G97/ 7

A robust measure of core inflation in New Zealand, 1949-96

Scott Roger[1]

March 1997

Abstract

This paper develops a stochastically-based method of measuring core inflation, extending earlier research by Bryan and Cecchetti (1993) and Roger (1995). The approach exploits the persistent tendency towards high kurtosis evident in the cross-sectional distribution of consumer price changes evident in New Zealand and elsewhere. High kurtosis makes the sample mean a less efficient and less robust estimator of the population, or underlying, mean price change than is an order statistic such as the median.

The quarterly cross-sectional distribution of price changes in New Zealand over the 1949-96 period also exhibits chronic right skewness. This tends to bias a median measure of inflation downwards relative to the population mean. It is found that a slightly higher percentile of the price change distribution reliably corrects for the asymmetry of the distribution, while maintaining its efficiency and robustness relative to the sample mean as an estimator of the population mean.

This percentile of the distribution, which corresponds on average to the mean, filters out the effects on the mean of relative price shocks, and is interpreted as a measure of core inflation. Testing of the measure suggests that the price movements being filtered out primarily reflect supply shocks having only a temporary impact on inflation.

The measure offers advantages over current approaches in terms of transparency and verifiability, and is also much better-suited to the filtering of supply shocks not directly affecting clearly identifiable components of the CPI in a readily measured way. Nonetheless, some further research on the stability of the distribution of price changes at a more disaggregated level is warranted before considering this kind of measure as a fully reliable indicator of core or underlying inflation.

JEL classifications: C43, E31, E52

Reserve Bank of New Zealand
PO Box 2498
Wellington, New Zealand
Fax: (+644) 473-1209
Internet: RogerS@RBNZ.GOVT.NZ

A robust measure of core inflation in New Zealand, 1949-96

Summary

Measures of 'core' or 'underlying' inflation seek to remove the distortionary influence of outlier price changes associated with supply shocks. The motivation for constructing such measures is basically that, in general, monetary policy ought not to react to the direct price level consequences of supply shocks.

Current standard methods of estimating 'core' or 'underlying' consumer price inflation are generally *ad hoc* in various respects and quite limited in their capacity to handle shocks that do not fit neatly into identifiable CPI categories. This is not very satisfactory, especially if policy performance and credibility is substantially judged on the basis of the core inflation measure.

In this paper well-established statistical methods are used to develop a measure of core inflation. Essentially, the method involves aggregating the movements in CPI component price changes in a way that is less influenced by exceptional price movements than is the mean. If exceptional relative price movements are typically associated with supply disturbances, the measure can provide a simple, robust and readily verified measure of core inflation.

The cross-sectional distribution of New Zealand CPI subgroup price changes over the 1949-96 period is found to be chronically right-skewed and kurtotic. This suggests that a central percentile of the cross-sectional distribution of price changes - though not the median percentile - should provide a substantially more efficient and robust estimator of the core inflation rate than does the sample mean.

Tests of this measure of core inflation are quite positive. They suggest in particular that, by filtering out the effect on inflation of relative price movements, the measure does not discard a significant useful source of leading information on future inflation. Indeed, the evidence points the other way; that the relative price shocks mainly add noise to inflation, and that by filtering out such noise, we can gain a more accurate picture of both current and future inflation.

I.   Overview

Since the end of the 1970s an increasing number of countries have recognised the control of inflation as the primary objective of monetary policy. In those countries, beginning with New Zealand, that have adopted explicit targets for inflation, the targets have been expressed in terms of the rate of change in the Consumers Price Index (CPI).

Yet it is also widely recognised that, at times, the CPI inflation rate will give a misleading impression of the general trend of prices. Consequently, central banks in many countries construct measures of 'core' or 'underlying' inflation that purport to more accurately represent the general trend of prices than does the official or 'headline' measure of inflation.

This paper extends and refines the basic proposition put forward in Bryan and Cecchetti (1993) and Roger (1995), that so-called *robust* measures of inflation, such as the trimmed-mean or median measure of inflation offer simple, reliable and transparent methods for estimating a measure of underlying inflation, and the impact of relative price disturbances of the CPI.

Section II reviews standard methods of measuring 'core' inflation. In each case, the aim is to remove the influence of 'unrepresentative' or outlier price movements on the aggregate, mean-based, measure of inflation. These methods are far from ideal, and essentially *ad hoc*.

Section III discusses a stochastic or probabilistic approach to the measurement of the central tendency of inflation, based on long-standing principles of statistics. If the distribution of price changes is typically characterised by high kurtosis (a high probability of exceptional price changes), the sample mean rate of inflation provides a less reliable estimate of the general trend of inflation than do robust estimators such as a trimmed mean or median.

In Section IV it is shown that the quarterly distribution of consumer price changes in New Zealand over the 1949 to 1996 period has typically shown high kurtosis and positive skewness. This points to the potential superiority - in terms of efficiency and robustness - of a trimmed mean or median-based measure over the sample mean as estimators of core inflation.

In Section V a median-based measure of core inflation is developed. Because the distribution of price changes is typically right-skewed, the median tends to understate the mean. This bias can be eliminated by taking a percentile of the price distribution slightly above the 50th. This measure is shown to be substantially more efficient than the sample mean as an estimator of the population or underlying mean rate of inflation.

In Section VI the ability of the median-based measure of core inflation to screen out supply shocks as opposed to demand shocks is examined. The evidence suggests that the relative price shocks screened out using a median-based measure have the characteristics that are normally associated with 'supply' shocks in the economics literature.

Section VII offers concluding comments.

II.   Current methods of measuring 'core' or 'underlying' inflation

It is common in many countries for central banks, finance ministries, statistical agencies or private sector economists to distinguish between inflation as measured by official price series, such as the Consumer Price Index (CPI), and some concept of inflation variously described as 'underlying', 'trend' or 'core' inflation.

However it is described, the basic idea is that, at times, exceptional movements of particular prices represented in the official aggregate price index will give a 'distorted' impression of the general rate or central tendency of price movement or inflation in the sense that the movement in the aggregate price index is quite different from the movement of most prices comprising the index.

The challenge is to define a measure of price movement or inflation that is free of or, at least, less prone to such distortions. Ideally, the measure chosen should be:

Timely. If the measure is not available for use in a timely manner either in the first instance, or is subject to revision over an extended period, its practical value will be severely impaired

Robust and unbiased. If the measure cannot be relied upon to remove the sorts of distortions that it ought to, or if it shows a systematically different trend than the series from which it is derived, it will provide false signals, lead to policy biases and fail to gain public credibility.

Verifiable. If the measure of 'core' inflation is not readily verifiable by anyone other than its creator, it is unlikely to have great credibility. As a result, it will have limited practical value either as a measure against which to assess monetary policy performance, or as a guide for inflation expectations and, thereby, wage and price determination.

The most common approaches to deriving measures of 'core' inflation from the CPI are described in Roger (1995). These include:

Adjustment by exclusion. This method involves modifying the domain of the CPI to exclude component price series judged likely to display perverse behaviour (e.g. interest rate components) or to be prone to exceptional or non-representative price (e.g. seasonal food and energy components). By excluding such series the modified index should be less subject to distortion than the original index.

This approach cannot be considered to be robust, unless one can be sure that distortionary price shocks will not affect the components that remain in the index. In this regard, past volatility of particular series may not be a reliable guide to future volatility. Moreover, even if prescient guesses are made as to what will turn out to be the most volatile price series, a cut-off point must inevitably be chosen arbitrarily, leaving the measure exposed to low probability but potentially large magnitude distortions to the 'core' price index.

Adjustment by smoothing. Typically this involves some form of time-series averaging, either at the level of the individual price series or at the aggregate level, to remove the effects of deterministic seasonality. By removing these effects, a clearer sense of the ongoing trend may be obtained.

Unfortunately, this approach also fails on robustness grounds, because it is unable to filter irregular price shocks or stochastic seasonality. Nor can this approach be considered to be timely. Smoothing procedures almost inevitably involve some form of averaging of current period price movements with those of earlier periods, so that the measure of 'core' inflation will almost inevitably be a lagging indicator of the true trend. This results in an awkward trade-off to be made between the degree of smoothing of inflation data (to minimize false signals about the general trend of inflation), and its timeliness (to minimize tardiness in policy adjustments).[2]

Specific adjustment. This involves modifying recorded price changes at specific times to eliminate the influence of specific developments on the measured aggregate inflation rate.

This approach has the advantage of allowing judgment to be brought to bear in determining which price movements are 'exceptional'. In most instances, however, it is also intrinsically the least systematic, transparent or verifiable precisely because it often depends on large elements of discretion or judgment.[3] The key element of judgment involved is typically in deciding which part of a price movement constitutes part of 'general' inflation, and which part is a relative price shift that should not be treated in the same way as other relative price shifts that occur all the time. In practice, it is very rare that such judgments can be made in any easy, consistent and defensible way.

In addition, the specific adjustment method is particularly handicapped in dealing with shocks to prices or costs below the retail level. For example, petroleum price shocks will have a direct impact on a fairly wide range of retail prices. Identifying and quantifying this impact, spread across several CPI categories is virtually impossible.

Each of the standard approaches outlined above suffers from important drawbacks either in terms of reliability or transparency. Both aspects, however, assume a particular importance in countries where the central banks are charged with achieving a well-defined inflation target.

III. A stochastic approach to measuring the central tendency of inflation

The elementary point that there may exist ... estimators superior to least squares for the non-Gaussian linear model is a well kept secret in most of the econometrics literature.
R. Koenker and G. Bassett, 1978

1. The stochastic approach

In this section the estimation of the central tendency of inflation is approached from an explicitly *stochastic* perspective.[4] The problem can be characterised as follows. In any given period, there is an array or distribution of price changes and some of these may be quite unrepresentative of the general trend. In general, it is not possible to pre-determine which particular prices will be

---

[2]   This point is noted by Cecchetti (1996).
[3]   See Roger (1995), McCallum (1995) and Spiegel (1995) for more detailed comment on the specific adjustment approach as applied in the New Zealand context.
[4]   The essentials of the stochastic approach to price measurement are spelled out by Theil (1967), and the approach is reviewed thoroughly by Diewert (1995b)

affected, or when, or by what the magnitude. These uncertainties constitute the essential weaknesses of measuring core inflation using the exclusion or smoothing methods.

One way to pre-specify a measure of the general or 'core' rate of change in prices in any given period is to base it on a measure of 'central tendency' of price changes which is fairly insensitive to extreme price changes, whatever their provenance and whenever they occur.

A key element of this approach is to think of the distribution of price changes in a particular price index such as the CPI as being a particular *sample* drawn from a characteristic *population* distribution of price changes. In each period we measure and observe a sample drawn from the aggregate *population* distribution of price changes.

The *sample* distribution will routinely differ from the *population* distribution for a number of reasons. One is due to the standard type of measurement error associated with the fact that the recorded price changes are only a sampling from the total number of price changes within the relevant domain of actual price changes, so 'bad' samples are always possible. In addition, errors may be introduced inadvertently at any of several stages of the CPI compilation process.

For the purposes of this analysis, the observed distribution of consumer prices in a particular period is still thought of as being a sample, even if measurement was comprehensive and error-free. To illustrate, suppose that petroleum prices rose very sharply in a particular quarter and that this price increase was measured and recorded with complete accuracy. The extreme price rise would result in a distribution of price changes making up the CPI that was quite different from the typical or characteristic distribution of CPI price changes (in this case, the sample distribution would be strongly right-skewed). The price distribution observed in the quarter would be considered as a 'bad' draw or sample in the sense that it was drawn from a distribution that was unrepresentative of the typical or population distribution of price changes.

The basic idea behind so-called *robust* estimators of central tendency is to define a measure of central tendency that is likely to be relatively unaffected by unusual or 'bad' sample distributions of price changes. The strength of such measures is not that they are tailor-made for every situation; their strength lies in their ability to work pretty reliably in even quite exceptional circumstances.

If sample price changes were being drawn from a well-known, stable population distribution, finding a reliable measure of central tendency would be fairly straight-forward. But the population distribution of price changes is not known, and may vary over time. The measurement of 'central tendency' of price changes in such circumstances becomes an exercise in statistical inference. This issue is discussed below.

2. Efficient and robust estimation of the population or 'underlying' mean

If we cannot observe the true or population distribution we are limited to an estimate of the underlying or population mean based on the sample price changes. In choosing an estimator of the population mean, three properties are highly desirable: unbiasedness, efficiency, and robustness.

If the population distribution of price changes can be assumed to approximately Normal, then the mean of samples from that distribution will be the best estimator of the true mean in the

sense of being unbiased and efficient.[5]   However, if the distribution is not Normal, or unknown, then the sample mean should still be an unbiased estimator of the population mean, but it may not be as efficient or robust as a variety of other estimators.

The relative efficiency of alternative estimators of the population mean is particularly sensitive to the kurtosis (the fourth moment) of the distribution.[6]   If the population distribution of price changes is characterised by high kurtosis, samples drawn from the distribution will include a higher proportion of extreme values than is characteristic of the Normal distribution.   It is precisely such extreme price movements that are regarded as distorting the sample mean.

Unfortunately, there is a common perception that the sample mean is the most efficient estimator for *all* distributions, not just for the Normal distribution. This is false. In fact, it does not require much change in the shape of the distribution for the sample mean to become a relatively inefficient estimator of the population mean.

In general, as the kurtosis of the distribution increases, the efficiency of estimators - like the sample mean - that place a high weight on observations in the tails of the distribution falls relative to estimators that place a low weight on observations in the tails.[7]

The Normal distribution, with a kurtosis of 3, occupies the middle ground between 'high' kurtosis (leptokurtic) and 'low' kurtosis (platykurtic) distributions. This is because the most efficient estimator for this distribution - the sample mean - places equal weight on all the observations. For distributions with a kurtosis of less than 3, the most efficient estimators place relatively high weight on observations in the tails, while for distributions with kurtosis greater than 3, the most efficient estimators place relatively low weight on observations in the tails. Such estimators are known as order statistics, because the weight attached to observations depends on their order or ranking in the distribution.

A common and particularly simple estimator that places a relatively low weight on observations in the tails of the distribution is the trimmed-mean. This measure involves zero-weighting of some (essentially arbitrary) proportion of the observations at each end of the distribution of observations. The trimming of the tails, it may be noted, need not be symmetric.

More complex, but also fundamentally arbitrary, weighting schemes are offered by the class of *L*-statistics (see, e.g., Judge *et al* (1988), Huber (1981) or David (1981)). These involve linear combinations of order statistics. Whereas the trimmed-mean assigns zero weight to, say, the top and bottom 5% of observations, and equally weights the central 90% of observations, a more complex *L*-statistic could have a gradual (and often non-linear) decrease in weights to outlying

---

[5]   In this paper, references to the 'Normal' distribution or 'Normality' use a capitalised 'N' to distinguish the technical term relating to that particular distribution from the everyday meaning of 'normal' and 'normality'.

[6]   The kurtosis is critical since, as Kendall and Stuart (1969) observe:
       "...the sampling variance of a moment depends on the population moment of twice the order..." (p.234). Thus the variance (second moment) of the sample mean depends on the kurtosis of the population distribution.

[7]   As is shown in Roger (1995), the mean is a least squared errors estimator, involving a 'loss' function that places a high 'penalty' on extreme price changes in determining the 'centre' of the distribution. In contrast, the median is a least absolute errors estimator, involving equal penalties on all price changes in determining the centre of the distribution. The median is, therefore, less affected by extreme price changes than is the mean.

observations.

The mean and the median can be thought of as very particular *L*-statistics. In the case of the mean, equal weights are placed on all of the observations, with the important consequence that the ordering of the observations ceases to have any effect on the value of the statistic. By contrast, the median (like any other individual percentile of the distribution) is an extreme order statistic: all but a single observation are zero-weighted.

Given the potentially infinite number of *L*-statistics to choose from, how should the 'best' estimator be selected? At this point it is useful to distinguish between the most efficient estimator for a particular distribution and the most reliable estimator for a variety of distributions.

If the fundamental distribution of price changes is known, then it might well be possible to find an estimator that is demonstrably more efficient than all others. When the underlying or population distribution is not known, however, it is appropriate to focus on the *robustness* of the estimator. A *robust* estimator may not be the most efficient estimator, but will rarely perform very poorly. In other words, in circumstances of uncertainty, dependable approximation is a more desirable property of an estimator than highly erratic precision. In general, the sample mean is not a very robust estimator, and declines rapidly in efficiency as the kurtosis of the distribution increases.

Hogg (1967) offers a simple scheme for selection of a robust estimator, based on extensive Monte Carlo testing of alternative measures applied to a wide range of frequency distributions:
- if the kurtosis of the distribution is between 2 and 4, the sample mean is the recommended estimator;
- if kurtosis is between 4 and 5.5, then the 25% trimmed-mean performs well;
- if kurtosis is above 5.5, the sample median is recommended.

Koenker and Bassett (1978) compare the variances of the sample mean, the median, the 10% and 25% trimmed-means and two slightly more complex *L*-statistics as estimators of the population mean of a number of specific distributions. Their results reinforce the essentials of Hogg's scheme: that the more kurtotic or 'fat-tailed' the distribution, the lower the weight placed on outlier observations by the most efficient estimator; that the mean is not very robust to departures from normality; and that estimators such as the trimmed-mean or median are robust for a wide range of (leptokurtic) distributions.

The thrust of these findings, in short, is that the most robust and efficient estimator of the population or underlying mean of the distribution cannot be specified *a priori*; it is sensible to look at the empirical distribution first. What can be said *a priori* is that even if the mean is the most efficient estimator, it is unlikely to be particularly robust.

IV.   The distribution of consumer price changes in New Zealand, 1949-96

It is illogical to construct a narrow model for the underlying distribution prior to sampling and then to make statistical inferences about the distribution characteristics from the sample, without worrying whether or not the model is appropriate
R. Hogg 1967

In this section, the distribution of CPI component price changes is examined. The examination extends and refines that of Roger (1995) in several respects.

Roger (1995) found that at a quite disaggregated level, the distribution of price changes in New Zealand showed significant right-skewness over the 1981-95 period. Two possible reasons for this were discussed. One was that fairly elevated rates of inflation through much of the period might be causing the right skewness. Another possibility was that substantial increases in various government charges (related to the economic reform process) might be the cause. An implication of both of these is that as inflation is stabilised at a low level and as the shift towards user charges and market prices for government services is completed, the right-skewness of the distribution of price changes should disappear.

One way of examining the validity of these hypotheses is to examine the distribution of price changes at a more aggregated level (minimising the government charge effect on the shape of the distribution), and over a longer time frame including previous periods of low inflation. That is what is done in this paper.

1.   The data

In this paper, analysis is restricted to data covering the period from 1949Q1-1996Q4 at the subgroup level of aggregation.[8] The subgroup level of aggregation as chosen partly because the characteristics of the distribution at this level are essentially similar to those of the data at lower levels of aggregation, and partly because only data at this level of aggregation are available as far back as 1949.

The period from 1949 to 1996 spans nine CPI regimens: 1949-55, 1956-65, 1966-74, 1975-77, 1978-80, 1981-83, 1984-88, 1989-93, 1994-98. Throughout, the CPI has been calculated as a Laspeyres, Dutot type index.[9] Prior to 1975, the CPI was basically a consumption price index, while since then it has been an expenditure price index. The main implication of this shift in methodology was to include house prices or construction costs, as well as mortgage and other credit costs directly into the CPI.

The New Zealand CPI is a quarterly series. For a number of component series, however, measurement has been at the semi-annual frequency. In most instances, the items affected have had a very low weight in the overall regimen. In the 1949-56 regimen, however, housing costs (having a large weight in the regimen) were measured only semi-annually. For the purposes of this analysis, the data were modified by interpolating (geometrically) between the quarters in which measurements were made.

---

[8]   Roger (1997a) examines the cross-sectional distributions of New Zealand CPI price changes at three levels of aggregation, of which the subgroup level is the most aggregated.

[9]   For a discussion of alternative indexes and their properties, see Diewert (1995a).

Two additional adjustments have been made to the data. First, credit service costs (mainly mortgage interest payments) have been systematically excluded from the data, from their introduction into the CPI from 1974Q4. Second, the direct impact on the CPI of the introduction of the Goods and Services Tax (GST) in 1986, and the subsequent increase in the GST rate in 1989, has been removed (as described in Roger (1996)).

2.   Moments of the distribution of price changes

The distribution of subgroup level price movements over the 1949-96 period corroborate and reinforce the earlier findings of Roger (1995) for much more disaggregated data over the 1981-95 period. Right skewness and high kurtosis are found to be persistent features of the distribution of changes in consumer prices. What is particularly striking is the apparent stability of the shape of the distribution, despite substantial shifts in its location, changes in the composition of the CPI and its method of calculation, and despite shifts in the monetary policy regime and the degree of openness of the economy and the extent of government intervention in price setting in the economy.

The moments of the cross-sectional distribution are calculated on three somewhat different bases:

- In the first method, weighted sample moments are calculated quarter by quarter.[10] The quarterly values are then used to calculate multi-period averages and higher moments, adjusting the moments of moments for the number of quarterly observations. By this method, we gain the precision of many quarterly observations of moments, but each quarterly moment is calculated from relatively few observations, particularly at the subgroup level of aggregation.

- The second method involves pooling of normalised quarterly distributions over a calendar year, and then calculating moments for the year as a whole. The normalisation necessarily eliminates information about changes in the means and variances, but, by pooling quarterly data, more precision is gained in estimating skewness and kurtosis. Partly offsetting this gain in precision will be the loss of precision from fewer (annual) observations over which to calculate long-term averages and higher moments of the moments.

- The third method pools the normalised quarterly data over multi-year periods and then calculates the higher moments.

Table 1 reports the means, medians and standard deviations of the first four (adjusted) sample moments of the cross-sectional distribution of price changes over the 1949-96 period. The 'quarterly', 'annual' and 'multi-year' figures refer to the averages, medians and standard deviations of the sample moments calculated according to the three methods described above.

---

[10]   The appropriate, unbiased estimates of population moments for an unequally weighted distribution are derived in Roger (1997a).

Table 1
Moments of the distribution of consumer price changes at the Subgroup level of aggregation, 1949Q2-96Q4. Sample size: 191 quarters / 48 years

Moments of sample values	Calculation basis	Adjusted sample moments of distribution			
		Mean	Std. Dev.	Skew	Kurtosis
Mean	**Quarterly**	1.6	2.0	0.7	7.2
	Annual	-	-	0.7	6.8
	Multi-year	-	-	0.6	6.7
Median	**Quarterly**	1.2	1.7	1.0	5.5
	Annual	-	-	0.7	5.8
	Multi-year	-	-	-	-
Standard Deviation	**Quarterly**	1.3	1.2	1.9	7.3
	Annual	-	-	0.8	3.9
	Multi-year	-	-	-	-

The sample moments shown in table 1 indicate that the distribution of price changes is not typically Normal, but right-skewed and leptokurtic (fat-tailed). A more striking impression of this is given by figure 1, showing the pooled normalised cross-sectional distribution of quarterly price changes at the subgroup level of aggregation over the 1949-96 period.

Figure 1
Frequency distribution of CPI subgroup level quarterly price changes 1949-96
(Pooled normalised percentage price changes)
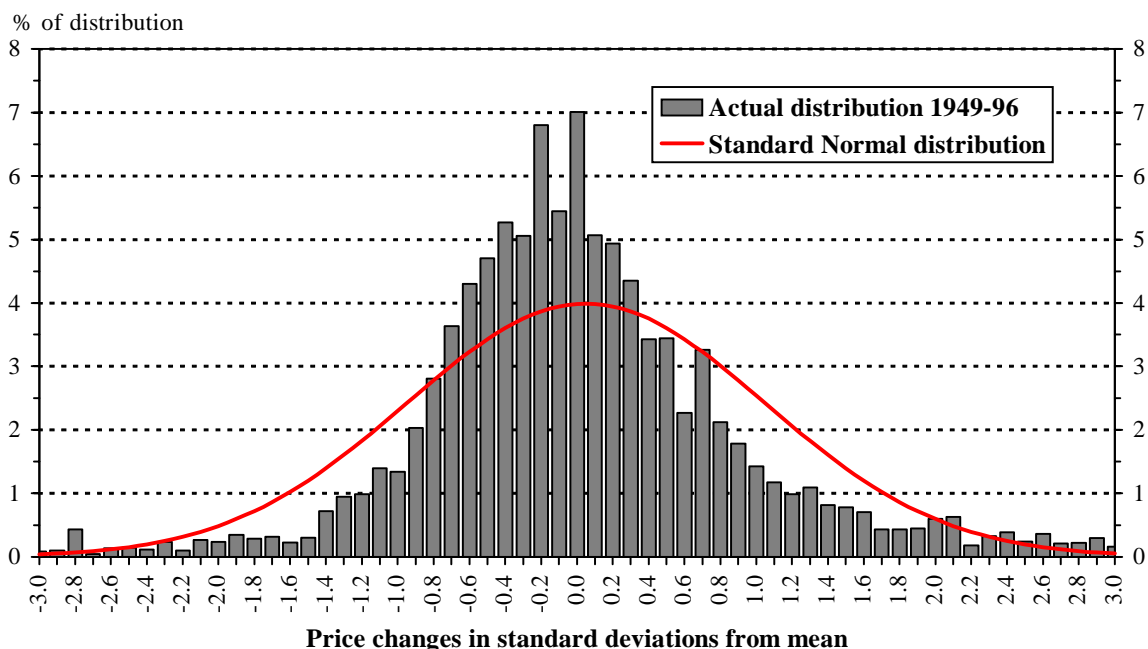
It may be noted that the tendency towards right-skewness together with excess kurtosis in the distribution of consumer price changes does not appear to be an idiosyncratic feature of New Zealand data. The same characteristics appear to be present in varying degrees in US, Canadian, British, French and Australian CPI data, to the best of this writer's knowledge. Bryan and Cecchetti (1996) report sample moments for US data. Their results for the CPI disaggregated into 36 components and based on quarterly averages are roughly comparable to New Zealand

subgroup level data. For the 1967Q1-96Q1 period, Bryan and Cecchetti report skewness of 0.23 (versus 0.79 for New Zealand) and kurtosis of 8.07 (versus 7.65 for New Zealand).

V.   A robust measure of inflation

The final practical conclusion, therefore, is that the weighted median serves the purposes of a practical barometer of prices ... as well as, if not better than, formulae theoretically superior. In spite, however, of the peculiar simplicity and ease of computation which characterises the median, and in spite of Edgeworth's strong endorsement, it remains still almost totally unused, if not unknown.
Irving Fisher, 1922

1.   The problem of skewness and a solution

The moments reported in table 1 indicate that, the distribution of price changes in New Zealand, even at a high degree of aggregation, has not been even close to Normal over the 1949-96 period. In particular, the evidence points to chronically high kurtosis, as well as moderate but also chronic right-skewness.

The discussion in Section III suggests that if the distribution is characterised by kurtosis on the order of that shown in table 1, then the sample mean is likely to be a much less robust or efficient estimator of the underlying or population mean of the distribution than would be an estimator placing less weight on extreme price changes.

There is, however, an important obstacle to overcome. The discussion in standard textbook analysis of the efficiency of the mean relative to other estimators of the population mean (such as the median or trimmed mean measures) is based on the assumption that each of the alternative measures is an unbiased or, at least, a consistent estimator of the population mean. This reflects a common assumption that the population distribution is symmetric or, at least, not skewed.

In the case at hand, however, the textbook assumption does not hold: right-skewness is a chronic feature of the empirical distribution. As a result, there appears to be a dilemma between using the sample mean - a relatively inefficient, but unbiased, estimator of the population mean - or using a relatively efficient, but biased, estimator based on an order statistic.[11]

The dilemma can be resolved quite simply, at least under certain circumstances. For distributions for which the mean exists, we know that the lowest ranked observation of the distribution will be a consistently downward biased estimator of the population mean, while the

---

[11]   The bias raises an interesting issue. The 'bias' is a relative concept: the mean is biased relative to the median and *vice versa*. It is not obvious that we should conclude that the mean is somehow 'less' biased as a measure of central tendency than the median. The geometric mean is also biased relative to the arithmetic mean, yet many statisticians would say that it is the arithmetic mean that is upward biased as opposed to saying that the geometric mean is downward biased. If the median is a less biased estimator of the geometric mean than is the arithmetic mean, perhaps we should regard the median not only as a relatively efficient measure of central tendency, but also as a less biased measure of 'true' central tendency.

highest ranked observation will be consistently upward biased. Somewhere in between will be an order statistic or percentile of the distribution that is, on average, an unbiased estimator of the population mean. In the case of any symmetric distribution (eg. the Normal distribution) the 50th percentile observation - the median - or any other order statistic centred on the 50th percentile, will be an unbiased estimator of the population mean.

If the population distribution is skewed, however, a different percentile will correspond to the population mean. In the case of a chronically right-skewed distribution, a percentile somewhat above the 50th percentile or median will be an unbiased estimator of the population mean. In this paper the percentile which corresponds to the *sample* mean of the distribution will be called the *mean percentile*, while the percentile corresponding to the *population* mean will be called the *population mean percentile*.

Now, although the sample mean may not be the most efficient estimator of the population mean, it should be an unbiased estimator. By transitivity, therefore, the percentile of the empirical distribution that, *on average*, corresponds to the sample mean should also be an unbiased estimator of the population mean.

An important potential difficulty with the approach outlined above is that, if the shape of the population distribution varies over time, the percentile of the distribution corresponding to the population mean (the *population mean percentile*) will also be time-varying. Of particular concern is the possibility that the shape of the population distribution may be systematically related to the average inflation rate (i.e. that the shape and 'location' of the distribution may not be independent).

Ball and Mankiw (1994) and Balke and Wynne (1996) present models in which the skewness of the distribution of price changes is expected to be positively correlated with the rate of inflation, at least in the short-term. Bryan and Cecchetti (1996) argue that positive correlation predicted by the Ball and Mankiw model will only hold in the short-term (i.e., the period over which a significant number of prices in the economy are 'sticky' in nominal terms), while in the Balke and Wynne model, the positive correlation will be more persistent (because it is rooted in 'stickiness' in the production structure of the economy rather than in 'menu' cost dynamics).

An implication of the positive correlation hypothesis is that the use of a time-invariant percentile price change, as an estimator of the population mean price change, will tend to understate the trend rate of inflation if the trend is rising, and overstate it when the trend rate is decreasing, at least over the short term.

2. The sample and population mean percentiles

(i) Sample mean percentiles

Figure 2, below, shows the evolution of the percentile of the quarterly price change distribution corresponding to the sample mean (the *sample mean* percentile) over the 1949-96 period. The main features to note are that:

- The series is highly volatile on a quarter-to-quarter basis. Essentially, this provides an indication of how unrepresentative the mean rate of inflation often is - at times the mean has been below all but about 15% of the (weighted) price changes in the CPI, while at other times it has been higher than over 90% of the (weighted) price changes in the regimen. The figure,

therefore, illustrates far more clearly than the coefficient of skewness the extent to which the mean is able to be pulled away from the central mass of price changes by price changes in the tails of the distribution.

- Although there is considerable quarter-to-quarter variation in the sample mean percentile, it shows no obvious cyclical or long-term trend, nor does it show clear signs of having risen during the inflationary surge from the mid-1970s to mid-1980s.

Figure 2

Sample mean percentile, subgroup level, 1949Q2-96Q4
(proportion of price changes below the sample mean)
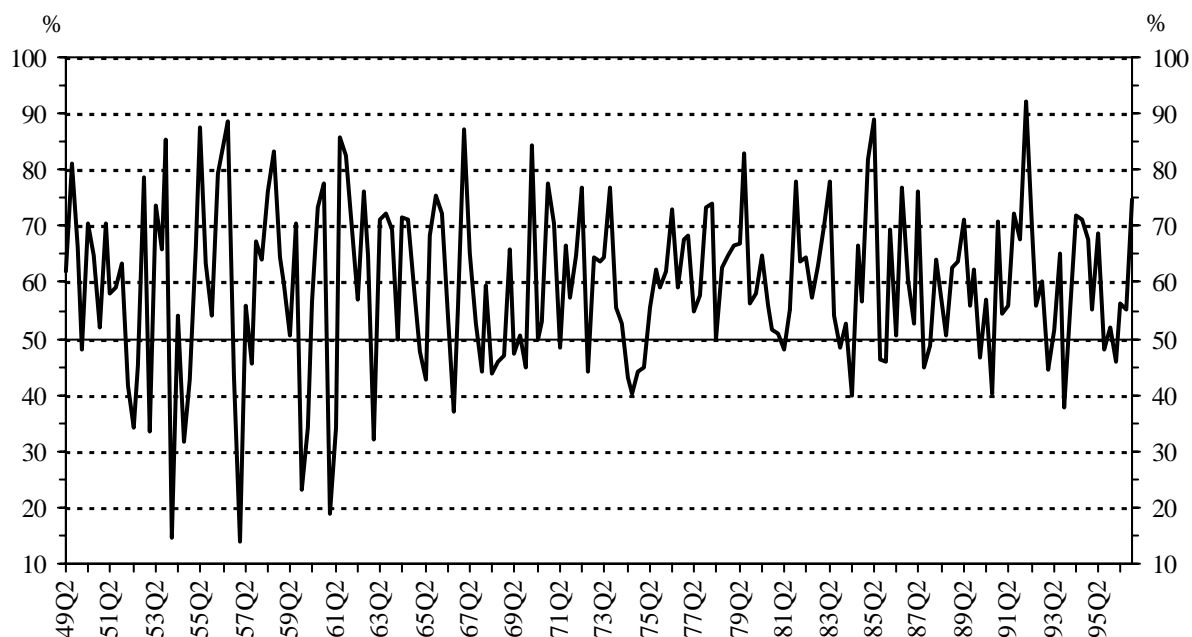
Additional evidence is provided in table 2, which shows sample correlations between the sample mean percentile, the sample mean inflation rate and the change in the sample mean, averaged over different time intervals.[12]

The table indicates a positive short-term correlation between the sample mean percentile and the mean and the change in the mean. Beyond the one year frequency, however, the correlation diminishes into insignificance, not only because the correlation coefficient falls but, also, because the standard errors of the estimates rise as the number of observations shrinks. The results are consistent with the notion that relative price disturbances may lead to temporary movements in the aggregate inflation rate *à la* Ball and Mankiw. But it is *not* consistent with the proposition that higher average or trend inflation will produce greater skewness.

Table 2

---

[12]   The change in the mean is included in these tables because it can be debated whether the inflation rate over the periods in question are I(0) or I(1).

Correlations between the sample mean percentile, the sample mean inflation rate and the change in the sample mean at the subgroup level of aggregation

Averaging period	Number of periods	Sample period	Correlation between sample mean percentile and:	
			Mean	**Change in mean**
Quarter	190	49Q3-96Q4	0.28	0.34
1 year	47	1950-96	0.18	0.33
2 years	23	1951-96	0.14	0.18
3 years	15	1952-96	0.10	0.11

(ii)   The population mean percentile

While the evidence discussed above may diminish concerns that the sample mean percentile might be substantially affected by the trend rate of inflation, or changes in the trend, it does not directly address the issue of whether the sample mean percentile is stable over time. Even if the average rate of inflation over time is not a significant determinant of the shape of the distribution of price changes, other factors (including the methodology for calculating the CPI) could be.

If the distribution of price changes in a given quarter is considered to be a particular 'draw' or sample from a characteristic 'underlying' or population distribution, then by pooling the sample distributions over many quarters, the population distribution will be approximated.

Figure 3 shows cumulative (normalised) frequency distributions pooled over 10-year sub-periods. The results are quite striking. For all of the sub-periods (consisting of 30-40 quarterly observations or 'samples'), it is apparent that the basic shape of the distribution is essentially similar - and substantially different from the Normal cumulative distribution - despite quite different average inflation rates and despite substantial changes in economic structure (including the degree of economic openness of the economy, the degree of government intervention in price setting) and the particular composition or construction of the CPI.

It may be noted, in particular, that for the cumulative distributions shown, the percentile of the distribution corresponding to the mean (i.e. zero standard deviations from the mean) lies somewhere between the 50th and 60th percentile, reflecting the chronic right-skewness of the distribution.

Table 3 seeks to pin down more precisely the percentile corresponding to the population mean (the *population mean percentile*). As discussed earlier, in the context of calculating moments of the distribution, approximations can be based on averaging of quarterly sample mean percentiles, or averaging of figures for normalised data pooled over longer periods such as calendar years or multi-year periods. Calculations based on each method are reported in table 3.

Figure 3
Cumulative frequency distribution of CPI subgroup quarterly price changes,1949-96
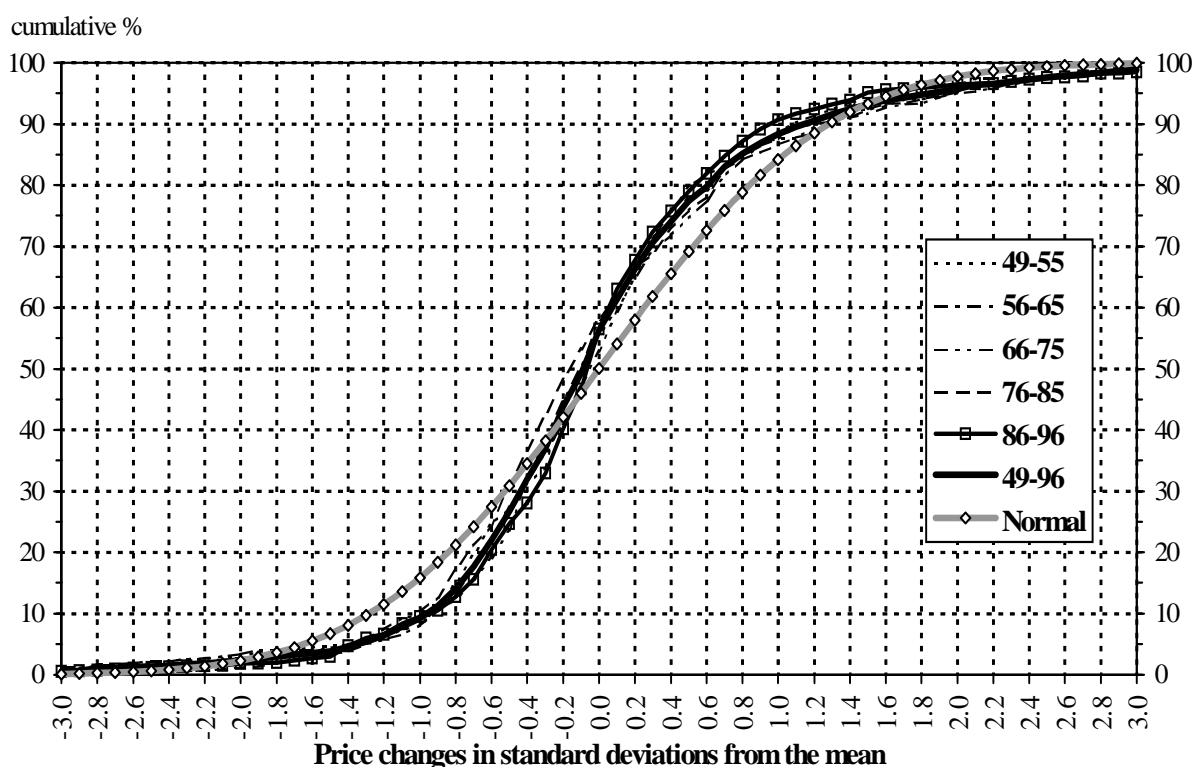(pooled, normalised price changes, in standard deviations from mean)

cumulative %

Table 3 also report median values of the sample mean percentiles, since the averages of the sample mean percentiles at the quarterly or pooled annual level are likely, in some cases, to have been significantly distorted by outliers in particular quarters.

Table 3
Estimates of the population mean percentile at the subgroup level of aggregation

Averaging period	Quarterly data		Annual pooled data		Multi-year pooled data
	Average of sample mean percentiles	Median of sample mean percentiles	Average of sample mean percentiles	Median of sample mean percentiles	Sample mean percentile
1949-55	58.1	61.8	56.7	60.9	56.1
1956-65	60.8	66.0	58.4	57.9	57.8
1966-75	57.3	55.4	53.6	54.2	52.8
1976-85	62.0	62.4	59.1	59.7	58.3
1986-96	59.8	56.7	57.4	56.8	56.4
1949-96	59.7	59.6	57.1	58.3	56.3
1975-96	60.6	59.1	58.0	58.5	57.1
1981-96	60.0	56.8	57.4	57.7	56.5

The table shows a surprising stability in the mean percentile over time. Once the standard errors (of 1-3 percentage points) are taken into account, it can be said that there has been no significant shift in the mean percentile over any sustained period the past 48 years.

Nonetheless, the results do show an average for the sample mean percentile that is fairly systematically higher for quarterly subgroup data than for the pooled annual data or the multi-year pooled data. In this writer's view, the pooled data are likely to provide a more precise

indication of the centre of the distribution of the mean percentile than is the average for the quarterly data. Whichever measure is regarded as most indicative of the population mean percentile, the evidence strongly points to a value well above the median or 50th percentile.

One way of examining the implications of choosing alternative assumptions about the population mean percentile is to calculate the implicit price levels associated with different percentiles. Using these, we can compare the rates of 'drift' or bias in the rates of change in the implicit price levels relative to the corresponding rate of change in the CPI price level, as shown in table 4.

Table 4
Rates of drift in implicit price levels associated with different percentiles of the quarterly subgroup price distribution relative to the CPI mean ex credit & GST

Price measure	1949-1996		1975-1996		1981-1996	
	Average annual % change	Average annual % drift vs. CPI	Average annual % change	Average annual % drift vs. CPI	Average annual % change	Average annual % drift vs. CPI
CPI ex credit & GST	6.67	0.0	8.57	0.0	6.30	0.0
Percentile						
56*th*	6.44	-0.21	8.42	-0.14	6.33	0.02
57*th*	6.53	-0.13	8.48	-0.09	6.48	0.07
58*th*	6.65	-0.02	8.57	-0.00	6.48	0.17

Table 4 shows that the 56th, 57th and 58th percentile price change estimators of the population mean all display very little drift relative to the sample mean throughout the 1949-96 period. Of the three, the 58th percentile measure appears to show the least bias over the entire 1949-96 period, but the most bias over the 1981-96 period. Conversely, the 56th percentile shows the most bias over the 1949-96 period, but the least over the 1981-96 period.

On balance, the 57th percentile appears to a reasonable approximation of the population mean percentile, displaying no significant drift either over the full 1949-96 period, or over the more recent period. On this basis, the 57th percentile price change will be assumed to be an essentially unbiased estimator of the mean for the remainder of this paper.

The validity of this assumption is sensitive to the stability of the degree of asymmetry in the distribution of price changes. If the degree of asymmetry were changing over time, then the percentile of the distribution corresponding to the population mean would also be time-varying. The evidence presented in tables 3 and 4 is not conclusive in this regard, and will be the subject of further research. Nonetheless, table 4 suggests that even if the population mean percentile is time-varying, assuming otherwise is not likely to result in a very substantial bias.

3.   The relative efficiency and robustness of the 57th percentile

The rules of thumb provided by Hogg (1967), and discussed earlier, together with the evidence of kurtosis of the distribution typically more than twice that of the Normal distribution, strongly points to the 57th percentile measure as being a more robust (as well as unbiased) measure of the underlying or population mean than is the sample mean.

There remains the question, however, of whether this robustness is purchased at a high cost in terms of a reduction in the efficiency of the 57th percentile relative to the sample mean. The relative efficiency of the population mean percentile can be measured by the standard error of that percentile relative to the standard error of the sample mean.

For a weighted distribution, the (adjusted) standard error of the sample mean is given by:[13]

$$\sigma_{\bar{x}(n)} = \sigma_{x(n)} \sqrt{\sum_{i=1}^{n} w_i^2}$$

where $\sigma_{\bar{x}(n)}$ is the standard error of the mean price change measured at the 'n'th level of aggregation.

$\sigma_{x(n)}$ is the standard deviation of price changes (x) at the 'n'th level of aggregation.

$w_i, i = 1...n,$ are the weights or empirical probabilities of price changes, x_i, at the 'n'th level of aggregation.

The standard error of the pth percentile (for a weighted distribution) is given by:

$$\sigma_{p(k)} = \frac{\sigma_{x(k)}}{f_k(p)} \sqrt{p(1-p) \sum_{i=1}^{k} w_i^2}$$

where p is the 'p'th percentile of the frequency distribution.

$\sigma_{p(k)}$ is the standard error of the 'p'th percentile price change at the 'k'th level of aggregation.

$\sigma_{x(k)}$ is the standard deviation of price changes, x, at the 'k'th level of aggregation.

$f_k(p)$ is the value of the density function of the 'p'th percentile, at the 'k'th level of aggregation.

The relative efficiency of the pth percentile is given by the ratio of the two standard errors:

$$\frac{\sigma_{p(k)}}{\sigma_{\bar{x}(n)}} = \frac{\sqrt{p(1-p)}}{f_k(p)} \bullet \frac{\sigma_{\bar{x}(k)}}{\sigma_{\bar{x}(n)}}$$

---

[13]   The formula for the standard error for weighted distributions is derived in Roger (1997a).

In the regular case, where the standard errors of the sample mean and the *p*th percentile are measured at the same level of aggregation of prices, the relative efficiency of the *p*th percentile will depend only on the value of the term $f(p)/\times[p(1-p)]$.

For a Normal distribution, $f(p) = 0.398$ and $\times[p(1-p)] = 0.5$, so that the relative efficiency of the median is about 0.8. In other words, if the population distribution is Normal, the sample mean will be approximately 1.25 times as efficient as the sample median

The relative efficiency of the *p*th percentile, measured at one level of aggregation, can also be compared with the standard error of the sample mean measured at a different level of aggregation. For the purposes of this analysis, the standard error of the CPI sample mean is based on the calculation at the most disaggregated level available: the item level. The standard error of the 57th percentile, however, is calculated at the subgroup level of aggregation.

The relative efficiency of the 57th percentile measure is estimated in two ways in this paper:

- The first method involves calculation of the standard errors of the sample mean and sample 57th percentile, on a quarter-to-quarter basis. The standard errors and the ratio of the two are then averaged across quarters. Because item level CPI data are not available prior to 1981, this method is applied only to the 1981-96 sample period. Median values of the standard errors and their ratios are also reported because the distributions of these statistics over the sample period display high kurtosis and right skewness. Consequently, the median values are probably more indicative than the mean values.

- The second method involves basing the calculations on multi-year pooled data. The pooled data should more accurately approximate the population distribution, providing a firmer basis for calculation of the true standard errors of the sampling errors for the mean and 57th percentile measures. Because the data is normalised prior to pooling, standard errors for the mean are normalised to unity. It is nonetheless possible to calculated the standard errors for the 57th percentile relative to the mean, and these are also reported in table 5.

Table 5
Relative efficiency of the 57th percentile as an estimator of the population mean

Sample period	Standard error of estimate (in %)		Relative efficiency of 57th percentile measure
	CPI sample mean	**Sample 57th percentile**	
Based on quarterly data			
1981-96 mean	0.41	0.36	1.97
1981-96 median	0.38	0.28	1.38
Based on pooled (normalised) annual data			
1949-96	1.0	0.82	1.22
1975-96	1.0	0.76	1.31
1981-96	1.0	0.70	1.43

Table 5 indicates that the distribution of quarterly price changes in New Zealand is sufficiently kurtotic that, *even on average*, the 57th percentile measures of price change is substantially more efficient as an estimator of the underlying or population mean rate of inflation than is the sample mean. In other words, achieving robustness in the estimate of core inflation is accompanied by greater efficiency rather than at the expense of efficiency.

VI. Tests of the measures

The 57th percentile measure, like any other similar order statistic down-weights outlier price changes relative to the mean. The issue addressed in this section is whether the price changes being down-weighted do generally represent supply shocks. If so, then the 57th percentile can be reasonably interpreted as a measure of 'core' inflation. For the purposes of this section, the 57th percentile measure of inflation will be described as the 'core' inflation measure.

Four sets of tests are involved. The first examines the degree of serial correlation in the differential between the mean and core inflation measures. This differential provides a measure of the impact of relative price shocks on the mean inflation rate. The second test examines the statistical independence or causality between the core measure and the relative price shocks filtered out by the measure. The third test examines whether, by excluding information on relative price shocks, the core measure discards useful information about the future movement of the CPI or, alternatively, discards 'noise'. Finally, the fourth test examines the question of whether the relative price shocks, based on the core measure, can be thought of as the kinds of 'supply' shocks normally associated with shifts in the short-run Phillips curve.[14]

1. Serial correlation of relative price shocks

The differential between the inflation rate as measured by the CPI (ex credit & GST) and the core measure can be viewed as providing an estimate of the impact of relative price shocks to inflation. Relative price shifts, either temporary or permanent, may stem from a variety of sources. These include (i) classic 'supply' disturbances (such as international commodity price shocks, or shifts in relative prices as a result of government policy), (ii) shifts in consumer preferences, (iii) seasonal shifts related to weather or regular re-pricing schedules for government and other producers, or to semi-annual or annual price sampling for some commodities by Statistics New Zealand (iv) pure error in sampling or processing of data.

For random shocks, serial correlation would not be expected, unless such shocks were typically spread over a number of periods. By contrast, with seasonal shocks, positive fourth-order (and, possibly, second-order for goods priced semi-annually) serial correlation might be expected.

Of particular concern is the possibility of positive first-order correlation, which might occur if the measure of core inflation was treating (i.e., down-weighting) some part of changes in generalised inflation as relative price disturbances.

---

[14] I would like to thank Catherine Connolly for performing a battery of causality tests, only a few of which are reported here. I would also like to thank Weshah Razzak for estimation of the various Phillips curves.

In fact, no significant first-, second- or fourth-order serial correlation in the relative price shocks was found at the 95% confidence interval over the 1949-96, 1975-96 or 1981-96 periods.[15] This suggests that past values of the relative price shocks provide no significant useful information about future shocks, and that ongoing inflation is not being mistaken as relative price movement.

2. Causality between relative price shocks and core inflation

The next line of enquiry involves testing the independence of the relative price shocks from changes in the core rate of inflation. Of particular concern are two plausible reasons for correlation between the relative price shocks and core inflation. The first is simply that, even if relative price shocks are driven by supply disturbances, they may feed into generalised inflation either through formal indexation of a variety of prices and wage or transfer payments to the aggregate CPI inclusive of the shocks, or through backward-looking inflation expectations formation. If these effects are important, causality would tend to run from relative price shocks to core inflation, but not *vice versa*.

Alternatively, if different prices display different degrees of inertia, a change in the generalised rate of inflation could be manifested in some prices moving much more rapidly towards the new equilibrium inflation rate than others. As a consequence, a change in the generalised inflation could generate skewness in the distribution that might be filtered out as a relative price disturbance. In this case, the true causality would run from generalised inflation to relative price shocks, though in a statistical sense one could conceivably find the causality to run the other way.

To test for such interdependence, two-way Granger causality tests were conducted between the relative price shocks and the core inflation rate, using 4 lags on the dependent and independent variable.[16] A potential complication with this procedure is that while the shocks are stationary, essentially by construction, the core inflation rate may not be. Unit root testing, however, suggests that a unit root in the quarterly inflation rate can be rejected at the 10% significance level and, on this basis, Granger causality tests were performed using the quarterly 57th percentile price changes.

The results of these Granger causality tests are not completely unequivocal. For the subgroup level data over the full 1950Q2-96Q4 period, no significant causality is found running from core inflation to the shocks or *vice versa*. Over shorter sample periods, however, causality cannot always be rejected.

Nonetheless, it is probably safe to say that while relative price shocks may, at times, spill over into core inflation, there is no evidence to support the view that ongoing inflation is mistakenly left out of the core measure.

3. Causality between CPI and core inflation

---

[15] The closest indication was for second-order serial correlation at the subgroup level, with a coefficient of -0.14 with a p-value of 0.058, for the 1949Q4-96Q4 period.

[16] See Engle and Granger (1987).

The third line of testing also examines whether the relative price shocks stripped out by the core measure eliminates useful information about future inflation. This time the test is whether the core measure is Granger-caused by the CPI ex credit services and GST, or whether it Granger-causes the CPI ex credit services and GST (hereafter just the 'CPI' for convenience).

If the relative price shocks purged from the core measure provide additional information on future inflation, the CPI should tend to Granger-cause the core measure, but not the reverse. Alternatively, if the relative price shocks are essentially noise, then the core measure should tend to Granger-cause the CPI, but not the reverse.

Testing in this case is complicated somewhat by weaker evidence of stationarity in the first difference of the CPI, and by evidence that the two inflation series are cointegrated. The procedure used was, first, to estimate the cointegration vector between the CPI and the core measure and test the residuals for stationarity (following Engle and Granger, 1987). The second step includes the lagged residuals in the Granger causality tests between the CPI and the core measure. Six lagged terms of the dependent and independent variable were included in the test, together with a single lagged value of the cointegration residual.

Once again, the results of the Granger causality tests cannot be described as unequivocal. Over the full 1950Q4-96Q4 period, significant causality is found to run only one way: from the core measure to the CPI. However, over the 1950Q4-73Q2 and 1973Q2-96Q4 sub-periods, no significant causality is found in either direction, while over the 1985Q1-96Q4 sub-period, causality is found to run in both directions.

The general thrust of the results suggests that, although the relative price shocks may sometimes be informative, more often than not, the core measure provides a clearer signal about future inflation than the CPI inclusive of the shocks.

4.   Core inflation and the Phillips curve

The final set of tests involves using the relative price shocks in a short-run Phillips curve specification for inflation. In principle, the identification and estimation of a short-run Phillips curve should be improved if shifts in the curve can be distinguished from movements along it. If the estimated relative price shocks arise primarily from supply-side developments, then including these shift terms in the Phillips curve specification should improve the estimation. Alternatively, if the relative price shocks are in fact picking up relative price changes associated with more generalised changes in inflation, including them in the specification should not improve the estimation.

Two non-linear expectations-augmented short-run Phillips curve specifications were estimated, with and without relative price shock terms:

(1) $(\Pi)^{\gamma}_t = a + b\,(\Pi^e)^{\gamma}_{t-1} + c_1\,(\Delta CU)_{t-1} + c_2\,(\Delta CU)_{t-2} + c_3\,(\Delta CU)_{t-3} + c_3\,(\Delta CU)_{t-3} + c_4\,(\Delta CU)_{t-4} + e_t$

(2) $(\Pi)^{\gamma}_t = a + b\,(\Pi^e)^{\gamma}_{t-1} + c_1\,(\Delta CU)_{t-1} + c_2\,(\Delta CU)_{t-2} + c_3\,(\Delta CU)_{t-3} + c_3\,(\Delta CU)_{t-3} + c_4\,(\Delta CU)_{t-4}$

$$+\, d\,(\Pi - \Pi_{57})_t + e_t$$

where:

Π_t is the quarterly rate of change of the CPI (ex credit services and GST) in period t.

Π^e_{t-1} is the National Bank survey measure of expected inflation.

(ΔCU) is the change in the rate of capacity utilisation, a measure of the change in excess demand pressures on inflation.[17]

$(\Pi - \Pi_{57})_t$ is the relative price shock term, measured by the differential between the CPI (ex credit services and GST) and the 57th percentile price change.

The equations were estimated over the period from 1983Q1-96Q3, due to the lack of inflation expectations data prior to 1983. The results are reported in table 6. Equation (1) and, to an even greater extent, equation (2) display noticeable positive serial correlation in the residuals, as shown by the autocorrelation coefficient ρ . The coefficients and diagnostic statistics for the two equations are reported for estimation after correcting for serial correlation.

Table 6
Relative price shocks and non-linear expectations-augmented Phillips curve estimates.
Sample period: 1983Q1-1996Q3

	Equation (1): no relative price shocks			Equation (2): with relative price shocks		
	Estimate	Std. error	t-stat	Estimate	Std. error	t-stat
a	-0.35	0.27	-1.32	-0.26	0.21	-1.24
b	0. 93	0.09	10.00*	0.92	0.10	9.79*
c_1	0.022	0.009	2.46**	0.014	0.005	3.14*
c_2	0.014	0.009	1.64***	0.009	0.004	2.18**
c_3	0.018	0.009	2.00**	0.011	0.004	2.60*
c_4	0.023	0.009	2.64*	0.015	0.004	3.59*
d	..	..	..	13.99	3.11	4.50*
γ	0.20			0.33		
ρ	0.21	0.13	1.56	0.37	0.13	2.90
LLF	221.3			228.8		
DW	1.96			2.03		
$Rbar^2$	0.73			0.80		
96Q4 forecast error	15.7%			7.5%		

Note: "*" indicates significance at the 1% level, "**" indicates significance at the 5% level, "***" indicates significance at the 10% level.

Table 6 shows that adding the relative price shock term improves not only the overall significance and fit of the equation, but also the significance of the coefficients on individual right-hand side variables, and out-of-sample forecast performance. These results are consistent with the view that the measure of relative price shocks derived in this paper are appropriately interpreted as primarily reflecting supply-side events.[18]

---

[17]   Equations (1) and (2) were also estimated using two other measures of excess demand: the difference between actual GDP and the level of 'potential' GDP as measured by (i) the common HP1600 two-sided smoother, and (ii) an HP1600 one-sided filter measure of 'potential' GDP. The results were essentially similar to those reported above, and so were not included.

[18]   These results corroborate those of Ball and Mankiw (1992) and Bryan and Cecchetti (1993), who also find that generalized inflation is better represented by median-based or trimmed-mean measures of price changes than by the mean.

VII.   Concluding comments

The analysis in this paper builds on previous research by Bryan and Cecchetti (1993) and Roger (1995) investigating a stochastic approach to the measurement of core inflation. The approach has significant potential advantages over other current standard approaches in terms of simplicity, statistical robustness, and verifiability.

The high kurtosis of the distribution of price changes in New Zealand since 1949 suggests that a median-based measure of price change is likely to be more efficient and more robust than alternative estimators of core inflation.

In addition, it is found that the distribution of price changes is chronically right-skewed. As a result, the median (50th percentile) price change typically understates the trend inflation rate. A slightly higher percentile price change eliminates this bias while retaining the attractive properties of the median.

The purely statistical basis of the stochastic approach is both a strength and a potential weakness. The strength of the approach is two-fold: it provides a sounder basis for a measure of core inflation both in terms of robustness and theoretical justification than current standard approaches; it also provides a measure of core inflation that is readily computed and verified by any independent observer with access to the same data. Both aspects are bound to be of importance for a central bank whose performance is assessed on the basis of the evolution of a measure of core inflation.

A potential weakness of the approach is that no explicit distinction is made between relative price movements arising from supply shocks as opposed to demand shocks. The evidence presented in this paper, however, suggest that most extreme movements in relative prices do in fact represent supply shocks.

A second weakness of the approach is that the degree of asymmetry in the distribution of price changes may change through time, in which case an estimator assuming an unchanging distribution will be biased. The evidence presented in this paper suggests that there has been surprisingly little change in the asymmetry of the distribution of price changes over the past 50 years. This finding may, however, be sensitive to the level of aggregation of prices. The degree of aggregation of prices examined in this paper might obscure significant changes in the shape of the distribution at a less aggregated level. Research on this question is in progress and will be published in due course.

References

Balke, N. and M. Wynne (1996), "Supply shocks and the distribution of price changes", *Federal Reserve Bank of Dallas Economic Review*, 1st Quarter 1996, pp. 10-18.

Ball, L., and N.G. Mankiw (1992), "Relative-price changes as aggregate supply shocks", *National Bureau of Economic Research Working Paper No. 4168*, Sept.

Ball, L., and N.G. Mankiw (1994), "Asymmetric price adjustment and economic fluctuations", *The Economic Journal*, Vol. 104, pp. 247-261.

Bryan, M., and S. Cecchetti (1993), "Measuring core inflation", *National Bureau of Economic Research Working Paper No. 4303*, March.

Bryan, M. and S. Cecchetti (1996), "Inflation and the distribution of price changes", *National Bureau of Economic Research Working Paper No. 5793*, Oct.

Cecchetti, S. (1996), "Measuring short-term inflation for central bankers", *mimeo*, Department of Economics, The Ohio State University, September.

David, H. (1981), *Order statistics*, 2nd ed., Wiley, New York.

Diewert, W. E. (1995a), "Axiomatic and economic approaches to elementary price indexes", *University of British Columbia, Department of Economics Discussion Paper* 95-01.

Diewert, W. E. (1995b), "On the stochastic approach to index numbers", *University of British Columbia, Department of Economics Discussion Paper* 95-31.

Engle, R. and C. Granger (1987), "Co-integration and error correction: representation, estimation, and testing", *Econometrica*, Vol. 55(2).

Fisher, I. (1922) *The purchasing power of money*, reprinted (1971) by A.M. Kelly Publishers, New York.

Hogg, R. (1967), "Some observations on robust estimation", *Journal of the American Statistical Association*, Vol. 62 (Dec.), pp. 1179-86.

Huber, P. (1981), *Robust statistics*, John Wiley & sons, New York.

Judge, G., R.C. Hill, W. Griffiths, H. Luetkepol and T.-C. Lee (1988), *Introduction to the theory and practice of Econometrics*, 2nd ed., John Wiley & sons, New York.

Kendall, M. and A. Stuart (1969), *The advanced theory of statistics*, Vol. 1 (distribution theory), 3rd ed., Charles Griffin & Co., London.

Koenker, R. and G. Bassett Jr. (1978), "Regression quantiles", *Econometrica*, Vol. 46(1), pp. 33-50.

McCallum, B. (1995), "New Zealand's monetary policy arrangements: some critical issues", *Reserve Bank of New Zealand Discussion Paper G95/4*.

Roger, S. (1995), "Measures of underlying inflation in New Zealand, 1981-95", *Reserve Bank of New Zealand Discussion Paper G95/5*, Sept.

Roger, S. (1996), "The 1986-87 and 1989-90 GST shocks: revised estimates of the impact on the CPI", *Research note 96/1*, Reserve Bank of New Zealand.

Roger, S. (1997a), "The distribution of consumer price changes in New Zealand, 1949-96", *Reserve Bank of New Zealand Discussion Paper*, forthcoming.

Spiegel, M. (1995), "Rules vs. discretion in New Zealand monetary policy", *Federal Reserve Bank of San Francisco Weekly Letter No.95-09*, March.

Theil, H. (1967), *Economics and information theory*, North-Holland, Amsterdam.