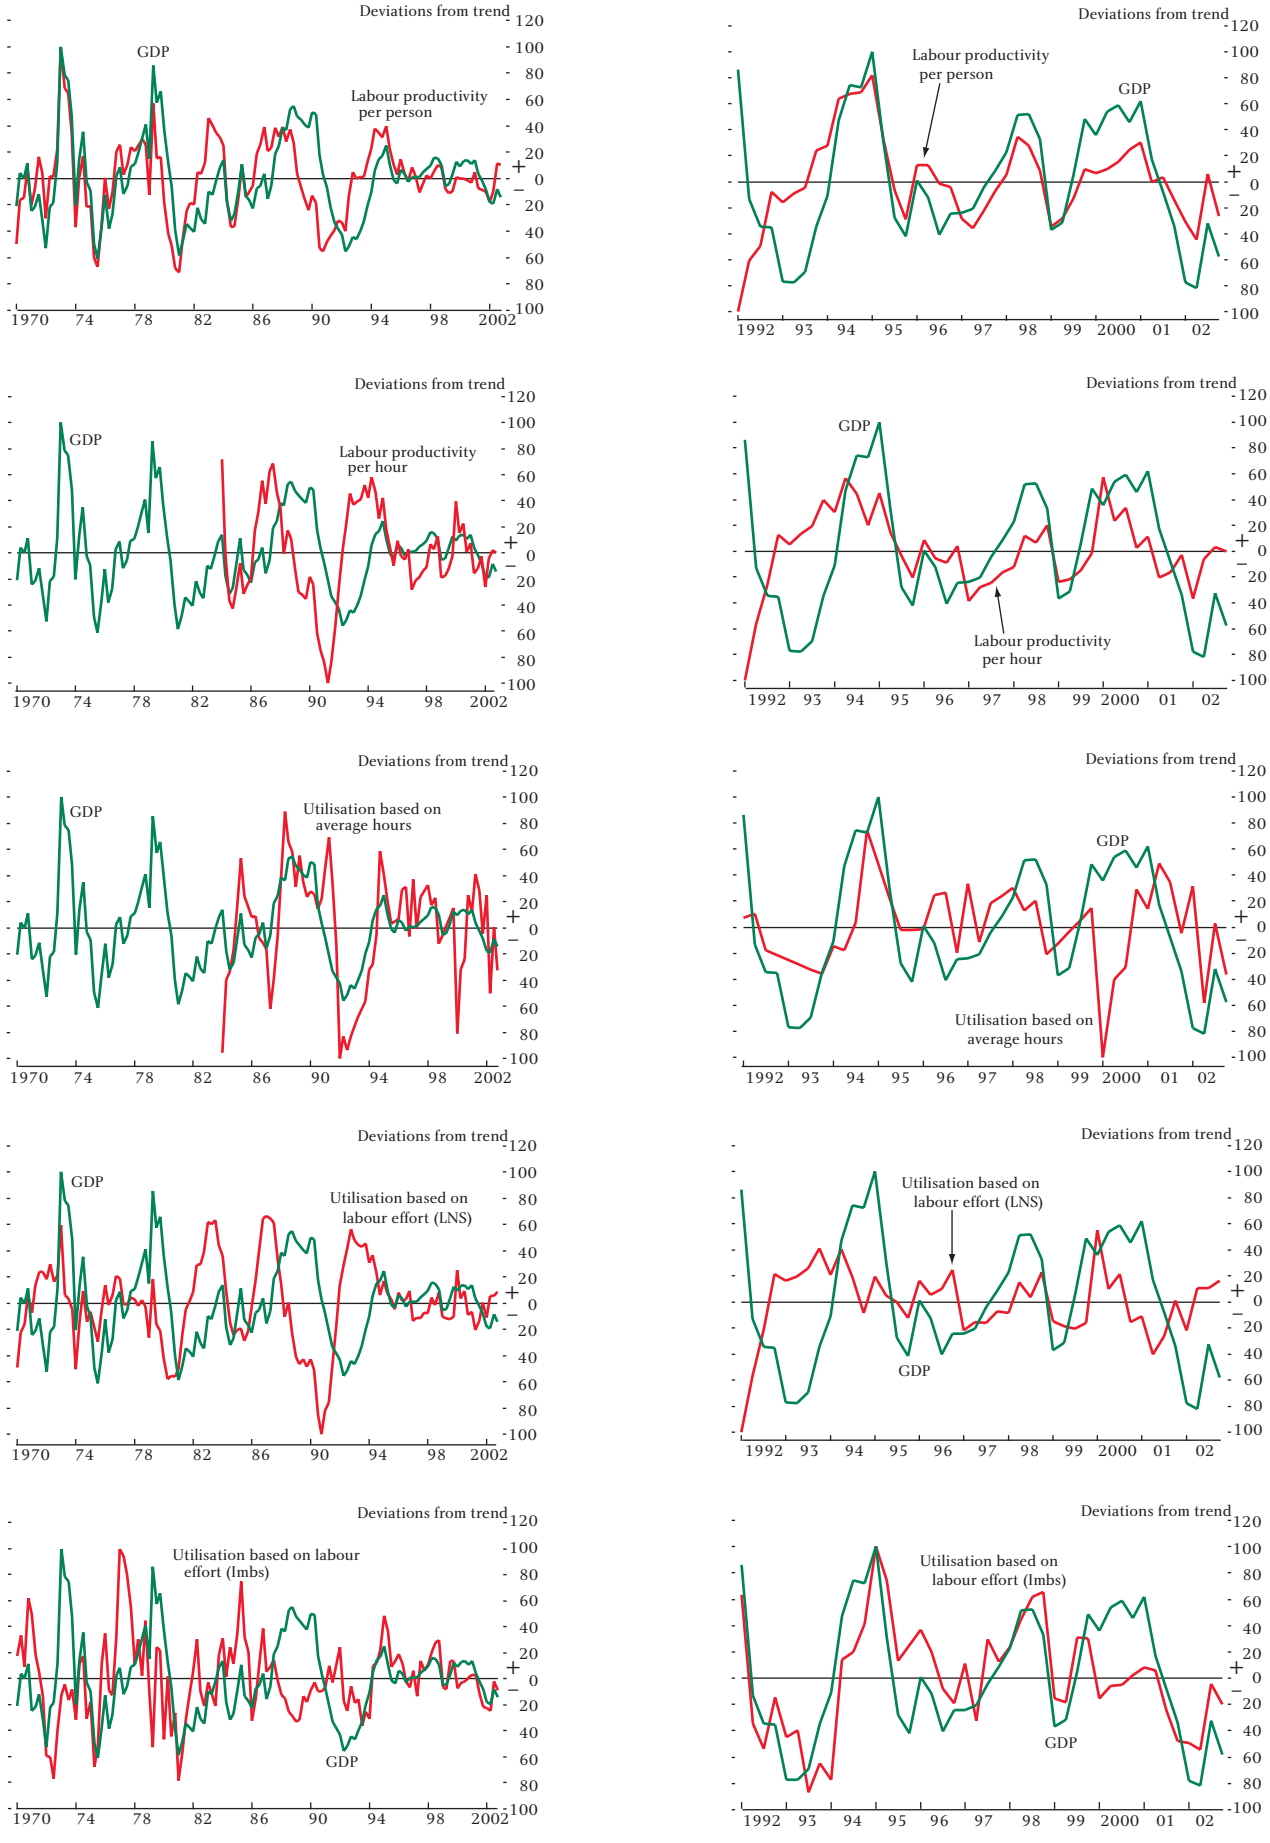

could be driving the behaviour of labour productivity per worker. Moreover, labour productivity is an imperfect measure of utilisation as it absorbs changes in the capital stock (and its utilisation), the rate of technological progress and the skill composition of the labour force. Labour productivity can also be affected by factors such as the degree of competition in the final goods market, and the quality and composition of the labour force.

Utilisation measure based on labour productivity in hours

As with the previous measure, this measure assumes that total hours worked are utilised more intensely when labour productivity in hours is higher. Hence, lower labour productivity implies higher labour hoarding. The second chart on the right-hand side of Chart 5 shows that utilisation decreased in the slowdown, though for a shorter period than in the case of labour productivity per worker. Furthermore, it appears to have reverted to trend in recent quarters. Although this measure factors in the behaviour of average hours (by dividing output by total hours worked), it still shares the other drawbacks of the previous measure. The chart on the left shows that recent deviations from trend are quite small compared with previous ones. One should bear in mind, however, that the hours data prior to 1992 are not absolutely reliable as they are based on interpolated annual data that have not yet been adjusted for the results of the Census.(1)

Utilisation measure based on average hours worked

Basu and Kimball (1997) argue that, in the presence of heads adjustment costs, the choice of average hours worked and effort by cost-minimising firms must be closely related.(2) This is because the cost of altering effort and hours is believed to be cheaper than changing heads employed. Labour effort is not directly observable, but average hours worked are; so they can be used as a reliable proxy for factor utilisation. The third chart on the right of Chart 5 shows a decline in utilisation in 2002 that lags GDP, unlike the previous measures. The left chart shows that the latest deviation from trend is small, relative to the previous data.

Detrending the average hours series allows one to control (imperfectly) for the downward trend observed in recent years. The pre-1992 data are particularly useful in this case because the filtering process captures the fall in average hours during the early-1990s' recession. The post-1992 data, however, only cover part of this recession, hence the dip is not fully captured by the filtering process. Another advantage of average hours over other measures is that they are not affected, at least directly, by changes in the capital stock (and its utilisation) and the rate of technological progress. These characteristics make average hours a more reliable proxy for labour utilisation, despite the difficulty of removing the downward trend observed in recent years.

Model-based measures of labour utilisation

Utilisation measure based on total hours worked

In a recent Bank of England working paper, Larsen, Neiss and Shortall (2002) (hereafter LNS) develop a version of the model of Burnside and Eichenbaum (1996) (hereafter BE) in order to measure factor utilisation. Both studies assume that firms hoard workers in the short run. They further assume that individuals work a fixed number of average hours, so changes in effort will capture movements in labour utilisation. The resulting effort series in LNS is mainly driven by total hours worked. It is also a function of other variables, such as the capital stock, government expenditure and a technology shock. The fourth chart on the right indicates that labour utilisation increased in 2002, contrary to all the other measures. The erratic behaviour of the series could be reflecting the effect of total hours, which are in turn determined by the opposite movement of heads and average hours. This is the main drawback of the series, despite the advantage of it being derived from a fully optimising model.

Utilisation measure based on ratio of output to consumption

Based on a model similar to BE, Imbs (1999) develops a model that allows for labour hoarding in hours, to construct series on input utilisation rates for ten OECD countries.(3) Unlike LNS and BE, his measure of labour effort is a function of the ratio of output to consumption and two estimated parameters of the optimisation problems of households and firms. The intuition behind this measure is as follows. Given that effort is chosen

(1) The utilisation measures and the GDP series on the second and third charts on the left have been filtered using different sample periods. The relative size of the measures compared with that of GDP could be affected by the filtering process. The qualitative properties of the series, which we rely on mostly for the analysis of these measures, will remain unchanged.
(2) This paper shows that variable capital and labour utilisation explain 40%–60% of the cyclicality of the Solow residual in US manufacturing. In a more recent study, Basu and Fernald (2000) decompose labour productivity into technology shocks, factor utilisation, imperfect competition, increasing returns, and resource reallocations. They find that variable utilisation and resource reallocations are particularly important in explaining procyclical productivity.
(3) The model also assumes variable capital utilisation.

optimally, the household's marginal loss of supplying effort (measured in units of consumption) has to be equal to the marginal output extracted by firms from this additional effort. Hence, movements of output relative to consumption (shaped by these key parameters) should proxy movements in the equilibrium level of effort. This formulation takes advantage of the consumption data by combining consumption and labour supply decisions of households with the profit maximisation decisions of firms.

The last row of Chart 5 displays the results. Labour intensity decreased between 2001 Q2 and 2002 Q2, showing signs of labour hoarding during the slowdown. The advantage of this measure is that it does not rely on hours data, therefore the left panel offers a consistent long-run utilisation series. The recent fall in utilisation could, therefore, be interpreted as the first indication of labour hoarding since the early-1990s' recession. Although this measure and the one based on total hours worked share the advantage of backing out labour utilisation series from optimising models, they also share the constraint of being dependent on parameter estimates and structural equations that have to be assumed to build these series.

Concluding remarks

This article has attempted to measure the extent of labour hoarding by comparing the cyclical behaviour of different measures of labour utilisation based on aggregate data. It compared five measures of labour utilisation, contrasting their relative merits and their ability to measure labour hoarding. Three of them are *ad hoc*, in the sense that they provide empirical measures of labour utilisation that are not derived from any optimisation problem. The other two measures are based on optimising models that back out labour utilisation as an effort variable that is part of the labour input.

Most of these measures indicate that labour was underutilised during the recent slowdown, implying that firms hoarded labour to some extent. However, the magnitude of the reduction differs between the measures, and the measures themselves are subject to various limitations. One other feature of the results was that the recent decrease in utilisation appeared to be quantitatively small compared with previous episodes where labour utilisation was significantly below trend, and the volatility of all the measures appears to have been lower in the past decade. These features could indicate that the labour market has become more flexible, allowing for more hiring and firing and hence less variation in utilisation. It might also be related to greater stability of inflation and output, associated with the new macroeconomic policy framework. The relative importance of these effects is, however, a matter for future research.

Technical appendix

Ad-hoc measures of labour utilisation

Utilisation measure based on labour productivity in heads

The utilisation (effort) measure based on labour productivity in heads takes the form:

$$e_t = Y_t / N_t$$

where Y is GDP at factor cost and N is heads employment.

Utilisation measure based on labour productivity in hours

The utilisation (effort) measure based on labour productivity in hours takes the form:

$$e_t = Y_t / N_t h_t$$

where Y is GDP at factor cost, N is heads employment, and h is average hours worked.

Utilisation measure based on average hours worked

The utilisation (effort) measure based on average hours worked takes the form:

$$e_t = k h_t$$

where h is average hours worked and k is a positive constant. For simplicity, the article assumes $k = 1$, as this value will not affect the calculations of the measures compared here.

Model-based measures of labour utilisation

Utilisation measure of Larsen, Neiss and Shortall

The utilisation measure of Larsen, Neiss and Shortall takes the following form:

$$e_t^* = \pi_0 \bar{H}_t^* + \pi_1 H_t^* + \pi_2 K_t^* + \pi_3 G_t^* + (\pi_3 - \pi_4) X_t^*$$

where an asterisk denotes the growth rates of the variables, and where $\bar{H}^*$ is effective total hours, H^* is total hours, K^* is capital, G^* is government expenditure, and X^* is total factor productivity. Using calibrated parameter values the authors find that $\pi_0 = -0.50$, $\pi_1 = -0.01$, $\pi_2 = -0.49$, $\pi_3 = 0.49$, $\pi_4 = 0.06$. The model assumes that total hours are costly to adjust over time. Effective total hours equal total hours minus this adjustment cost. The authors estimate this cost to be quite small. As a result, the behaviour of effective total hours and total hours is almost the same. The high and negative value of π_0 and the small value of π_1 therefore imply that total hours drive the effort series inversely.

Utilisation measure of Imbs

The utilisation measure of Imbs takes the form:

$$e_t = \left(\alpha \frac{Y_t}{C_t} \right)^{1/(1+\theta)}$$

where Y is the level of GDP, C is private consumption, α is the share of labour in output, and θ measures a representative household's disutility associated with providing effort. As in Imbs (1999), this article assumes $\alpha = 0.793$ and $\theta = 0.231$.

References

Basu, S and Fernald, J (2000), 'Why is productivity procyclical? Why do we care?', *NBER Working Paper Series*, No. 7040.

Basu, S, Fernald, J and Shapiro, M (2001), 'Productivity growth in the 1990s: technology utilisation or adjustment?', *NBER Working Paper Series*, No. 8359.

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

Becker, G (1975), *Human capital*, 2nd edition, Chicago: University of Chicago Press.

Burnside, C and Eichenbaum, M (1996), 'Factor hoarding and the propagation of business cycle shocks', *American Economic Review*, Vol. 86(5), pages 1,154–74.

Darby, J, Hart, R and Vecchi, M (2001), 'Wages, work intensity and unemployment in Japan, UK and USA', *Labour Economics*, Vol. 8, pages 243–58.

Fair, R (1985), 'Excess labour and the business cycle', *American Economic Review*, Vol. 75, Issue 1, pages 239–45.

Fay, J and Medoff, J L (1985), 'Labour and output over the business cycle: some direct evidence', *American Economic Review*, Vol. 75, pages 638–55.

Hamermesh, D (1993), *Labor demand*, Princeton: Princeton University Press.

Horning, B (1994), 'Labor hoarding and the business cycle', *International Economic Review*, Vol. 35, No. 1.

Imbs, J (1999), 'Technology, growth and the business cycle', *Journal of Monetary Economics*, Vol. 44, pages 65–80.

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

Mincer, J (1962), 'On the job training: costs, returns and some implications', *Journal of Political Economy*, Vol. 70, pages 50–79.

Oi, W Y (1962), 'Labor as a quasi-fixed factor', *Journal of Political Economy*, Vol. 70, pages 538–55.

Parsons, D O (1972), 'Specific human capital: an application to quit rates and layoff rates', *Journal of Political Economy*, Vol. 80, pages 1,120–43.

Summers, L (1986), 'Some skeptical observations on the real business cycle theory', *Quarterly Review: Federal Reserve Bank of Minneapolis*, Vol. 10, pages 23–27.