What do measures of core inflation really tell us?

By Alan Mankikar and Jo Paisley of the Bank's Conjunctural Assessment and Projections Division.

Despite the widespread use of the term 'core inflation', there is neither a widely accepted theoretical definition, nor an agreed method of measuring it. The wide range of conceptual bases is potentially confusing, particularly when the measures display different trends. This article offers an overview of some of the issues. It examines how core inflation has been defined, sets out to what extent the concept might be useful for policy-makers and assesses the wide range of available measures in the United Kingdom.

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. 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'. The wide range of conceptual bases is potentially confusing, and can make the resulting large number of measures of core inflation hard to interpret, particularly when they display different trends. This article sets out how core inflation might be useful to monetary policy makers and provides a conceptual and empirical evaluation of different measures of core inflation in the United Kingdom.

Core inflation and monetary policy

In the United Kingdom, the Monetary Policy Committee's (MPC) remit 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'. Price stability is currently defined as keeping the annual inflation rate of the retail price index excluding mortgage interest payments (RPIX) at 2½% at all times.

To achieve its remit, the MPC needs to have an understanding of how the economy works, how shocks are transmitted and how monetary policy affects the economy. Importantly, because changes in policy affect activity and inflation with a lag, monetary policy makers need to 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 so policy-makers are most interested in the *outlook* for inflation, typically over the one to two-year horizon where monetary policy can have most of its influence. In making judgments about the outlook for inflation, policy-makers employ a variety of economic models and monitor a wide range of economic variables and indicators, which potentially reveal information about the shocks affecting the economy.

Inflation itself is one of the numerous variables that policy-makers monitor in order to make judgments about the outlook for inflation. Chart 1 shows, however, that month-to-month movements in annual RPIX inflation can be volatile, making outturns potentially hard to interpret. A key task for policy-makers, as with all the variables that they monitor, is to read through the volatility or 'noise' in the data to extract as much information as possible.

An important question is whether an outturn for inflation—or indeed any other economic variable—changes the outlook and merits a change in policy. As Zeldes (1994) notes, 'presumably the answer depends on the persistence of the inflation innovation in the absence of any change in monetary policy'. Blinder, a former central banker, suggests that '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?' (Blinder (1997)). That is, when policy-makers

Chart 1
RPIX

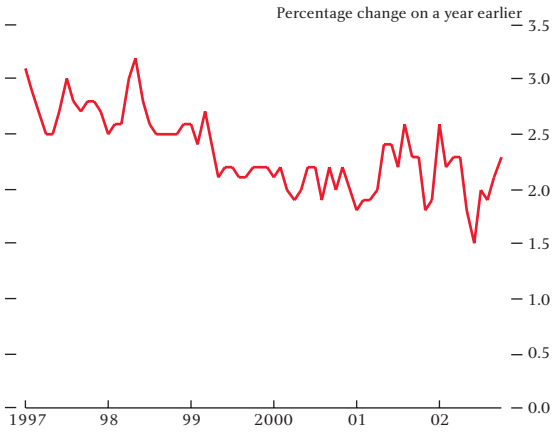

see changes in inflation, they are interested in the 'news' for the outlook for inflation.

In making these judgments, policy-makers do not view outturns for inflation in isolation. They must understand current movements in inflation and the shocks affecting the economy in the context of the other variables that they monitor and the models of the economy that they employ. Measures of core inflation may be useful if they help policy-makers see through the 'noise' in inflation outturns and provide a better indication of underlying inflationary developments. But what do we mean by 'noise'?

What induces 'noise' in measured inflation?

There are two main reasons why annual RPIX inflation may be 'noisy'. The first is economic, while the second is a function of the focus on annual inflation rates. Both are discussed below.

Movements in relative prices and aggregate inflation

In a world with fully flexible prices and unchanged monetary policy, a shock to a particular sector (such as a change in tastes or technology) would lead to instantaneous changes in relative prices, which would, other things being equal, leave the aggregate price level, and therefore the aggregate inflation rate, unchanged.

In practice, however, movements in relative prices do affect the aggregate price level and therefore the aggregate 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 of menu costs associated with changing prices, or perhaps because of 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, which mean that movements in relative prices can affect the aggregate price level. Consumer price indices cover only a subset of prices in the economy.(1) Relative price movements between two goods, one included in the RPI basket and the other not, would 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 changes in relative prices, and so are affected by relative price movements.

In theory, since relative price changes should have no long-run effect on the price level or inflation, they should not require a monetary policy response. So policy-makers would like to be able to distinguish between relative price movements and changes in prices that reflect underlying inflationary pressures. A measure of core inflation that is free from the noise induced by changes in relative prices may therefore be useful.

Interpreting changes in annual inflation rates

Inflation targets around the world are exclusively framed in terms of annual inflation rates, so that the price level in the latest month is compared with the price level twelve months earlier. Focusing on annual inflation rates is a simple way of trying to overcome the problem of seasonal price changes, ie that prices are changed at similar times each year. But while annual inflation rates are less volatile than monthly or even quarterly inflation rates, they can still be quite noisy. A one-off change in the price level, for example, will affect the annual inflation rate for a whole year before it drops out of the annual comparison. A key question when interpreting movements in the annual rate is to what extent they reflect price changes this year 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 rate.

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 between those that reflect a one-off or temporary price level change. Indeed, it may only be

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

some time after the event that one can be confident how to interpret a given change in an annual inflation rate. Nevertheless, measures of core inflation that attempt to smooth the volatile movements in inflation may help in this regard.

Uses of measures of 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 policy mean that monetary policy makers are most interested in the outlook for inflation.

Measuring core inflation

Since core inflation is unobservable, there is no 'right' answer as to how to measure it, and therefore no single agreed method. In the literature, there have been two main ways in which core inflation is measured. 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 also typically incorporate some prior view about the time-series properties of inflation to help distinguish between core and non-core inflation. The measures calculated under this approach use past relationships between aggregate inflation and its determinants to distinguish movements in inflation that reflect underlying pressures from those that reflect transitory shocks.

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

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 among the components of the index—up to 50% from each tail of the distribution in the case of the (weighted) median. The trimmed mean does not require *a priori* judgment about 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 and the weighted median for the United Kingdom are shown in Charts 2 and 3, together with RPIX inflation.

Chart 2
RPIX and trimmed mean

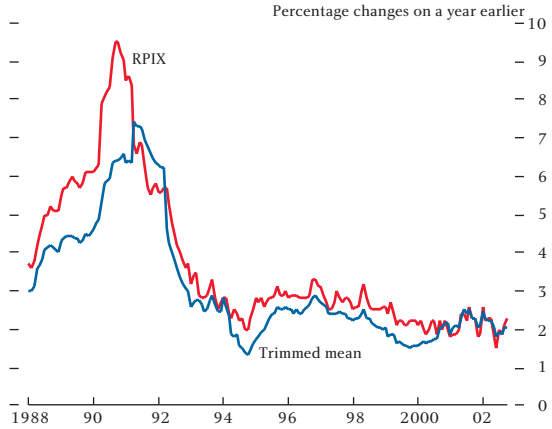

Chart 3
RPIX and weighted median

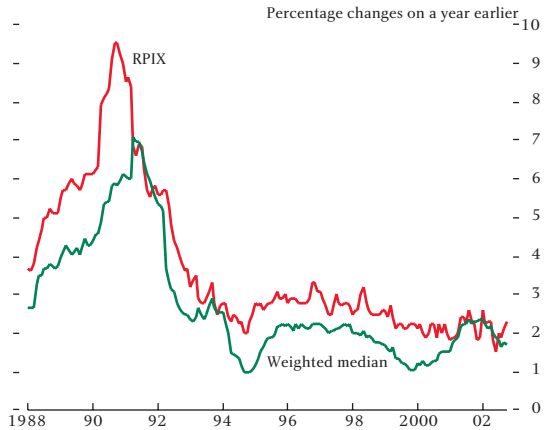

The ability of the trimmed mean to exclude relative price movements, but retain price movements associated with aggregate shocks, depends on the former being at the extremes of the price distribution. A recent UK case where trimming might have been appropriate is the outbreak of foot-and-mouth disease. A restriction on domestic meat supplies led to a sharp increase 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 affecting primarily one sector of the economy, which would be expected to lead to an adjustment in relative prices. In this case, trimming out these sharp price increases might have provided a better indication of underlying inflation in those particular months. But the question then arises how those meat prices and other prices adjust back to their equilibrium level over the 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 not be seriously misleading, there are instances when trimming would unambiguously misinform policy-makers. 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 following the shock, 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.

In practice, the validity of the trimmed mean as a measure of core inflation hinges on the premise that price fluctuations beyond some (albeit arbitrary) threshold are associated mainly with movements in relative prices and temporary price level effects. That is, these price changes must be generally larger in absolute magnitude than those price changes associated with aggregate shocks. But it is not clear that the magnitude of a price change is, in itself, necessarily a reliable signal of the cause of the shock. It seems more likely that both the trimmed mean and the excluded tails will contain a mixture of the effects of aggregate and relative price shocks. Indeed, this would be entirely consistent with theoretical arguments based on menu cost and staggered-price setting models.

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% trimmed mean measure monitored at the Bank of England.(1) Of the 21 components which are excluded more than 50% of the time between 1975 and 2002, five are seasonal food, 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 these 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.

One 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, it is unlikely that one would want to place much weight on the inflationary signals given by the trimmed mean. Furthermore, it is not clear how much of the distribution of price changes should be trimmed, so there is still a large degree of judgment needed. Some have decided this by considering how well measures with different degrees of 'trim' approximate a particular 'reference measure', with the 37-month centred moving average of inflation being a popular benchmark (see Bryan and Cecchetti (1993) for example). But it is difficult to determine whether the benchmark is sensible. One argument for using a reference measure is that it is 'smooth'. But if underlying aggregate shocks affecting 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 (see also the section on model-based measures).

Measures based on 'exclusion'

Some measures of core inflation are derived by permanently excluding certain components from the price index, *a priori*. In the United Kingdom, there are two prominent examples of measures of inflation in which certain items are permanently excluded. First, mortgage interest payments (MIPs) are excluded from

(1) At the Bank of England, the 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 arranged according to their weight, to give a string of 1,000–n numbers, where n is the current weight of 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.

the all-items retail price index to give RPIX, the target measure. MIPs are 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.(1) These exclusions may be useful for monetary policy purposes. Although indirect taxes are important components of a cost-of-living index, they do not constitute 'core' inflation: changes in indirect taxes may reflect headline consumer price inflation (duties are often raised in line with the rate of RPI inflation) but are independent of the underlying inflationary process.

Other components of aggregate consumer price indices 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 pressures in the targeted rate of inflation. Two examples of such measures for the United Kingdom are shown in Charts 4 and 5.

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 large changes in their prices and consequently in aggregate inflation. The argument for excluding energy prices is less clear cut. To the extent that petrol prices are driven by global oil supply conditions, this may be a valid reason for exclusion. But, it is likely that global demand conditions will also have a significant influence on the prices of these commodities, implying that energy prices contain useful information about underlying inflation.

Like the trimmed mean, an advantage of measures based on excluding components is that they are timely and easy to compute and explain. However, their downside is that they require a once-and-for-all (subjective) judgment about what the least informative price components are for estimating core inflation. And in a sense, these types of measure add little to the information set of monetary policy makers. They are just another way of representing certain components' contributions to the annual aggregate inflation rate, which are monitored as a matter of routine already in the Bank.

Chart 4
RPIX and RPIX excluding seasonal food and petrol

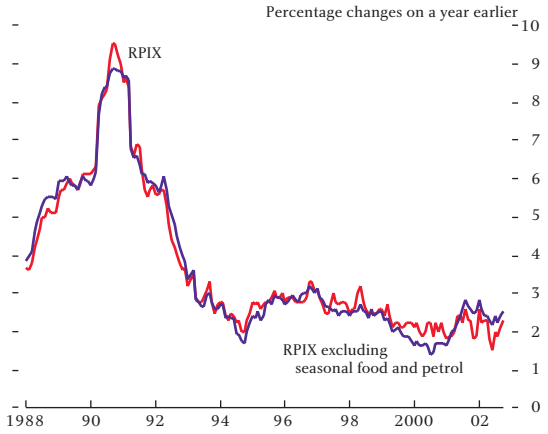

Chart 5
RPIX and RPIX excluding food, drink, petrol and tobacco

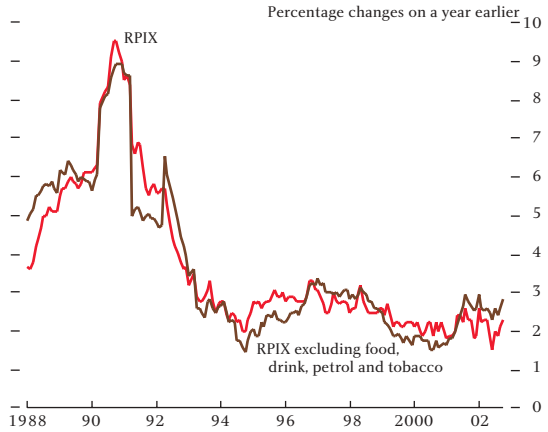

Chart 6
RPIX and 'persistence-weighted' RPIX

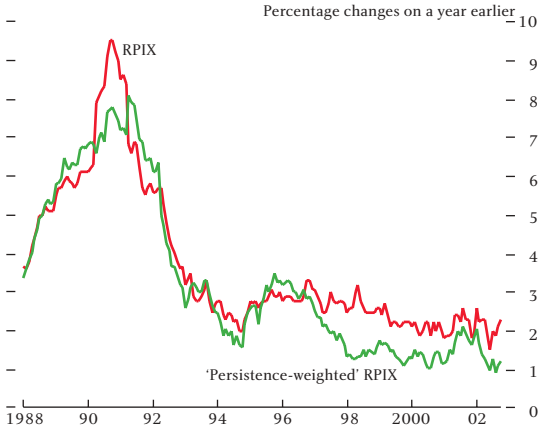

(1) Stripping out the effects of indirect taxes from consumer prices is not straightforward, since it involves making behavioural assumptions about the extent to which duty changes are passed on to consumers. For a description of how RPIY in the United Kingdom is constructed see Beaton and Fisher (1995).

Measures based on the whole price distribution

Other measures of core inflation use all available (disaggregated) 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 weights 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).

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 by expenditure weights. Underlying this particular approach is the view that some 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 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 take into account such weights, 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.(1) 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.(2)

Model-based approaches

Model-based approaches 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 of core inflation, which he 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 sustain the wage-price momentum, and in the tax system'.

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'.(3) The non-core element is essentially unanticipated inflation—and this is the component of measured inflation that does have a medium to long-run

(1) 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.
(2) 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.
(3) A shock that raises output permanently (and so raises actual and potential output) is assumed to have no long-run effect on inflation.

impact on output.(1) This definition clearly hinges on how one defines 'medium to long run' as opposed to short run. Quah and Vahey are trying to capture inflationary pressures that feed into or reflect inflation expectations. 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 aggregate supply shocks).

The two definitions seem to differ according to the effect of cyclical influences on core inflation. In Eckstein's world, core inflation should not be cyclical; in Quah/Vahey's world, core inflation should be strongly correlated with output in the short run. Roger (1998) suggests that we should not overdo the differences: the difference between a transient influence on inflation and cyclical and long-term influences is an artificial construct. This distinction should really be drawn in reference to the policy-maker's horizon. If the policy-maker is focusing on the medium run, then the Quah/Vahey definition is appropriate. 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. 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. These use different variables, and therefore identification schemes, which result in different estimates of core inflation. This non-robustness to the precise specification of the model is a limitation to their practical and routine use by policy-makers.

Domestically generated inflation

Domestically generated inflation (DGI) may be viewed as a particular type of core inflation measure that aims to exclude the one-off price level effects of external shocks on the aggregate rate of inflation. 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 know the extent and duration of such effects. Once the effects have worked through the economy, inflation should 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 are constructed and monitored: the GDP deflator excluding export prices; RPIX excluding import prices; and a measure based on unit labour costs (ULC). These are shown in Chart 7. Even if the conceptual basis for DGI is attractive, there are some practical concerns about how well the measures of DGI achieve their objective. First, the measures are only likely to strip out the direct, but not the indirect effects of external shocks, which should ideally be excluded as well. Second, the measures are sensitive to the precise assumptions underlying their construction. Third, and perhaps most worryingly, the different measures have shown very different trends over the past couple of years.

Chart 7
Measures of domestically generated inflation

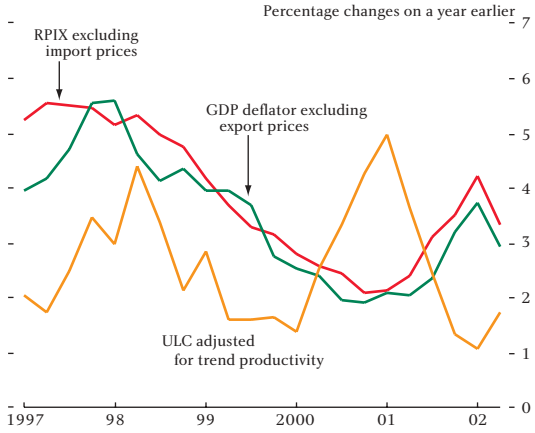

Evaluating measures of core inflation

There are several ways in which we might assess the usefulness of measures of core inflation. In the

(1) Quah and Vahey estimate a structural vector autoregressive (SVAR) model containing RPI inflation and output, on which they impose long-run identifying restrictions. But in their model the level of core inflation is not determined—since their VAR consists of just output and inflation, 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. But the precise nature of identifying restrictions and the data used will affect the estimates of SVARS.

literature, many 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 by the public. In addition, 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.

And though 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. The following sections, use a cointegration framework 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 movements in that measure of core inflation will not be informative about the future path of RPIX inflation.(1) 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 policy-maker's horizon, typically one to two years, then cointegration itself is not sufficient to render a measure of core inflation useful.

The next section sets out some tests proposed by Marques *et al* (2000) that 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.

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 a 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 strongly exogenous).

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 inflation and the targeted rate of inflation move one-for-one in the long run, and that the impact of relative price movements 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 was 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, or to use Marques *et al*'s terminology, 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. If it did, it would be extremely difficult to infer anything about the future path of targeted inflation by looking at core inflation, as the relationship would run both ways.

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 for the United Kingdom are shown in Table A.

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,

(1) 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 that surprising given that inflation has fallen over the sample period of the tests.

Table A
Replicating Marques *et al* tests for measures of core inflation in the United Kingdom[a]

	Condition (i)		Condition (ii)	Condition (iii)
	($\pi_t - \pi_t^*$) stationary	…and mean zero	Core inflation (π^*) should be an 'attractor' of targeted inflation (π)	Targeted inflation (π) should not be an attractor of core inflation (π^*)
RPIX excluding food	✔	✘	✔	✘
RPIX excluding seasonal food	✔	✘	✔	✘
RPIX excluding food and fuel	✔	✘	✘	✘
RPIX excluding food, alcohol, tobacco and petrol	✔	✔	✔	✔
RPIX excluding seasonal food and petrol	✔	✔	✔	✔
Trimmed mean	✔	✘	✘	✘
Weighted median	✔	✘	✘	✘
'Persistence-weighted' RPIX	✔	✘	✔	✔
'Quah and Vahey' measure	✔	✘	✔	✔
RPIY	✔	✘	✔	✔
DGI: RPIX excluding import prices	✔	✘	✘	✘
DGI: ULC measure	✔	✔	✔	✔
DGI: GDP deflator excluding export prices	✔	✘	✔	✘

n.a. = not available.

(a) A tick indicates that a measure passes the test at the 10% significance level.

tobacco and petrol, and the DGI measure based on ULC.

How should these results be interpreted? On the face of it, they suggest that RPIX excluding seasonal food and petrol and RPIX excluding food, alcohol, tobacco and petrol are potentially the two 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 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 RPIX inflation to some prescribed path, for example 2.5% at all times. And suppose that, over the sample period, policy had been used actively, and set optimally to achieve the target. Then, RPIX inflation would simply follow the prescribed path, save perhaps 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, RPIX 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 RPIX inflation.[1] This finding would cause us to reject this measure of core inflation as useful in providing forward-looking information about the future path of RPIX 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 in making that judgment.

But even if the results of the tests were not affected by the policy followed over the sample period, there still may be problems. For example, condition (iii) implies that lagged inflation contains no information about core inflation, which led Marques *et al* to reject a large number of measures of core inflation in Portugal. But there is a risk that this is overly stringent. There may be circumstances when core inflation might lag targeted inflation: for example, when movements in relative prices (temporarily) affect the aggregate inflation rate and inflation expectations. In this case, targeted inflation might lead core inflation to the extent that monetary policy is accommodative in allowing the relative price movements to affect inflation expectations.

Indeed, 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.

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 policy-makers monitor.

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

Conclusion

When policy-makers see a change in measured inflation, a key question is how much news there is 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 these 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 on the usefulness of measures of core inflation is provided by Hogan *et al* (2001) who suggest, in the Canadian context, that each one can provide a different insight into the inflation process. As this article has also found, no single measure performs well across the board. Hogan *et al* suggest that there can be value in looking at a range of measures, as long as it is clear 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 display different trends that they need to take a much closer look at the reasons behind those divergences.

References

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*.

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 and Cecchetti, S G (1993), 'The consumer price index as a measure of inflation', *Federal Reserve Bank of Cleveland Economic Review*, Vol. 4, pages 15–24.

Bryan, M 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.

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.

Hogan, S, Johnson, M and Laflèche, T (2001), 'Core inflation', *Bank of Canada Technical Report*, No. 89.

Marques, C R, Neves, P D and Sarmento, L M (2000), 'Evaluating core inflation measures', *Banco de Portugal Working Paper,* No. 3–2000.

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

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

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.