

Stephen G. Cecchetti is executive vice president and director of research at the Federal Reserve Bank of New York and a research associate with the National Bureau of Economic Research. The author is grateful to the Department of Economics, University of Melbourne, for their generous hospitality; to Alan Blinder, Michael Bryan, Margaret Mary McConnell, Mark Wynne, and seminar participants at the European Monetary Institute, the Innocenzo Gasparini Institute for Economic Research, the Reserve Bank of Australia, and the Università di Brescia for their comments; and to the Federal Reserve Bank of Cleveland for financial and research support.

Measuring Short-Run Inflation for Central Bankers

Stephen G. Cecchetti

In recent years, many 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. There is active debate over whether the United States will 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 article is to examine the difficulties surrounding the measurement of the changes in the aggregate price level.

Measuring inflation is a surprisingly difficult task. Although it is conceptually easy to survey the prices of individual commodities at any given time, using them 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, two problems are associated with measuring inflation. The first concerns transitory phenomena, or *noise*, that should not affect policymakers' 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 indexes 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 the noise in the targeted price

index. Second, so long as inflation has some weight in a policymaker's objective function, it is important to know how to interpret monthly movements in aggregate prices. This argues for the development of indexes aimed at minimizing the presence of noise.

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 indexes. These biases can be divided into the following two broad categories:

- Those biases related to the way in which individual prices are weighted together to form an aggregate index (weighting bias). An example is substitution bias.
- Those biases that result from actual errors in measuring the individual prices themselves (measurement bias), such as quality or new goods bias.

Although noise is by definition temporary, bias is not. The importance of bias for monetary policy is twofold. 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 this longer run issue, it seems likely that biases themselves will be time varying. Again, this suggests that it is difficult to interpret movements in measured price indexes.

Bias might vary through time for many reasons. For example, if the rate of technological progress varies, then bias resulting from quality adjustment may not be time invariant.¹ Substitution bias may also be time varying. In the consumer price index (CPI), this bias is caused by the fixed expenditure-based weighting scheme used in its construction. This occurs because, as relative prices change,

¹ Gordon (1992) measures quality bias in consumer durable goods and finds large differences through time.

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, however, the extent of this bias will depend on the degree to which such substitution occurs. If during a particular period very little variation occurs in price changes across commodities, then one would expect the substitution bias to be small. By contrast, during times in which a large spread in cross-sectional inflation occurs, substitution bias might be large.

Numerous researchers 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 (forthcoming) all provide a catalog of estimates of the biases in the CPI. But all these estimates come from studies of product-specific microeconomic data and so lack the generality necessary to help gauge the overall bias in the aggregate index.

A number of solutions have been suggested to remedy price-index measurement problems. Solutions to the high-frequency noise problem 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 after the initial shock and move down after 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) attempt a general treatment of the weighting-bias problem. There, weighting bias is treated as a statistical problem that can be overcome (at least in part) by 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 manner) commodities based on the strength of a common inflation signal.

My purpose here is to examine the severity of the noise- and weighting-bias problems and propose some partial solutions. I begin with a simple framework intended to clarify the issues associated with noise and weighting bias. I then 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 the averaging over 3-, 6-, and 12-month horizons. A simple series of tests suggests that the most accurate estimator is a 10 percent trimmed mean averaged over three to six months. I go on to examine the problem of weighting bias, suggesting the extent to which the DFI is a reduced-bias estimator of aggregate inflation. I then present my conclusions.

GENERAL CONSIDERATIONS

It is useful to begin with a general discussion of noise and bias. First, I will present a simple accounting framework

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

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.

that illuminates 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, I describe the DFI. 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.

A Framework for Analysis

During any given 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 p_{it} as the change in the log of individual price i at time t , \dot{P}_t as the common factor, and \dot{x}_{it} as relative price inflation, this can be written as

$$(1) \quad \dot{p}_{it} = \dot{P}_t + \dot{x}_{it}.$$

\dot{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. For example, defining the weights as w_{it} , then,

$$(2) \quad \pi_t \equiv \sum_i w_{it} \dot{p}_{it}.$$

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

$$(3) \quad \sum_i w_{it} = 1, \forall t.$$

Combining Equations 1 and 2, and using Equation 3 yields

$$(4) \quad \pi_t = \dot{P}_t + \sum_i w_{it} \dot{x}_{it}.$$

This formulation clearly shows that the deviation of measured price indexes from the common inflation is given by the second term on the right side of Equation 4. The problem with standard indexes is that this term is nonzero—both at any given time and in expectation.

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

$$(5) \quad \pi_t - \dot{P}_t = \sum_i w_{it} \dot{x}_{it} = n_t + b_t.$$

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

$$(6) \quad \lim_{k \rightarrow \infty} \frac{1}{k} \sum_{j=1}^k n_{t+j} = 0.$$

This has the important implication that lengthening the observation interval eliminates the transitory noise. In other words, even if data are available monthly, it may be wise to look at changes over 3, 6, or even 12 months.

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

$$(7) \quad b_t = \mu_b + \omega_t,$$

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 occur when a relatively large number of new products are introduced. Quality and new goods biases might therefore be large.

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

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

that imply a measure of aggregate price inflation like

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

The immediate implication is that length-

ening 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 policymakers. As an alternative, core inflation measures, such as the CPI excluding food, energy, and the weighted median, are designed to minimize the observed noise without sacrificing the high frequency of measurement. As Bryan and Cecchetti (1994) discuss, the weighted median reduces noise in a number of ways. First, it downweights the importance of sector-specific shocks likely to only eventually average to zero across all prices. And second, the weighted median 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 \hat{p}_{it} s, together with their weights, and orders them from largest to smallest. The weighted median is the \hat{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 percent trimmed mean as the measure constructed by removing 25 percent of the upper and lower tails of the cross-sectional distribution of the \hat{p}_{it} s, and then averaging the remaining 50 percent.

The Dynamic Factor Index

In this section I estimate a measure 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 \hat{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. The objective here is to compute a reduced-bias estimate of inflation from consumer price data.⁴

Bryan and Cecchetti (1993) propose an alternative to the expenditure-weighting schemes that is based on the strength of the inflation signal, \hat{P}_t , relative to the noise, \hat{x}_{it} , in each observed time series, \hat{p}_{it} . To do this, they assume the following model:

$$(10) \quad \hat{p}_t = \hat{P}_t + \hat{x}_{it},$$

$$(11) \quad \psi(L)\hat{P}_t = \delta + \xi_t, \text{ and}$$

$$(12) \quad \theta(L)\hat{x}_t = \eta_t,$$

where \hat{p}_t and \hat{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, \hat{P}_t , is identified by assuming it to be uncorrelated with relative price disturbances at all leads and lags. Logically, if one of the \hat{x}_{it} s were correlated with \hat{P}_t , it would not be idiosyncratic.

Bryan and Cecchetti (1993) estimate the model using a Kalman filtering algorithm, assuming all lag polynomials are AR(2). The result is an estimate of \hat{P}_t the DFI,

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

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

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

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

Use of the DFI has one clear advantage over other methods of bias reduction. Because 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,

⁴ See Bryan and Cecchetti (1993).

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 indexes (e.g., difficulties in measuring quality changes), the DFI will be biased as well.⁵ Second, the DFI still contains transitory noise. While $E(\hat{P}) = \hat{P}$, there will be deviations of the estimates from the true value. But even more important, since the DFI is constructed using an econometric procedure, the entire history can change as new data are added and the parameters of the model are reestimated.

NOISE

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 crucial to monetary policymakers' decisions is timely (i.e., high-frequency) estimates of movements in the long-term trend in inflation. As an approximation, I will define this trend to be the 36-month centered moving average of actual inflation. The choice of a 3-year window is not crucial; it is simply meant as an illustration.⁶

This section is devoted to measuring different procedures' ability 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 an estimator's small-sample variance. To compare different estimators for a parameter, such as the population mean of this distribution, it is natural to calculate the variance of each candidate estimator. The estimator with the lowest variance is then the most efficient.

The second standard for comparison is to review the distribution of deviations of each estimator from the 36-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 a result of seasonality that can be easily removed? Do any of the candidate measures have less seasonality than the others? How quickly does the seasonality disappear as the observation interval is increased?

Statistical Efficiency

All 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 the 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 same disaggregated component price series.

To understand the efficiency issue, compare the sample mean and the sample median. Both attempt to measure the population mean of a distribution. But for any given sample, one would not expect the two measures 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.

Although the sample mean is 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 easy to demonstrate that the sample median is more efficient than the sample mean when the data are drawn from leptokurtic distributions (i.e., distributions with fat tails relative to the norm). The reason for this is that, with a fat-tailed distri-

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

⁶ Changing to a 3-year *forward*-moving average does not change the qualitative results I will report.

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 bootstrap 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 0.05867 percent per year.

bution, it is more likely 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. In other words, 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 increasingly more of the sample. For 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 four to five causes the variance of the sample 10 percent trimmed mean to fall below that of the sample mean. As the kurtosis rises further, it is optimal (in this sense) to trim increasingly 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 (in 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. First, for each of the 36 components of the CPI, I take the deviation

of monthly inflation, p_{it} , from the 36-month centered moving average of inflation in the CPI-U. I refer to this as π_t^* . This yields a relative price-change matrix 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 through which 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 (RMSE) of each of the five following estimators: (1) the 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 percent trimmed mean, (4) the 25 percent 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, a slight efficiency loss occurs from removing food and energy in the common way. The standard deviation of the CPI excluding food and energy is 2 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 percent lower). The sample 10 percent trimmed mean is clearly the most efficient estimator, with an RMSE that is more than 15 percent below that of the sample mean, with a 30 percent lower variance!

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

⁸ A simple Monte Carlo experiment establishes that this sample value implies a fat-tailed population. 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 2

Estimates of Noise in Price Indexes

	Sample: 1967.01 to 1996.04							
	$\dot{P}^* = \text{CPI-U}$				$\dot{P}^* = \text{DFI}$			
	<i>k</i> = 1	<i>k</i> = 3	<i>k</i> = 6	<i>k</i> = 12	<i>k</i> = 1	<i>k</i> = 3	<i>k</i> = 6	<i>k</i> = 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							
	$\dot{P}^* = \text{CPI-U}$				$\dot{P}^* = \text{DFI}$			
	<i>k</i> = 1	<i>k</i> = 3	<i>k</i> = 6	<i>k</i> = 12	<i>k</i> = 1	<i>k</i> = 3	<i>k</i> = 6	<i>k</i> = 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							
	$\dot{P}^* = \text{CPI-U}$				$\dot{P}^* = \text{DFI}$			
	<i>k</i> = 1	<i>k</i> = 3	<i>k</i> = 6	<i>k</i> = 12	<i>k</i> = 1	<i>k</i> = 3	<i>k</i> = 6	<i>k</i> = 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_t^* is the 36-month centered moving average of inflation in either the CPI-U or the DFI. All data are seasonally adjusted, monthly at annual rates. Boldfaced numbers are the minimum for each *k* during each sample and π_t^* .

Reducing Noise

A second way to examine noise is to compare a sequence of inflation measures at increasingly longer horizons (i.e., measures of 1-month, 3-month, 6-month, and 12-month inflation). In all cases, I continue to study deviations from a 36-month centered moving-average objective such as π_t^* . Defining

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

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

$$(16) \quad \pi_t^K - \pi_t^*.$$

Table 3

Percentage Deviations of Inflation Measures from π^*


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

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

Figure 1


Seasonally Adjusted CPI

Monthly with 36-month centered moving average



Unadjusted CPI

Monthly with 36-month centered moving average



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

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

Table 2 reports the RMSE of the deviations of π_t^k from π_t^* for three sample periods. Not surprisingly, if the objective is to follow the DFI's moving average, then the DFI 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 DFI is only best at short horizons. Using higher values for k , the limited influence estimators clearly outperform the DFI in the most recent period. Again, the worst results are those for the CPI excluding food and energy.

Note that the RMSE criteria assume one has a symmetric aversion to overestimating and underestimating π^* . However, this may not be the case. It seems likely that when inflation is high, a policymaker would be more averse to underestimating the inflation trend than to overestimating it. In other words, 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, policymakers 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 percent of the distribution for each measure of inflation. Table 3 shows that using monthly inflation in the CPI-U, a 75 percent 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 that would be required for one to change one's estimate of the level of π^* .

The alternative estimators show that the 10 percent trimmed mean is the best on average. In fact, the 3-month change in the 10 percent trimmed mean has a smaller 75 percent confidence interval than the 6-month change in the CPI-U. The same comparison holds for the 10 percent trimmed mean averaged over 6 months and the CPI-U averaged over 12 months.¹⁰

Table 3 indicates 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 the numbers in Table 3 are used, take a case where the Federal Reserve announces it is targeting the 3-year moving average inflation rate and using a 12-month average in a particular inflation measure to monitor their performance. If the Fed chooses the CPI-U to measure effectiveness, a specified target band that was 1 percentage point wide, historical experience implies that the results would be outside this band fully one-quarter of the time.

Seasonality

When setting policy, central bankers would like to avoid responding to seasonal fluctuations in price data. Although seasonality may be easy to understand in theory, it is extremely difficult to remove. Figure 1 plots inflation (seasonally adjusted and unadjusted) as measured by the CPI-U, together with the 36-month centered moving average, using the same scale for both. 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 presents a set of statistics to emphasize this point. I have calculated three measures through three separate sample periods for both seasonally adjusted

Table 4

Seasonality in Inflation Data

	Seasonally Adjusted			Unadjusted		
	St. Dev.	R ²	Noise	St. Dev.	R ²	Noise
Sample: 1967.01 to 1996.04						
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	—	—	—

NOTES: Standard deviation is that of inflation at an annual rate. R² is a multiple correlation coefficient from a regression on monthly seasonal dummy variables. Noise is one minus the ratio of the variance in the 12-month inflation to the variance in 1-month inflation.

and unadjusted data: the standard deviation of monthly inflation within the sample; the R² from regressing the series on a set of seasonal dummy variables; and the proportion of the variance in the series attributable to fluctuations of fewer than 12 months, one minus the ratio of the variance of monthly series to variance of the 12-month change, both at annual rates (noise).¹¹

Only 7 percent of the variation in the unadjusted CPI-U, over the full sample, is accounted for by seasonality. But the R²'s rise substantially for the more recent

⁹ This conclusion may be a bit unfair because the DFI estimates are constructed using full-sample information.

¹⁰ Although it is not the case in the data under study here, an additional problem arises if, when calculating average inflation over the past 30 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

Table 5

Seasonality in Inflation Data at Various Horizons

	Sample: 1982.01 to 1996.04					
	Seasonally Adjusted			Unadjusted		
	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	—	—	—

NOTES: Standard deviation is that of inflation at an annual rate. R² is a multiple correlation coefficient from a regression on monthly seasonal dummy variables. Noise is one minus the ratio of the variance in the 12-month inflation to the variance in 1-month inflation.

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. I would have 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 57th percentile.

period, with more than one-third of the variance in the unadjusted CPI-U explained by seasonal dummy variables. It is important to note that 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.

Tremendous time variation occurs in the pattern of seasonals in prices.¹² Bryan and Cecchetti (1995) make the point that since 1982, seasonality in inflation has been much

more pronounced. The reason is that trend inflation has been very stable. To extract the seasonal pattern from the observed monthly fluctuations is therefore easier.

Bryan and Cecchetti (1995) also express the view that it would be 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. Bryan and Cecchetti argue that such a procedure suffers from a technical difficulty because of the variation in relative price movements 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, relative price inflation reduces the statistical power of the pretests used for the initial decision of whether to seasonally adjust. Bryan and Cecchetti conclude 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 (1996), I examined the question of whether the seasonality in inflation varies during the business cycle and found that—for the CPI-U in 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 50 percent from a typical recession to a typical boom.

What happens to seasonality as data are averaged over 3 or 6 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 because they again suggest that substantial gains are associated with moving from 1-month to 3-month inflation measures. For the 10 percent 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 3- to 6-month changes 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 1-month change in the median has approximately the same noise as the 3-month change in the CPI-U. And the three-month change in the 10 percent trimmed mean looks like the 6-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 3-month changes. No obvious reason exists to go all the way to 12-month changes. Second, I believe clear evidence favors the use of limited influence estimators. In the horse race I have conducted here using recent U.S. data, the 3-to 6-month changes in the 10 percent trimmed mean are the clear winners.

BIAS

Earlier I argued that by measuring the common inflation trend in a broad cross-section of prices, the DFI eliminated one source of bias in inflation statistics. Therefore, by combining the DFI 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 to reduce the noise.¹³

I have suggested that the time variation in the weighting bias might be correlated using 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:


$$(17) \quad \begin{aligned} \pi_t^{12} - \hat{p}_t^{12} = & \alpha + \beta_0 \sigma_t(\dot{p}_i) \\ & + \beta_1 S_t(\dot{p}_i) + \beta_2 \lambda_t + u_t, \end{aligned}$$

where $\sigma_t(\dot{p})$ and $S_t(\dot{p})$ are the standard deviation and the skewness, respectively, of the distribution of 12-month inflation measured over the 36 components in the CPI, and λ_t is

Figure 2


CPI-U, Median CPI, and DFI

Seasonally adjusted, 12-month change



Estimated Bias in CPI-U and Median CPI

Seasonally adjusted, CPI-DFI, Med-DFI 12-month change



the level of capacity utilization, all averaged during the same 12-month period.

Table 6 presents 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. Note that the average weighting bias ranges from 0.40 percentage points, for the 10 percent trimmed mean during the 1982 to 1996 period, to 0.85 for the CPI excluding food and energy during the same sample. Focusing again on the more recent sample period, the estimated bias clearly increases with the cross-sectional variance in relative price changes.

CONCLUSION

How can policymakers obtain timely measures of movements in long-run infla-

¹¹ This is equivalent to one minus the R^2 of a regression of the monthly series on lagged 12-month changes, with the coefficient restricted to equal one.

¹² See Bryan and Cecchetti (1995).

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.

Table 6

Measuring Time Variance in Weighting Bias

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

NOTE: Results are 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 in which the coefficient equals zero. Standard errors are computed using the Newey and West (1994) automatic lag selection procedure.

tion trends? First, monthly percentage changes in virtually any inflation measure contain so much noise that they are 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 percent 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 are actually useful. Moving from 1-month to 3-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 the monthly data seem 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 DFI, are generally around 1/2 of 1 percentage point at an annual rate. The weighting bias estimates themselves have substantial time variation—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* (February 1995), pp. 161-93.
- Bell, William R., and Steven C. Hillmer. "Issues Involved With the Seasonal Adjustment of Economic Time Series," *Journal of Business and Economic Statistics* (October 1984), pp. 291-320.
- Bryan, Michael F., and Stephen G. Cecchetti. "The Consumer Price Index as a Measure of Inflation," *Economic Review*, Federal Reserve Bank of Cleveland (1993 : 4), pp. 15-24.

_____ and _____. "Measuring Core Inflation," *Monetary Policy*, N. Gregory Mankiw, ed., University of Chicago Press, 1994, pp. 195-215.

_____ and _____. "The Seasonality of Inflation," *Economic Review*, Federal Reserve Bank of Cleveland (1995 : 2) pp. 12-23.

_____ and _____. "Inflation and the Distribution of Price Changes," NBER working paper No. 5793, October 1996.

Cecchetti, Stephen G., Anil K. Kashyap, and David W. Wilcox. "Interaction of Seasonal and Business Cycles in Production," unpublished manuscript, November 1996.

Gordon, Robert J. "Measuring the Aggregate Price Level: Implications for Economics Performance and Policy," NBER working paper No. 3969, January 1992.

Haldane, Andrew G. "Introduction," *Targeting Inflation*, A. G. Haldane, ed., Bank of England, 1995, pp. 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 Series 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* (October 1994), pp. 631-53.

Roger, Scott. "Robust Measurement of Core Inflation: New Zealand, 1949-1996," Reserve Bank of New Zealand Working Paper, 1996.

Shapiro, Matthew D., and David W. Wilcox. "Bias in the Consumer Price Index," *NBER Macroeconomics Annual*, B. Bernanke and J. Rotemberg, eds., M.I.T. Press, forthcoming.

Stock, James H., and Mark W. Watson. "A Probability Model of the Coincident Economic Indicators," *Leading Economic Indicators: New Approaches and Forecasting Records*, K. Lahiri and G. H. Moore, eds., 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.