NBER WORKING PAPER SERIES

MEASURING SHORT-RUN INFLATION FOR CENTRAL BANKERS

Stephen G. Cecchetti

Working Paper 5786

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 1996

This paper was prepared for the Federal Reserve Bank of St. Louis's Economic Policy Conference, October 17-18, 1996. I am grateful to the Department of Economics, University of Melbourne for their generous hospitality, to Michael Bryan, Margaret Mary McConnell and seminar participants at the Reserve Bank of Australia for their comments, and to the National Science Foundation and the Federal Reserve Bank of Cleveland for financial and research support. This paper is part of NBER's research program in Monetary Economics. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

NBER Working Paper 5786
October 1996

MEASURING SHORT-RUN INFLATION FOR CENTRAL BANKERS

ABSTRACT

As central bankers intensify their focus on inflation as the primary goal of monetary policy, it becomes increasingly important to have accurate and reliable measures of changes in the aggregate price level. Measuring inflation is surprisingly difficult, involving two types of problems. Commonly used indices, such as the Consumer Price Index (CPI), contain both transitory noise and bias. Noise causes short-run changes in measured inflation to inaccurately reflect movements in long-run trends, while bias leads the long-run average change in the CPI to be too high.

In this paper I propose methods of reducing both the noise and the bias in the CPI. Noise reduction is achieved by average monthly inflation in measures called 'trimmed means' over longer horizons. Trimmed means are statistics similar to the median that are calculated by ignoring the CPI components with extreme high and low changes each month, and averaging the rest. I find that using three month averages halves the noise, while removing the highest and lowest ten percent of the cross-sectional distribution of inflation reduces the monthly variation in inflation by one-fifth.

Stephen G. Cecchetti
Department of Economics
Ohio State University
1945 N High Street
Columbus, OH 43210-1172
and NBER
cecchetti.1@osu.edu

1 Introduction

In recent years, a large number of central banks have moved toward explicit inflation targeting. Haldane (1995) provides a list that includes Australia, Canada, Finland, Israel, New Zealand, Spain, Sweden and the United Kingdom. It seems likely that the United States will soon join this group. As the focus of monetary policy shifts, it has become increasingly important to have accurate, reliable measures of inflation. The purpose of this paper is to examine the difficulties surrounding the measurement of the changes in the aggregate price level.

Measuring inflation is a surprisingly difficult task. While it is conceptually easy to survey the prices of individual commodities at any given time, using these to produce a measure appropriate for monetary policy is far from straightforward. Gauging movements in aggregate prices is neither theoretically nor practically easy.

Broadly speaking, there are two problems associated with measuring inflation. The first concerns transitory phenomena, or *noise*, that should not affect policy makers' actions. Sources of such noise include changing seasonal patterns, broad-based resource shocks, exchange rate changes, changes in indirect taxes, and asynchronous price adjustment. Knowledge of the extent to which noise is present in measured aggregate price indices is important for two reasons. First, for those central bankers who target inflation or prices alone, the width of a credible target band depends on how the noise in the targeted price index. Second, so long as inflation has some weight in a policy maker's objective function, it is important to know how to interpret monthly movements in aggregate prices. This argues for the development of indices aimed at minimizing this problem.

The second potentially severe difficulty associated with measurement involves biases that are a consequence of weighting schemes sampling techniques, and quality adjustments employed in the calculation of price indices. Shapiro and Wilcox (1996) provide a catalog of the long run biases in the Consumer Price Index (CPI), together with estimates of their size. These bias estimates can be divided into two broad categories: those related to the way in which individual prices are weighted together to form an aggregate index (weighting bias), an example is substitution bias; and those that result from actual errors in measuring the individual prices themselves (measurement bias), such as quality or new goods bias.

While noise is by definition temporary, bias is not. The importance of bias for monetary policy is two-fold. First, if the central bank is setting a numerical target for a particular inflation index, the extent of the bias will dictate whether, for example, the objective of price stability implies zero measured inflation. But in addition to the this longer-run issue, it seems likely that the biases themselves will be time-varying. Again, this suggests that it is difficult to interpret movements in measured price indices.

There are many reasons bias might vary over time. For example, if the rate technological progress varies, then bias resulting from quality adjustment may not be time-invariant.[1] Substitution bias may also be time-varying In the CPI, this bias is caused by the fixed expenditure-based weighting scheme used in its construction. The reason for this is that as relative prices change, expenditures on relatively more costly goods fall. But since the CPI has weights that are invariant to relative price changes, it will show an increase in the aggregate price level even when one has not occurred.[2] Clearly, though, the extent of this bias will depend on the degree to which such substitution is occurs. If, during a particular period, there was very little variation in price changes across commodities, then one would expect the substitution bias to be small. By contrast, during times in which there is a large spread in cross-sectional inflation, substitution bias might be large.

Numerous studies address the issue of bias in price measurement by performing very careful calculations based on highly disaggregated information on prices, quantity, product quality and the like. Lebow, Roberts and Stockton (1992), Wynne and Sigalla (1993) and Shapiro and Wilcox (1996) all provide a catalog of estimates of the biases in the CPI. But all of these estimates come from studies of product specific microeconomic data, and so lack the generality necessary to help gauge the overall

[1] Gordon (1992) measures quality bias in consumer durable goods, and finds large differences over time.

[2] Bryan and Cecchetti (1993) discuss and estimate time variation in this and all other weighting biases.

bias in the aggregate index.

There have been a number of suggested solutions to price-index measurement problems. Solutions to the problem of high-frequency noise in price data include calculation of low-frequency trends over which this noise is reduced. But from a policymaker's perspective, this greatly reduces the timeliness, and therefore the relevance, of the incoming data. Another common technique for measuring the underlying or *core* component of inflation excludes certain prices from the computation of the index on the assumption that these are the ones with high-variance noise components. This is the "ex. food and energy" strategy, where the existing index is reweighted by placing zero weights on some components, and rescaling the remaining weights.

Finally, Bryan and Cecchetti (1994) suggest removing these transitory elements from the aggregate index by calculating the weighted-median CPI. The median addresses the difficulty caused by the fact that a large and sudden increase in the price of one good may not be matched immediately by an equivalent decrease in the price of some other good. Instead, the offsetting adjustment will take time. Such a price shock will cause standard measures of inflation, based on the mean of inflation in the prices of individual goods, to move up following the initial shock, and down following the compensating adjustment. These temporary movements will not be present in the median, because it eliminates the undesirable effects of temporarily high or low prices in specific sectors.[3]

Bryan and Cecchetti (1993) is an attempt at a general treatment of the weighting bias problem. In that paper, weighting bias is treated as a statistical problem that can be overcome, at least in part, using a technique developed by Stock and Watson (1991) in their construction of a coincident index of real activity. The result is a 'Dynamic Factor Index' (DFI) in which a measure of the aggregate price level is constructed by weighting (in a time-varying way) commodities based on the strength of a common inflation signal.

The purpose of this paper is to examine the severity of the noise and wieghting

[3]The intuition behind the use of limited-influence estimators, such as the median, is based on Ball and Mankiw's (1995) model of costly price adjustment.

bias problems, and propose some partial solutions. I begin in Section 2 with a simple framework that helps to clarify the issues associated with noise and weighting bias, followed by a presentation of the Dynamic Factor Model. In Section 3, I present results that examine the extent of noise and ways in which it might be reduced. Two general methods of noise reduction are introduced: the computation of limited influence estimators, such as the median; and averaging over three, six and twelve month horizons. A simple series of tests suggests that the most accurate estimator is a 10% Trimmed Mean averaged over three to six months. I examine the problem of weighting bias in Section 4, where I present results suggesting the extent to which the DFI is a reduced-bias estimator of aggregate inflation. Section 5 concludes.

2 General Considerations

It is useful to begin with a general discussion of noise and bias. I begin in Section 2.1 with a simple accounting framework that helps illuminate the sources and consequences of noise and bias. The result here is that the simplest way to reduce noise is to lengthen the observation interval over which inflation is measured.

Building on the general discussion, Section 2.2 proceeds with the description of the Dynamic Factor Index. Here I show why the DFI is likely to be an estimator of common trend in prices that eliminates one source of the bias in the inflation statistics.

2.1 A Framework for Analysis

During any given time period, inflation in a particular product's price can be decomposed into two elements, one common to all price changes and the other idiosyncratic. Defining $\dot{p}_{it}$ as the change in the log of an individual price i at time t, $\dot{\mathcal{P}}_t$ as the common factor and $\dot{x}_{it}$ as relative price inflation, this can be written as

$$\dot{p}_{it} = \dot{\mathcal{P}}_t + \dot{x}_{it} \ . \tag{1}$$

$\dot{\mathcal{P}}_t$ is the quantity of interest. This is the common trend in prices, and is therefore the proper analog to changes in the aggregate price level of macroeconomic theory.

An aggregate price index can be constructed from a set of these product prices using a set of weights. Defining the weights as w_{it}, for example then,

$$\pi_t \equiv \sum_i w_{it}\dot{p}_{it} \ . \tag{2}$$

It is useful to normalize the weights to sum to one,

$$\sum_i w_{it} = 1 \quad \forall t. \tag{3}$$

Combining (1) and (2), and using (3), yields

$$\pi_t = \dot{\mathcal{P}}_t + \sum_i w_{it}\dot{x}_{it} \ . \tag{4}$$

It is clear from this formulation that the deviation of measured price indices from the common inflation is given by the second term on the right hand side of (4). The problem with standard indices is that this term is nonzero — both at any given time, and in expectation.

As the introduction suggestions, the difference between π_t and $\dot{\mathcal{P}}_t$ is the sum of two parts: noise plus bias. Writing these as b_t and n_t respectively, yields

$$\pi_t - \dot{\mathcal{P}}_t = \sum_i w_{it}\dot{x}_{it} = n_t + b_t \tag{5}$$

I will now discuss each of these in turn.

The noise, n_t, represents both transitory and permanent shocks to the price level. It is stationary and mean zero, with the property that

$$\lim_{k\to\infty} \frac{1}{k} \sum_{j=1}^{k} n_{t+i} = 0 \ . \tag{6}$$

This has the important implication that lengthening the observation interval eliminates the transitory noise. In other words, even if data is available at monthly frequency, it may be wise to look at changes over three, six or even twelve months.

The bias might be thought of as having a mean and a transitory component. That is to say

$$b_t = \mu_b + \omega_t \ , \tag{7}$$

where ω_t is mean zero, but may be serially correlated. Theory predicts that ω_t might have substantial unconditional variance. In thinking about bias, it is clear that high relative price dispersion might lead to higher levels of commodity substitution bias. Furthermore, periods of high aggregate growth might be times when a relatively large number of new products are introduced, and so quality and new goods biases might be large.

All of this leads one to consider measures of individual price change such as

$$\dot{p}_{it}^k = \frac{1}{k} ln \left(\frac{p_{it+k}}{p_{it}} \right) \tag{8}$$

which imply a measure of aggregate price inflation like

$$\pi_t^k = \dot{\mathcal{P}}_t^k + \mu_b + \frac{1}{k} \sum_{j=1}^{k} (\omega_{t+j} + n_{t+j}) \ . \tag{9}$$

The immediate implication is that lengthening the observation interval over which inflation is measured will gradually eliminate the noise.

But measuring inflation as a moving average over many months or even years, clearly reduces its usefulness to policy makers. As an alternative, core inflation measures, such as the CPI excluding food and energy and the weighted median, are designed to minimize the observed noise without sacrificing the high frequency of measurement. As Bryan and I discuss in Bryan and Cecchetti (1994), the weighted median reduces noise in a number of ways. First, it downweights the importance of sector-specific shocks that are likely to may only eventually average to zero across all prices. And second, it reduces the impact of errors in price-setting or measurement when any of these is far from the central tendency of the cross-sectional distribution of price changes.

The weighted median is only one in a class of limited-influence estimators, called *trimmed means*, that have the potential to reduce the noise in price statistics. To construct a trimmed mean, one simply takes the $\dot{p}_{it}$'s, together with their weights, and orders them from largest to smallest. The weighted median is the $\dot{p}_{it}$ such that one-half the weight is above and below it in this ordering. A more general alternative is to specify a percentage of the weight to remove, and then average the remaining mass of the distribution. So, for example, I will define the 25% Trimmed Mean as the measure constructed by removing 25% of the upper and lower tail of the cross-sectional distribution of the $\dot{p}_{it}$'s, and then averaging the remaining 50%.

2.2 The Dynamic Factor Index

In this section I describe the computation of an estimate of consumer price inflation that eliminates one source of bias. It is important to distinguish bias in a price statistic as a measure of inflation from bias in the CPI as a utility-based measure of welfare. Bias in the CPI as a measure of inflation is simply the deviation of measured π_t from $\dot{\mathcal{P}}_t$, whereas bias in the CPI as a measure of the cost-of-living is the deviation of the CPI from a constant utility price index. As in my early work with Michael Bryan, the objective here is to compute a reduced-bias estimate of inflation from consumer price data.[4]

Bryan and I propose an alternative to the expenditure-weighting schemes that is based on the strength of the inflation signal $\dot{\mathcal{P}}_t$, relative to the noise, $\dot{x}_{it}$, in each observed time series, $\dot{p}_{it}$. To do this, we assume the following model:

$$\dot{p}_t = \dot{\mathcal{P}}_t + \dot{x}_{it} \tag{10}$$

$$\psi(L)\dot{\mathcal{P}}_t = \delta + \xi_t \tag{11}$$

$$\theta(L)\dot{x}_t = \eta_t \tag{12}$$

[4]See Bryan and Cecchetti (1993).

where $\dot{p}_t$ and $\dot{x}_t$ are vectors; ψ and θ are a vector and matrix, respectively, of lag polynomials with stationary roots; δ is a scalar constant; and ξ_t and η_t are a scalar and a vector i.i.d. random process, respectively. The common element, $\dot{\mathcal{P}}_t$, is identified by assuming it to be uncorrelated with relative price disturbances at all leads and lags. Logically, if one of the $\dot{x}_{it}$'s were correlated with $\dot{\mathcal{P}}_t$, it would not be idiosyncratic.

Bryan and I estimate the model using a Kalman filtering algorithm, assuming all of the lag polynomials are AR(2). The result is an estimate of $\dot{\mathcal{P}}_t$ the Dynamic Factor Index (DFI),

$$\hat{\mathcal{P}}_t = \sum_j \hat{w}_j(L)\dot{p}_{jt} \ . \tag{13}$$

That is to say, the DFI is an estimate of the common trend in individual inflation series such that

$$E\left(\sum_j \hat{w}_j(L)\dot{x}_{jt}\right) = E(b_t + n_t) = 0 \ . \tag{14}$$

$\hat{P}_t$ is free of weighting bias.

Use of the Dynamic Factor Index has one clear advantage over other methods of bias reduction. Since its estimation is based on maximum likelihood methods, it allows for the computation of well-defined standard errors. This is something others often do informally.

But the DFI also has limitations. First, it eliminates only one source of bias — that associated with correlations between relative price changes and the weights. To the extent that bias arises from the mismeasurement of individual component price indices — e.g., difficulties in measuring quality changes — the Dynamic Factor Index will be biased as well.[5] Secondly, the DFI still contains transitory noise. While $E(\hat{\mathcal{P}}_t) = \dot{\mathcal{P}}_t$, there will be deviations of the estimates from the true value. But even more importantly, Finally, since the Dynamic Factor Index is constructed using an econometric procedure, the entire history can change as new data is added and the parameters of the model are reestimated.

[5]This problem is discussed in Bryan and Cecchetti (1993). There we consider the consequences of dropping the presumably less reliable service price measures and recalculating the index. This does have some impact on the results.

3 Noise

In order to assess the extent of transitory noise in inflation measures, as well as to evaluate proposed solutions, it is important to specify what we would ideally like to measure. My sense is that the information that is crucial for monetary policymakers' decisions is timely (high-frequency) estimates of movements in the long-term trend in inflation. As an approximation, I will define this trend to be the thirty-six-month centered moving average of actual inflation. The choice of a three year window is not crucial, it is simply meant as an illustration.

This section is devoted to measuring the ability of different procedures to reduce noise. The evaluation process will proceed in a series of steps. First, I will examine the *efficiency* of various estimators of inflation. This is a standard statistical criterion related to the small-sample variance of an estimator. To understand it, take an example of a set draws from mean zero distribution, such as the standard-normal distribution. In order to compare the sample mean and the sample median as estimators of the population mean of this distribution, it is natural to calculate the variance of each candidate estimator of the mean based on the number of draws. The estimator with the lowest variance is then the most efficient.

The second standard for comparison is to look at the distribution of deviations of each estimator from the thirty-six-month centered moving average of inflation. The purpose of this is to determine how large a move in inflation measured by a candidate estimator is required before one could confidently infer that the trend has moved.

Finally, I will examine seasonal fluctuations. How much of the fluctuation in inflation is due to seasonality that can be easily removed? Do any of the candidate measures have less seasonality that the others? How quickly does the seasonality disappear as the observation interval is increased?

3.1 Statistical Efficiency

All of the candidate estimators for inflation that I study are based on alternative ways of combining the various component price series. For example, the CPI itself is simply the average of the components computed using the expenditure weights. The 'CPI excluding food and energy' and limited-influence estimators, such as the weighted-median CPI, each combine the component series in a slightly different way. But regardless of the particular method used, the aggregate inflation measure is always based on the disaggregated component price series.

To understand the issue of efficiency, compare the sample mean and the sample median. Both of these attempt to measure the population mean of a distribution. But for any given sample, one would not expect that the two to yield the same result. What determines which measure is preferred? One possible answer is to choose the estimator with the lowest small-sample variance — i.e. the most efficient. As an example, compare the small-sample variance of the sample mean and the sample median, obtained from a sample of 15 $N(0,1)$ draws. The sample mean has variance equal to $(1/15) = 0.067$ while the sample median has sampling variance equal to 0.103.

While the sample mean is the more efficient than the sample median when the data are drawn from a normal distribution, this will not be true in general. In fact, it is possible to show that the sample median is more efficient than the sample mean when the data are drawn from leptokurtic distributions — that is, distributions with fat tails relative to the normal.[6] The reason for this is that with a fat-tailed distribution, it is more likely that one will obtain a draw of an observation in one tail of the distribution that is not balanced by an equally extreme observation in the opposite tail. That is, the sample has a higher probability of being skewed even though the underlying distribution is symmetrical. This increased chance of drawing a skewed sample causes the distribution of the sample mean itself to spread out. With high kurtosis, the sample mean is a higher variance estimator of the population mean than is the sample median.

A similar result holds for the entire class of trimmed means. As the kurtosis of the distribution increases, it is efficient to trim more and more of the sample. For

[6] Roger (1996) notes that for highly leptokurtic distributions, the mean is not the most efficient estimator of the central tendency of the distribution.

example, in a simple experiment with data drawn from a mixture of a normal and a uniform distribution, an increase in the kurtosis of the mixed distribution from 4 to 5 causes the variance of the sample 10% Trimmed Mean to fall below that of the sample mean. As the kurtosis rises further, it is optimal (in this sense) to trim more and more of the sample.[7]

As noted in Bryan and Cecchetti (1996), price-change distributions are highly leptokurtic. In fact, the average sample kurtosis of monthly price changes across the 36 major components of the CPI, over the 1967 to 1996 sample, exceeds nine.[8] This suggests that the sample mean may be a very poor estimator of the mean of the cross-sectional distribution of inflation.

Judging the relative efficiency of candidate estimators is straightforward. To do so, I use a Monte Carlo experiment based on the actual data. It proceed as follows. First, for each of the thirty-six components of the CPI, I take the deviation of monthly inflation, $\dot{p}_{it}$, from the thirty-six month centered moving average of inflation in the CPI-U. I refer to this as π_t^*. This yields a matrix of relative price changes that is 36 components by approximately 360 months. I draw a series of samples from this distribution of actual price changes by taking one randomly drawn observation for each of the 36 time-series — one draw from each column of the matrix.

This is a bootstrap procedure, and I generate 10,000 samples, each with 36 relative price changes. For each of these, I compute the mean, standard deviation and root-mean-square error of each of five estimators: (1) this CPI itself (the sample mean), (2) the CPI excluding food and energy (the mean of a sample with certain elements systematically excluded), (3) the 10% Trimmed Mean, (4) the 25% Trimmed Mean, and (5) the median. In all cases, the 1985 CPI expenditure weights are used. The results are reported in Table 1.

The results of this experiment are quite striking. First, there is a slight efficiency

[7]This experiment is described in Bryan and Cecchetti (1996).

[8]A simple Monte Carlo experiment, establishes that this sample value implies population that is fat-tailed. I drew 10,000 sample of 36 draws from a standard normal, and computed the kurtosis of each, using the CPI weights. Ninety-five percent of the resulting empirical distribution lies between 1.67 and 5.57.

Table 1: The Efficiency of Limited-Influence Estimators of Inflation

	CPI-U	CPI ex F&E	10% Trim	25% Trim	Median
Mean	0.057	0.124	-0.087	-0.004	0.027
St.Dev.	1.926	1.958	1.612	1.671	1.736
RMSE	1.926	1.959	1.618	1.672	1.736

Results of boot-strap experiment, 10,000 draws, deviation of inflation in 36 components from the 36 month centered moving average in the CPI-U. All estimators are constructed using the CPI weights. Mean of the actual distribution is 0.0587 percent per year.

loss from removing food and energy in the common way. The standard deviation of the CPI excluding food and energy is two-percent higher than that of the sample mean. But the real improvement comes from moving to trimmed means and the median. The RMSE of the sample median is 10 percent lower than that of the sample mean (the variance is 20% lower). The sample 10% Trimmed Mean is clearly the most efficient estimator, with an RMSE that is over 15% below that of the sample mean, with a 30% lower variance!

3.2 Reducing Noise

A second way to examine noise is to compare a sequence of measures of inflation at increasingly longer horizons. That is, monthly measures of one-month, three-month, six-month and twelve month inflation. In all cases, I continue to study deviations from a thirty-six month centered-moving-average objective such as π_t^*.

Defining

$$\pi_t^K = \frac{1}{K}[\ln(P_{t+K}) - \ln(P_t)] \ , \tag{15}$$

where, P_t is candidate measure of the price level, then I will study

$$\pi_t^K - \pi_t^*. \tag{16}$$

For the median and trimmed means, I calculate the index of price levels from monthly inflation estimates, and then use these to construct the multi-period inflation averages.

While it is computationally infeasible to evaluate the statistical efficiency of the Dynamic Factor Index, it can still be included in the analysis. I use it in two ways. First, it will be in the set of P_t's I examine. Second, I use it as an alternative to the CPI-U in constructing a measure of π_t^*.

Table 2 reports the root-mean-square error of the deviations of π_t^K from π_t^*, for each of three sample periods. Not surprisingly, if the objective is to follow the moving average of the Dynamic Factor Index, then the Dynamic Factor Index itself does the best job, regardless of the sample period and the horizon used.[9] If, however, the objective is to follow movements in the moving average of the CPI-U, then the Dynamic Factor Index is only best at short horizons. Using higher values for k, the limited influence estimators clearly out perform the Dynamic Factor Index in the most recent period. Again, the CPI excluding food and energy clearly worst.

It is worth noting that the RMSE criteria assumes that one has a symmetric aversion to over and underestimating π^*. This may not be the case. It seems likely that when inflation is high, a policy maker would be more averse to underestimating the inflation trend than to overestimating it. That is, the cost of missing an increase in inflation would be higher than the cost of missing an equivalent decline. When inflation is near zero, this asymmetry may flip because of the potentially high costs incurred when deflations force real interest rates to rise. In this case, policy makers might be more concerned about missing inflation declines than increases. The results presented in Table 2 are unaffected by moving to an asymmetric loss function.

A second way to present this information is to examine the distribution of the different measures of $(\pi_t^K - \pi_t^*)$. Table 3 reports the interval that contains 75% of the distribution for each measure of inflation. The table shows that using monthly inflation in the CPI-U, a 75% confidence interval for no change in π_t^* is about 3 percentage points. Even for a 12 month moving average, the confidence interval is still about 1 full percentage point wide. This is the sustained increase in the index

[9]This conclusion may be a bit unfair, as the Dynamic Factor Index estimates are constructed using full-sample information.

Table 2: Estimates of Noise in Price Indices

	$\dot{P}^*$= CPI-U				$\dot{P}^*$= DFI			
Sample: 1967.01 TO 1996.04								
	$k=1$	$k=3$	$k=6$	$k=12$	$k=1$	$k=3$	$k=6$	$k=12$
CPI-U	2.01	1.46	1.13	1.01	2.20	1.70	1.42	1.27
CPI ex F&E	2.13	1.65	1.37	1.29	2.30	1.85	1.60	1.47
10% Trim	1.80	**1.38**	**1.11**	1.02	1.99	1.61	1.37	1.24
25% Trim	1.91	1.47	1.17	1.03	2.12	1.72	1.46	1.30
Median	2.03	1.49	1.17	**1.01**	2.21	1.72	1.44	1.26
DFI	**1.66**	1.40	1.30	1.25	**1.54**	**1.25**	**1.11**	**0.98**
Sample: 1967.01 TO 1981.12								
	$\dot{P}^*$= CPI-U				$\dot{P}^*$= DFI			
	$k=1$	$k=3$	$k=6$	$k=12$	$k=1$	$k=3$	$k=6$	$k=12$
CPI-U	2.15	**1.62**	**1.36**	**1.30**	2.46	2.00	1.76	1.60
CPI ex F&E	2.52	2.04	1.76	1.65	2.70	2.24	1.95	1.74
10% Trim	2.09	1.68	1.41	1.34	2.38	2.01	1.76	1.58
25% Trim	2.20	1.78	1.48	1.36	2.51	2.14	1.87	1.67
Median	2.33	1.81	1.47	1.31	2.59	2.12	1.81	1.58
DFI	**1.91**	1.77	1.70	1.66	**1.79**	**1.62**	**1.49**	**1.31**
Sample: 1982.01 TO 1996.04								
	$\dot{P}^*$= CPI-U				$\dot{P}^*$= DFI			
	$k=1$	$k=3$	$k=6$	$k=12$	$k=1$	$k=3$	$k=6$	$k=12$
CPI-U	1.85	1.24	0.82	0.52	1.90	1.31	0.94	0.73
CPI ex F&E	1.64	1.07	0.78	0.63	1.80	1.31	1.11	1.02
10% Trim	1.43	0.95	**0.64**	**0.40**	1.47	1.03	0.77	0.61
25% Trim	1.54	1.00	0.68	0.43	1.60	1.10	0.84	0.68
Median	1.64	1.03	0.69	0.48	1.72	1.16	0.88	0.73
DFI	**1.36**	**0.88**	0.67	0.57	**1.22**	**0.68**	**0.46**	**0.35**

Notes: Calculations are the mean of $(\pi_t^k - \dot{P}_t^*)^2$. π^* is the 36 month centered moving average of inflation in either the CPI-U or the Dynamic Factor Index (DFI). All data is seasonally adjusted, monthly at annual rates. Bold-face numbers are the minimum for each k during each sample and π^*.

Table 3: Percentile Deviations of Inflation Measures from π^*

	$k=1$	$k=3$	$k=6$	$k=12$
CPI-U	1.56	1.07	0.70	0.42
	−1.91	−1.22	−0.81	−0.51
CPI ex F&E	1.71	1.13	0.98	0.90
	−0.94	−0.38	−0.23	−0.18
10% Trim	1.07	0.52	0.45	0.37
	−1.21	−0.73	−0.56	−0.53
25% Trim	1.33	0.70	0.57	0.52
	−1.26	−0.77	−0.50	−0.50
Median	1.59	0.94	0.65	0.65
	−1.72	−0.99	−0.65	−0.57
DFI	0.80	0.35	−0.02	0.05
	−2.07	−1.28	−0.99	−0.88

Sample period is 1982:01 to 1996:04. These are the 12.5 and 87.5 percentiles of the distribution of deviations in the index from the 36 month centered moving average of inflation in the CPI-U.

that would be required for one to change one's estimate of the level of π^*.

Moving to the alternative estimators shows that the 10% Trimmed Mean is the best on average. In fact, the 3 month change in the 10% Trimmed Mean has a smaller 75% confidence interval than the 6 month change in the CPI-U. The same comparison holds for the 10% Trimmed Mean averaged over 6 months and the CPI-U averaged over 12 months.[10]

The numbers in this table have implications for the size of a credible band that could be announced by a central banker interested in setting an inflation target. To see how they are used, take a case where the Federal Reserve announces that they are targeting the three-year moving average inflation rate and using a twelve-month

[10]While it is not the case in the data under study here, an additional problem arises if, when calculating average inflation over the past thirty years, the sample mean yields an answer above that of the sample trimmed means or median. This would be the result if the distribution of relative-price changes over the long-run were skewed. Should this have happened here, I would have adopted a procedure based on Roger's (1996) study of New Zealand, and centered the sample trimmed mean not on the 50th percentile of the distribution, but on the percentile that yields the same average inflation in the sample as the mean. For the New Zealand case, Roger reports that this is the 60th percentile.

average in a particular inflation measure to monitor their performance. If they chose the CPI-U to measure effectiveness, a specified a target band that was one-percentage point wide, historical experience implies that they would be outside of this band fully one-quarter of the time.

3.3 Seasonality

When making policy, central bankers would like to avoid responding to seasonal fluctuations in price data. While seasonality may be easy to understand in theory, it is extremely difficult to actually remove from most economic time-series. Figure 1 plots the seasonally adjusted and not seasonally adjusted inflation as measured by the CPI-U, together with the thirty-six month centered moving average. Note how little seasonal adjustment actually reduces high-frequency noise. Seasonal adjustment does very little to reduce short-run, monthly, variation in inflation data.

Table 4 reports a set of statistics which emphasize this point. I have calculated three measures over three separate sample periods, for both seasonally adjusted and unadjusted data. The first is the standard deviation of monthly inflation within the sample. The second is the R^2 from regressing the series on a set of seasonal dummy variables. And the third, labeled *Noise*, is the proportion of the variance in the series attributable to fluctuations of less than twelve months, one minus the ratio of the variance of monthly series to variance of the twelve month change, both at annual rates.[11]

Only 7% of the variation in the not-seasonally-adjusted CPI-U, over the full sample, is accounted for by seasonality. But the R^2's rise substantially for the more recent period, with over one-third of the variance in the unadjusted explained by seasonal dummy variables. Importantly, the high-frequency noise accounts for between one-third and three-quarters of the variation in the monthly *seasonally adjusted* inflation series, with substantially higher values in the more recent sample.

As Bryan and I have emphasized in a paper on seasonality in inflation, there is

[11]This is equivalent to one minus the R^2 of a 'regression' of the monthly series on lagged twelve-month changes, with the coefficient restricted to equal one.

FIGURE 1

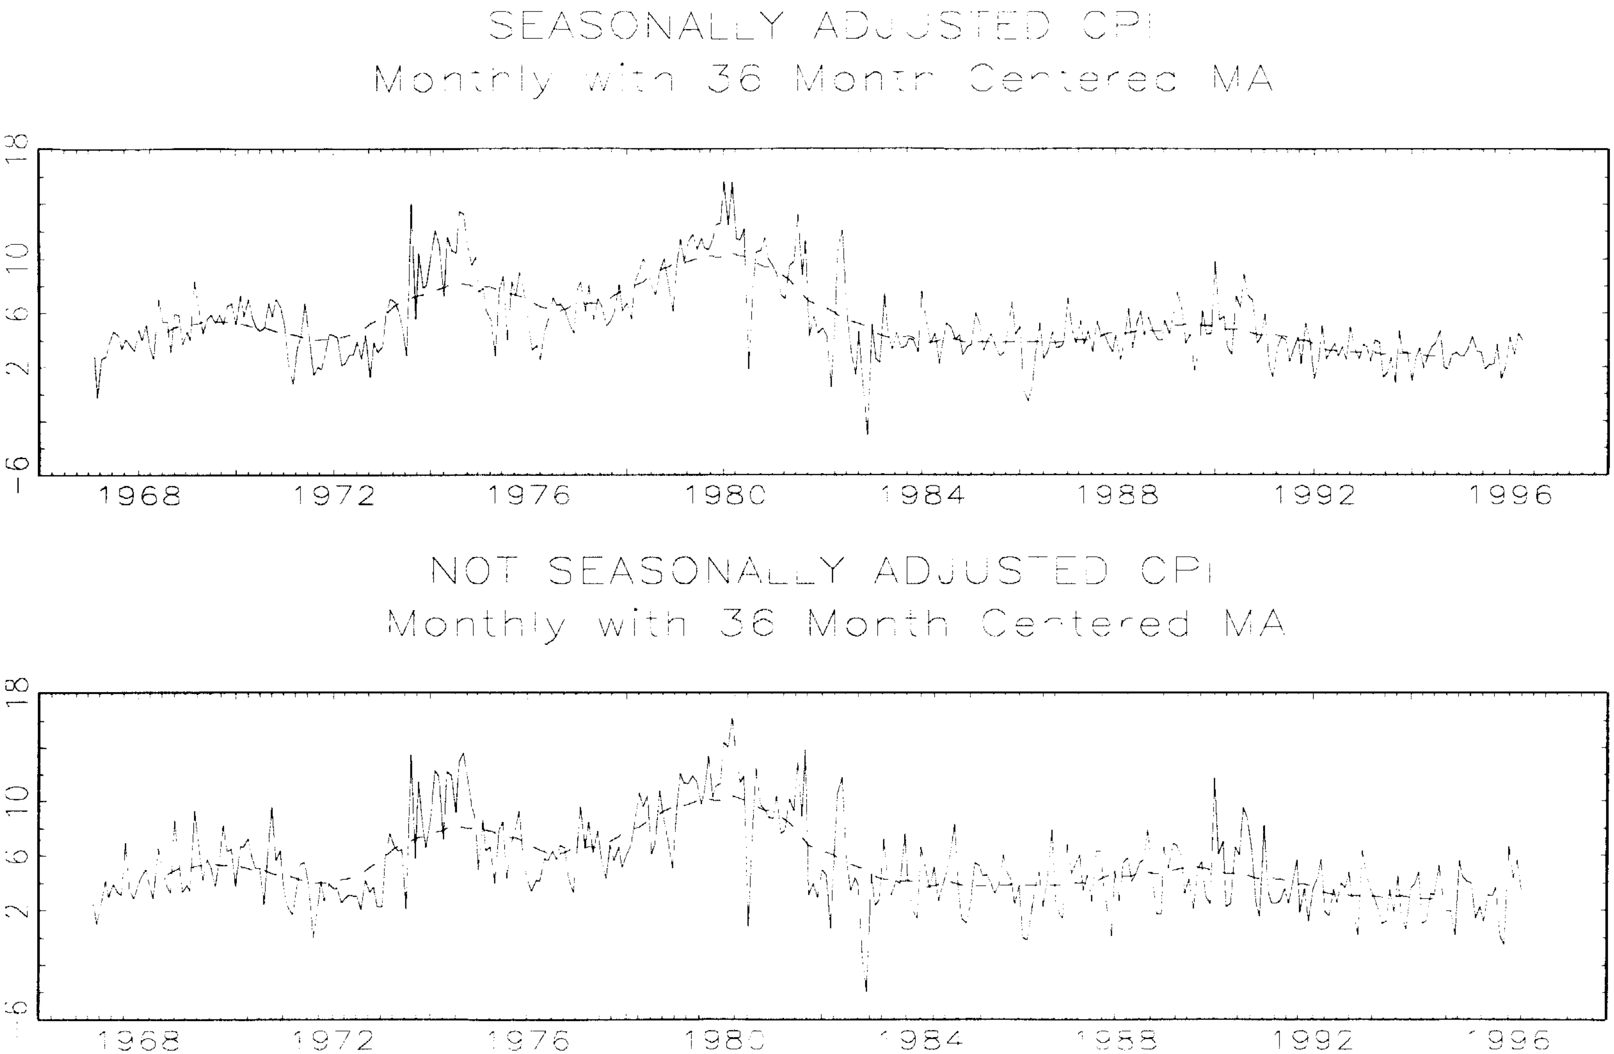
SEASONALLY ADJUSTED CPI
Monthly with 36 Month Centered MA
18
10
6
2
-6
1968
1972
1976
1980
1984
1988
1992
1996
NOT SEASONALLY ADJUSTED CPI
Monthly with 36 Month Centered MA
18
10
6
2
-6
1968
1972
1976
1980
1984
1988
1992
1996

Table 4: Seasonality in Inflation Data

	Seasonally Adjusted			Not Seasonally Adjusted		
Sample 1967.01 to 1996.04						
	St.Dev.	R^2	Noise	St.Dev.	R^2	Noise
CPI-U	3.00	0.01	0.37	3.22	0.07	0.45
CPI ex F&E	2.90	0.02	0.43	3.41	0.18	0.59
10% Trim	2.92	0.01	0.34	3.15	0.05	0.42
25% Trim	2.92	0.01	0.39	3.21	0.05	0.45
Median	2.85	0.01	0.32	3.01	0.05	0.39
DFI	2.47	0.01	0.28	—	—	—
Sample 1967.01 to 1981.12						
CPI-U	3.19	0.01	0.35	3.35	0.05	0.42
CPI ex F&E	3.33	0.02	0.43	3.72	0.18	0.55
10% Trim	3.22	0.02	0.35	3.35	0.03	0.39
25% Trim	3.19	0.02	0.41	3.38	0.02	0.40
Median	3.12	0.02	0.34	3.19	0.03	0.36
DFI	2.52	0.01	0.29	—	—	—
Sample 1982.01 to 1996.04						
CPI-U	1.89	0.07	0.72	2.33	0.34	0.83
CPI ex F&E	1.70	0.08	0.66	2.58	0.47	0.90
10% Trim	1.61	0.06	0.64	2.08	0.36	0.84
25% Trim	1.70	0.06	0.76	2.26	0.34	0.87
Median	1.52	0.07	0.57	1.92	0.36	0.81
DFI	1.49	0.09	0.64	—	—	—

Standard deviation is that of inflation at an annual rate. R^2 is multiple correlation coefficient from a regression on monthly seasonal dummy variables. 'Noise' is one minus the ratio of the variance in the twelve month inflation to the variance in one-month inflation.

tremendous time-variation in the pattern of seasonals in prices.[12] We make the point that since 1982, seasonality in inflation has been much more pronounced. The reason is that trend inflation has been very stable, and so it is easier to extract the seasonal pattern from the observed monthly fluctuations.

In that same paper, We also express our view that it would advantageous to seasonally adjust inflation data at the aggregate level. The current procedure is to test disaggregated series for the presence of seasonality, and then to seasonally adjust only those components of the price index that show sufficient statistical evidence of seasonal fluctuations. We argue that such a procedure suffers from a technical difficulty that arises because of the variation in relative price movements that is present in the disaggregated price series. Relative price inflation is a form of statistical pollution in these series that reduces the ability of an econometrician to measure the presence of seasonality. In other words, it reduces the statistical power of the pre-tests used for the initial decision of whether or not to seasonally adjust. Our conclusion is that if one is interested in a seasonally adjusted series, as monetary policymakers presumably are, then adjustment should be done at the aggregate level.

Yet another caution is in order here. Using the same techniques described in Cecchetti, Kashyap and Wilcox (1995), I have examined the question of whether the seasonality in inflation varies over the business cycle and found that for the CPI-U over the 1982 to 1996 sample, it very clearly does. The variance in the seasonals, i.e. the coefficients in a seasonal dummy variable regression, shrinks by fifty percent from a typical recession to a typical boom.

What happens to seasonality as data is averaged over three or six months? Obviously, the problem will decline in importance, but how quickly? Table 5 reports results for the 1982 to 1996 sample period. The first panel replicates the last panel of Table 4. The results are interesting in that they again suggest that there are substantial gains associated with moving from one-month to three-month inflation measures. For the 10% Trimmed Mean, for example, the amount of noise is cut by more than half, from 0.64 for $k = 1$ to 0.28 for $k = 3$. Moving from three- to six-month changes

[12]See Bryan and Cecchetti (1995).

Table 5: Seasonality in Inflation Data at Various Horizons

Sample 1982.01 to 1996.04

	Seasonally Adjusted			Not Seasonally Adjusted		
	$k=1$					
	St.Dev.	R^2	Noise	St.Dev.	R^2	Noise
CPI-U	1.89	0.07	0.72	2.33	0.34	0.83
CPI ex F&E	1.70	0.08	0.66	2.58	0.47	0.90
10% Trim	1.61	0.06	0.64	2.08	0.36	0.84
25% Trim	1.70	0.06	0.76	2.26	0.34	0.87
Median	1.52	0.07	0.57	1.92	0.36	0.81
DFI	1.49	0.09	0.64	—	—	—
	$k=3$					
CPI-U	1.40	0.02	0.51	1.63	0.21	0.66
CPI ex F&E	1.22	0.03	0.36	1.76	0.39	0.78
10% Trim	1.18	0.04	0.28	1.38	0.20	0.59
25% Trim	1.19	0.02	0.39	1.46	0.22	0.65
Median	1.16	0.04	0.28	1.29	0.16	0.55
DFI	1.11	0.02	0.33	—	—	—
	$k=6$					
CPI-U	1.08	0.01	0.29	1.16	0.07	0.37
CPI ex F&E	0.97	0.02	0.22	1.08	0.07	0.40
10% Trim	0.93	0.03	0.18	0.97	0.06	0.27
25% Trim	0.90	0.02	0.19	0.99	0.05	0.31
Median	0.92	0.03	0.17	0.93	0.05	0.26
DFI	0.96	0.01	0.11	—	—	—

Standard deviation is for inflation at an annual rate. R^2 is multiple correlation coefficient for a regression on monthly seasonal dummy variables. "Noise" is one minus the ratio of the variance in the twelve month inflation to the variance in one-month inflation.

reduces the noise to 0.18.

Comparing the noise in the CPI-U with that in the other estimators yields similar results to those gleaned from Table 3. For example, the one month change in the median has approximately the same noise as the three month change in the CPI-U. And the three month change in the 10% Trimmed Mean looks like the six-month change in the CPI-U.

I draw several overall conclusions from this exercise. First, the use of monthly changes in monthly data is clearly unwise. Instead, one should focus on at least three-month changes. There is no obvious reason to go all the way to twelve month changes. Second, I believe that there is clear evidence favoring the use of limited influence estimators. In the horse race I have conducted here, using recent U.S. data, the three- to six- month changes in the 10% Trimmed Mean is the clear winner.

4 Bias

In Section 2.2 I argued that by measuring the common inflation trend in a broad cross-section of prices, the Dynamic Factor Index eliminated one source of bias in inflation statistics. Therefore, by combining the Dynamic Factor Index with another measure of inflation I can begin to gauge the importance of time-variation in the weighting bias. Figure 2 plots a 12 month inflation in the DFI, the Median and the CPI-U. I smooth the data over 12 months in order to reduce the noise.[13]

I suggested earlier that the time-variation in the weighting bias might be correlated with two easily measurable quantities: the stage of the business cycle and the spread of the price change distribution. To examine this hypothesis, I use the following regression

$$\pi_t^{12} - \hat{\mathcal{P}}_t^{12} = \alpha + \beta_0 \sigma_t(\dot{p}_i) + \beta_1 \mathcal{S}_t(\dot{p}_i) + \beta_2 \lambda_t + u_t \ , \tag{17}$$

where $\sigma_t(\dot{p}_i)$ and $\mathcal{S}_t(\dot{p}_i)$ are the standard deviation and the skewness, respectively, of the distribution of 12 month inflation measured over the 36 components in the CPI,

[13]It would be possible in principle to construct confidence bands on the bias from the implied distributions of the coefficients estimated to construct the DFI. I leave this for future work.

FIGURE 2

CPI--U, MEDIAN CPI and DFI, sa
12 Month Change

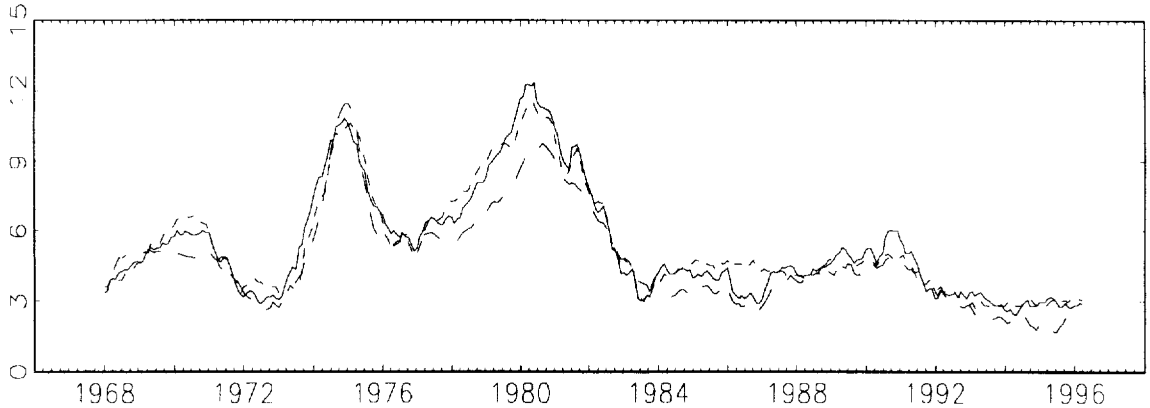

ESTIMATED BIAS in CPI--U and MEDIAN CPI, sa
CPI--DFI, Med--DFI 12 Month Changes

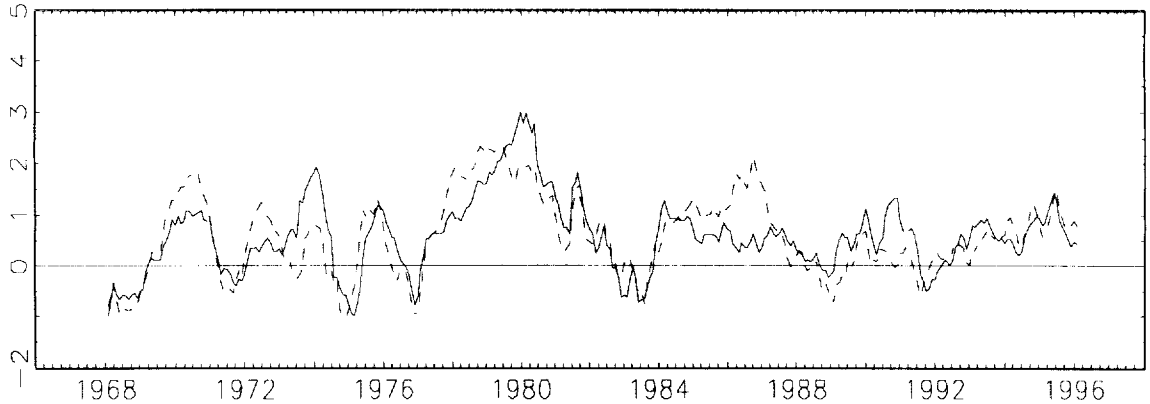

and λ_t is the level of capacity utilization, all averaged over the same 12 month period.

Table 6 reports the results of this regression, as well as the summary statistics for the estimated weighting bias over various samples using a number of inflation measures. First, note that the average weighting bias ranges from 0.40 percentage points, for the 10% Trimmed Mean over the 1982 to 1996 period, to 0.85 for the CPI excluding Food and Energy over the same sample. Focusing again on the more recent sample period, there is clear evidence that the estimated bias increases with the cross-sectional variance in relative price changes.

5 Conclusion

How can policymakers obtain timely measures of movements in long-run inflation trends? First, monthly percentage changes in virtually any inflation measure contain so much noise as to be virtually useless. Second, the CPI excluding food and energy is an extremely poor measure of any underlying trend, or core component of the CPI. It is *not* less volatile than the CPI-U itself. In fact, it usually fares worse than the overall price index.

After examining alternatives to the standard measures, I conclude that limited influence estimators are more efficient estimators of the central tendency of the price change distribution than is the overall mean. In particular, given the properties of U.S. price data, the 10% Trimmed Mean provides the measure of the changes in long-run trend inflation.

The results also lead to a conclusion regarding the frequency at which data is actually useful. Moving from one-month to three-month changes reduces the noise in the data so much that it is difficult to see why someone would look at monthly data. In fact, it may actually be that the costs of collecting monthly data exceed the benefits, given how little information it seems to contain.

Finally, I have presented a set of results that concern the size of the weighting bias in inflation measures and its variation over time. My reduced-bias estimates, which are constructed using the Dynamic Factor Index, are generally around one-half of

Table 6: Measuring Time-Variance in Weighting Bias

Sample	Mean Bias	$\sigma_t(\dot{p}_i)$	$\mathcal{S}_t(\dot{p}_i)$	λ_t	R^2
Weighting Bias in CPI-U					
67:01 to 96:04	0.59	0.145	0.0010	0.051	0.24
		(0.08)	(0.04)	(0.08)	
67:01 to 81:12	0.71	0.242	0.0002	0.040	0.27
		(0.04)	(0.85)	(0.30)	
82:01 to 96:04	0.46	0.062	0.0008	0.075	0.23
		(0.09)	(0.09)	(0.00)	
Weighting Bias in CPI ex F&E					
67:01 to 96:04	0.67	0.047	-0.0017	-0.019	0.13
		(0.70)	(0.01)	(0.59)	
67:01 to 81:12	0.48	0.126	-0.0024	-0.023	0.08
		(0.45)	(0.05)	(0.53)	
82:01 to 96:04	0.85	0.082	-0.0013	0.034	0.28
		(0.02)	(0.00)	(0.29)	
Weighting Bias in 10% Trim					
67:01 to 96:04	0.50	0.165	-0.0009	0.043	0.16
		(0.09)	(0.09)	(0.14)	
67:01 to 81:12	0.61	0.218	-0.0014	0.039	0.13
		(0.11)	(0.25)	(0.31)	
82:01 to 96:04	0.40	0.075	-0.0013	0.034	0.29
		(0.01)	(0.00)	(0.17)	
Weighting Bias in 25% Trim					
67:01 to 96:04	0.56	0.168	-0.0012	0.052	0.15
		(0.14)	(0.04)	(0.14)	
67:01 to 81:12	0.67	0.223	-0.0019	0.042	0.10
		(0.16)	(0.19)	(0.32)	
82:01 to 96:04	0.45	0.098	-0.0015	0.057	0.35
		(0.00)	(0.00)	(0.02)	
Weighting Bias in Median CPI					
67:01 to 96:04	0.59	0.135	-0.0016	0.042	0.14
		(0.10)	(0.00)	(0.22)	
67:01 to 81:12	0.68	0.142	-0.0017	0.032	0.05
		(0.20)	(0.25)	(0.46)	
82:01 to 96:04	0.50	0.086	-0.0019	0.038	0.36
		(0.03)	(0.00)	(0.12)	

Results from regression bias on distributional characteristics of relative price changes and a measure of the stage of the business cycle. Numbers in parentheses are p-values for the two-sided test that the coefficient equals zero. Standard errors are computed using the West and Newey (1994) automatic lag selection procedure.

one percentage point at an annual rate. The weighing bias estimates themselves have substantial time-variation, and this is a further source of noise in inflation statistics.

References

Ball, Laurence and N. Gregory Mankiw, "Relative-Price Changes as Aggregate Supply Shocks," *Quarterly Journal of Economics* 110 (February 1995) 161-193.

Bryan, Michael F. and Stephen G. Cecchetti, 'The Consumer Price Index as a Measure of Inflation,' *Economic Review of the Federal Reserve Bank of Cleveland* 29 (1993 Quarter 4) 15–24.

______ and ______, 'Measuring Core Inflation,' in *Monetary Policy*, N. Gregory Mankiw, ed., Chicago: University of Chicago Press for NBER, 1994, 195–215.

______ and ______, 'The Seasonality of Inflation,' *Economic Review of the Federal Reserve Bank of Cleveland*, (1995 Quarter 2) 12–23.

______ and ______, 'Inflation and Skewness in the Distribution of Price Changes: Theoretical Issues and Small-Sample Inference,' mimeo., August 1996.

Cecchetti, Stephen G., Anil K Kashyap and David W. Wilcox, 'Do Firms Smooth the Seasonal in Production in a Boom? Theory and Evidence,' mimeo., October 1995.

Gordon, Robert J., 'Measuring the Aggregate Price Level: Implications for Economics Performance and Policy,' N.B.E.R. Working Paper No. 3969, January 1992.

Haldane, Andrew G., 'Introduction,' in A. G. Haldane, ed., *Targeting Inflation*, Bank of England, November 1995, pg. 1-12.

Lebow, David E., John M. Roberts and David J. Stockton, 'Economic Performance under Price Stability,' Board of Governors of the Federal Reserve System, Working Paper No. 125, April 1992.

Newey, Whitney K. and Kenneth D. West, 'Automatic Lag Selection and Robust Covariance Matrix Estimation with Prewhitening' *Review of Economic Studies* 61 (October 1994) 631-53.

Roger, Scott, ' ' Working Paper, Reserve Bank of New Zealand, forthcoming.

Shapiro, Matthew D. and David W. Wilcox, 'Bias in the Consumer Price Index,' in B. Bernanke and J. Rotemberg, eds. *NBER Macroeconomics Annual*, Cambridge, MA.: M.I.T. Press, 1996, forthcoming.

Stock, James H. and Mark W. Watson, "A Probability Model of the Coincident Economic Indicators," in K. Lahiri and G. H. Moore, ed., *Leading Economic Indicators: New Approaches and Forecasting Records.* Cambridge: Cambridge University Press, 1991, pp. 63-89.

Wynne, Mark and Fiona Sigalla, 'A Survey of Measurement Biases in Price Indexes,' Federal Reserve Bank of Dallas, Research Paper No. 9340, October 1993.