

Core inflation: a critical guide

Alan Mankikar

and

Jo Paisley

Working Paper no. 242

Monetary Assessment and Strategy Division and Monetary and Financial Statistics Division,
Bank of England, Threadneedle Street, London, EC2R 8AH.

alan.mankikar@bankofengland.co.uk

jo.paisley@bankofengland.co.uk

This working paper is a revised and extended version of a paper presented at the Working Group on Econometric Modelling: Seminar on Core Inflation, held at the European Central Bank in May 2001. We would like to thank Hasan Bakhshi, Charlie Bean, Ian Bond, Alex Brazier, Spencer Dale, Simon Price, Iain de Weymarn and Tony Yates for their helpful comments. The views in this paper are those of the authors and not necessarily of the Bank of England.

Copies of working papers may be obtained from Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH; telephone 020 7601 4030, fax 020 7601 3298, e-mail mapublications@bankofengland.co.uk.

Working papers are also available at www.bankofengland.co.uk/wp/index.html.

The Bank of England's working paper series is externally refereed.

Contents

Abstract	5
Summary	7
1 Introduction	9
2 Monetary policy in theory and practice	9
3 What induces ‘noise’ in measured inflation?	12
(i) Problems associated with measuring ‘actual’ inflation	12
(ii) Movements in relative prices and aggregate inflation	12
(iii) Interpreting changes in annual inflation rates	13
4 Measuring core inflation	15
(i) Measures based on trimming	16
(ii) Measures based on ‘exclusion’	18
(iii) Measures based on the whole price distribution	19
(iv) Model-based measures	21
(v) Domestically generated inflation	24
5 Evaluating measures of core inflation	25
(i) Tests proposed by Marques <i>et al</i> (2000)	26
(ii) How do measures of core inflation in the United Kingdom perform in the tests?	26
6 Conclusion	28
Appendix A: Core inflation as a statistically robust measure of actual inflation	29
Appendix B: Frequency of RPIX components’ exclusion from the 15% symmetric trimmed mean	32
Appendix C: Marques <i>et al</i> ’s (2000) tests for measures of core inflation	33
References	35

Abstract

The term ‘core inflation’ is widely used by academics and central bankers. But despite its prevalence, there is neither a commonly accepted theoretical definition nor an agreed method of measuring it. The range of conceptual bases is potentially confusing, and can make the large number of available measures of core inflation difficult to interpret, particularly when they display different trends. Nevertheless, measures of core inflation can be helpful if they increase the signal to noise ratio in measured inflation. This paper examines a range of measures of core inflation for the United Kingdom, both conceptually and empirically, setting out their motivation and highlighting their potential limitations. No single measure performs well across the board, but a compromise conclusion on the usefulness of measures of core inflation is that each one may provide a different insight into the inflation process. There can be value in looking at a range of measures, as long as one bears in mind what information each type of indicator is best at providing. When all measures are giving the same message then, in a sense, monetary policy makers can reasonably consider that these measures are providing a reliable guide to inflationary pressures. It is when the measures start to diverge that policymakers need to take a much closer look at the reasons for those divergences.

Summary

The term ‘core inflation’ is widely used by academics and central bankers. But despite its prevalence, there is neither a commonly accepted theoretical definition nor an agreed method of measuring it. Some researchers, for example, have suggested that core inflation relates to the growth rate of the money supply. Others identify core inflation with the ‘durable’ part of inflation, while others define the term as that component of measured inflation that has no medium to long-run impact on real output. The range of conceptual bases is potentially confusing, and can make the large number of available measures of core inflation difficult to interpret, particularly when they display different trends. This paper sets out how the concept of core inflation might be useful to monetary policy makers and provides a conceptual and empirical evaluation of various measures of core inflation in the United Kingdom.

Month-to-month movements in inflation can be volatile, making outturns potentially difficult to interpret. The ‘noise’ might be a reflection of movements in relative prices, or it may reflect one-off price level effects that will affect the annual inflation rate for a year. A key task for policymakers, as with all economic variables they monitor, is to read through the volatility or ‘noise’ in the data to extract as much information as possible. Measures of core inflation can be helpful if they increase the signal to noise ratio in measured inflation.

This paper examines a range of measures of core inflation for the United Kingdom, setting out their motivation and highlighting their potential limitations. The literature has distinguished two main approaches to measuring core inflation. First, there is the statistical approach, where some researchers take an existing price index and either remove certain items from it or reweight the components of that index, or use statistical methods to try to extract the ‘persistent’ or underlying trend component. These measures can be thought of as summary statistics of the large amount of component data in the aggregate price index. Second, there are model-based measures, which are usually based on multivariate econometric analysis in which some structure has been imposed that is explicitly grounded in economic theory. These measures use past relationships between aggregate inflation and its determinants to distinguish movements in inflation that reflect underlying pressures from those that reflect transitory shocks. They also typically incorporate some prior view about the ‘smoothness’ of core inflation.

Because one can define core inflation in a number of ways, this can unfortunately create confusion and there is no single ‘right’ answer. One needs to be aware of the pros and cons of different measures. Measures that simply strip out the most volatile elements or reduce the weight of extreme observations in the price index raise the possibility of losing potentially useful information. Measures based on times-series models are sensitive to prior beliefs about the time-series properties of core inflation.

It would clearly be unwise to be too ambitious about what a measure of inflation can hope to capture. Core inflation certainly cannot act as a summary statistic for inflationary pressures that are relevant to the monetary policy decision. Some measures may be useful at certain times. And since there is no ‘right’ answer, the main test is whether policymakers find them useful for

helping to understand the inflation process better. Most important is to be clear how the term core inflation is used and what the concept underlying it is.

A compromise conclusion on the usefulness of measures of core inflation is that each one may provide a different insight into the inflation process. This paper finds that no single measure performs well across the board. Nevertheless, there may be value in looking at a range of measures, as long as one clearly bears in mind what information each type of indicator is best at providing. When all measures are giving the same message then, in a sense, monetary policy makers can reasonably consider that these measures are providing a reliable guide to inflationary pressures. It is when the measures start to diverge that policymakers need to take a much closer look at the reasons for those divergences.

1 Introduction

The term ‘core inflation’ is widely used by academics, central bankers and economic commentators. But despite its prevalence, there is neither a widely accepted theoretical definition nor an agreed method of measuring it. This may make it difficult to get a clear message from the measures, particularly when they have different trends.

There are two quite distinct uses of the term ‘core inflation’. First, it is used in the sense of a ‘statistically robust measure of *actual* inflation’. Such measures seek to exclude extreme price or erratic movements from the sample of price changes to obtain a more efficient and less biased estimate of inflation. This is very different to the second concept of core inflation, which refers to that component of actual inflation that reflects underlying economic pressures in the economy. It is the second that is of most direct significance for monetary policy makers who, like the Monetary Policy Committee (MPC), face a given inflation target. It is this use of the term, which is the main focus of this paper.

There is no agreed definition of core inflation in this second sense, but all are motivated by the observation that measures of inflation can be ‘noisy’, making it difficult to read underlying inflationary developments. Bryan and Cecchetti (1993), for example, have suggested that core inflation relates to the growth rate of the money supply. Blinder (1997) identifies core inflation with the ‘durable’ part of inflation, while Quah and Vahey (1995) define core inflation as ‘...that component of measured inflation that has no medium to long-run impact on real output’.

This paper tries to make sense of these sometimes competing, sometimes complementary concepts. It first looks at monetary policy in theory, and then in practice, highlighting where core inflation might play a role. It then considers what makes measured inflation ‘noisy’. A range of measures of core inflation in the United Kingdom is then examined, setting out the motivation for each and highlighting their potential advantages and limitations, before offering some evaluation of the competing measures based on statistical tests suggested in the literature. Section 6 concludes.

2 Monetary policy in theory and practice

A common approach to formalising the task facing monetary policymakers is to express their objectives in terms of a loss function, which they aim to minimise. Woodford (1999), for example, sets out a framework in which the monetary authority aims to minimise a loss function, derived from a model founded on optimising behaviour by private sector agents. The resultant loss function constitutes the expected discounted sum of the weighted average of the squared deviations of current and future expected inflation (π_t) from the target rate (π_t^*) and of real activity (y_t) from its potential (y_t^*). This is characterised by expression (1):

$$L = E \left[\sum_{i=0}^{\infty} \beta^{t+i} (w_p (\pi_{t+i} - \pi^*)^2 + w_y (y_{t+i} - y_{t+i}^*)^2) \right] \quad (1)$$

where w_p and w_y are the weights on inflation and output stabilisation, respectively, and $0 < \beta \leq 1$ is the discount rate. Although this highly stylised framework provides a useful way of thinking about the monetary policy problem in theory and a means of *evaluating* outcomes for inflation and activity, it does not tell policymakers *how* to achieve those outcomes. To do this, they need to have an understanding of how the economy works, how shocks are transmitted through the economy, and how monetary policy impacts on the economy. Policymakers typically employ a range of economic models to help them form their judgments about the appropriate monetary policy strategy to minimise their loss function.

Importantly, because changes in policy affect activity and inflation with a lag, monetary policy must be forward-looking. Given the lags in the transmission mechanism, monetary policy can do little to affect activity and inflation in the short run, and policymakers are most interested in the *outlook* for inflation, typically over the one to two-year horizon over which monetary policy can have most of its influence.

Following the Bank of England Act (1998), the MPC's remit, determined by the Chancellor each year, states that it must aim 'to maintain price stability, and subject to that, to support the economic policy of Her Majesty's Government, including its objectives for growth and employment'.⁽¹⁾ Initially, price stability was defined as keeping the annual rate of RPIX inflation at 2½% at all times, but in December 2003, the Chancellor announced a switch in the inflation target to keeping the annual rate of CPI inflation at 2% at all times.⁽²⁾ Differences arise between CPI and RPIX inflation because of differing composition and coverage, and the methods for aggregating the constituent prices.⁽³⁾ Owing to the relatively short back-run of CPI data, the analysis below is done in terms of RPIX inflation, the previous targeted measure.

The MPC's remit could be characterised by the type of loss function in **(1)**. But in practice, the monetary policy problem is not as straightforward as theory might suggest and there is a range of issues about which they will be uncertain. For example, economic data, which are themselves measured with sampling error, often get revised or may not correspond precisely to the economic concept most relevant to economic theory. Policymakers may also be uncertain about the underlying theory, or the responsiveness of the economy to policy changes.

Policymakers therefore tend to look at a wide range of data, aiming to build up a picture of how the economy is evolving relative to their expectations. Outturns for an economic variable may provide potentially useful information about the shocks currently affecting the economy and how they are being propagated through the economy. And when data corroborate one another, policymakers will tend to place more weight on those data outturns.

Outturns for inflation are of interest to policymakers because it is the target against which they will be judged. But also inflation outturns may themselves reveal potentially useful information

⁽¹⁾ The United Kingdom is not unique in having this type of remit with a focus also on real variables, like output and employment growth.

⁽²⁾ The CPI was formerly known as the HICP.

⁽³⁾ For more information, see the box on pages 38-39 on the May 2003 *Inflation Report*.

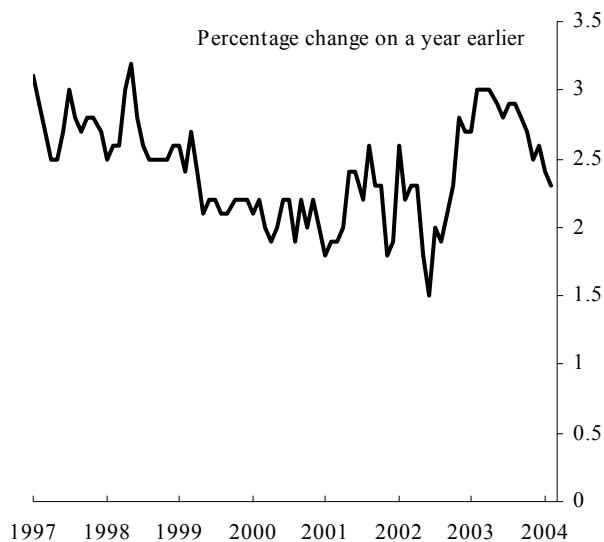
about the shocks currently affecting the economy. They may also be important to the extent that they feed into wage bargaining and affect inflation expectations.

There are a number of reasons, however, why interpreting movements in inflation can be problematic. One difficulty is that any given observation of inflation is consistent with a number of different types of economic shocks affecting the economy. Policymakers need to look at inflation outturns in the context of what else is happening in the economy and in particular how those outturns relate to the fundamental determinants of inflation.

A second difficulty with interpreting movements in inflation is that month-to-month changes can often be volatile or ‘noisy’ (see Chart 1) making outturns potentially hard to interpret. How should policymakers interpret such sharp movements in the annual inflation rate of the target variable?

Chart 1

Annual RPIX inflation



As Blinder (1997) noted, ‘The name of the game was distinguishing the signal from the noise, which was often difficult. What part of each monthly observation on inflation is durable and which part is fleeting?’ That is, when policymakers see large changes in measured inflation, they are interested in the ‘news’ for the outlook for inflation. Do these outturns merit a change in policy? Zeldes (1994) suggests that, ‘presumably the answer depends on the persistence of the inflation innovation in the absence of any change in monetary policy’.

Measures of core inflation are motivated by the observation that inflation can be ‘noisy’, making it difficult to read true underlying inflationary developments. Their usefulness stems from the extent to which they increase the signal to noise ratio in measured inflation. The next section sets out what is meant by ‘noise’ in measured inflation.

3 What induces ‘noise’ in measured inflation?

Focusing on annual RPIX inflation, there are three main reasons why this can be a ‘noisy’ indicator. The first is statistical, the second economic, while the third is a function of the focus on annual inflation rates. These are discussed below.

(i) *Problems associated with measuring ‘actual’ inflation*

There are many different possible conceptual bases for measuring ‘actual’ inflation, depending on the purpose. For any given conceptual basis, there is a range of *measurement* issues, such as the appropriate weights, coverage, the sample of prices collected and the best estimator of the unobserved population mean of price changes. One use of the term ‘core inflation’ in the literature is trying to obtain the ‘statistically best’ measure of the population mean price change (or of ‘actual’ inflation). This differs from the economic concept of core inflation, which is related to demand pressures in the economy.

The precise method used to try to obtain a better estimate of ‘actual’ inflation will depend, in large part, on the *shape* (ie the moments) of the population distribution of price changes. If the population distribution is approximately Normal, then the mean of samples taken from that distribution will be the best estimator of the population mean (in the sense of being efficient and unbiased). The stylised fact, however, is that the sample distribution of price changes appears to be non-Normal, which suggests that the population distribution is likewise non-Normal. There are two ways in which the sample distribution may differ from a Normal distribution. First, it may be ‘fat-tailed’ or leptokurtotic; second, it may be asymmetric or skewed. The statistical response to a leptokurtotic (but symmetric) distribution is to trim symmetrically; the response to a skewed distribution is to trim asymmetrically. Appendix A goes through these issues in more detail.

(ii) *Movements in relative prices and aggregate inflation*

Inflation is ultimately a monetary phenomenon determined by the stance of monetary policy. In a world with fully flexible prices and an unchanged monetary policy stance, a shock to a particular sector (such as a change in tastes or technology) would lead to instantaneous changes in *relative* prices, which would leave the *aggregate* price level, and therefore the aggregate inflation rate, unchanged.

In practice, relative price movements can affect the aggregate price level and therefore the rate of inflation, and sometimes for a considerable period. Why is this? For a start, prices are not fully flexible in the short run. This may be because there are menu costs associated with changing prices, or perhaps because there is staggered price-setting across firms. In these situations, a temporary wedge may open up between firms’ desired and actual prices—in other words, relative prices take time to adjust. There are also more practical reasons relating to the construction of the price index. Consumer price indices cover only a subset of prices in the economy.⁽⁴⁾ Relative price movements between two goods, one included in the RPI basket and the other not, would

⁽⁴⁾ The GDP deflator comes closer to a whole-economy price index.

change the level and therefore the inflation rate of the RPI. Also, many consumer price indices (including the RPI) do not allow for the substitution effects that would normally follow relative price shocks, and so are affected by relative price movements.

In theory, since relative price changes should have no long-run effect on either the aggregate price level or the aggregate inflation rate, they should not require a monetary policy response. So policymakers would like to be able to distinguish between movements in aggregate inflation which reflect relative price movements from those which reflect underlying inflationary pressures. A measure of core inflation that was free from the noise induced by relative price changes would be particularly helpful.

(iii) *Interpreting changes in annual inflation rates*

Inflation targets around the world are exclusively framed in terms of *annual* inflation rates. While annual inflation rates are less volatile than monthly or even quarterly inflation rates, they too may be relatively noisy (see Charts 2 and 3).

Chart 2

Annual RPIX inflation

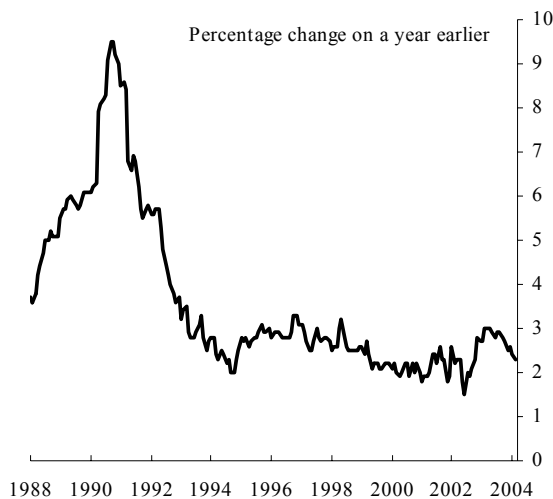
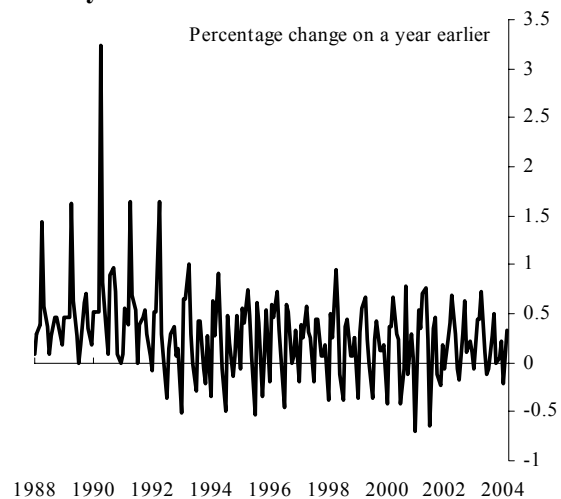


Chart 3

Monthly RPIX inflation



An annual inflation rate compares the current price level with the price level twelve months earlier. A one-off change in the price level will affect the annual inflation for a whole year before it drops out of the annual comparison. A key issue when interpreting movements in the annual inflation rate is to assess to what extent it reflects price changes that are occurring now and/or price changes last year (so-called ‘base’ effects). Changes in the seasonal pattern of price changes from year to year can also induce noise into the annual inflation rate. The following example illustrates these issues.

As a thought experiment, consider how temporary price level shocks may affect the annual inflation rate. In the charts below, there is a benchmark price index (the solid line) that rises by 2% each year, the ‘core’ inflation rate. Chart 4 plots the price level over 7 years. The dotted line uses the same underlying price data but has a temporary price level shock (+5) added in the first

quarter in each year, with the exception of the first year. Chart 5 indicates that inflation rises in the first quarter of year 2 when the first price level shock occurs. But it also indicates that as long as shocks of the same magnitude occur in the same quarters of each year, the annual inflation rate is thereafter unaffected.

Chart 4

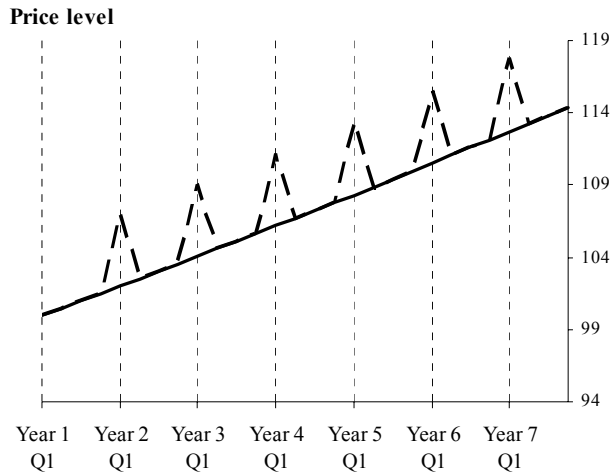
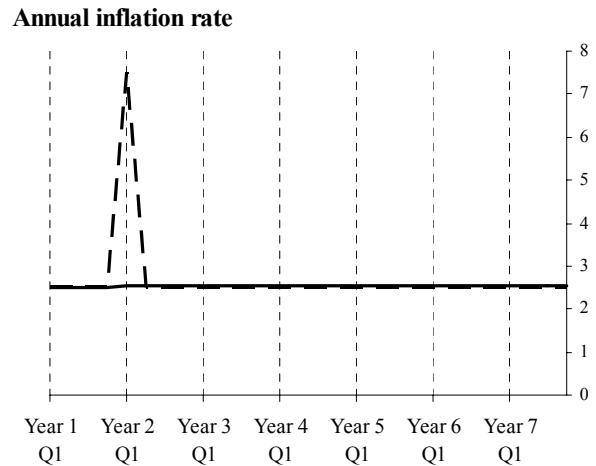


Chart 5



Next consider the case when the size of the temporary price level shock is the same as before except that the timing of the shock in year 3 moves from the first quarter to the second quarter, as shown in Chart 6. Chart 7 shows the corresponding annual inflation, which is clearly volatile, illustrating that a measure that could abstract from such temporary price level effects would give a far better read on ‘underlying’ or ‘core’ inflation.

Chart 6

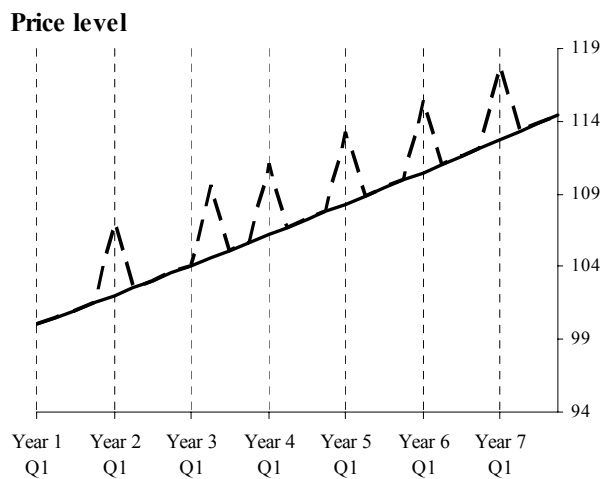
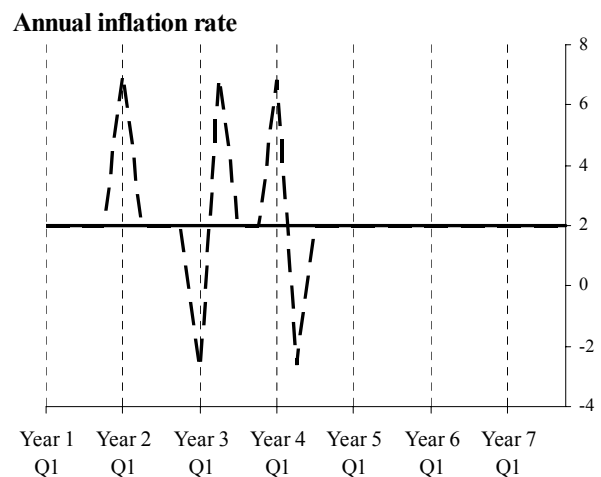


Chart 7



The difficulty is that it is virtually impossible in real time to distinguish between price changes that contain news about inflation and those that simply reflect a change in seasonality or that are the result of a one-off/temporary price level change. Indeed, it may only be some time after the event that one can be confident how to interpret a given change in an annual inflation rate.

Measures of core inflation that attempt to smooth volatile movements in inflation, like moving averages, may help in this regard. Several authors consider other measures besides annual inflation rates (eg three or six-month annualised rates). But these too are noisy and there is the added problem that one needs to seasonally adjust the data, which itself raises additional issues. An alternative way to reduce the short-term noise is to calculate annual inflation rates based on quarterly rather than monthly price data.

4 Measuring core inflation

Given the ‘noise’ associated with measured inflation, there are two natural uses of measures of core inflation. First, they might provide a ‘clean’ measure of *current* inflation: for example, the targeted inflation rate without the ‘noise’ induced by relative price movements. Second, measures of core inflation might be indicative of the outlook for inflation, providing information on the likely course of the targeted rate of inflation over the next few months or so, as relative prices continue to adjust to shocks affecting the economy. This may be particularly useful since the lags in the effects of monetary policy mean that monetary policy makers are most interested in the outlook for inflation.

Given that inflation is ultimately a monetary phenomenon, a potential way to measure core inflation is to link it somehow to growth in the ‘monetary’ aggregates. Unfortunately, as MPC (1999) note, ‘...the relationship between the monetary aggregates and nominal GDP in the United Kingdom appears to be insufficiently stable (partly owing to financial innovation) for the monetary aggregates to provide a robust indicator of likely future inflation developments in the near term’. This means that it is difficult to use monetary data to give a reliable indication of inflationary pressures in the short to medium run.

Even if the relationship between monetary growth and inflation was stable, there is a conceptual problem with trying to isolate that part of observed inflation that is purely the result of monetary growth, as opposed to permanent and transitory ‘shocks’. Bryan and Cecchetti (1994) point out the inherent problem: ‘If money were truly exogenous, one could measure core inflation by estimating this reduced form and then looking only at the portion of inflation that is due to past money growth and the permanent component of the shocks. But in reality, money growth responds to the shocks themselves, so measuring the long-run trend in prices requires estimating the monetary reaction function. In fact, this suggests that measuring core inflation necessitates that we identify monetary shocks as well as the shocks to which money is responding.’ In other words, any particular measure of core inflation is likely to reflect some mixture of current shocks as well as past policy.

Since core inflation is unobservable, there is no ‘right’ way to measure it, and therefore no single agreed method. There have been two main approaches to measuring core inflation. First, there is the statistical approach. Within this, there are those that take an existing price index and either remove certain items from it, or reweight the components of that index, or use statistical methods to try to extract the ‘persistent’ or underlying trend component. These measures can be thought of as summary statistics of the large amount of component data in the aggregate price index.

Second, there is the model-based approach. These usually involve multivariate econometric analysis in which some structure has been imposed that is explicitly grounded in economic theory. They typically use some prior view about the time-series properties of core inflation to help distinguish between core and non-core inflation. The measures calculated under this approach use past relationships between inflation and its determinants to distinguish movements in inflation that reflect underlying inflationary pressures from those that reflect transitory shocks.

The next section examines a range of available measures, setting out the motivation for each, and highlighting their potential advantages and limitations.

(i) *Measures based on trimming*

Trimmed mean measures of core inflation are calculated by excluding a certain percentage of the largest and smallest (weighted) price changes in the components of the index—up to 50% from each tail of the distribution in the case of the (weighted) median.⁽⁵⁾ Some authors, for example Bryan and Cecchetti (1993), have justified trimming on economic grounds. The motivation is that price changes at the extremes of the distribution may contain less information about current underlying price pressures than those further towards the centre of the distribution. Therefore, it is argued that it may be more informative to strip out extreme price movements than to look solely at the aggregate inflation rate (for a fuller explanation see Bakhshi and Yates (1999)).

Chart 8

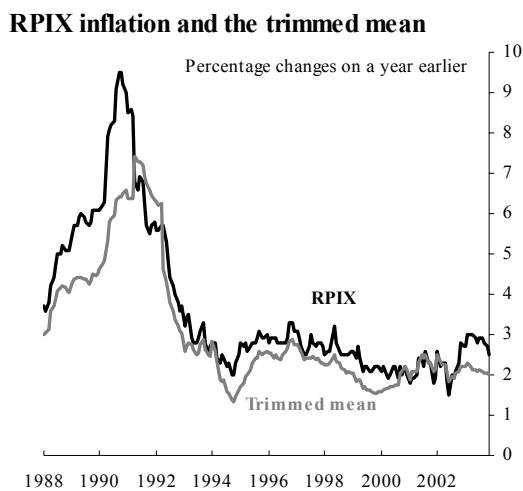
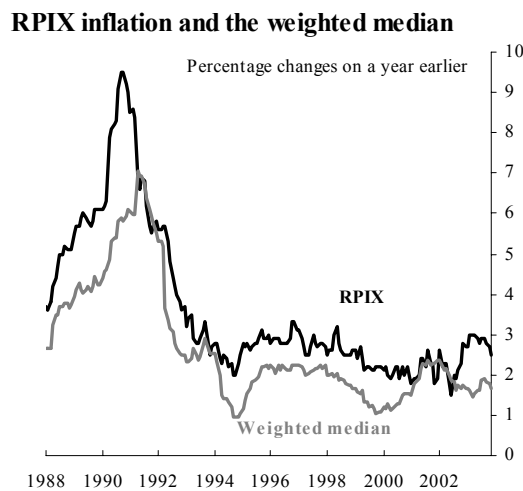


Chart 9



To justify using the trimmed mean as an informative measure of core inflation in an economic sense requires a model of the economy in which prices are sticky. To see why, consider an economy with fully flexible prices. As explained in Section 4 (ii), in such an economy, and with an unchanged monetary policy stance, a shock to a particular sector would lead to instantaneous changes in *relative* prices, which would leave the *aggregate* price level, and therefore the

⁽⁵⁾ At the Bank of England, one version of a trimmed mean inflation rate is calculated as follows. First, one-month percentage changes of the 81 subcomponents of the RPI are calculated. They are then ordered according to their weight, giving a string of 1,000 - *n* numbers, where *n* is the current weight on mortgage interest payments (MIPs). Second, these 1,000 - *n* numbers are sorted into ascending order. Third, the smallest 15% and largest 15% from these 1,000 - *n* numbers are excluded. Fourth, an average is taken over the remaining 0.7 * (1,000 - *n*) numbers. This gives the one-month change in the 70% trimmed mean of RPIX. This series of one-month percentage changes is used to create an index from which annual inflation rates can be calculated.

aggregate inflation rate, unchanged. In this type of world, it would make no sense to trim out large price movements (see Zeldes (1994)).

In the real world, of course, prices are *not* completely flexible and relative price changes can affect the inflation rate, at least in the short to medium run. The trimmed mean does not require *a priori* judgment concerning which components to include or exclude permanently. Rather, components' price changes are included or excluded on the basis of their relative magnitudes. The trimmed mean's ability to exclude relative price movements, but retain those price movements associated with say aggregate demand shocks, depends on the former being at the extremes of the price distribution. This implies that relative price movements must be generally larger in absolute magnitude than those price changes associated with aggregate demand shocks.

One UK case where trimming might have been appropriate was the outbreak of foot-and-mouth disease in 2001. A restriction on domestic meat supplies and costly imported substitute supplies led to a sharp rise in the retail prices of directly affected meat (eg pork, beef and lamb). These rises were unlikely to be related to underlying inflation because the *source* of the shock was known—a supply shock, impacting primarily on one sector of the economy. In this case, trimming out these sharp price increases might have provided a better indication of underlying inflation in that particular month. But the question then arises of how those meat prices and other prices adjust back to their equilibrium level over following months. If these subsequent relative price adjustments are not large enough to qualify for trimming, they would be included in the trimmed mean in the next few periods.

Though in this example trimming might have been helpful, there are examples where trimming would unambiguously misinform policymakers. For example, take an aggregate demand shock, such as an exogenous increase in world demand, which raises all firms' 'desired' prices. Say only a few firms change their prices in the first period, while the other firms leave their prices unchanged. As noted by Bakhshi and Yates (1999), trimming out the few price rises would yield a zero trimmed mean inflation rate, giving a misleading picture of underlying inflation. In this case, the information in the tails of the price distribution would be of more use to monetary policy makers than that in the centre of the distribution. Thus, knowing the source of the shock is crucial in determining whether it is wise to trim.

An informal way of gauging the usefulness of the trimmed mean is to look at the frequency with which price changes of each of the components of RPIX are excluded in the calculation of the 15% symmetric trimmed mean in the United Kingdom. Of the 21 components which are excluded more than 50% of the time between 1975 and 2002, five are seasonal food components, three are non-seasonal food and two are energy. Of the other eleven, four are components whose prices are regularly heavily discounted in the January and summer sales. It may therefore be sensible to exclude their price movements in those months. This limited evidence does at least suggest that the trimmed mean in the United Kingdom has predominantly excluded those items that are most subject to shocks affecting particular sectors and to short-term volatility. Appendix B provides more details.

An advantage of the trimmed mean is that it is timely and can be easily computed (so people outside the central bank can easily verify the measure). But overall, given other concerns highlighted above, it is unlikely that one would want to place much weight on the inflationary signals given by trimmed mean. There is still a large degree of judgment needed, for example on how much of the distribution of price changes should be trimmed.⁽⁶⁾ Some have decided this by considering how well measures of differing degrees of ‘trim’ approximate a particular ‘reference measure’, with the 37-month centred moving average of headline inflation being a popular benchmark. The difficulty with this is that one has no idea of whether the benchmark is sensible. One argument for using such a benchmark is that it is ‘smooth’. But if underlying aggregate demand shocks hitting the economy are not smooth, and/or the transmission of the effects of these shocks onto prices is changing, then a measure of core inflation would not be expected to be smooth either.⁽⁷⁾ The issue of the ‘smoothness’ of core inflation is returned to in the section on model-based measures.

(ii) *Measures based on ‘exclusion’*

Some measures of core inflation are derived by permanently excluding certain components from the price index, *a priori*. In the case of the RPI in the United Kingdom, there are two prominent examples of items which are permanently excluded. First, mortgage interest payments (MIPs) are excluded from the all-items retail price index to give RPIX, the former target measure. MIPs were excluded from the targeted measure, since otherwise changes in interest rates would have, at least in the short run, perverse effects on the targeted inflation rate. The second prominent measure of this kind is RPIY, which also excludes all indirect taxes.⁽⁸⁾ These exclusions may be useful for monetary policy purposes. For although indirect taxes are important components of a cost-of-living index, they do not constitute ‘core’ inflation under most definitions of that term.

Other components are often excluded on the grounds that their prices are considered to be too volatile—adding ‘noise’ to the measured inflation rate—and obscure the signal of underlying demand pressures in aggregate inflation. Seasonal food and energy prices are often excluded on this basis. Two examples of such measures for the United Kingdom are shown in Charts 10 and 11.

The case for excluding seasonal food prices is clearest. Since their supply is heavily influenced by changes in weather conditions, and given their relatively low elasticity of demand, shifts in supply can cause relatively large changes in prices and consequently in aggregate inflation. The argument for excluding energy prices is less clear cut. To the extent that energy prices are driven by temporary global oil supply conditions, this may be a valid reason for exclusion. But, it is

⁽⁶⁾ There are also some issues regarding the precise procedure of how to trim. Should monthly or annual inflation rates be trimmed and should seasonally adjusted or non seasonally adjusted price data be used? The trimmed mean measure described in footnote 3 is based on monthly, non seasonally adjusted price data from which an index is calculated for the basis of the annual calculation.

⁽⁷⁾ There is a question whether the root mean squared error (RMSE) or the mean absolute deviation (MAD) should be minimised. That is important since the results can be sensitive to the choice between the two (see Bakhshi and Yates (1999)).

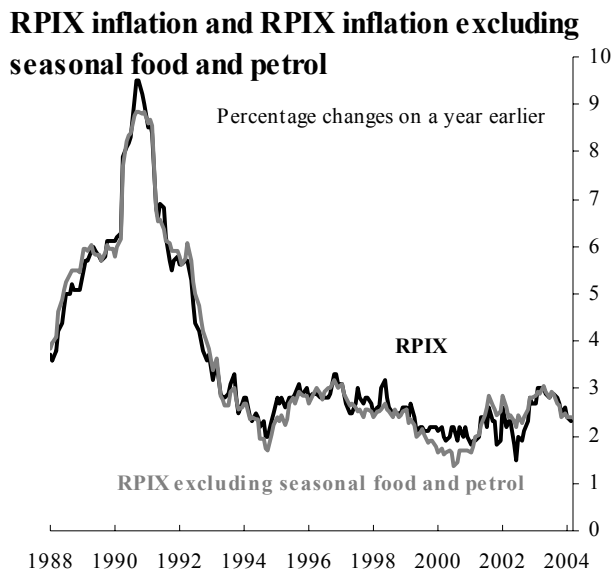
⁽⁸⁾ Stripping out the effects of indirect taxes from consumer prices is not straightforward, since it involves making behavioural assumptions about the extent to which duty changes are passed on to consumers. For a description of how RPIY in the United Kingdom is constructed see Beaton and Fisher (1995).

likely that more persistent global demand conditions will also have a significant influence on the prices of these commodities, implying that energy prices might contain useful information about underlying inflation.

On the other hand, the results from the previous section showed that seasonal food and energy prices are two of the components that are excluded from the trimmed mean in the majority of months. This suggests that these are reasonable items to exclude if one wanted to strip out the most volatile components.

Like the trimmed mean, the advantage of measures based on excluding components is that they are timely and easy to calculate and explain. Also the composition of the underlying basket is unchanged in each period, unlike the trimmed mean, so one is comparing like with like over time. However, their downside is that they require a once-and-for-all (subjective) judgment about the least informative price components for estimating core inflation. It does seem a little unrealistic to assume that some components' price changes contain no information at all for the measurement of core inflation.⁽⁹⁾ And in a sense, these types of measure add nothing to the information set of monetary policy makers. They are just another way of representing certain components' contribution to the annual aggregate inflation rate, which are monitored as a matter of routine in the Bank.

Chart 10



(iii) Measures based on the whole price distribution

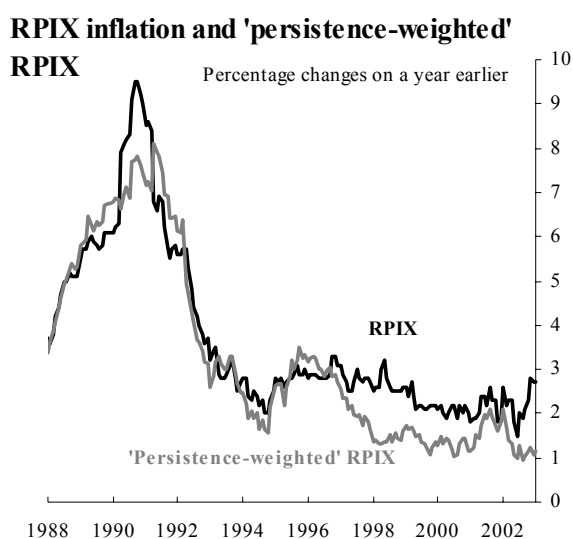
Other measures of core inflation use all available (disaggregated) prices information from the consumer price index. One such approach is to reweight the disaggregated price indices to maximise the 'signal' in the data, however that might be defined. For instance, sectors in which supply conditions are believed to be relatively important in determining prices might have their

⁽⁹⁾ Some countries have also made specific adjustments to cope with particular price shocks considered to be of a one-off nature and so will have only a temporary impact on the measured inflation rate: eg in New Zealand for the impact of international trade price shocks and in Sweden for exceptional exchange rate movements.

weight reduced, whereas prices in the remaining sectors would be assigned higher weights. Some authors have argued for components to be weighted according to the inverse of their volatility.

Blinder (1997) identifies 'core' inflation with the 'durable' part of inflation. In trying to estimate this component, he advocates constructing an index by weighting together individual price changes 'according to their usefulness in forecasting future inflation'. This idea is operationalised for the United Kingdom by Cutler (2001), who reweights the components of RPIX according to the 'persistence' of their annual inflation rates. The weights are obtained by estimating coefficients in a first-order autoregressive model for each component of RPIX in order to derive a 'persistence-weighted' RPIX measure (ie components with a more 'persistent' annual inflation rate are given a higher weight). This measure is shown in Chart 11.

Chart 11



Bryan and Cecchetti (1993) adopt an alternative approach based on dynamic factor analysis. They assume that individual inflation series share a component that is subject to common disturbances. The disturbance to the common inflation component is assumed to be uncorrelated with idiosyncratic (or relative) price shocks, either contemporaneously or serially, at all leads and lags. In the core inflation measure, prices are weighted according to their determination by common, as opposed to idiosyncratic, shocks rather than expenditure weights. Underlying this particular approach is the view that relative price changes are driven primarily by supply disturbances that are uncorrelated with the persistent or general tendency of inflation.

One concern with the reliability of measures based purely on statistical criteria is that they may be more vulnerable to the Lucas critique. For example, in 'persistence-weighted' RPIX, the coefficients in component price autoregressions will depend in part on past policy. If future policy were to factor such weights into its decisions, the weights would change, and the measure would become misleading. Problems with the stability of these types of measures would be more acute when the economy is undergoing significant structural change and, as in the United

Kingdom, when the definitions and classifications of the subcomponents of the RPI change.⁽¹⁰⁾ Another more general problem with any particular reweighted price index is that its inflation rate can have a different trend to that of the target measure, depending on the relative trends in the individual reweighted price series. If so, these types of core measure will exclude not only temporary disturbances to inflation but also a part of trend inflation.⁽¹¹⁾

(iv) *Model-based measures*

Model-based measures are attractive in that they are multivariate and use econometric techniques, in which some structure is imposed explicitly, grounded in economic theory. They typically derive measures of core inflation from aggregate inflation data and tend to rely on some prior belief about the time-series properties of core inflation—for example, how cyclical the measures should be. The difficulty with discriminating between them is that they are all based on slightly different definitions.

Eckstein (1981) is commonly attributed with the original definition, in which core inflation is identified as: ‘...the trend increase of the cost of the factors of production’. This ‘...originates in the long-term expectations of inflation in the minds of households and businesses, in the contractual arrangements which sustains the wage-price momentum, and in the tax system’.

For expositional purposes and following Roger (1998), suppose the short-run aggregate supply curve is given by:

$$\pi_t = [\pi_t^{LR} + g(x_{t+1}) + v_t]$$

where:

- π_t is the aggregate inflation rate in period t
- π_t^{LR} is the long-run or trend inflation rate
- x_{t-1} is a measure of cyclical excess demand pressure
- v_t is a measure of transient disturbances to inflation.

Then Eckstein’s notion of core inflation is given by:

$$\pi_t^c = [\pi_t - g(x_{t-1}) - v_t] = \pi_t^{LR}$$

This is the long-run component, while non-core inflation is given by:

$$\pi_t^{nc} = g(x_{t-1}) + v_t$$

Under this definition, core inflation is not affected by any cyclical influences and so should *not* be cyclical.

⁽¹⁰⁾ Redefinition of price series, through reweighting at low levels of aggregation, recategorisation of particular prices, or the addition/removal of various prices, means that the time-series properties of particular RPI components may change markedly.

⁽¹¹⁾ Treatment of ‘non-market’ prices, such as utility prices, is also problematic. These prices show persistent, non-cyclical trends together with infrequent (typically annual) jumps.

The definition used by Quah and Vahey (1995) is that core inflation is ‘...that component of measured inflation that has no medium to long-term impact on real output’.⁽¹²⁾ In their model with a vertical long-run Phillips curve, these are aggregate demand shocks and inflation is neutral in its effects on the real economy in the long run. The remainder is the part of inflation caused by shocks that have a permanent effect on output (ie permanent aggregate supply shocks). Quah and Vahey estimated a structural vector autoregressive (SVAR) model containing RPI inflation and output, on which they impose long-run identifying restrictions. Their measure of core inflation is close to a notion of persistent inflation. Again, following Roger (1998), the Quah/Vahey definition of core inflation is given by the following expression:

$$\pi_t^c = [\pi_t - v_t] = \pi_t^{LR} + g(x_{t-1})$$

This corresponds to the long-run component plus any cyclical movements associated with excess demand. So non-core inflation is simply that part that is due to the temporary disturbances:

$$\pi_t^{nc} = v_t$$

The definition of core inflation hinges on how one defines ‘medium to long run’. Quah and Vahey are trying to capture inflationary pressures that feed into or reflect inflation expectations. The non-core element is essentially unanticipated inflation—and this is the component of measured inflation that does have a medium to long-run impact on output.⁽¹³⁾

The two definitions seem to differ according to the effect of cyclical influences on core inflation. Under Eckstein’s definition, core inflation should not be cyclical; whereas, using Quah/Vahey’s definition, core inflation should be strongly correlated with output in the short run. Roger suggests we should not over-do the differences: the difference between a transient influence on inflation (v_t) and cyclical ($g(x_{t-1})$) and long-term influences (π_t^{LR}) is an artificial construct. This distinction should really be drawn in reference to the policymaker’s horizon. If the policymaker is focusing on the medium run, then the Quah/Vahey definition is appealing. If the policy horizon is longer, then Eckstein’s definition may be more relevant.

One attraction of the model-based approach is that the measures are more deeply based on economic theory. They also benefit from being the product of multivariate analysis, in that they use non-price variables in calculating core inflation. This seems more likely to produce a superior ‘economic’ measure of core inflation than one produced from univariate analysis.

The downside, however, is that the restrictions imposed are rarely uncontroversial. These models are also sensitive to their exact specification and identification scheme. For example, Folkertsma and Hubrich (2000) suggest that at least five different SVAR models have been proposed in the literature. They examine how reliably these various identifications recover the ‘true’ inflation

⁽¹²⁾ A shock that raises output permanently (and so raises actual and potential output) is assumed to have no long-run effect on inflation. Note one drawback of the Quah and Vahey model, overcome in other models, is that core inflation is only identified up to a constant—the *level* is undetermined.

⁽¹³⁾ One question this raises is what the level of core inflation is, since in Quah/Vahey this is not determined – their VAR consists of just output and inflation so there is no nominal anchor. Blix (1995) adds money to the Quah/Vahey two-variable VAR. In this case, the system is identified by assuming that changes in the level of the money stock, rather than changes in the growth rate of money, are output neutral in the long run. One of the problems with the SVAR approaches is that the precise nature of identifying restrictions and the data used will affect the estimates.

process, which they define as core inflation implied by their monetary general equilibrium model. As they point out, the validity or relevance of the experiment crucially depends on the empirical realism of their general equilibrium model. But the interesting point is that they find that ‘none of the schemes were able to produce a core inflation estimate that can be considered sufficiently accurate to qualify for monetary policy purposes’. They argue that this is not particularly surprising, given—as other authors have pointed out—the well-known problems associated with SVAR models. That said, this non-robustness to the precise specification of the model is a limitation to their practical and routine use by policymakers.⁽¹⁴⁾

Evaluating how well model-based measures isolate that part of measured inflation relating to underlying aggregate demand pressures is virtually impossible and will depend heavily on one’s prior beliefs about the characteristics that core inflation should display. In particular, many authors expect core inflation to be relatively ‘smooth’. But it is not self-evident that smoothness *per se* is a desirable property for a core inflation measure. If underlying demand and supply shocks are not smooth or persistent then presumably neither will core inflation.

Part of the difficulty on agreeing whether smoothness is a sensible property is that there is no agreed theory of inflation. For example, there are competing hypotheses concerning the observed skewness in the cross-sectional distribution of prices. This could be generated equally plausibly by sluggish adjustment of prices, (eg within a New Keynesian framework as in Ball and Mankiw (1995) or by a skewed distribution of shocks (as in Balke and Wynne (2000)). And it may well be that the ‘smoothness’ which we observe may simply be a product of monetary policy itself. This sort of consideration means that these types of measures are also not immune to Lucas critique type concerns.

There are parallels with Kalman filter estimates of the non-accelerating inflation rate of unemployment (NAIRU), over which there has also been debate concerning how smooth these estimates should be. Like core inflation, the NAIRU is unobservable. Nevertheless, the NAIRU can help to explain slow-moving structural changes in the labour market.

⁽¹⁴⁾ For example, Faust and Leeper (1997) question the reliability of SVAR-based measures to capture long-run effects sufficiently well in short samples.

Chart 12

RPIX inflation and the Quah-Vahey measure of core inflation

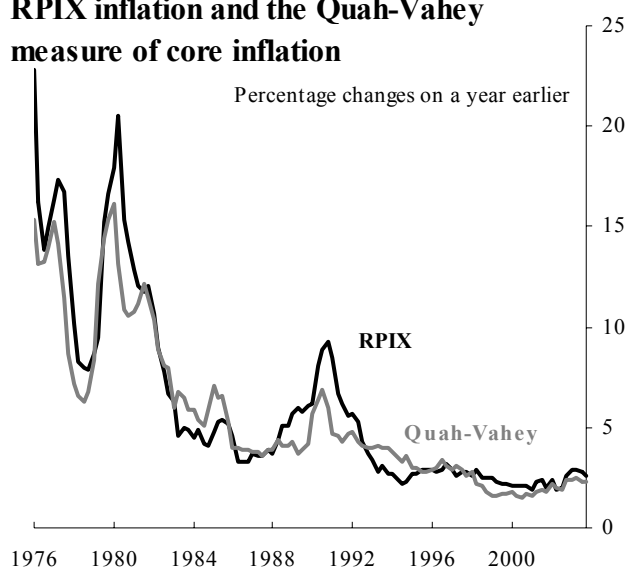


Chart 12 shows the Quah and Vahey measure of core inflation⁽¹⁵⁾ together with RPIX inflation. During the oil price shocks in 1979, 1990 and 1999/2000, RPIX inflation peaked at a higher rate than the Quah and Vahey measure of core inflation. Equally significantly, the converse is true for the sharp fall in the oil price in 1986. This is what we would expect if those oil price rises were driven by supply shocks, after which we would expect that other relative prices would adjust downwards. Interestingly, RPIX inflation peaked later than this measure of core inflation on both occasions. This perhaps suggests that the Quah and Vahey measure of core inflation might lead targeted inflation.

(v) *Domestically generated inflation*

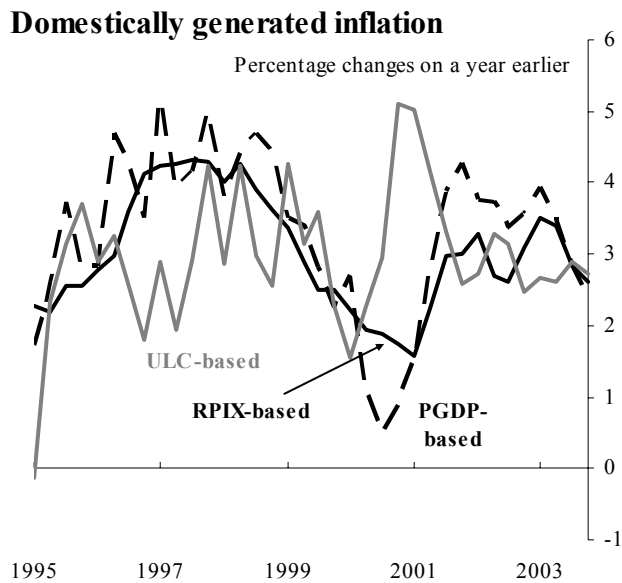
Domestically generated inflation (DGI) may be viewed as a particular type of measure of core inflation that aims to exclude the one-off price level effects of movements in the exchange rate. Since RPIX inflation is a weighted average of DGI and imported inflation, DGI may be useful in providing information on the pressure being exerted on prices by domestic conditions. The effects of an external shock on actual inflation will be temporary, though it may be hard to assess the extent and duration of such effects. Once the effects have worked through the economy, inflation will revert to DGI. If DGI had strong inertia then it would be a leading indicator of actual inflation during an external shock.

There is no unique definition of DGI, and so no single way of measuring it. It could be model or statistically based. At the Bank of England, three measures of DGI have been constructed and monitored: the GDP deflator excluding export prices; RPIX excluding import prices; and a measure based on unit labour costs. These are shown in Chart 13. Even if the conceptual case for DGI is attractive, there are some practical concerns about how well the measures of DGI achieve their objective. In particular, the measures are sensitive to the precise assumptions

⁽¹⁵⁾ Estimated using quarterly RPIX and GDP data, sample 1975-2000.

underlying their construction. And perhaps most worryingly, the measures have at times shown very different trends.

Chart 13



5 Evaluating measures of core inflation

There are several ways in which we might assess the usefulness of measures of core inflation. For example, several authors have put forward properties that they believe measures of core inflation should ideally possess. Roger (1998), for instance, suggests that measures of core inflation should be timely, credible, verifiable and easily understood to the public. In addition to these, Wynne (1999) argues that measures of core inflation should be computable in real time, forward-looking in some sense, have a track record and have an economic theoretical basis.

Although it may be helpful for measures of core inflation to possess some of these properties, a more useful method of evaluation is to assess how well the measures achieve what they were constructed to do. As already highlighted, one potential use of measures of core inflation is to provide information on the outlook for inflation. Below, we use cointegration analysis to try to determine which measures of core inflation in the United Kingdom are most informative about the future short-term path of annual RPIX inflation.

If a measure of core inflation does not cointegrate with RPIX inflation, then the two series will diverge over time, meaning that the long-run level of that measure of core inflation will not be informative about the future level of RPIX inflation.⁽¹⁶⁾ At the same time, the presence of cointegration does not eliminate the possibility that the two may diverge for considerable periods of time. If the period of adjustment is longer than the policymaker's horizon, typically one to two years, then cointegration itself is not sufficient to render a measure of core inflation useful.

⁽¹⁶⁾ Cointegration techniques should only be applicable to series that are I(1). The use of cointegration tests to evaluate measures of core inflation is valid, at least statistically, because RPIX inflation and the measures themselves are found to be I(1) in standard unit root tests. The finding that RPIX inflation and the various measures of core inflation are not I(0) is not surprising given that inflation has fallen over the sample of the tests.

The next section sets out some tests proposed by Marques *et al* (2000) which they argue measures of core inflation should satisfy, if they are to be useful in providing forward-looking information about the targeted rate of inflation. Like the measures of core inflation themselves, the tests are not without their problems, as discussed below. The following section applies the tests to the available measures of core inflation in the United Kingdom, before drawing inferences from the results.

(i) *Tests proposed by Marques et al (2000)*

Marques *et al* (2000) propose the following testable conditions when the targeted and candidate core inflation rates are found to be non-stationary:

- (i) Targeted (π_t) and core inflation (π_t^*) should be cointegrated with unit coefficient.
- (ii) Core inflation should be an ‘attractor’ of targeted inflation.
- (iii) Targeted inflation should not be an ‘attractor’ of core inflation (ie core inflation should be weakly exogenous).⁽¹⁷⁾

The attraction of the tests is that they attempt to formalise the relationship between targeted and core inflation by exploiting information contained in the differential between the two. The conditions essentially imply that the targeted rate of inflation should converge to core inflation in the long run, but not *vice versa*. The first condition ensures that core and the targeted rate of inflation move one-for-one in the long run, and that the impact of relative price shocks on the targeted inflation rate should have a zero mean once all relative prices have adjusted. A unit coefficient on core inflation ensures that targeted and core inflation do not display a permanently diverging trend. If this were not the case, it would suggest that the measure of core inflation is not fully capturing some part of the trend rate of inflation. Also, it would make it harder for the central bank to use the measure of core inflation in its communication of its actions to the public.

The second condition formalises the assumption that the targeted rate of inflation converges to core inflation in the long run; in other words, core inflation should be an ‘attractor’ of the targeted rate of inflation. If condition (ii) holds, then when π_t is above (below) π_t^* , π_t will at some point decrease (increase) and converge to π_t^* . The third condition says that core inflation should not converge to targeted inflation (ie core inflation should be weakly exogenous). Marques *et al* (2000) argue that if targeted inflation was an ‘attractor’ of core inflation, it would be more difficult to infer anything about the future path of targeted inflation by looking at core inflation, as the relationship would run both ways.

(ii) *How do measures of core inflation in the United Kingdom perform in the tests?*

The key results for a range of measures of core inflation are shown in Table A. Full details are provided in Appendix C.

⁽¹⁷⁾ The third test put forward by Marques *et al* also demands that targeted inflation does not Granger cause core inflation, so that, strictly, their test is for strong, rather than weak, exogeneity.

Table A: Replicating Marques *et al* tests for core inflation measures in the United Kingdom

	Condition (i)		Condition (ii)	Condition (iii)
	$(\pi_t - \pi_t^*)$ stationary	...and mean zero	Core inflation (π_t^*) should be an 'attractor' of target (π_t)	Target inflation (π_t) should not be an attractor of core inflation (π_t^*)
RPIX excl. food	✓	✗	✓	✗
RPIX excl. seasonal food	✓	✗	✓	✗
RPIX excl. food and fuel	✓	✗	✗	✗
RPIX excl. food, alcohol, tobacco and petrol	✓	✓	✓	✓
RPIX excl. seasonal food and petrol	✓	✓	✓	✓
Trimmed mean	✓	✗	✗	✗
Weighted median	✓	✗	✗	✗
'Persistence-weighted' RPIX	✓	✗	✓	✓
'Quah and Vahey' measure	✓	✗	✓	✓
RPIY	✓	✗	✓	✓
DGI: RPIX excl. import prices	✓	✗	✗	✗
DGI: ULC measure	✓	✓	✓	✓
DGI: GDP deflator excl. export prices	✓	✗	✓	✗

The results are mixed. Only three of the measures of core inflation pass all three tests: RPIX excluding seasonal food and petrol, RPIX excluding food, alcohol, tobacco and petrol, and the DGI measure based on ULC. On the face of it, the results suggest that these three are potentially the most useful measures of core inflation. However, care needs to be taken in interpreting the results. For a start, there are some problems with the tests.

In particular, the regressions in the tests are essentially reduced-form representations of the inflation process and the results will therefore be affected by past monetary policy. The following argument highlights the problem. Suppose the target for monetary policy was to keep annual CPI inflation to some prescribed path, for example 2% at all times. And suppose that, over the sample period, policy had been used actively, and set optimally to achieve the target. Then, CPI inflation would simply follow the prescribed path, save perhaps for some unavoidable and unforecastable error. If we were to perform Marques *et al*'s tests on a measure of core inflation, it would fail conditions (ii) and (iii). That is, CPI inflation would not be attracted to the measure of core inflation since it follows the exogenously prescribed path, but core inflation would be attracted to CPI inflation.⁽¹⁸⁾ This finding would cause us to reject this measure of core inflation as useful in providing forward-looking information about the future path of CPI inflation, even though it might well be useful in setting policy. Thus, failure in the tests does not necessarily mean that a measure of core inflation is not informative—it may just be that the effects of past policy mean that Marques *et al*'s tests do not help us make that judgment. In addition, because the differential between targeted and core inflation is likely to be some function of the stance of monetary policy, at least in the short run, the tests may be vulnerable to the Lucas critique. That is, if policy were to be based on some estimated relationship between core and targeted inflation, that relationship may change and become misleading as a guide to the future.

⁽¹⁸⁾ The Granger representation theorem implies that if two series are cointegrated, then one of them at most is weakly exogenous.

The tests put forward by Marques *et al* seem attractive and may be indicative of the relative usefulness of different measures of core inflation. However, the problems with the tests outlined above mean that the results are in no way conclusive, like the measures themselves.

So how useful are measures of core inflation? Bearing in mind what information each type of indicator is best at providing, it can be valuable to look at a range of measures. Measures of core inflation can then provide a different perspective on the inflationary process in the context of the other variables that policymakers monitor.

6 Conclusion

When policymakers see a change in measured inflation, a key question is how much news is there for the outlook for inflation. Does it reflect movements in the fundamental determinants of inflation? How persistent is the change likely to be? Measures of core inflation are potentially useful in answering these questions, but as summary statistics, they are no substitute for understanding the sources of shocks affecting the economy and how they are likely to evolve over the future. Moreover, the large number of available measures, based on a wide range of different conceptual bases, is potentially confusing.

A compromise conclusion is that each one can provide a different insight into the inflation process. This paper has found that no one measure performs well across the board, but there can be value in looking at a range of measures, as long as one clearly bears in mind what information each type of indicator is best at providing. When all measures are giving the same message then, in a sense, monetary policy makers can reasonably consider that they are providing a reliable guide to inflationary pressures. It is when the measures start to diverge that they need to take a much closer look at the reasons for those divergences.

Appendix A: Core inflation as a statistically robust measure of actual inflation

The precise method used to obtain a better estimate of the population mean will depend in large part on the *shape* of the population distribution of price changes. If the population distribution is approximately Normal, then the mean of samples taken from that distribution will be the best estimator of the population mean (in the sense of being efficient and unbiased). The stylised fact, however, is that the sample distribution of price changes appears to be non-Normal, which may indicate that the population distribution is likewise non-Normal. There are two important dimensions along which this is the case, both with different implications for the ‘statistically best’ estimator of the population mean. First, the population sample may be ‘fat-tailed’ or leptokurtotic. Second, the distribution may be asymmetric or skewed. The statistical response to the first issue is to trim symmetrically; the response to the second is to trim asymmetrically. For a recent discussion of these issues in the UK context, see Andrade and O’Brien (2001).

(a) *Dealing with leptokurtosis*

When the population distribution is fat-tailed, there is a greater probability than in a Normal distribution that a price change will be sampled from the tails. Though the sample mean will not necessarily be *biased*, it will be *inefficient*. In this case, ‘limited-influence’ estimators such as the trimmed mean, which downplay the influence of the tails, can provide a more efficient estimate of the central tendency of price changes than the sample mean. These estimators vary according to how much of the distribution is trimmed, though for symmetric distributions the trim should also be symmetric. At one extreme, the weighted median trims 50% of the distribution on either side.⁽¹⁹⁾ For symmetric leptokurtotic distributions, the robustness and efficiency of alternative estimators largely depends on the degree of kurtosis of the distribution.⁽²⁰⁾

Economic theory provides no direct guide on the optimal trim. Rather, this is a purely statistical issue. One commonly used method of trying to identify the optimal trim is by examining how well trimmed measures based on different degrees of trimming approximate some ‘reference measure’. Bakhshi and Yates (1999), for example, use the 37-month moving average of RPIX inflation as their benchmark.⁽²¹⁾ On that basis, they tentatively concluded that it was optimal to trim 15% from each tail, though they note that there are clearly issues regarding the suitability of the reference measure. Also ‘[t]he trimming point seems unduly sensitive to the exact method used’, as noted by Andrade and O’Brien (2001), who suggest that ‘... we are some way from a clear and undisputed measure of core inflation’.

(b) *Dealing with skewness*

When a distribution is skewed or asymmetric, estimates of central tendency that ignore this will be biased. In this case asymmetric trimming has been suggested. For a positively skewed distribution, the mean will correspond to a percentile between 50 and 100. If this percentile can be determined, and we assume that this holds for the population distribution also, then the sample

⁽¹⁹⁾ The simple mean is at the other extreme, where none of the distribution has been trimmed.

⁽²⁰⁾ See Roger (1997) for a useful overview of the literature. He draws the distinction between robustness and efficiency of different estimators. As he puts it, ‘a robust estimator may not be the most efficient estimator, but will rarely perform poorly’.

⁽²¹⁾ In this, they follow Bryan and Cecchetti (1993) who choose the 37-month moving average on the grounds that this is a suitably long time horizon over which relative prices will have adjusted to shocks.

value of that percentile can be considered to be the best estimator of the population mean. Roger (1997), for example, found that the distribution of measured quarterly price changes in New Zealand between 1949-96 was both chronically positively skewed *and* highly leptokurtotic. He proposed the 57th percentile as a good measure of core inflation (that is, a good measure of the true population mean) since it was unbiased and robust in the face of the leptokurtosis *and* positive skewness in the New Zealand data. This is essentially a case of trimming asymmetrically.

In the United Kingdom, there is evidence that the distribution of monthly retail price changes displays both high kurtosis and positive skewness (see Charts A1-4). This is the case whether we are looking at the distribution of monthly (see Charts A1 and A2) or annual price changes (see Charts A3 and A4).⁽²²⁾ By itself, excess kurtosis would support using a symmetric trimmed mean, but the evidence of prolonged positive skewness in the cross-sectional distribution of price changes would favour asymmetric trimming. More generally, the charts suggest that the moments of the price distribution are changing, which might suggest that the optimal trim—both in terms of how much of each tail is trimmed and the degree of asymmetry—is time varying.

⁽²²⁾ One would expect the statistical properties of price changes to differ according to the length of the period over which the price change is calculated (eg quarterly changes will tend to include offsetting monthly movements in volatile components and so will tend to be smoother).

Chart A1

Moments of weighted monthly RPIX inflation

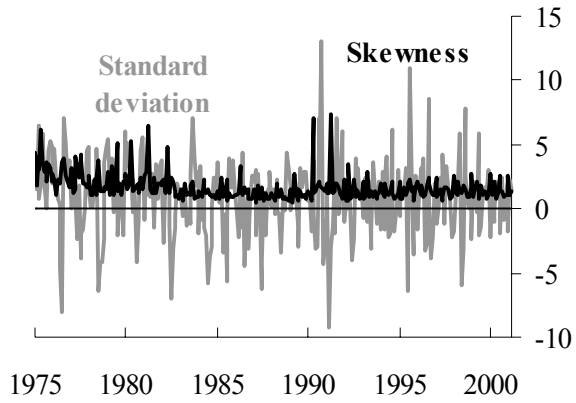


Chart A2

Kurtosis of weighted monthly RPIX inflation

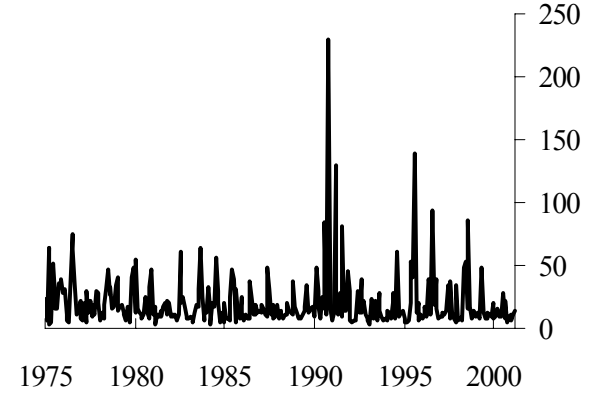


Chart A3

Moments of weighted annual RPIX inflation

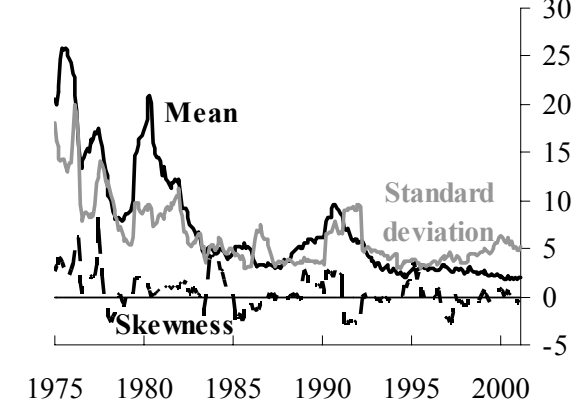
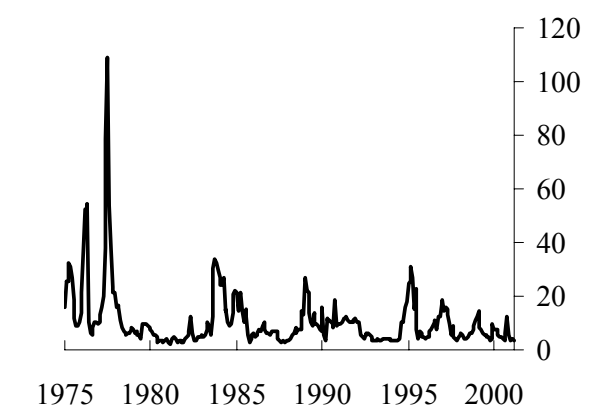


Chart A4

Kurtosis of weighted annual RPIX inflation



Appendix B: Frequency of RPIX components' exclusion from the 15% symmetric trimmed mean

% on months excluded from the trimmed mean

Housing depreciation	87.8	Bread	31.3
Fresh vegetables	86.9	Other clothing	30.3
Poultry	81.6	Sugar	29.4
Potatoes	78.1	Electrical appliances	29.4
UK holidays	77.5	Garden products	29.4
Lamb	75.6	Electricity	26.6
Eggs	75.3	Cereals	25.6
Foreign holidays	74.4	Fees and subscriptions	25.6
Fruit	69.1	Biscuits	25.3
Other meat	64.7	Other equipment	25.0
Insurance and ground rent	64.4	Other foods	23.4
CDs and tapes	63.8	Other travel	23.4
Personal articles	61.6	Telephones	22.5
Petrol and oil	60.3	Vehicle tax and insurance	21.6
Pork	59.4	Other tobacco	21.3
Women's clothing	58.4	Council tax	21.3
Oil and other fuel	55.3	Sweets and chocolates	20.9
Toys and sports goods	55.0	Wine	20.6
Take aways	51.6	Cigarettes	20.6
Tea	50.6	Maintenance of motorcars	20.6
Restaurant meals	50.6	Books and newspapers	20.0
Coffee	49.7	Milk	19.1
Furnishings	46.9	Chemist goods	18.8
Oil and fat	46.6	Bus and coach fares	17.8
Beef	46.3	Gas	17.2
Bacon	46.3	Water	16.9
Men's clothing	46.3	Entertainment and other recreation	16.6
clothing children	45.9	Rent	16.3
Audio visual equipment	45.6	Repairs and maintenance	15.9
Butter	45.3	DIY	15.9
Furniture	45.3	Rail fares	15.6
Fish	45.0	Personal services	14.7
Milk products	39.1	Consumables	14.4
Cheese	36.6	Canteen meals	11.6
Coal	35.6	Beer	11.3
Pet care	35.0	TV licences	10.6
Footwear	34.1	Postage	10.3
Soft drink	33.8	Domestic services	10.3
(Used) cars	32.5		

Appendix C: Marques *et al*'s (2000) tests for measures of core inflation

Marques *et al* (2000) put forward the following necessary conditions that measures of core inflation, π^* , should meet for them to be useful indicators of the future path of target inflation.⁽²³⁾

(i) Target and core inflation should be cointegrated with unit coefficient (ie the difference between the two inflation rates should be a stationary series with mean zero). This is tested in two stages. First, we perform an ADF test to establish the stationarity of $(\pi_t - \pi_t^*)$. Second, given that $(\pi_t - \pi_t^*)$ is stationary, we test the null hypothesis that $\alpha = 0$ (ie that the series $(\pi_t - \pi_t^*)$ is mean zero) in the static regression:

$$(\pi_t - \pi_t^*) = \alpha_t + \varepsilon_t \quad (\text{C1})$$

(ii) Core inflation should be an 'attractor' of target inflation. To test this condition we estimate the following error-correction model:

$$\Delta\pi_t = \sum_i \alpha_i \Delta\pi_{t-i} + \sum_i \beta_i \Delta\pi_{t-i}^* - \gamma(\pi_{t-1} - \pi_{t-1}^*) + \varepsilon_t \quad (\text{C2})$$

The hypothesis that $\gamma = 0$ is then tested using conventional t -statistics. The rejection of the hypothesis suggests that π^* is an 'attractor' of π , and that target inflation does converge to core inflation in the long run.

(iii) Target inflation should not be an 'attractor' of core inflation (ie core inflation should be weakly exogenous). To test this condition we estimate the following error-correction model:

$$\Delta\pi_t^* = \sum_i \delta_i \Delta\pi_{t-j}^* + \sum_i \theta_i \Delta\pi_{t-j} - \lambda(\pi_{t-1} - \pi_{t-1}^*) + \varepsilon_t \quad (\text{C3})$$

The hypothesis that $\lambda = 0$ is then tested using conventional t -statistics. The failure to reject the hypothesis suggests that π is not an 'attractor' of π^* , and that core inflation does not converge to target inflation in the long run.

⁽²³⁾ A problem with this analysis is that we are looking at the relationship between two very imperfect measures of inflation, namely the targeted inflation rate and the various measures of core inflation.

Table C1 – Results for the United Kingdom of Marques *et al*'s tests

Column	Stationarity of $(\pi_t - \pi_t^*)$ $\alpha = 0$ given			
	1	2	3	4
RPIX excluding seasonal food	Yes (-3.11)	No (0.00)	No (0.00)	No (0.03)
RPIX excluding food	Yes (-2.89)	No (0.00)	No (0.09)	No (0.00)
RPIX excl. seas. food and petrol	Yes (-2.75)	Yes (0.22)	No (0.00)	Yes (0.35)
RPIX excluding food and fuel	Yes (-2.49)	No (0.00)	Yes (0.13)	No (0.00)
RPIX excluding food, alcohol, tobacco and petrol	Yes (-3.51)	Yes (0.27)	No (0.00)	Yes (0.10)
Trimmed mean	Yes (-2.27)	No (0.00)	Yes (0.12)	No (0.00)
Weighted median	Yes (-1.89)	No (0.00)	Yes (0.11)	No (0.00)
'Persistence-weighted' RPIX	Yes (-2.48)	No (0.00)	No (0.02)	Yes (0.13)
Quah and Vahey measure	Yes (-2.71)	No (0.00)	No (0.00)	Yes (0.54)
RPIY	Yes (-2.67)	No (0.00)	No (0.00)	No (0.32)
DGI: RPIX excl. import prices	Yes (-2.36)	No (0.00)	Yes (0.31)	No (0.00)
DGI: ULC measure	No (-1.14)	n/a	n/a	n/a
DGI: PGDP excl. export prices	Yes (-3.11)	No (0.00)	No (0.00)	No (0.00)

Columns 1 and 2 report the results of the tests of condition (i); column 3 reports the results of the tests of condition (ii) and column 4 shows the results of the tests of condition (iii).

Column 1: Reports the ADF statistics of unit root tests of the stationarity of $(\pi_t - \pi_t^*)$. Yes/No indicates the series is/is not stationary.

Column 2: Reports the p-value on the null hypothesis that $\alpha = 0$ in the regression (C1). Yes/No indicates that α is not/is significantly different from zero.

Column 3: Reports the p-value on the null hypothesis that $\gamma = 0$ in regression (C2). Yes/No indicates that γ is not/is significantly different from zero.

Column 4: Reports the p-value on the null hypothesis that $\lambda = 0$ in regression (C3). Yes/No indicates that λ is not/is significantly different from zero.

References

- Andrade, I C and O'Brien, R J (2001)**, 'A measure of core inflation in the UK', *Department of Economics, Institute for Economics and Business Administration (ISEG), Technical University of Lisbon Working Paper*, No. 2001/05.
- Bakhshi, H and Yates, A (1999)**, 'To trim or not to trim? An application of a trimmed mean inflation estimator to the United Kingdom', *Bank of England Working Paper no. 97*.
- Balke, N S and Wynne, M A (2000)**, 'An equilibrium analysis of relative price changes and aggregate inflation', *Journal of Monetary Economics*, Vol. 45, pages 269–92.
- Ball, L and Mankiw, N G (1995)**, 'Relative-price changes as aggregate supply shocks', *Quarterly Journal of Economics*, February, pages 161–93.
- Beaton, R and Fisher, P G (1995)**, 'The construction of RPIY', *Bank of England Working Paper no. 28*.
- Blinder, A (1997)**, 'Commentary', *Federal Reserve Bank of St Louis Review*, May/June.
- Blix, M (1995)**, 'Underlying inflation: a common trends approach', *Sveriges Riksbank Arbetsrapport*, No. 23.
- Bryan, M F and Cecchetti, S G (1993)**, 'The consumer price index as a measure of inflation', *Federal Reserve Bank of Cleveland Economic Review*, No. 4, pages 15–24.
- Bryan, M F and Cecchetti, S G (1994)**, 'Measuring core inflation', in Mankiw, N G (ed), *Monetary policy*, NBER Studies in Business Cycles.
- Cutler, J (2001)**, 'A new measure of core inflation in the UK', *MPC Unit Discussion Paper no. 3*.
- Eckstein, O (1981)**, *Core inflation*, Prentice Hall.
- Faust, J and Leeper, E M (1997)**, 'When do long-run identifying assumptions give reliable results?', *Journal of Business and Economic Statistics*, Vol. 15, pages 345–53.
- Folkertsma, C K and Hubrich, K S E M (2000)**, 'Performance of core inflation measures', *De Nederlandsche Bank NV*, Research Memorandum WO&E No. 639/0034.
- Marques, C R, Neves, P D and Sarmiento, L M (2000)**, 'Evaluating core inflation measures', *Banco de Portugal Working Paper*, No. 3–00.
- Monetary Policy Committee (1999)**, *The transmission mechanism of monetary policy*, Bank of England.

Quah, D and Vahey, S P (1995), 'Measuring core inflation', *Economic Journal*, Vol. 105, September, pages 1,130–44.

Roger, S (1997), 'A robust measure of core inflation in New Zealand, 1949-96', *Reserve Bank of New Zealand Discussion Paper*, G97/7.

Roger, S (1998), 'Core inflation: concepts, uses and measurement', *Reserve Bank of New Zealand Discussion Paper*, G98/9.

Woodford, M (1999), *Inflation stabilisation and welfare*, Princeton University.

Wynne, M A (1999), 'Core inflation: a review of some conceptual issues', *ECB Working Paper Series*, No. 5.

Zeldes, S (1994), 'Comment' on Bryan and Cecchetti, in Mankiw, N G (ed), *Monetary policy*, NBER Studies in Business Cycles.