UNDERLYING INFLATION: CONCEPTS, MEASUREMENT AND PERFORMANCE

Ivan Roberts

Research Discussion Paper
2005-05

July 2005

Economic Analysis Department
Reserve Bank of Australia

I would like to thank Tony Richards, Christopher Kent, Alex Heath, Tim Bulman and Adam Cagliarini for helpful comments and suggestions. Of course, the responsibility for any errors lies with the author. The views expressed in this paper are those of the author and do not necessarily reflect those of the Reserve Bank of Australia.

Abstract

This paper explores the concept of underlying inflation and the properties of various measures of underlying inflation in the Australian context. Underlying inflation measures are routinely calculated and monitored by central banks in many countries, including the Reserve Bank of Australia. Alternative measurement concepts are explored, and a range of measures that have been calculated for Australia are discussed and evaluated on the basis of statistical criteria. These criteria capture the intuition that a good measure of underlying inflation should be less volatile than CPI (or headline) inflation, be unbiased with respect to CPI inflation, and capture the 'trend' in CPI inflation so that, on average, CPI inflation will tend to adjust towards the measure of underlying inflation. In the Australian context, statistical measures of underlying inflation, such as the trimmed mean or weighted median, perform fairly satisfactorily against these criteria. The performance of these measures can be further improved by seasonally adjusting prices at the CPI component level. Although underlying inflation measures have become less necessary in the past decade as inflation itself has become less volatile, these findings suggest that underlying inflation measures can still add value to the analysis of inflationary trends.

JEL Classification Numbers: C43, E31

Keywords: underlying inflation, core inflation, trimmed mean, weighted median, volatility-weighted measures

Table of Contents

UNDERLYING INFLATION: CONCEPTS, MEASUREMENT AND PERFORMANCE

Ivan Roberts

1. Introduction

In principle, underlying inflation, also known as 'core' inflation, can be regarded as the medium-term trend in inflation. That is, it is the inflation rate that would be recorded if one were to abstract from (or down-weight) sharp, quickly reversed, movements in prices or one-off shocks that create short-term volatility in measured inflation.[1] It is particularly important to have an accurate measure of current underlying inflation in the context of an inflation-targeting regime. In addition to providing a less noisy indicator of inflationary pressure than targeted inflation, it is possible that an accurate measure of underlying inflation could enable better forecasts of targeted inflation than would be achieved using the targeted measure alone.

Since 1993, monetary policy in Australia has targeted inflation of between 2 and 3 per cent, on average, over the course of the business cycle. Because of the flexible nature of this target, monetary policy is not required to respond to all movements in the published consumer price index (CPI), and consequently has focused on medium-term inflationary trends. Accordingly, various measures of underlying inflation have been monitored by the Reserve Bank of Australia (RBA) and other economists in Australia in recent years. Indeed, from 1993 to late 1998 the targeted inflation rate was a measure of underlying inflation which excluded a fixed set of items from the CPI.[2] This series removed the influence of volatile components of the CPI, such as fresh fruit and vegetables, and items whose prices

---

[1] Although widely used in the literature on underlying inflation, this idea has not been without critics. Keynes (1930, p 78) argued that 'the hypothetical change in the price level, which would have occurred if there had been no change in relative prices, is no longer relevant if relative prices have in fact changed – for the change in relative prices has in itself affected the price level' (cited in Diewert 1995).

[2] This series was referred to as 'Treasury underlying' inflation. It is no longer published by the Australian Bureau of Statistics.

were partly determined outside the market, such as tobacco products. More importantly, the series removed interest rate charges (included in the CPI over that period), thereby precluding a mechanical relationship between changes in monetary policy and targeted inflation. After interest charges were removed from the published CPI in 1998, it became the targeted measure of inflation (Reserve Bank of Australia 1998).

The Australian CPI is a transparent, publicly recognised gauge of price inflation, published quarterly. However, like consumer price indices in other countries, it is subject to volatility, which can make it difficult to distinguish between noise and the inflationary trend. For this reason, since the mid 1990s the RBA has incorporated a selection of underlying inflation measures into its analysis of the economy. These include both simple exclusion-based measures (such as the CPI excluding fruit, vegetables and automotive fuel), which are commonly used in many countries, and statistical measures (such as the weighted median) that exclude various items from the CPI on a time-varying basis. Table 1 lists the statistical measures published by five countries and some relevant research.

Table 1: Statistical Measures of Underlying Inflation in Selected Countries

Country	Published statistical underlying inflation measures	Relevant empirical research
Australia	Weighted median, 30 per cent trimmed mean	Kearns (1998); Heath, Roberts and Bulman (2004)
Canada	Double-weighted measure	Laflèche (1997)
NZ	Weighted median	Roger (1997, 1998)
Switzerland	15 per cent trimmed mean	Faber and Fischer (2000)
US	Weighted median, 15 per cent trimmed mean	Bryan and Cecchetti (1994); Bryan, Cecchetti and Wiggins (1997)

Heath *et al* (2004) calculated a range of statistical underlying inflation measures for Australia and found that after 1993 none of these were better predictors of CPI inflation than a constant forecast of around 0.6 per cent per quarter (2½ per cent annualised). This suggests that underlying inflation measures have become less necessary for forecasting inflationary trends as inflation has become lower and more stable. The aim of this paper is to extend this work by providing a more detailed discussion of underlying inflation concepts and the statistical problems that can hamper accurate measurement, and by subjecting the data to a simple econometric test to determine whether or not CPI inflation adjusts towards its underlying trend.

The rest of the paper proceeds as follows. Section 2 outlines the theoretical underpinnings of several underlying inflation measures, and proposes a test to gauge the extent to which these measures capture the underlying trend in published inflation. In Section 3, the properties of a selection of underlying inflation measures are considered. Seasonal adjustment at the disaggregated component level is shown to eliminate the average bias (with respect to CPI inflation) observed for many of these measures. Calculating underlying inflation measures based on the distribution of four-quarter-ended price changes is shown to achieve a similar result. Evidence is then presented in favour of the proposition that CPI inflation tends to adjust, at least partially, towards statistical measures of underlying inflation. Section 4 concludes.

2. Underlying Inflation Concepts and Alternative Measures

Underlying inflation is a difficult concept to pin down. As Vega and Wynne (2003) note, the existing literature lacks an agreed-upon theoretical framework, notwithstanding efforts by some authors (for example, Bryan and Cecchetti 1994). Thus, in practice, assessments of the spectrum of proposed measures often focus on their usefulness as unbiased medium-run predictors of CPI inflation, and not their approximation of ideal properties suggested by the literature.

There are two possible approaches to the measurement of underlying inflation. The first is to use a theoretical model to generate an underlying inflation series (see, for example, Quah and Vahey 1995). The second is to construct measures of underlying inflation based on the characteristics of the cross-section of individual price changes in each period and/or over time. This is the approach followed in this paper. In the context of an inflation-targeting regime, it is proposed that measures of underlying inflation should be unbiased estimators of targeted inflation over the medium term and capture the systematic component of inflation, if they are to add value to the analysis of inflationary trends.

2.1 Volatility-weighted Measures

The concept of underlying or core inflation has been surveyed in detail by Wynne (1999). The earliest notion of underlying inflation is often attributed to Stanley Jevons and Francis Edgeworth, who regarded it as the systematic component of aggregate inflation (Diewert 1995). Without loss of generality, the growth rate of the price of each good or service can be decomposed into a systematic component and an independent random component:

$$\pi_{it} = \pi_t^* + \nu_{it} \tag{1}$$

where π_t^* is the systematic component (underlying inflation) that is common to all items in the economy, and ν_{it} is the non-systematic component reflecting a relative price movement specific to individual item *i*.

If the ν_{it} are independently distributed across components and have a common variance, then the maximum likelihood estimator of π_t^* is the unweighted average of the π_{it}. But the idea that all prices in the economy are equally informative about underlying inflation trends is counter-intuitive; we know that some prices are affected by one-off shocks (such as changes to the tax and welfare systems), and that other prices can be relatively volatile due to temporary supply or demand shocks. The assumption of a common variance is thus clearly untenable. Diewert (1995) shows that if $E(\nu_{it}) = 0$ and $Var(\nu_{it}) = \sigma_i^2$ (that is, the variance is unique to each price), then the maximum likelihood estimator of the systematic component of inflation (π_t^*) is the 'neo-Edgeworthian' measure:

$$\hat{\pi}_t^* = \frac{\sum_{i=1}^{N} \frac{\pi_{it}}{\hat{\sigma}_i^2}}{\sum_{i=1}^{N} \frac{1}{\hat{\sigma}_i^2}}$$

where $\hat{\sigma}_i^2$ is the estimated variability in each component *i*.[3] This measure weights observed inflation in each item by the reciprocal of its volatility, so that more

[3] In theory, the underlying inflation and the variance estimates should be solved for simultaneously, using an iterative algorithm, if π_t^* is to be the maximum likelihood estimate. However, an *ad hoc* approximation is used by several authors, including Aucremanne (2000) and Marques, Neves and Sarmento (2000). See Diewert (1995) for details.

volatile items, which may give a less informative signal about underlying inflation, are given smaller weights.

However, the neo-Edgeworthian measure itself is not free from criticism. Diewert (1995) argues that inflation rates ought to be weighted according to their economic importance (for example, their share in expenditure) rather than by reference to a purely statistical criterion. Wynne (1999) notes that the neo-Edgeworthian measure may appeal to policy-makers if the increase in the 'cost of living' is not considered the most relevant macroeconomic inflation concept. However, there are reasons to be suspicious of a measure which entirely discards information about consumers' expenditure patterns, especially when the inflation target is framed in terms of the CPI, which is weighted by expenditure shares.

One approach that combines the cost of living and purely statistical approaches is the 'double-weighted' measure described by Laflèche (1997). This measure multiplies the neo-Edgeworthian weights by effective expenditure weights w_{it} drawn from the published CPI, as follows:

$$\hat{\pi}_t^* = \frac{\sum_{i=1}^{N} \frac{w_{it}}{\hat{\sigma}_i^2} \pi_{it}}{\sum_{i=1}^{N} \frac{w_{it}}{\hat{\sigma}_i^2}} .$$

This measure can potentially provide a compromise between the economic significance of a component and the clarity of the inflationary signal it provides. Good examples of this are food and automotive fuel, which are often considered too volatile to be included in core inflation measures, but are relatively important items in consumer expenditure.[4]

[4] Although the intuition behind the double-weighted measure seems reasonable, this measure is somewhat *ad hoc* on statistical grounds. In practice, it tends to be similar to the neo-Edgeworthian measure. The double-weighted measure can be derived as a maximum likelihood estimator of the systematic component of inflation, along similar lines to Diewert (1995), assuming that $E(\nu_{it}) = 0$ and $Var(\nu_{it}) = \sigma_i^2 w_{it}^{-1}$. The assumption that the variance in each item's relative price change should be inversely proportional to that item's weight in expenditure is questionable (Diewert 1995). See the Appendix for more details.

2.2 Trimmed Means and their Variants

The decomposition of Equation (1) has been used to support another way of measuring underlying inflation. Bryan and Cecchetti (1994) combine this decomposition with the Ball and Mankiw (1995) interpretation of relative price changes as aggregate supply shocks. Ball and Mankiw observe that according to classical theory, in which nominal prices are perfectly flexible, real factors such as productivity determine relative prices and monetary factors determine the overall price level. They argue that the assumption that firms face menu costs implies that changes in the price level are positively related to the skewness of relative price changes. This suggests that relative price shocks may seriously distort the underlying signal provided by standard price indices.

Bryan and Cecchetti (1994) observe that if the Ball and Mankiw (1995) model is correct, extracting the signal π_t^* from the individual price changes π_{it} is complicated by distributional issues. If the set of price changes is not normally distributed (a reasonable assumption), then published aggregate inflation π_t will not necessarily be a robust estimator of underlying inflation. They propose trimmed weighted means, using CPI weights and components, as a suitable alternative.

A trimmed weighted mean can be calculated by removing a certain proportion of the weight from each tail of the distribution of price changes, rescaling the remaining weights to sum to one, and calculating the weighted mean of the remaining distribution. The weighted median is calculated as the price change in the middle of the distribution, and is equivalent to a trimmed mean calculated such that 50 per cent of the distribution above and below the central observation is trimmed. Formally, following Vega and Wynne (2003), items in the CPI are ranked from smallest to largest price change. We define the cumulative weight for items labelled 1 to *i* as $W_{it} = \sum_{j=1}^{i} w_{(j)t}$, where $w_{(j)t}$ denotes the sorted *j*th weight, and define a set $I_\alpha = \{i : \alpha < W_{it} < 1-\alpha\}$. The (symmetric) trimmed mean formula is then:

$$\hat{\pi}_t^* = \frac{1}{1-\frac{2\alpha}{100}} \sum_{i \in I_\alpha} w_{(i)t} \pi_{(i)t}$$

where α is the percentage trimmed from each tail and $\pi_{(i)t}$ is the sorted *i*'th price change. The weighted median is the limiting case of the trimmed mean when α tends to 50. The formula can easily be adjusted to allow for asymmetric trimming, whereby different amounts are removed from the upper and lower tails of the distribution (Roger 1997; Kearns 1998).

Bryan and Cecchetti (1994) recommend using trimmed means for two reasons. First, according to their model, trimming the tails of the distribution of price changes should help to identify the body of price changes that are influenced by monetary rather than by real factors. Second, a trimmed mean provides a more robust measure of central tendency than the standard CPI inflation rate, by reducing the influence of 'outliers' in the consumer basket that exhibit transitory price movements and thus distort the underlying inflationary impulse. Bryan *et al* (1997) argue that if the distribution of price changes exhibits chronic excess skewness or excess kurtosis, then trimming the tails of that distribution so that it more closely approximates a normal distribution will also yield a more efficient estimate of underlying inflation than the standard CPI inflation rate.

Diewert (1997) has criticised trimmed means on the grounds that core inflation measures justified using the decomposition (1) and the assumption $E(\nu_{it}) = 0$ are flawed because they assume that, *ex ante*, all price changes have the same mean (π_t^*). Against this argument it should be noted that the use of trimmed means as underlying inflation measures need only assume that those prices which have not been 'trimmed' have the same mean.[5] Furthermore, one can accept the robustness argument of Bryan and Cecchetti (1994) without reference to the hypothesis that all prices have the same mean. One need not assume that all prices have a common mean and variance to predict that an estimator such as the trimmed mean will better capture the underlying trend in inflation than the published CPI. Whether or not trimmed means identify 'monetary' inflation *per se*, or for that matter the π_t^* in Equation (1), a trimmed mean based on the CPI has an advantage over some other measures of underlying inflation. Specifically, trimmed means reduce the influence

5 Assuming that the trimmed prices are genuine outliers that contain no relevant information, one could derive the weighted mean of the trimmed distribution as a maximum likelihood estimator of π_t^* from decomposition (1). But as Diewert (1995) has shown, this derivation, due to Clements and Izan (1987), makes an implausible assumption about the variance of the ν_{it}. See the Appendix for more details.

of extreme price movements in the distribution of price changes on a time-varying basis – that is, according to the characteristics of the cross-section in each period. Moreover, they weight price changes together on the basis of their economic significance, even if the latter is balanced in the calculation by a concern with statistical robustness.

A comparison with simple exclusion-based measures (for example, the CPI excluding fruit, vegetables and automotive fuel) is instructive. An exclusion-based measure gives a zero weight to items that are thought to contribute excessively, on average, to volatility in measured inflation, and therefore removes these items from the distribution of price changes in every period, regardless of their position in the distribution. Given that these items can record close-to-average inflation rates, and that other items sometimes record more excessive inflation, it is arbitrary to remove only these 'volatile' items all the time. In some periods genuine outliers are excluded, while in others they are not; the extent to which an exclusion-based measure is an accurate measure of central tendency depends on the period in which the measurement is taken. In contrast, a well-selected trimmed mean will be a robust measure of central tendency at any point in time. Furthermore, by trimming the tails of an ordered distribution of price changes, a trimmed mean – regardless of how much of the cross-section is trimmed – does not entirely *exclude* outlying information, but rather *limits* its influence on the inflation calculation. The trimmed mean will be affected even by elements which are 'excluded' because their magnitude will determine the location of the centre of the ordered distribution, but will be less affected by outliers than published inflation. In theory, the weighted median is the measure least likely to be affected by outlying price movements.

The items that are typically given a zero weight by trimmed means are often unsurprising. Table 2 shows the items in the Australian CPI which the 30 per cent symmetric trimmed mean published by the RBA trims most consistently from the quarterly distribution of price changes. This is done for two sample periods, from 1987:Q1 to 2004:Q4 and from 1993:Q1 to 2004:Q4. Only the top ten most frequently trimmed items for each sample period are shown; virtually all items get trimmed from the distribution at one time or another. Unsurprisingly, fruit, vegetables and automotive fuel are all regularly trimmed from the distribution of price changes, but a large number of other items are removed as well, including many non-market goods and services. It is worth noting that the most frequently trimmed items are basically the same in both sample periods, and that vegetable prices are removed from the distribution at almost every point in time. But automotive fuel prices are removed only two-thirds of the time, suggesting that these prices are not always ‘outliers’ in the distribution of price changes.

Table 2: Consistently Trimmed Prices

Symmetric 30 per cent trimmed mean CPI[(a)]

CPI expenditure class	**Proportion of time removed from distribution (per cent)**	**CPI expenditure class**	**Proportion of time removed from distribution (per cent)**
1987:Q1–2004:Q4 (72 observations)		1993:Q1–2004:Q4 (48 observations)	
Vegetables	97	Vegetables	100
Fruit	90	Fruit	94
Pharmaceuticals	88	Pharmaceuticals	85
Audio, visual and computing equipment	75	Audio, visual and computing equipment	77
Lamb and mutton	74	Lamb and mutton	75
Automotive fuel	68	Automotive fuel	67
Overseas holiday travel and accommodation	65	Glassware, tableware and household utensils	63
Glassware, tableware and household utensils	60	Pets, pet foods and supplies	63
Domestic holiday travel and accommodation	60	Domestic holiday travel and accommodation	63
Poultry	58	Overseas holiday travel and accommodation	63

Note: (a) Excluding interest charges and tax effects.

Although the weighted median is theoretically the most robust estimator of underlying inflation, it is unlikely to be the most efficient, as it ‘trims’ potentially informative observations. Thus there exists a trade-off between robustness and efficiency that may be exploited by varying the trimming percentage.[6] Bryan and Cecchetti (1994) suggest that the optimal trim can be found by searching across different trims and choosing the time-invariant trim that minimises the root mean squared error (RMSE) when the underlying measure is compared to a ‘benchmark’

[6] For thorough discussions concerning this trade-off, see Hampel *et al* (1986, chapter 1), and Aucremanne (2000).

underlying series. The benchmark that is often chosen is a moving average. However, choosing such a benchmark series requires that the benchmark itself is a good measure of underlying inflation, and a moving average or similar measure may not always be suitable. Heath *et al* (2004) report that the optimal trim chosen when using Australian data proved sensitive to the smoothness of the benchmark series chosen as well as the sample periods used in the calculation of the RMSE and mean absolute deviation statistics, suggesting that few firm conclusions could be drawn from this procedure in the Australian context.[7] Aucremanne (2000) also reports that this method proved unstable in an application to Belgian data.

The usefulness of a trimmed mean or weighted median measure may be called into question if the average over time of such a measure is significantly biased with respect to that of the weighted mean (that is, 'headline' CPI inflation). As indicated by Roger (1997) and Kearns (1998), rather than centreing the trim on the 50th percentile, one way to correct for such bias would be to centre it on the percentile that ensures that the average of quarterly changes in the underlying variable lines up with that corresponding to the target variable. This issue proved particularly problematic in New Zealand, where strong and persistent right-hand skewness in the distribution of price changes resulted in a large difference between the weighted mean and weighted median, implying that the 57th percentile was a more appropriate centre (Roger 1997). A discussion of this problem in the Australian context is provided by Kearns (1998). This paper will focus on alternative methods, such as seasonally adjusting the disaggregated price data, as a means of eliminating bias in trimmed means.

As noted above, one of the main reasons Bryan and Cecchetti recommend trimmed means as measures of core inflation is that they reduce the effect of departures from normality (skewness and kurtosis) on the distribution of price changes, and thus allow core inflation to be estimated in a more robust fashion. An alternative to trimming a constant percentage of price changes from the distribution (as is the standard practice) would be to choose the least amount of trim necessary to accept the hypothesis that the trimmed distribution has skewness and kurtosis statistics

7 They also discuss the drawbacks of various series that have commonly been used as benchmark inflation series in the literature.

equivalent to those of a normal distribution. Aucremanne (2000) follows this approach, using the Jarque-Bera statistic to test normality in the cross-sectional distribution for each time period. This procedure allows the degree of trim to vary for each cross-section distribution being considered, and thus has the potential to exploit the trade-off between robustness and efficiency more effectively than standard trimmed means.[8] Heath *et al* (2004) extend this approach by relaxing the implicit assumption that the central trimming percentage remains constant over time and allowing it to vary in line with the characteristics of the cross-section.

Several other measures of underlying inflation have been proposed in the literature, of which at least three deserve brief mention. The first is the 'dynamic factor index' measure proposed by Bryan and Cecchetti (1993), which effectively weights individual price changes by the strength of their signal-to-noise ratio, and is thus designed to avoid 'bias' due to any particular form of production or expenditure weighting scheme. The second is what Cutler (2001) describes as a 'persistence-weighted' measure of core inflation, constructed by estimating first-order autoregressive models of disaggregated inflation series in a recursive manner and collecting the autoregressive coefficients, which are designated time-varying 'persistence weights' if positive and multiplied by zero if negative. Individual price change data are then aggregated using the persistence weights to create a measure of core inflation. The third approach, which has been used by several authors (for example, Machado *et al* 2001, Maria 2004, Shu and Tsang 2004), is to employ the first principal component of disaggregated price change data as an estimator of underlying inflation. The first principal component of a set of data is a linear combination of the data that explains as much variation in the data as possible, and

8 Although the aim of this approach is to strike a balance between robustness and efficiency, the Jarque-Bera statistic used to test for normality may not itself be a particularly powerful estimator if the data are sufficiently non-normal. The extent to which this is likely to be a problem for this paper's results is unclear. Aucremanne (2000), and more recently OECD (2005), have experimented with the Huber 'skipped' mean (Huber 1964) in an attempt to avoid the problems that non-normality poses for robust estimators, though it appears that the well-established robustness properties of this estimator have yet to be proved for weighted data.

can be interpreted as the ‘common trend’ in a series of data.[9] While all three measures have attractions and disadvantages, they are difficult to implement in practice without a sufficiently long run of data for each price series. In the Australian context, a key limitation of the CPI data is the regular introduction of new items and cancellation of old items in the CPI, which restricts the length of time series for many items. This precludes the effective application of these methods to Australian data, and consequently they are not addressed further in this paper.

2.3 Assessing the Alternatives

It remains for us to determine which, if any, of the possible underlying inflation measures are of use in assessing the divergence of current inflation from its trend, and possibly also predicting future inflation. The literature has produced a number of specific criteria by which to judge different underlying inflation measures. Roger (1998) argues that, in an inflation-targeting context, an appropriate measure of underlying inflation should be timely, credible (verifiable by agents independent of the central bank), easily understood by the public, and not significantly biased with respect to targeted inflation. Wynne (1999) suggests that such a measure should also be computable in real time, have some predictive power relative to future inflation, have a track record of some sort, and not be subject to substantial revisions. As Wynne (1999) emphasises, these features are important insofar as the central bank seeks to use a measure of underlying inflation as an important part of its routine communications with the public to explain policy decisions.

While many of these criteria are sensible, they do little to clarify the statistical conditions that a suitable underlying inflation indicator should satisfy. Heath *et al* (2004) have argued that two properties are particularly desirable. The first is unbiasedness with respect to CPI inflation. Bias can be assessed informally by comparing the average of underlying inflation with that of CPI inflation over a

---

9 In general, it is difficult to give an economic interpretation to the weights assigned to the disaggregated price data by this procedure. As the first principal component is scale-dependent, the usual practice is to standardise the price data so each price series has a zero mean and a unit variance. But this implies that if the first principal component is rewritten as a weighted average of disaggregated price changes, the weight of a given price change will be inversely proportional to its standard deviation. Thus, this formulation has a property in common with the volatility-weighted measures discussed earlier.

given period. Of course, the existence or extent of bias observed may depend on the specific period over which the calculation is performed. But it is nonetheless a useful indicator of which measures have bias properties that merit further investigation. We can test whether the bias is statistically significant by estimating the equation:

$$\pi_t = \beta_0 + \beta_1 \pi_t^* + v_t \qquad (2)$$

and testing the joint null hypothesis that $\beta_0 = 0$ and $\beta_1 = 1$. This restriction reduces Equation (2) to Equation (1).

A second property that Heath *et al* (2004) consider desirable is that underlying inflation contain information about future trends in CPI inflation over and above the information provided by the CPI itself. This condition, originally suggested by Bryan and Cecchetti (1994), can be formalised by stipulating that underlying inflation should Granger cause CPI inflation and that Granger causality should not run in the opposite direction. Heath *et al* (2004) applied this test to Australian data and found that for a sample beginning in 1987 a large number of underlying inflation measures performed satisfactorily on this criterion. However, they also found that for a sample beginning in 1993 none of the same measures performed satisfactorily; in regressions of CPI inflation on a constant and lags of CPI and underlying inflation, only the constant proved to be significantly different from zero.

The Granger causality test provides one test of predictive ability, but there exists an alternative, albeit related, test which may be closer in spirit to the problem faced by analysts when they seek to interpret an inflation figure and various competing measures of underlying inflation at a given point in time. In particular, at times when there is a divergence between CPI inflation and a measure of underlying inflation, an important issue is how the gap might be closed in the next period. It could be that CPI inflation will move towards the current level of underlying inflation. Alternatively, underlying inflation could move towards the current level of CPI inflation.

Formally, this can be tested by performing the following two regressions:

$$\Delta\pi_t = \gamma_{10} + \gamma_{11}(\pi_{t-1} - \pi^*_{t-1}) + \varepsilon_{1t} \tag{3}$$

$$\Delta\pi^*_t = \gamma_{20} + \gamma_{21}(\pi_{t-1} - \pi^*_{t-1}) + \varepsilon_{2t} \,. \tag{4}$$

If $\hat{\gamma}_{11}$ is significantly less than zero and $\hat{\gamma}_{21}$ is equal to zero, then it can be concluded that CPI inflation tends to adjust towards underlying inflation, but not vice versa. Alternatively, if $\hat{\gamma}_{21}$ is significantly less than zero and $\hat{\gamma}_{11}$ is equal to zero, then it can be concluded that it is underlying inflation that tends to adjust towards CPI inflation. Finally, if the parameters $\hat{\gamma}_{11}$ and $\hat{\gamma}_{21}$ are both significantly less than zero, then it is likely that the underlying inflation measure being tested is only a rough approximation to the underlying trend in inflation that it is supposed to estimate. For the sake of brevity, this test is referred to as the 'gap' test. Though admittedly simple, the gap test provides a gauge of whether prospective underlying inflation measures satisfy a minimum property we would expect to see in a suitable measure. Similar tests have been applied to UK and OECD data, respectively, by Cutler (2001) and OECD (2005). Assuming that inflation is a stationary variable, the gap test can be conducted by estimating Equations (3) and (4) separately by ordinary least squares and comparing *t*-statistics for the two parameters with relevant critical values.[10]

However, Equations (3) and (4) can be conceived as nested within more general specifications. For example, Equation (3) can be generalised to the following equation:

$$\Delta\pi_t = \delta_{10} + \delta_{11}\pi_{t-1} + \delta_{12}\pi^*_{t-1} + \varepsilon_{1t} \,. \tag{5}$$

[10] The interpretation of inflation as a stationary process finds empirical support in Cecchetti and Debelle (2004). On the assumption that CPI and underlying inflation are both I(1) processes, Marques *et al* (2000) and Marques, Neves and da Silva (2002) implement a test similar to the gap test described here. They suggest that CPI and underlying inflation should be cointegrated and that only CPI inflation should respond to deviations from this cointegrating relationship. Essentially, this approach requires the estimation of Equations (3) and (4) in a system, augmented for dynamics. Making similar assumptions, Dixon and Lim (2004) apply this procedure to four-quarter-ended Australian inflation data.

In other words, Equation (3) constitutes a restriction on Equation (5), such that $\delta_{11} = -\delta_{12} (= \gamma_{11})$. Equation (5) can also be interpreted as a Granger causality test with a lag order of 1.[11]

A modification to Equation (5), which may make sense in an environment where inflation is relatively stable, is:

$$\Delta\pi_t = \alpha_{10} + \alpha_{11}(\pi_{t-1} - \pi^*_{t-1}) + \alpha_{12}(\pi_{t-1} - \overline{\pi}) + \varepsilon_{1t} \tag{6}$$

where $\overline{\pi}$ refers to the mean of CPI inflation over the sample period under consideration. This equation can be used to test whether CPI inflation tends to adjust towards a given measure of underlying inflation, or towards a particular rate of inflation, or some combination of both. We may refer to this test as the augmented 'gap' test. If the term $\hat{\alpha}_{12}$ is significantly less than zero, then it can be inferred that CPI inflation reverts towards some constant rate, $\overline{\pi}$. If $\hat{\alpha}_{11}$ is significantly less than zero, then it can be concluded that CPI inflation adjusts towards a particular measure of underlying inflation.

3. Properties of Underlying Inflation Measures

This section examines the features of a range of underlying inflation measures calculated for Australia, assesses any bias with respect to CPI inflation, and considers their performance as indicated by the 'gap' test described earlier. The

[11] Formally, Granger causality is tested by estimating the following two equations:

$$\pi_t = \gamma_{10} + \sum_{i=1}^{P} \phi_{1i}\pi_{t-i} + \sum_{i=1}^{P} \eta_{1i}\pi^*_{t-1} + \varepsilon_{1t}$$

$$\pi^*_t = \gamma_{20} + \sum_{i=1}^{P} \phi_{2i}\pi_{t-i} + \sum_{i=1}^{P} \eta_{2i}\pi^*_{t-1} + \varepsilon_{2t} .$$

If $\sum_{i=1}^{P} \eta_{1i} \neq 0$ and $\sum_{i=1}^{P} \phi_{2i} = 0$ then it can be concluded that underlying inflation Granger causes CPI inflation, and not vice versa. Assuming a lag order of P = 1 and restricting the coefficients such that $\phi_{11} + \eta_{11} = 1$ and $\phi_{21} + \eta_{21} = 1$, where $\eta_{i1} = -\gamma_{i1}$, results in Equations (3) and (4).

measures considered are: the double-weighted measure;[12] the 30 per cent symmetric trimmed mean and the symmetric weighted median; the optimal Jarque-Bera trimmed mean discussed in the previous section; and two exclusion-based measures published by the Australian Bureau of Statistics (ABS) – the CPI excluding volatile items and market prices excluding volatile items. The trimmed means are calculated in three ways: first, using unadjusted quarterly price changes; second, using seasonally adjusted quarterly price changes; and third, using year-ended price changes. The double-weighted measure is calculated using unadjusted and seasonally adjusted quarterly price change data.

Measures based on the distributions of seasonally adjusted quarterly price changes and annual price changes are considered because the prices of certain items display identifiable seasonality in a yearly cycle.[13] Although in aggregate the Australian CPI is not particularly seasonal (Australian Bureau of Statistics 2004), several items in the distribution of price changes (for example, health, education, utilities and property rates & charges) do display significant seasonality. Seasonal analysis of the disaggregated quarterly price changes can be performed using packages such as the US Census Bureau's X12-ARIMA program.[14] A seasonal analysis of the Australian data revealed that in total, around one-third of the CPI expenditure classes display significant seasonality.[15] Sliding spans analysis was applied to all series with a sufficiently long run of data, and only those series with seasonal factors found to be robust on this basis were accepted as seasonal. Applying X12-ARIMA seasonal factors to the associated CPI components enabled the construction of a new distribution of price changes, from which 'seasonally adjusted' measures of core inflation could be calculated.

---

12 The weights used in the construction of the double-weighted measure are CPI effective expenditure weights multiplied by reciprocals of the variances of disaggregated price changes (calculated for the entire sample). This is the standard *ad hoc* construction. See the Appendix for further details on this measure.

13 Seasonal adjustment has been employed by several authors to improve empirical estimates of core inflation. See Bryan and Cecchetti (1999) and Aucremanne (2000), for applications to Japanese and Belgian data, respectively.

14 See the Census Bureau's website for more information: <http://www.census.gov/srd/www/x12a/>.

15 The number of CPI expenditure classes varies between 89 and 107 depending on which period is being considered.

An alternative method of accounting for seasonality in the disaggregated price data is to calculate underlying inflation rates on the basis of the distribution of four-quarter-ended price changes. Since seasonal items often record large increases only once a year, seasonality observed in the quarterly data often ‘washes out’ of the year-ended calculation. For the aggregate CPI and the exclusion-based measures, the four-quarter-ended price change is the same whether it is calculated directly or from compounding the quarterly movements, because the weighted mean calculation is linear. Likewise, the double-weighted measure presented in this paper will only differ for year-ended inflation data if the weights are recalculated on a four-quarter-ended basis. But the weighted median and trimmed mean can be quite different when calculated on a different basis.

The way in which the Australian CPI is constructed creates practical problems for utilising the four-quarter-ended distribution. Changes to the CPI basket, which occur roughly at five-year intervals, but which have been slightly more frequent in the past ten years, mean that while a continuous time series of quarterly price-change distributions can be obtained, there are gaps in the series of four-quarter-ended distributions. This prevents a four-quarter-ended weighted median, trimmed mean and optimal Jarque-Bera trim from being directly calculated in a transition period. The statistics corresponding to these measures shown in this paper are therefore based on interpolated data.[16]

Table 3 considers the average variability of the different underlying inflation measures. Two alternative sample periods are shown: 1987:Q1 to 2004:Q4 and

[16] The interpolation method used is as follows. For each break period there are three quarters for which four-quarter-ended price changes are not available. Denote as quarter t the last quarter for which four-quarter-ended data for the ‘old’ series are available (that is, the quarter prior to the break). The first quarter of the break period is then denoted $t+1$, the second $t+2$ and the third $t+3$. The annual rate to quarter $t+1$ is calculated by compounding the (aggregate) three-quarter-ended statistical measure for t with the quarter-on-quarter statistical measure for $t+1$. The second quarter is interpolated by compounding the two-quarter-ended statistical measure for t with the two-quarter-ended statistical measure for $t+2$. The third quarter is interpolated by compounding the quarter-on-quarter statistical measure for t with the three-quarter-ended statistical measure for $t+3$. As the interpolated figures may be biased downwards due to differences between higher- and lower-frequency distributions, a bias adjustment is made to each figure, estimated by averaging the bias for the calculation of the corresponding period (for example, the March quarter) over the previous four equivalent periods (that is, the previous four March quarters).

Table 3: Variability of Underlying Inflation Measures[(a)]

Standard deviation (per cent)

Indicator	1987:Q1–2004:Q4	1993:Q1–2004:Q4
	Quarterly	
Consumer Price Index	**0.57**	**0.30**
Quarterly measures: not seasonally adjusted		
Weighted median	0.49	0.19
30 per cent trimmed mean	0.46	0.17
Market prices excluding volatile items	0.53	0.26
CPI excluding volatile items	0.52	0.24
Optimal Jarque-Bera trim	0.47	0.19
Double-weighted measure	0.48	0.21
Quarterly measures: seasonally adjusted		
Weighted median	0.48	0.17
30 per cent trimmed mean	0.47	0.17
Optimal Jarque-Bera trim	0.47	0.18
Double-weighted measure	0.46	0.20
	Year-ended	
Consumer Price Index	**2.09**	**0.67**
Measures based on the annual distribution of price changes[(b)]		
Weighted median	2.05	0.53
30 per cent trimmed mean	2.02	0.48
Optimal Jarque-Bera trim	2.02	0.53

Notes: (a) All series exclude interest charges and tax effects.
(b) Interpolated.

1993:Q1 to 2004:Q4. The earlier period is the longest over which data are available for all measures, while the inflation-targeting period, which starts in 1993, is a logical alternative for analysis, given that it corresponds to a period of low and stable inflation. Unsurprisingly, all of the underlying inflation measures shown are less volatile than CPI inflation for both sample periods. The exclusion-based measures, which remove specified items whose prices are deemed particularly volatile in every period, are the most volatile of all the underlying inflation measures. It also bears noting that the volatility-weighted measures themselves do not necessarily display significantly lower volatility than the other measures. This is not surprising, as trimmed means with substantial trims weight the least volatile components more heavily and give a zero weight to the most volatile components. It is also apparent that seasonally adjusting the disaggregated data makes little difference to the variability of a given statistical underlying inflation measure. Finally, the measures based on the distribution of annual price changes are less volatile than year-ended CPI inflation. Differences between the three types of underlying inflation measure become more distinct when the question of bias is considered.

3.1 Average Inflation and Bias

Table 4 shows summary statistics for this selection of measures, calculated on a year-ended basis. Looking first at the unadjusted quarterly measures, it is apparent that the weighted median, trimmed mean and market prices excluding volatile items measures all show an average negative bias of 0.2 percentage points or more with respect to CPI inflation. The optimal Jarque-Bera trimmed mean also displays a (smaller) downward bias. Compared with more standard trimmed means, the optimal Jarque-Bera trim on average excludes a substantially lower percentage of the price change distribution (around 8 per cent) in addition to being – at least in principle – more robust to departures from normality in the original distribution of price changes. The double-weighted measure and the CPI excluding volatile items measure display a lower degree of bias than the other measures. Only in the cases of the trimmed mean and the market prices excluding volatile items is the bias statistically significant.

Turning to the seasonally adjusted measures in Table 4, it is apparent that the observed negative bias is negligible for the weighted median, the trimmed mean and the optimal Jarque-Bera trim. The double-weighted measure, indeed, has a positive bias when calculated using seasonally adjusted price changes. None of these biases is statistically significant.[17] It is also apparent from Table 4 that none of the measures based on the distribution of annual price changes is substantially biased. These findings suggest that, in general, seasonally adjusted quarterly

---

[17] Nonetheless, the bias for the double-weighted measure is sufficiently large to warrant caution interpreting this measure in practice. The fact that the measure based on unadjusted data exhibits negligible bias suggests that seasonal adjustment is probably not necessary for this measure.

measures and measures based on the distribution of annual price changes may be more appropriate than the unadjusted quarterly measures for assessing the underlying inflation rate.

Table 4: Underlying Inflation – Assessment of Bias[(a)]

Year-ended percentage change

Indicator	Average	
	1987:Q1–2004:Q4	1993:Q1–2004:Q4
Consumer Price Index	**3.62**	**2.48**
Quarterly measures: not seasonally adjusted		
Weighted median	3.40	2.26
30 per cent trimmed mean	3.37*	2.25**
Market prices excluding volatile items	3.12*	2.20**
CPI excluding volatile items	3.38	2.42
Optimal Jarque-Bera trim	3.36	2.24
Double-weighted measure	3.58	2.44
Quarterly measures: seasonally adjusted		
Weighted median	3.51	2.36
30 per cent trimmed mean	3.57	2.44
Optimal Jarque-Bera trim	3.53	2.37
Double-weighted measure	3.86	2.80
Measures based on the annual distribution of price changes[(b)]		
Weighted median	3.56	2.42
30 per cent trimmed mean	3.56	2.40
Optimal Jarque-Bera trim	3.54	2.38

Notes: ** and * denote statistically significant bias at the 5 and 10 per cent levels, respectively.
(a) All series exclude interest charges and tax effects.
(b) Interpolated. These measures are not tested for bias.

3.2 Explaining the Presence or Absence of Bias

Why is it that underlying inflation measures based on unadjusted quarterly data display a downward bias while those based on seasonally adjusted data or year-ended price changes do not? As far as (unadjusted) quarterly trimmed means are concerned, it is often the case that items displaying large increases once or twice a year are trimmed from the distribution of price changes in the quarters when they record their largest increases, resulting in a lower average inflation rate over the year for trimmed measures than for the published CPI. Seasonally adjusting price changes at the disaggregated level can eliminate biases caused in this way, by smoothing price changes over the course of a year. This reduces the chance that a highly seasonal item will be trimmed from the distribution of price changes, providing that inflation over the year in that item is not significantly greater than overall CPI inflation.[18] As shown in Section 2.2, the pharmaceuticals component is one of the items most frequently removed by the symmetric 30 per cent trimmed mean. But when these prices are seasonally adjusted it is removed far less frequently. The same is true of many other items, including domestic holiday travel & accommodation, education, electricity and property rates & charges, all of which are removed roughly half as frequently when seasonally adjusted.

A similar explanation can be given for the lack of bias observed in measures based on the distribution of year-ended price changes. By smoothing seasonal price spikes over the course of a year, the four-quarter-ended calculation eliminates the effect of regular seasonality in determining which items are trimmed. Several items that display significant seasonality are removed from the annual distribution much less frequently than they are from the quarterly distribution (for example, automotive fuel, pharmaceuticals and domestic holiday travel & accommodation).[19]

---

[18] A similar point is made by Aucremanne (2000), in reference to Belgian CPI data.

[19] On a more abstract level, it is worth noting that measures based on distributions of price changes calculated for very high frequencies are in general likely to yield downwardly biased estimates of underlying inflation. As an extreme example, at daily frequencies, the prices of many goods and services would be unchanged. Hence, a hypothetical daily inflation rate would be much more volatile than, say, a quarter-on-quarter inflation rate. If a fixed percentage of extreme movements were trimmed on a daily basis, a trimmed mean might well show zero inflation on most, if not all, days and the cumulated rate over a longer (say, quarterly) horizon would thus generally be lower than CPI inflation calculated at a quarterly frequency. This suggests that, up to some optimal point, incrementally increasing the horizon over which price changes are calculated should reduce the bias of trimmed mean measures with respect to CPI inflation. There is some empirical support for the idea that trimmed means calculated using high frequency price changes tend to display lower average inflation than comparable measures calculated using lower frequency distributions. Shiratsuka (1997) shows that, in the Japanese case, there is a noticeable difference between annualised monthly and year-ended observations of underlying inflation. A weighted median based on the annual distribution of price changes is used in New Zealand for similar reasons.

While the seasonally adjusted underlying inflation measures and the measures based on the distribution of year-ended price changes do not display significant bias with respect to CPI inflation, a few caveats are worth noting. First, measures based on seasonally adjusted price changes are subject to revision as new inflation data become available. For a number of highly seasonal component series, such as those in the education category, only a short sample of data is available (with as few as 19 observations). Hence, the corresponding seasonal factors may not prove stable in the longer term.[20] Second, any changes in observed seasonal patterns, or the appearance of temporary atypical seasonality in specific items, could potentially affect the signal provided by these measures. The extent to which this is a first- or second-order empirical issue has yet to be established for Australian data.

While there is no strong *a priori* reason to prefer measures calculated using the quarterly rather than the four-quarter-ended distribution, as discussed previously the fact that three observations are lost in each transition period means that interpolated data must be used. It would also be difficult to derive an informative quarterly series from the annual data using mechanical techniques, although this need not diminish the value of the annual interpolated series as an analytical tool. In the absence of a corresponding quarterly series, this paper does not attempt to use econometric criteria to assess measures based on the distribution of annual price changes. But to the extent that these measures are less biased with respect to CPI inflation than the unadjusted quarterly measures, they still have a useful role to play in the assessment of underlying inflationary trends.

3.3 Econometric Assessment of Underlying Inflation Measures

We first present the results of the 'gap' test described in Section 2.3 to determine whether CPI inflation adjusts towards underlying inflation measures. The results are displayed in Table 5. For each sample the estimates of γ_{11} and γ_{21} in Equations (3) and (4) are presented. It is apparent that, for the 1987:Q1–2004:Q4 sample, the coefficient γ_{11} is statistically significant for all of the underlying inflation measures considered, irrespective of whether or not they are based on seasonally adjusted data, suggesting that CPI inflation tends to move towards each

[20] This may explain why underlying inflation measures based on seasonally adjusted price data have not previously been calculated for Australia.

underlying inflation measure. In addition, in all but two cases it is clear that underlying inflation does not adjust towards the CPI inflation rate – that is, $\hat{\gamma}_{21}$ is statistically insignificant. The exceptions are the unadjusted and seasonally adjusted trimmed means, for which $\hat{\gamma}_{21}$ is significantly less than zero.

Table 5: 'Gap' Tests

	1987:Q1–2004:Q4 (72 observations)		1993:Q1–2004:Q4 (48 observations)	
	$\hat{\gamma}_{11}$	$\hat{\gamma}_{21}$	$\hat{\gamma}_{11}$	$\hat{\gamma}_{21}$
Quarterly measures: not seasonally adjusted				
Weighted median	–1.20 (0.00)***	–0.08 (0.43)	–0.85 (0.00)***	0.16 (0.27)
30 per cent trimmed mean	–1.48 (0.00)***	–0.35 (0.01)***	–1.32 (0.00)***	–0.22 (0.12)
Market prices excluding volatile items	–1.01 (0.00)***	0.15 (0.29)	–0.85 (0.00)***	0.27 (0.04)**
CPI excluding volatile items	–1.26 (0.00)***	–0.15 (0.39)	–1.12 (0.00)***	0.03 (0.85)
Optimal Jarque-Bera trim	–1.24 (0.00)***	–0.10 (0.27)	–1.10 (0.00)***	–0.01 (0.92)
Double-weighted measure	–1.18 (0.00)***	–0.07 (0.47)	–0.86 (0.00)***	0.16 (0.14)
Quarterly measures: seasonally adjusted				
Weighted median	–1.28 (0.00)***	–0.12 (0.16)	–0.92 (0.00)***	0.08 (0.51)
30 per cent trimmed mean	–1.48 (0.00)***	–0.30 (0.00)***	–1.20 (0.00)***	–0.09 (0.44)
Optimal Jarque-Bera trim	–1.22 (0.00)***	–0.05 (0.52)	–1.00 (0.00)***	0.06 (0.61)
Double-weighted measure	–1.19 (0.00)***	–0.09 (0.59)	–0.90 (0.00)***	0.07 (0.20)

Notes: *** and ** denote significance at the 1 and 5 per cent levels, respectively; *p*-values are given in parentheses. The estimated coefficients refer to the parameters in Equations (3) and (4).

At first glance, the results for the 1993:Q1–2004:Q4 sample appear to suggest that the same proposition holds – i.e., that there is strong evidence that CPI inflation moves towards the various underlying inflation measures. These results should be treated with some caution, however, because the implicit restriction underlying the gap test referred to in Section 2.3, which broadly holds for the longer sample, does not hold for this shorter sample period.[21] To obtain a clearer picture of what is going on in the shorter sample, we estimate Equation (6) to perform the augmented gap test, as discussed in Section 2.3. It will be recalled that this equation simply augments the gap regression with the term $(\pi_{t-1} - \bar{\pi})$. In this case, $\bar{\pi}$ refers to the mean of CPI inflation for the 1993:Q1–2004:Q4 sample, which is 0.6 per cent per quarter, or 2½ per cent in annualised terms. Table 6 shows the results of estimating Equation (6) for the 1993:Q1–2004:Q4 sample.

It is apparent from Table 6 that in all cases the coefficient ($\hat{\alpha}_{12}$) on the difference between CPI inflation and $\bar{\pi}$ is significantly less than zero. This implies that when CPI inflation has departed from its average of 2½ per cent in annualised terms, it has tended to adjust towards that average rate in the following quarter. This finding is consistent with the results of Heath *et al* (2004) indicating that CPI inflation is, statistically, explained reasonably well by a constant over this sample period.

The coefficient on the underlying inflation term ($\hat{\alpha}_{11}$) is also significantly less than zero in the cases of the unadjusted trimmed mean, the market prices excluding volatile items measure, and the optimal Jarque-Bera trim. However, the coefficient is insignificant at ordinary levels of confidence for the four seasonally adjusted measures. Nonetheless, for a number of cases in which $\hat{\alpha}_{11}$ is statistically insignificant, it appears that this coefficient may still be economically significant. A good example is the regression involving the seasonally adjusted trimmed mean, which has an estimated coefficient of $\hat{\alpha}_{11} = -0.5$. This coefficient suggests that when a gap emerges between CPI inflation and the trimmed mean inflation rate, CPI inflation typically adjusts half of the way towards the trimmed mean inflation rate in the following quarter. Broadly similar results hold for the seasonally adjusted optimal Jarque-Bera trim and several of the unadjusted measures.

---

[21] The restriction is accepted by the data at borderline levels of significance for the 1987:Q1–2004:Q4 sample. This is discussed by Heath *et al* (2004) in reference to the results of Granger causality tests. For the 1987:Q1–2004:Q4 sample, the data fail to reject the restriction at around the 5 per cent significance level for all of the seasonally adjusted measures considered except the Jarque-Bera optimal trim. The data reject the restriction for the unadjusted measures at the same significance level. However, this is probably due to the limitation of the gap equation to a lag order of 1. As shown in Heath *et al* (2004), all of these underlying inflation measures are found to Granger cause CPI inflation over the 1987:Q1–2003:Q4 period when multiple lags (of an order chosen using a general-to-specific methodology) are admitted into the specification.

Table 6: Augmented 'Gap' Tests

	1993:Q1–2004:Q4 (48 observations)		
	$\hat{\alpha}_{10}$	$\hat{\alpha}_{11}$	$\hat{\alpha}_{12}$
Quarterly measures: not seasonally adjusted			
Weighted median	0.00	0.08	–0.99
	(0.89)	(0.78)	(0.00)***
30 per cent trimmed mean	0.00	–0.64	–0.56
	(0.47)	(0.10)*	(0.04)**
Market prices excluding volatile items	0.00	–0.40	–0.70
	(0.54)	(0.04)**	(0.00)***
CPI excluding volatile items	0.00	–0.37	–0.80
	(0.93)	(0.20)	(0.00)***
Optimal Jarque-Bera	0.00	–0.57	–0.56
	(0.49)	(0.04)**	(0.02)**
Double-weighted measure	0.00	–0.28	–0.72
	(0.91)	(0.24)	(0.00)***
Quarterly measures: seasonally adjusted			
Weighted median	0.00	–0.17	–0.82
	(0.94)	(0.57)	(0.06)*
30 per cent trimmed mean	0.00	–0.50	–0.62
	(0.93)	(0.17)	(0.03)**
Optimal Jarque-Bera trim	0.00	–0.40	–0.65
	(0.84)	(0.16)	(0.01)**
Double-weighted measure	0.00	–0.28	–0.71
	(0.62)	(0.24)	(0.01)***

Notes: ***, ** and * denote significance at the 1, 5 and 10 per cent levels, respectively; *p*-values are given in parentheses. The estimated coefficients refer to the parameters in Equation (6).

This economic interpretation of the coefficient $\hat{\alpha}_{11}$ suggests that CPI inflation partially adjusts towards measures of underlying inflation when a gap opens up between CPI inflation and the underlying inflation measure in question. The fact that several of the quarterly underlying inflation measures display this property implies that they can be viewed as providing useful supplementary information concerning the 'trend' in inflation in the 1993:Q1–2004:Q4 sample period. As noted above, the results from the longer sample, which contains considerably more variation, suggest that all of the measures considered satisfy this minimum property. Therefore, it seems that the analysis of inflationary trends could benefit from the application of a number of these measures.

In practice, however, it is clear that some measures will be more useful than others for assessing the current trend in inflation. Trimmed means calculated using disaggregated price data that have been seasonally adjusted and measures based on the annual distribution of price changes both tend to be unbiased with respect to CPI inflation. The observed bias in the market prices excluding volatile items measure cannot be remedied in this fashion. In addition, as explained in Section 2, the fact that exclusion-based measures cannot be relied upon to exclude appropriate components from the CPI on a regular basis suggests more generally that they should be used with caution.

Even when bias has been corrected, however, different measures give different signals from time to time. Figure 1 displays the time-series properties of the seasonally adjusted trimmed mean, the trimmed mean based on the annual distribution of price changes, and CPI inflation.

Figure 1: CPI and Underlying Inflation[(a)]

Year-ended

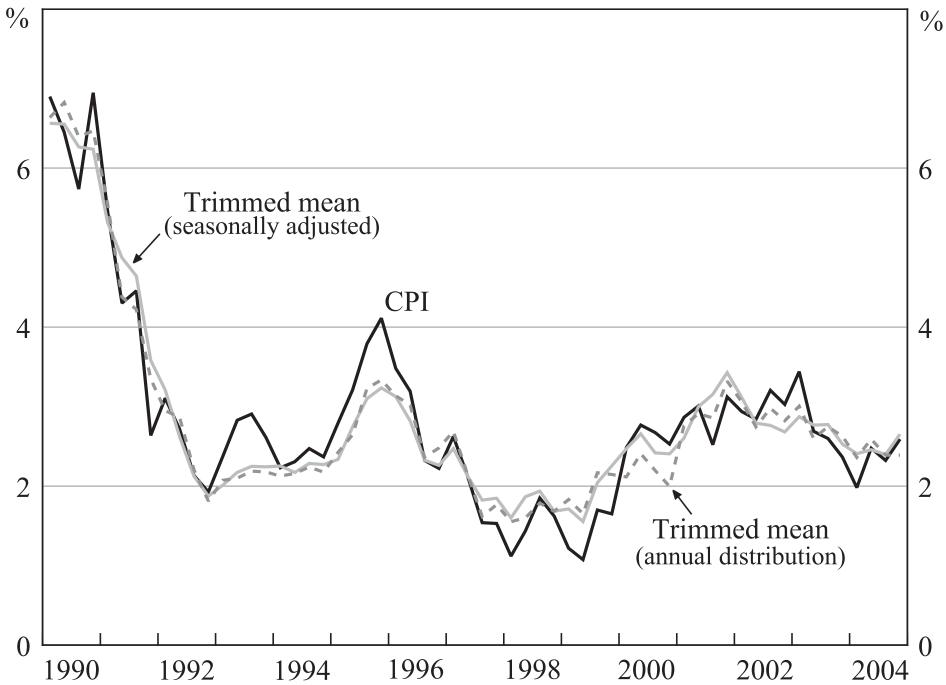

Note: (a) RBA estimates, excluding interest charges prior to the September quarter 1998 and adjusted for the tax changes of 1999/2000.

It is apparent that all three series tend to co-move. However, at times the two underlying inflation measures give somewhat different signals regarding the trend in inflation (for example, in early 2000 and late 2001 the two measures move briefly in opposite directions). The same is true of other underlying inflation measures considered in this paper. As individual measures may be inaccurate guides to the underlying trend in a given quarter, there is thus justification for an approach to analysing underlying inflation that relies more on a 'suite' of underlying inflation measures than on any single measure.[22]

4. Conclusion

Several inferences can be drawn from the results of this paper. While statistical measures of underlying inflation may not have an agreed-upon theoretical foundation, it is clear that in practice they have an advantage over simpler exclusion-based measures. By removing elements from the distribution of price changes on a time-varying basis (rather than applying a fixed rule for all periods), trimmed means or weighted medians exploit the trade-off between efficiency and robustness better than some other core inflation measures. Moreover, they are relatively intuitive, making them easy to interpret in practice.

In some cases, these procedures can result in measures of underlying inflation that are biased with respect to CPI inflation. However, as this paper has shown, there is no evidence of bias for measures calculated using the distribution of annual (rather than the quarterly) price changes, or the distribution of seasonally adjusted quarterly price changes. Of these two techniques, the latter provides a more up-to-date reading of inflation, although it may be subject to revision as seasonal factors are updated over time.

[22] Similar conclusions have been drawn for the UK (Mankikar and Paisley 2002) and Canada (Hogan, Johnson and Laflèche 2001). The benefits of an approach that combines the information provided by a variety of measures has long been recognised by the RBA (for example, Reserve Bank of Australia 1994).

Econometric testing suggests that a variety of measures provide a satisfactory gauge of the underlying trend in inflation. A number of measures assessed in this paper have the desirable property that CPI inflation tends to move towards underlying inflation over time, and thus have the potential to add value to the analysis of inflationary trends in Australia. Nonetheless, to ensure that such analysis is robust, it is desirable to consider a range of measures, and to interpret them in the light of other available information about broader economic developments.

Appendix: Derivation of Underlying Inflation Estimators

First, we show how to derive the double-weighted measure described by Laflèche (1997) as a maximum likelihood estimator of underlying inflation, along similar lines to Diewert's (1995) derivation of the neo-Edgeworthian measure. The model is defined by the equation:

$$\pi_{it} = \pi_t^* + \nu_{it} . \tag{A1}$$

Assume that $E(\nu_{it}) = 0$ and $Var(\nu_{it}) = \sigma_i^2 w_{it}^{-1}$, for $i = 1,\ldots,N$ and $t = 1,\ldots,T$. Assuming further that we are sampling from a normal distribution, the log-likelihood function for this model, aside from a constant, is:

$$\ln L = -\frac{1}{2}\sum_{i=1}^{N}\sum_{t=1}^{T}\ln\left(\frac{\sigma_i^2}{w_{it}}\right) - \frac{1}{2}\sum_{i=1}^{N}\sum_{t=1}^{T}\frac{w_{it}}{\sigma_i^2}\left(\pi_{it} - \pi_t^*\right)^2 .$$

Partially differentiating with respect to π_t^* and σ_i^2, and setting the resulting expressions equal to zero, yields the following system of $T + N$ simultaneous equations that can be used to determine the maximum likelihood estimators for the model:

$$\hat{\pi}_t^* = \frac{\sum_{i=1}^{N} \frac{w_{it}\pi_{it}}{\hat{\sigma}_i^2}}{\sum_{i=1}^{N} \frac{w_{it}}{\hat{\sigma}_i^2}} \qquad \text{for } t = 1,\ldots,T\text{; and}$$

$$\hat{\sigma}_i^2 = \frac{1}{T}\sum_{t=1}^{T} w_{it}(\pi_{it} - \hat{\pi}_t^*)^2 \qquad \text{for } i = 1,\ldots,N.$$

Note that the formula for $\hat{\sigma}_i^2$ multiplies relative price inflation at each point in time by the effective expenditure weight, and thus does not correspond to the usual notion of a variance.

Similar measures have been proposed before (for example, Laflèche 1997, Marques *et al* 2000, Aucremanne 2000), but by using the reciprocal of the sample standard deviation they only approximate this estimator. The neo-Edgeworthian measure is usually constructed in a similar *ad hoc* manner, although it has been more rigorously estimated as the solution to a full system of equations for the US and for the euro area by, respectively, Wynne (1997) and Vega and Wynne (2003). Regular breaks in the Australian CPI would make precise estimation of either the double-weighted or neo-Edgworthian measure difficult in the Australian context.

Second, we derive the weighted mean of the distribution of price changes, in the spirit of Clements and Izan (1987). One could also do this for the trimmed distribution of price changes, assuming that the trimmed price changes are simply outliers that have been 'cleaned' from the data. In the Australian case, this could make the assumption that we are sampling from a normal distribution more defensible. The model is comprised of Equation (A1) and the assumptions that $E(v_{it}) = 0$ and $Var(v_{it}) = \sigma_t^2 w_{it}^{-1}$. The only difference compared with the previous model is that now σ_t^2 is assumed to be common to all goods and services. This time, the log likelihood function is:

$$\ln L = -\frac{1}{2}\sum_{i=1}^{N}\sum_{t=1}^{T}\ln\left(\frac{\sigma_t^2}{w_{it}}\right) - \frac{1}{2}\sum_{i=1}^{N}\sum_{t=1}^{T}\frac{w_{it}}{\sigma_t^2}(\pi_{it} - \pi_t^*)^2 .$$

Differentiating with respect to π_t^* and setting the partial derivative equal to zero, we have:

$$\frac{1}{\hat{\sigma}_t^2}\sum_{i=1}^{N} w_{it}(\pi_{it} - \hat{\pi}_t^*) = 0 \qquad \text{for } t = 1,\ldots,T.$$

Since $\hat{\sigma}_t^2$ cancels out, the maximum likelihood estimator for the systematic component of inflation is:

$$\hat{\pi}_t^* = \sum_{i=1}^{N} w_{it}\pi_{it} \qquad \text{for } t = 1,\ldots,T.$$

References

Aucremanne L (2000), 'The use of robust estimators as measures of core inflation', National Bank of Belgium Working Paper No 2.

Australian Bureau of Statistics (ABS) (2004), 'Seasonally adjusted and trend estimates for the consumer price index (CPI)', *Australian Economic Indicators*, ABS Cat No 1350.0, December, pp 3–8.

Ball L and NG Mankiw (1995), 'Relative-price changes as aggregate supply shocks', *Quarterly Journal of Economics*, 110(1), pp 161–193.

Bryan MF and SG Cecchetti (1993), 'The consumer price index as a measure of inflation', Federal Reserve Bank of Cleveland *Economic Review*, 29(4), pp 15–24.

Bryan MF and SG Cecchetti (1994), 'Measuring core inflation', in NG Mankiw (ed), *Monetary policy*, University of Chicago Press, Chicago, pp 195–215.

Bryan MF, SG Cecchetti and RL Wiggins II (1997), 'Efficient inflation estimation', NBER Working Paper No 6183.

Bryan MF and SG Cecchetti (1999), 'The monthly measurement of core inflation in Japan', *Bank of Japan Monetary and Economic Studies*, 17(1), pp 77–101.

Cecchetti SG and G Debelle (2004), 'Has the inflation process changed?', prepared for the Bank for International Settlements annual conference 'Understanding low inflation and deflation', Brunnen, Switzerland, 18–19 June.

Clements KW and HY Izan (1987), 'The measurement of inflation: a stochastic approach', *Journal of Business and Economics Statistics*, 5(3), pp 339–350.

Cutler J (2001), 'Core inflation in the UK', External MPC Unit Discussion Paper No 3, Bank of England.

Diewert WE (1995), 'On the stochastic approach to index numbers', University of British Columbia Discussion Paper No 95-31.

Diewert WE (1997), 'Commentary for "Alternative strategies for aggregating prices in the Consumer Price Index"', Federal Reserve Bank of St. Louis *Review*, 79(3), pp 126–138.

Dixon R and GC Lim (2004), 'Underlying inflation in Australia: are the existing measures satisfactory?', *Economic Record*, 80(251), pp 373–386.

Faber SM and AM Fischer (2000), 'Core inflation in Switzerland', Swiss National Bank *Quarterly Bulletin*, December, pp 66–75.

Hampel FR, EM Ronchetti, PJ Rousseeuw and WA Stahel (1986), *Robust statistics: the approach based on influence functions*, Wiley, New York.

Heath A, ID Roberts and TJ Bulman (2004), 'Inflation in Australia: measurement and modelling', in C Kent and S Guttmann (eds), *The future of inflation targeting*, Proceedings of a Conference, Reserve Bank of Australia, Sydney, pp 167–207.

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

Huber PJ (1964), 'Robust estimation of a location parameter', *Annals of Mathematical Statistics*, 35(1), pp 73–101.

Kearns J (1998), 'The distribution and measurement of inflation', Reserve Bank of Australia Research Discussion Paper No 9810.

Keynes JM (1930), *A treatise on money*, Harcourt Brace and Co, New York.

Laflèche T (1997), 'Statistical measures of the trend rate of inflation', Bank of Canada *Review*, Autumn, pp 29–47.

Machado JF, CR Marques, PD Neves and AG da Silva (2001), 'Using the first principal component as a core inflation indicator', Banco de Portugal Working Paper No 9-01.

Mankikar A and J Paisley (2002), 'What do measures of core inflation really tell us?', Bank of England *Quarterly Bulletin*, 42(4), pp 373–383.

Maria JR (2004), 'On the use of the first principal component as a core inflation indicator', Banco de Portugal Working Paper No 3-04.

Marques CR, PD Neves and AG da Silva (2002), 'Why should central banks avoid the use of the underlying inflation indicator?', *Economics Letters*, 75(1), pp 17–23.

Marques CR, PD Neves and LM Sarmento (2000), 'Evaluating core inflation indicators', Banco de Portugal Working Paper No 3-00.

OECD (Organisation for Economic Co-operation and Development) (2005), 'Measuring and assessing underlying inflation', OECD *Economic Outlook*, 77(1), pp 322–343.

Quah D and SP Vahey (1995), 'Measuring core inflation', *Economic Journal*, 105(432), pp 1130–1144.

Reserve Bank of Australia (1994), 'Measuring "underlying" inflation', Reserve Bank of Australia *Bulletin*, August, pp 1–6.

Reserve Bank of Australia (1998), 'The implications of recent changes to the consumer price index for monetary policy and the inflation target', Reserve Bank of Australia *Bulletin*, October, pp 1–5.

Roger S (1997), 'A robust measure of core inflation in New Zealand, 1949–96', Reserve Bank of New Zealand Discussion Paper No G97/7.

Roger S (1998), 'Core inflation: concepts, uses and measurement', Reserve Bank of New Zealand Discussion Paper No G98/9.

Shiratsuka S (1997), 'Inflation measures for monetary policy: measuring the underlying inflation trend and its implication for monetary policy implementation', *Bank of Japan Monetary and Economic Studies*, 15(2), pp 1–26.

Shu C and A Tsang (2004), 'Alternative measures of core inflation on the mainland', Research Memorandum, Hong Kong Monetary Authority, May.

Wynne MA (1997), 'Commentary for "Measuring short-run inflation for central bankers"', Federal Reserve Bank of St. Louis *Review*, 79(2), pp 163–168.

Wynne MA (1999), 'Core inflation: a review of some conceptual issues', Federal Reserve Bank of Dallas Working Paper No 9903.

Vega J-L and MA Wynne (2003), 'A first assessment of some measures of core inflation for the euro area', *German Economic Review*, 4, pp 269–306.