

This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: Monetary Policy

Volume Author/Editor: N. Gregory Mankiw, ed.

Volume Publisher: The University of Chicago Press

Volume ISBN: 0-226-50308-9

Volume URL: <http://www.nber.org/books/greg94-1>

Conference Date: January 21-24, 1993

Publication Date: January 1994

Chapter Title: Measuring Core Inflation

Chapter Author: Michael F. Bryan, Stephen G. Cecchetti

Chapter URL: <http://www.nber.org/chapters/c8333>

Chapter pages in book: (p. 195 - 219)

6 Measuring Core Inflation

Michael F. Bryan and Stephen G. Cecchetti

Discussions of the goals of monetary policy generally focus on the benefits of price and output stabilization. After formulating a loss function that weights these two objectives, the next step is to examine different policy programs and operating procedures in order to achieve the desired outcomes.

But these discussions take for granted our ability to measure the objects of interest, namely, aggregate price inflation and the level of output. Unfortunately, the measurement of aggregate inflation as a monetary phenomenon is difficult, as nonmonetary events, such as sector-specific shocks and measurement errors, can temporarily produce noise in the price data that substantially affects the aggregate price indices at higher frequencies. During periods of poor weather, for example, food prices may rise to reflect decreased supply, thereby producing transitory increases in the aggregate index. Because these price changes do not constitute underlying monetary inflation, the monetary authorities should avoid basing their decisions on them.

Solutions to the problem of high-frequency noise in the price data include calculating low-frequency trends over which this noise is reduced. But from a policymaker's perspective, this greatly reduces the timeliness, and therefore

Michael F. Bryan is economic advisor at the Federal Reserve Bank of Cleveland. Stephen G. Cecchetti is professor of economics at The Ohio State University and a research associate of the National Bureau of Economic Research.

The authors thank Laurence Ball, Ben Bernanke, Giuseppe Bertola, Alan Blinder, John Campbell, William Gavin, Robin Lumsdaine, N. Gregory Mankiw, James Powell, John Roberts, Alan Stockman, Alan Viard, Stephen Zeldes, the conference participants, and anonymous referees and seminar participants at Boston College, New York University, Princeton University, the University of Pennsylvania, and Wilfrid Laurier University, for comments and suggestions. In addition, they thank Edward Bryden and Christopher Pike for research assistance, Robin Ratliff for editorial assistance, and Michael Galka for the artwork. The views stated herein are those of the authors and not necessarily those of the Federal Reserve Bank of Cleveland or of the Board of Governors of the Federal Reserve System.

the relevance, of the incoming data. Another common technique for measuring the underlying or *core* component of inflation excludes certain prices in the computation of the index based 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 the remaining weights are rescaled.

As an alternative to the consumer price index (CPI) excluding food and energy, Bryan and Pike (1991) suggest computing median inflation across a number of individual prices. This approach is motivated by their observation that individual price series (components of the CPI) tend to exhibit substantial skewness, a fact also noted by Ball and Mankiw (1992), among others.¹

In this paper, we show that a version of Ball and Mankiw’s (1992) model of price-setting implies that core inflation can be measured by a limited-influence estimator, such as the median of the cross-sectional distribution of individual product price inflation first suggested by Bryan and Pike (1991). In the simplest form of the model, price setters face a one-time cross-sectional shock and can pay a menu cost to adjust their price to it immediately. Those firms that choose not to change prices in response to the shock can do so at the beginning of the next “period.” Only those price setters whose shocks were large will choose to change, and as a result, when the distribution of shocks is skewed, the mean price level will move temporarily—for example, positive skewness results in a transitory increase in inflation. This structure captures the intuition that the types of shocks that cause problems with price measurement are infrequent and that these shocks tend to be concentrated, at least initially, in certain sectors of the economy.

Removing these transitory elements from the aggregate index can be done easily. The problem is that when the distribution of sector-specific shocks is skewed, the tails of the distribution of resulting price changes will no longer average out properly. This implies that we should not use the mean of price changes to calculate the persistent component of aggregate inflation. Instead, a more accurate measure of the central tendency of the inflation distribution can be calculated by removing the tails of the cross-sectional distribution. This leads us to calculate trimmed means, which are limited-influence estimators that average only the central part of a distribution after truncating the outlying points. The median, which is the focus of much of our work below, is one estimator in this class.

The remainder of this paper is divided into four parts. Section 6.1 provides a brief discussion of the conceptual issues surrounding the measurement of core inflation. We describe a simple model and examine some evidence suggesting that shocks of the type discussed in Ball and Mankiw are likely to affect measured inflation at short horizons of one year or less. Section 6.2 reports estimates of the (weighted) median and a trimmed mean, both calcu-

1. Vining and Elwertowski (1976) discuss this fact at some length.

lated from thirty-six components of the CPI over a sample beginning in February 1967 and ending in December 1992. Section 6.3 presents evidence as to whether our measures conform to a key implication of Ball and Mankiw's view. Differences between core inflation and movements in the CPI should reflect aggregate supply shocks and, to the extent that they are accommodated, should be related to future growth in output. By contrast, core inflation itself should not forecast money growth. We find that these predictions are borne out for the median CPI.

In section 6.4, we examine some additional properties of our estimates, including their ability to forecast inflation at horizons of three to five years. While inflation is very difficult to predict, we find that the core measure based on the weighted-median forecasts future inflation better than either the CPI excluding food and energy or the all items CPI. We conclude this section with the presentation of actual predictions of future inflation. Using our preferred specification, we find that inflation is expected to average approximately $2\frac{2}{3}$ percent per year for the five years ending in December 1997.

The final section of the paper offers our conclusions. Briefly, we are encouraged by the performance of the weighted median. Because it is both easy to calculate and simple to explain, we believe that it can be a useful and timely guide for inflation policy.

6.1 Defining Core Inflation

While the term *core inflation* enjoys widespread common use, it appears to have no clear definition.² In general, when people use the term they seem to have in mind the long-run, or persistent, component of the measured price index, which is tied in some way to money growth. But a clear definition of core inflation necessarily requires a model of how prices and money are determined in the economy. Any such formal structure is difficult to formulate and easy to criticize, so we will proceed with a simple example that we believe captures much of what underlies existing discussions.³

Our goal here is to use existing data on prices to extract a measure of money-induced inflation: that is, the component of price changes that is expected to persist over medium-run horizons of several years. To see how this might be done, assume that we can think of the economy as being composed of two

2. Early attempts to define core inflation can be found in Eckstein (1981) and Blinder (1982).

3. The main conceptual problem in defining core inflation can be described as follows. Any macroeconomic model will imply some quasi-reduced form in which inflation depends on a weighted average of past money growth and past permanent and transitory "shocks." If money were truly exogenous, one could measure core inflation by estimating this reduced form and then looking only at the portion of inflation that is due to past money growth and the permanent component of the shocks. But in reality, money growth responds to the shocks themselves, so measuring the long-run trend in prices requires estimating the monetary reaction function. In fact, this suggests that measuring core inflation necessitates that we identify monetary shocks as well as the shocks to which money is responding.

kinds of price setters. The first have flexible prices in the sense that they set their prices every period in response to realized changes in the economy. The second group of price setters set their prices infrequently, and face potentially high costs of readjustment.⁴ These price setters are the familiar contracting agents of the New Keynesian theory, who set their prices both to correct for past unexpected events and in anticipation of future trends in the economy. From the point of view of measuring inflation, we might think of the first group, the realization-based price setters, as creating noise in inflation measured using existing price indices, as their price paths can exhibit large transitory fluctuations. Because they can change their prices quickly and often, these firms have little reason to care about the long-run trends in aggregate inflation or money growth.

By comparison, the expectations-based price setters have substantially smoother price paths, since they cannot correct mistakes quickly and at low cost. Our view is that the expectations-based price setters actually have information about the quantity we want to measure. If we knew who these people were, we could just go out and measure their prices. But since we do not, we must adopt a strategy in which we try to infer core inflation from the data we have.

A simple model of our view of price-setting behavior draws on Ball and Mankiw's study of the skewness of the distribution of price changes and its relationship to aggregate supply shocks. They examine price-setting as a single-period problem that can be described as follows. Each firm in the economy adjusts its price at the beginning of each period, taking into account anticipated future developments. Following this initial adjustment, each firm is then subjected to a mean zero shock and can pay a menu cost to change its price a second time. Only some firms will experience shocks that are large enough to make the second adjustment worthwhile. As a result, the observed change in the aggregate price level will depend on the shape of the distribution of idiosyncratic shocks. In particular, if the shock distribution is skewed, the aggregate price level will move up or down temporarily.

We concentrate here on a single-period problem in order to highlight the fact that we are interested in the impact of infrequent shocks. In effect, we are presuming that at the beginning of the single period under study, all price setters have completed their responses to the last disturbance of this type. This is really an assumption about the calendar time length of the model's "period." Some evidence of this is provided below.

To make the model a bit more specific, assume that the economy is composed of a large number of firms, that trend output growth is normalized to zero, and that velocity is constant.⁵ Furthermore, take money growth (\dot{m}) to be

4. Different firms will fall into these two groups for a number of reasons. We would expect, for example, that the flexible-price group will be composed of firms with some combination of low costs of price adjustment and high variance of shocks.

5. In this simple framework, we are not able to address the problems created by transitory velocity shocks.




Fig. 6.1 Distribution of relative price shocks

exogenously determined and given by a known constant (although this is not necessary). Under these conditions, each firm will initially choose to change its price by \dot{m} , and aggregate inflation will equal monetary inflation. It follows that we can define core inflation as

$$(1) \quad \pi^c = \dot{m}.$$

If we were to further assume that money growth follows a random walk, then π^c would be the best forecast of future inflation.⁶

Following this initial price-setting exercise, each firm experiences a shock, ε_i , to either its production costs or its product demand. The distribution of these shocks, $f(\varepsilon_i)$, has some arbitrary shape, such as the one drawn in figure 6.1. If each firm were to reset its prices following the realization of the ε_i 's instead of before, they would have changed them by

$$(2) \quad \pi_i = \dot{m} + \varepsilon_i.$$

But this is no longer possible without paying a menu cost. As a consequence, only firms with large $|\varepsilon_i|$ will choose to change again. With further structure on the problem, it would be possible to calculate the critical values of ε_i that lead to this action.⁷ For purposes of exposition, we assume that all firms face the same menu costs, and thus will all have the same threshold values for ε . These are labeled $\underline{\varepsilon}$ and $\bar{\varepsilon}$ in figure 6.1. It is only those firms with $\bar{\varepsilon} < \varepsilon_i < \underline{\varepsilon}$ that will change their prices. (These thresholds will differ with the cost of price adjustment, and so, in general, they will differ across firms.)

We can now examine the resulting distribution of observed price changes. First, all of the firms that chose not to act based on the realized shocks will

6. The level of core inflation will also be the level of inflation at which actual output, y , equals the natural rate, y^* . Any deviations of inflation from π^c will result in changes in real money balances and move y away from y^* . A simple interpretation of this definition is that we are attempting to measure the point at which the current level of aggregate demand intersects the long-run (vertical) aggregate supply curve.

7. See Ball and Mankiw (1992), section III, for an example.




Fig. 6.2 Distribution of nominal price changes

have changed their prices by \bar{m} . This results in a spike in the cross-sectional price-change distribution. On the other hand, the firms that did pay the menu cost and adjusted to the shock will have nominal price changes that are in the tails above and below this spike. The result is pictured in figure 6.2.

In computing aggregate observed inflation, π , we would naturally average over all of the prices in the economy. When the distribution of ε_i is symmetrical, this yields $\pi = \pi^c = \bar{m}$. But when the distribution of shocks is skewed, observed inflation is not going to equal π^c . In fact, π will be greater than or less than π^c depending on whether $f(\varepsilon_i)$ is positive or negatively skewed.⁸

Because our goal is to measure π^c from the available price data, this simple analysis leads us to an estimate that can be computed directly from the data. Instead of averaging over the entire cross-sectional distribution of price changes, consider trimming the distribution by averaging only the central part of the density. From figure 6.2 it is clear that if we average the central portion of the distribution—in the example this is the spike at \bar{m} —then we obtain an accurate estimate of π^c . As a result, we are led to compute limited-influence estimates of inflation, such as the median. These estimators are calculated by trimming the outlying portions of the cross-sectional distribution of the component parts of aggregate price indices.

The results of this simple example suggest that we examine the median, but the model is extremely specific. The implications of the analysis certainly remain valid if we assume that the shocks under consideration are infrequent and

8. The impact of the shape of $f(\varepsilon_i)$ lasts for at least two periods. To see this, note that at the beginning of the period following a shock, when all of the Ball-Mankiw price setters have the opportunity to adjust again, the relationship between measured and core inflation will depend on the distribution of shocks in the past period. When $f(\varepsilon_i)$ is positively skewed, current-period inflation will be above core inflation, while in the period following the shock, measured inflation will be below core inflation.

that the economy has fully adjusted to the last one by the time the next one arrives. But if shocks of this type arrive every period, then we need to consider a multiple-period dynamic model, one that is substantially more difficult.⁹ A completely satisfactory presentation would incorporate staggered price-setting explicitly, and the results are likely to imply more complex time-dependent and parametric measures of core inflation.¹⁰ Nevertheless, we feel that the intuition we gain from this exercise is useful, and that it guides us to explore a new estimator for inflation that is easy to calculate.

There is a way to use the available price information to estimate the frequency—that is, every month, once per year, etc.—at which these difficulties are likely to arise. To see how this can be done, rewrite equation (2) with time subscripts, and replace money growth with average aggregate inflation:

$$(3) \quad \pi_{it} = \pi_t + \varepsilon_{it}.$$

Now consider measuring average per period inflation in each sector over a horizon of K periods. Using (3), we can write this as

$$(4) \quad \pi_{it}^K = \sum_{j=1}^K \pi_{i,t+j} = \pi_t^K + \frac{1}{K} \sum_{j=1}^K \varepsilon_{i,t+j},$$

where π_t^K is average aggregate inflation per-period over the K -period horizon. Next, examine the distribution of π_{it}^K , computed cross-sectionally over the sectors. If the skewness disappears as K increases, this suggests that there is a horizon at which the problems caused by the asymmetric shocks disappear.

Using data on thirty-six components of the all urban consumers CPI (seasonally adjusted by the Bureau of Labor Statistics) from February 1967 through December 1992, and measuring inflation as the change in the natural log of the price level, we have computed the cross-sectional skewness in the price change distribution using overlapping data for K going from one to forty-eight months.¹¹ Throughout, we define inflation as the change in the log of the price index level. The results are reported in table 6.1. We have conducted a Monte Carlo experiment in order to determine if a particular level of skewness is surprising. Using the null that each sector's relative price change is drawn from a normal distribution with mean zero, and variance equal to the uncondi-

9. We have examined a simple multiple-period version of the Ball and Mankiw model, and find that as long as the shocks are temporally independent, the price-change distribution remains bunched at \bar{m} , but the bias in the mean depends on the change in the skewness, rather than its level—for example, the bias is positive when the skewness increases. While this may seem disappointing at first, there is empirical evidence that skewness changes substantially over time (see, for example, Ball and Mankiw's table II).

10. While such measures will have the advantage of being grounded in a more realistic structural model, they are likely to have the disadvantage of requiring imposition of a time-invariant stochastic structure on the data. Such methods are always vulnerable to the standard critiques.

11. The data set was chosen so that there would be a reasonably large number of component series, and at the same time we retain complete coverage of the components in the index. Skewness is calculated using the 1985 fixed expenditure weights.

Table 6.1 Frequency Distribution of Skewness (computed using overlapping observations of K months)

K	Average of Skewness	Percent Rejected at			
		1%	5%	10%	20%
1	0.346	0.09	0.15	0.19	0.39
3	0.348	0.05	0.09	0.16	0.34
6	0.230	0.03	0.09	0.14	0.30
9	0.253	0.02	0.08	0.13	0.24
12	0.239	0.02	0.06	0.10	0.28
24	0.171	0.00	0.03	0.07	0.26
36	0.029	0.00	0.01	0.03	0.14
48	-0.070	0.00	0.00	0.02	0.14

Note: Frequency distribution is computed from the percentiles of the skewness distribution based on ten thousand draws of 36 $N(0, \sigma_i^2)$ variates, weighted by the 1985 consumer price index weights. The variance, σ_i^2 , is set equal to the unconditional time-series variance of inflation in each of the components in the data computed for each value of K .

tional variance of that sector's K -period price changes over the entire sample, we compute the empirical distribution of the skewness for ten thousand draws. The results are then used to evaluate the observed skewness in the data. The calculations in table 6.1 show clearly that at frequencies of twelve months or shorter, some periods have substantial skewness in the price change distribution.¹² From this we conclude that the problems in inflation measures that we wish to eliminate exist at frequencies of one year and perhaps longer.¹³

It is important to note that the definition of core inflation as the rate of money growth presumes that there is no monetary accommodation. In order to derive these very simple results, we have assumed that m does not depend in any way on the ε 's. As such, we are proposing a measure of core inflation that forecasts the level of future inflation in the absence of monetary response to supply shocks.

We conclude this discussion with two additional remarks about the median and related estimators. First, the computation of a limited-influence estimator from the cross-sectional distribution of price changes each period has a number of potential advantages over standard methods. In particular, measures such as the median are robust to the presence of many types of noise. For example, to the extent that some price-change observations contain a combination of sampling measurement errors and actual price-setting mistakes, both of which are likely to be short-lived, this noise creates misleading movements in the

12. These results are sensitive to the specific assumptions about the heteroskedasticity of the shock distributions. For example, if we were to assume that all the relative price shocks were drawn from the same normal distribution, then we would continue to observe a substantial number of large values for the skewness until K equals seventy-two months.

13. We note, but cannot explain, the fact that as K becomes large, the observed distribution is becoming more concentrated than the empirical distribution would suggest.

aggregate index only when it is far from the central tendency of the distribution. Estimating a trimmed mean or a median will downweight the importance of this effect and result in a more robust measure of inflation. In addition, calculation of the median is a natural way to protect against problems such as the energy price increases of the 1970s—we do not need to know which sector will be subjected to the next large shock.

Finally, calculation of the median can give us additional information about price-setting behavior in the sectors covered by the indices we study. In particular, we can count how often a particular good is in the middle of the cross-sectional distribution. If the median good were selected randomly each month, this sample frequency would equal the unconditional probability of the good being the median good. Sample probabilities above or below the unconditional probability suggest that a sector is dominated by either expectations-based, inertial behavior or by realization-based, auction behavior.

6.2 Estimates of Core Inflation

Using the data on the CPI described above, we now examine various measures of inflation. We compute two trimmed means using the fixed 1985 CPI weights as measures of the number of prices in each category. In other words, in computing the histogram for inflation in each month, we assume that the weight represents the percentage of the distribution of all prices that experienced that amount of inflation. We report results for the weighted median and for a 15 percent trimmed mean. The median is measured as the central point, as implied by the CPI expenditure weights, in the cross-sectional histogram of inflation each month. The 15 percent trimmed mean is computed by averaging the central 85 percent of the price-change distribution each month. Obviously, we could report results for an index computed by trimming any arbitrary percentage of the tails of the distribution. We have chosen 15 percent because it has the smallest monthly variance of all trimmed estimators of this type.¹⁴

Table 6.2 reports the summary statistics for the all items CPI, the CPI excluding food and energy, the weighted median, and the 15 percent trimmed mean. As one would expect, the variance of the core measures is substantially lower than that of the CPI measures. In fact, the standard deviation of the median and the 15 percent trimmed mean are both on the order of 25 percent less than that of the total CPI, and 10 percent less than that of the CPI excluding food and energy. Furthermore, all of the series show substantial persistence, although the standard Dickey-Fuller test fails to reject stationarity in all of the series (the 10 percent critical value is -3.12).¹⁵

14. The 15 percent trimmed mean also has the highest first-order autocorrelation of all the trimmed estimators.

15. The results of testing for a unit root in inflation are extremely sensitive to the sample period chosen. Using data from 1960 to 1989, for example, Ball and Cecchetti (1990) model inflation as nonstationary.

Table 6.2 Comparison of Various Measures of Inflation, 1967:2 to 1992:12
(computed from monthly data measured at annual rates)

	All Items Consumer Price Index	Consumer Price Index Excluding Food and Energy	Weighted Median	15% Trimmed Mean
Mean	5.67	5.71	5.64	5.56
Standard deviation	3.79	3.25	2.95	2.86
1st-order autocorrelation	0.64	0.60	0.68	0.76
Dickey-Fuller (24) ^a	-2.85	-2.59	-2.44	-2.46
Correlation Matrix				
All items Consumer price index	1.00	0.73	0.75	0.84
Consumer price index excluding food and energy	0.73	1.00	0.80	0.87
Weighted median	0.75	0.80	1.00	0.93
15% trimmed mean	0.84	0.87	0.93	1.00

^aDickey-Fuller tests are based on examining the coefficient on the lagged price level in the regression given by $\Delta p_t = a + bp_{t-1} + \sum_{i=1}^k c_i \Delta p_{t-i} + v_t$, where p is the log of the price level, Δp is the first difference in the log of the price (inflation), v is a random error, and the remaining terms are parameters. The null hypothesis is that the log of the price level (p_t) has a unit root, namely, that $b = 0$. (See Dickey and Fuller 1981.) The reported results, for $k = 24$, are unaffected by setting $k = 12$.

Figure 6.3 presents a plot of the twelve-month lagged moving average of each of the series—an observation plotted at date t is the sum of monthly inflation from $t - 11$ months through t . The graph reveals a number of interesting patterns in the core measures, in addition to demonstrating how they are less variable. First, both the median and the 15 percent trimmed mean show lower peaks. Furthermore, the core measures display substantially lower inflation than either the all items CPI or the CPI excluding food and energy during the high-inflation period of 1979 to 1981. Finally, the results clearly demonstrate that the low inflation of 1986 was largely the consequence of transitory shocks to relative prices.

As we mentioned at the end of section 6.1, the sample frequency a good is at the median provides us with interesting information about the nature of price-setting in various sectors. Table 6.3 reports the unconditional probability of each good being the median together with the sample probability that the good is the median. The unconditional probability of a good being the median cannot be computed in a simple analytic way. Instead, we calculate the quantity of interest using a Monte Carlo experiment in which we draw 1.5 million random sample orderings and tabulate the frequency each good is at the median.




Fig. 6.3 Comparison of inflation measures (twelve-month moving averages, 1968 to 1992)

The results show several intriguing properties.¹⁶ The most striking is that the shelter component of the index, with an unconditional probability of 37.01 percent (the CPI weight is 27.89), is the median good 47.04 percent of the time. Food away from home (unconditional probability = 5.42, sample frequency = 9.65) and medical care services plus commodities (unconditional probability = 5.89, sample frequency = 9.65) are also in the center of the distribution more often than random chance suggests they should be. All of these

16. The results are the same for both the 1967 to 1979 sample period and the 1980 to 1992 sample period.

Table 6.3 Unconditional Probability and Sample Frequency for Median Good (%)

Description	1985 Consumer Price Index Weight	Unconditional Probability Good Is at Median	Frequency Good Is at Median
Cereals and bakery products	1.43	1.22	0.96
Meats, poultry, fish, and eggs	3.03	2.63	0.96
Dairy products	1.23	1.05	0.32
Fruits and vegetables	1.85	1.61	0.32
Other food at home	<u>2.38</u>	<u>2.06</u>	<u>1.93</u>
Total food at home	9.92	8.57	4.49
Food away from home	6.08	5.42	9.65
Fuel oil and other household fuels	0.42	0.36	0.00
Gas and electricity (energy services)	3.64	3.20	1.93
Motor fuel	<u>3.30</u>	<u>2.88</u>	<u>0.32</u>
Total energy	7.36	6.44	2.25
Shelter	27.89	37.01	47.91
Medical care commodities	1.26	1.09	3.22
Medical care services	<u>5.43</u>	<u>4.80</u>	<u>6.43</u>
Total medical care	6.69	5.89	9.65
Men's and boys' apparel	1.45	1.27	0.96
Women's and girls' apparel	2.52	2.21	2.25
Infant and toddler apparel	0.22	0.19	0.00
Other apparel commodities	0.55	0.47	0.32
Alcoholic beverages	1.62	1.40	0.64
House furnishings	3.70	3.26	1.61
Housekeeping supplies	1.15	0.99	1.29
Footwear	0.80	0.69	0.64
New vehicles	5.03	4.44	4.18
Used cars	1.13	0.98	1.29
Other private transportation commodities	0.68	0.59	0.00
Entertainment commodities	2.03	1.76	0.64
Tobacco and smoking products	1.66	1.43	0.96
Toilet goods and personal care appliances	0.63	0.55	0.32
Schoolbooks and supplies	<u>0.24</u>	<u>0.21</u>	<u>0.00</u>
Total other commodities	23.41	20.44	15.10
Apparel services	0.56	0.48	0.32
Housekeeping services	1.47	1.26	1.29
Auto maintenance and repair	1.52	1.33	1.29
Other private transportation services	3.85	3.41	1.93
Public transportation	1.49	1.27	0.00

Table 6.3 (continued)

Description	1985 Consumer Price Index Weight	Unconditional Probability Good Is at Median	Frequency Good Is at Median
Entertainment services	2.33	2.02	1.61
Personal care services	0.55	0.49	0.64
Personal and education services	3.58	3.13	2.89
Other utilities and public services	<u>3.27</u>	<u>2.85</u>	<u>0.96</u>
Total other services	18.62	16.24	10.93

Note: Calculations use thirty-six components of the consumer price index (CPI), monthly from February 1967 to December 1992. Unconditional Probability Good Is at Median is calculated from a Monte Carlo experiment with 1.5 million draws using the 1985 CPI weights. Frequency Good Is at Median simply counts the number of months a particular good is the median good, and divides by the total number of months. Sums of unconditional probabilities are estimates assuming independence.

are markets in which long-term contracts or customer relationships are important.¹⁷

The results in table 6.3 also shed some light on the practice of excluding food and energy arbitrarily. We find that both food at home and energy are in the center of the distribution much less frequently than their weights would suggest. If we assume independence and simply sum the probabilities and the sample frequencies, then food at home plus energy has an unconditional probability of 15.01, but the goods in these groups are at the median a total of only 6.74 percent of the time. From this we conclude that if we were to construct an index that removed food and energy components, it would retain “food away from home” and be an “excluding food at home and energy” index.

Our next task is to demonstrate the usefulness of these proposed measures of core inflation and to show that they are superior estimates of money-induced inflation. In the following sections, we examine a series of characteristics of the core measures. First, we study the relationship between money and inflation directly. Then we consider the ability of the alternative price measures to

17. We have also computed the percentage of the months in which each commodity lies in the central half of the cross-sectional distribution, and they are consistent with those reported in the table. If goods were ordered randomly, then a component with weight w_i should appear in the middle of the distribution approximately $50 + 2w_i$ percent of the time. Food away from home, with a weight of 6.08, is in the middle 50 percent of the distribution in 82.5 percent of the months of the sample, far more than the 62 percent that would result from random chance. By contrast, the inflation in energy prices and the prices of food at home appear in the center of the distribution far less than half of the time—motor fuel, for example, has a weight of 3.30 and is present in only 14.2 percent of the months.

forecast CPI inflation over long horizons under the assumption that, since supply disturbances affect measured CPI inflation only in the short run, current core inflation should provide useful information about future aggregate price increases.

6.3 Core Inflation and Money Growth

A primary motivation for our study of core inflation is to find a measure that is highly correlated with money growth. To test our success in this endeavor, we first consider the ability of money growth to forecast each of the alternative inflation measures in simple regressions.

A straightforward way of evaluating the relationship between money growth and various measures of inflation employs the following simple regression:

$$(5) \quad \frac{1}{K}(\ln p'_{t+K}) - \ln p'_t = \alpha + \sum_{i=1}^m \beta_i [\ln M_t - \ln M_{t-i}] + \varepsilon_t,$$

where M is a measure of money.

We look at the ability of the monetary base, M1, and M2 to forecast the average level of inflation over the next one to five years. The results for $m = 24$ months, which are representative, are presented in table 6.4, where we report the R^2 's of the regressions (5). The table shows that the past year's money growth is most highly correlated with changes in the weighted median, with the 15 percent trimmed mean a close second.

Next, we conduct a series of Granger-style tests to establish where changes in money growth actually forecast changes in inflation, once we take account of the ability of past inflation to forecast itself. Curiously, previous research has found that tests of this type show that the forecasting relationship (and the direction of causality) between the CPI and money operates in the opposite direction—from inflation to money growth. A recent study by Hoover (1991), for example, provides substantial evidence for this counterintuitive result. We might interpret Hoover's conclusions as suggesting that movements in standard aggregate price indices are dominated by supply disturbances that influence both prices and money. Purging the price statistics of these distortions should reveal the money-to-inflation relationship that is otherwise obscured.

In order to test whether a candidate variable y forecasts x in the Granger sense, we examine the coefficients on x in the regression

$$(6) \quad y_t = \alpha + \sum_{i=1}^m b_i y_{t-i} + \sum_{i=1}^m c_i x_{t-i} + u_t$$

We report results for testing whether all of the c_i 's are zero simultaneously. This can be interpreted as a test for whether x forecasts y , once lagged y is taken into account.

Table 6.4 Forecasting Long-Horizon Inflation with Money Growth (1967:02 to 1992:12)

Horizon (K)	All Items Consumer Price Index	Consumer Price Index Excluding Food and Energy	Weighted Median	15% Trimmed Mean
<i>M</i> = Monetary Base				
12	0.16	0.19	0.19	0.18
24	0.16	0.22	0.22	0.21
36	0.13	0.19	0.17	0.17
48	0.07	0.14	0.10	0.10
60	0.03	0.11	0.06	0.06
<i>M</i> = M1				
12	0.03	0.03	0.04	0.03
24	0.01	0.01	0.01	0.01
36	0.02	0.01	0.03	0.02
48	0.10	0.07	0.11	0.10
60	0.17	0.12	0.18	0.16
<i>M</i> = M2				
12	0.19	0.12	0.22	0.18
24	0.24	0.23	0.32	0.28
36	0.24	0.27	0.33	0.30
48	0.18	0.25	0.26	0.24
60	0.10	0.16	0.15	0.14

Note: The table reports the R^2 from a regression of twenty-four lags of money growth on inflation over the next K months. See equation (5).

Results for $m = 12$ are presented in table 6.5.¹⁸ These clearly suggest that both M1 and M2 growth forecast core inflation as measured by the weighted median and the 15 percent trimmed mean. But, as we expect, the deviations of actual inflation from the 15 percent trimmed mean and the weighted median both forecast M1, while the weighted median forecasts M2 growth. Unfortunately, the results for the monetary base are less compelling.

While these tests have a number of well-known problems that prevent us from interpreting them as evidence of true causality, we find the results tend to confirm our measures and interpretation. Specifically, the reason that others have found that inflation forecasts money growth appears to be a sign of the monetary accommodation of the aggregate supply shocks that we measure as the deviations of the all items CPI from the core measure. Furthermore, the

18. The results are unaffected by increasing m to 24.

Table 6.5 Tests of Granger-Style Forecasting Ability: Money and Inflation (1967:03 to 1992:06)

Inflation Measure	Monetary Base		M1		M2	
	M to π	π to M	M to π	π to M	M to π	π to M
All items consumer price index	0.13	0.57	0.72	0.03	0.34	0.00
Consumer price index excluding food and energy	0.05	0.53	0.01	0.03	0.03	0.04
Weighted median	0.11	0.34	0.01	0.18	0.00	0.30
15% trimmed mean	0.35	0.36	0.01	0.05	0.00	0.11
Consumer price index-2	0.15	0.62	0.18	0.24	0.36	0.06
Consumer price index-3	0.19	0.91	0.39	0.00	0.25	0.05
Consumer price index-4	0.28	0.96	0.37	0.04	0.52	0.15

Note: Values are p -values for Granger F -tests.

fact that core inflation is forecast by money growth, but does not itself forecast money growth, suggests a measure of inflation that is in some sense tied to monetary policy.

6.4 Forecasting CPI Inflation

It is typically difficult to forecast medium- and long-term inflation in either a univariate or multivariate setting. Nevertheless, we set out to examine the ability of these different price measures to forecast actual inflation over horizons of one to five years. In this section we proceed in two related directions. First, we study univariate forecasts of CPI inflation over horizons of one to five years. The univariate forecasts reported in section 6.4.1 show that recent core inflation does a slightly better job than inflation in either the all items CPI or the CPI excluding food and energy. Section 6.4.2 examines the marginal forecasting power of core inflation when it is added to a multivariate equation including money, output, and interest rates with essentially the same result.

6.4.1 Univariate Methods

The results in section 6.1, table 6.1 suggest that short-run problems in measurement of aggregate inflation are likely to disappear over horizons of one year or more. This suggests that the all items CPI provides an accurate measure of inflation over longer horizons and thus is useful as a benchmark for forecasting inflation. Using equation (4), we identify the object of interest as

$$(7) \quad \Pi_r^K = \frac{1}{K} [\ln(\text{CPI}_{t+k}) - \ln(\text{CPI}_t)],$$

where K might indicate one, two, or three years.

If we were simply interested in constructing the best estimate of Π_r^K pos-

sible, then we could continue in a number of directions, such as constructing a multivariate vector autoregression. But since our main interest is in the informativeness of the measures of core inflation, we proceed slightly differently.¹⁹ Restricting ourselves to price data alone, we examine our alternative measures of inflation and see which of them forecasts Π_t^K best. To do this, consider the following simple regression of the average CPI inflation at horizon K on inflation in a candidate index over the previous year:

$$(8) \quad \Pi_t^K = \alpha + \beta(\ln p_t' - \ln p_{t-12}') + \varepsilon_t^K,$$

where p' is the one of the four indices: all items CPI, CPI excluding food and energy, the weighted median, and the 15 percent trimmed mean. We provide two sets of comparisons. In the first, we estimate (8) from monthly data through December 1979 and then use the fitted regression to forecast from January 1980 through the end of the sample (which will vary depending on the choice of the horizon K).²⁰ The second exercise examines the forecast error when the forecast is simply cumulative inflation over the prior twelve months. We report results for this naive rule over the entire available sample.

Table 6.6 reports the root mean square errors for each of these forecasting exercises, along with summary statistics for Π_t^K .²¹ The results suggest two conclusions. First, we confirm the general impression that it is difficult to forecast inflation. For horizons of two years or longer over the sample beginning in 1980, the root mean square errors of the forecasts are more than half the mean of the series being forecast.

Second, the core measures provide the best forecasts at long horizons. Among the alternatives, the weighted median yields the best forecast of long-horizon CPI inflation. One view of core inflation, then, is that it is a forecast of future inflation over the next three to five years.²²

6.4.2 Multivariate Methods

An alternative to the univariate forecasting equation (8) is to examine the marginal forecasting power of prices in an equation that includes a set of variables Z :

$$(9) \quad \Pi_t^K = \alpha + \sum_{i=0}^{11} \beta_i (\ln p_t' - \ln p_{t-i}') + \gamma Z_t + \varepsilon_t^K.$$

19. Yet another alternative would be to define core inflation as the optimal forecast of Π_t^K . This has the disadvantage that it is difficult to calculate in real time. In addition, such a definition would force revision of the entire history of estimates with the arrival of each new month's data.

20. The estimates of β in (8) are very close to one for most of the cases, implying that the current twelve-month moving average of the index is the best forecast of long-horizon inflation.

21. We have restricted the constant in equation (8) to zero, as this reduces the root mean square forecast errors. This is consistent with the general notion that inflation is highly persistent.

22. All of our results are robust to either adding lags of the right-hand-side variable to the forecasting regression (8), or including many lags of single-period inflation rather than twelve-month averages.

Table 6.6 Comparison of Forecasts of Long-Horizon Consumer Price Index Inflation (1967:02 to 1992:12)

Candidate Index	Root Mean Square Error Horizon K in Months				
	12	24	36	48	60
	Forecasts Beginning 1980:01				
All items consumer price index	2.25	2.72	3.07	3.40	3.82
Consumer price index excluding food and energy	2.58	3.02	3.41	3.79	4.25
Weighted median	2.08	2.48	2.80	3.10	3.49
15% trimmed mean	2.21	2.62	2.99	3.32	3.74
	Summary Statistics for Π_t^k during Forecasting Period				
Mean	4.51	4.32	4.18	4.10	4.01
Standard deviation	2.14	1.52	1.06	0.89	0.71
	Full Sample				
All items consumer price index	2.17	2.71	2.95	3.05	3.08
Consumer price index excluding food and energy	2.64	2.94	3.02	3.00	3.00
Weighted median	2.30	2.58	2.67	2.67	2.66
15% trimmed mean	2.32	2.64	2.76	2.76	2.75
	Summary Statistics for Π_t^k during Forecasting Period				
Mean	5.83	5.90	5.96	6.02	6.10
Standard Deviation	2.86	2.63	2.40	2.21	2.05

Note: The top panel of the table reports the root mean square error of forecasts of inflation beginning in 1980:01 constructed from an equation estimated over the period from 1967:02 through 1979:12. The bottom panel reports the root mean square error of forecasts of inflation over the entire available sample, based on the previous twelve months.

We examine the case in which the Z 's are twelve monthly lags of money growth, the growth in industrial production, the nominal interest rate on a constant K -month maturity U.S. government bond, and inflation in the CPI itself. To test the proposition of interest, we compare the P -values from F -tests that all the β 's are zero simultaneously when the equation is estimated over the entire available sample period.

The results are reported in table 6.7. As the table clearly shows, the weighted median is consistently informative about future changes in the CPI, *over and above* the information contained in the past changes in the CPI itself. The result is robust to both the horizon and the choice of how money is measured. Interestingly, the CPI excluding food and energy appears to contain little additional information useful in predicting future inflation.

As a final exercise, we use the estimated multivariate forecasting equation (9) to compute actual forecasts of inflation from 1993 to 1997. Table 6.8 reports the fitted values for regressions over various horizons, with different measures of money and core inflation, using actual data through December 1992. We also present estimates of the standard errors of these forecasts.

The estimated forecasts vary substantially depending on the definition of money and the measures of inflation included in the simple linear regression. But the weight of the evidence thus far suggests that we should focus on results for M2 and the weighted median. Using this preferred combination, we find that inflation is forecast to average 3.76 percent for 1993, 3.02 percent for the three years ending December 1995, and 2.68 percent over the five years ending December 1997. The standard errors of all of these estimates are a bit over 1 percent, so a 95 percent confidence interval for the five-year horizon would be (0.3,5.1). Thus, in the absence of accommodation of any future shocks, current monetary policy will result in inflation that is roughly comparable to that of the past decade (1983 to 1992), when price increases averaged approximately 4 percent per year.

Table 6.7 Multivariate Forecasts of Inflation: The Marginal Contribution of Past Inflation (1967:02 to 1992:12)

	<i>K</i>		
	12	36	60
	Monetary Base		
Consumer price index excluding food and energy	0.61	0.03	0.48
Weighted median	0.07	0.00	0.09
15% trimmed mean	0.37	0.33	0.58
	M1		
Consumer price index excluding food and energy	0.01	0.16	0.56
Weighted median	0.63	0.00	0.00
15% trimmed mean	0.14	0.71	0.05
	M2		
Consumer price index excluding food and energy	0.50	0.93	1.00
Weighted median	0.02	0.00	0.00
15% trimmed mean	0.22	0.40	0.50

Note: The table reports the *p*-values for the *F*-tests associated with adding twelve lags of the candidate index to a regression of average inflation *K* months into the future on twelve monthly lags of the nominal interest and the growth rates of either the monetary base, M1 or M2, industrial production, and the CPI.

Table 6.8 Forecasts of Inflation: 1993 to 1997 (average annual rates, standard errors in parentheses)

	Annual Average from December 1992 to		
	December 1993	December 1995	December 1997
	Monetary Base		
Consumer price index excluding food and energy	4.62 (1.31)	6.22 (1.32)	5.95 (1.32)
Weighted median	4.30 (1.25)	5.30 (1.22)	5.47 (1.27)
15% trimmed mean	4.69 (1.29)	6.01 (1.28)	5.85 (1.33)
	M1		
Consumer price index excluding food and energy	4.98 (1.14)	5.66 (1.25)	4.56 (1.19)
Weighted median	4.49 (1.20)	4.60 (1.31)	3.79 (1.22)
15% trimmed mean	4.91 (1.17)	5.38 (1.27)	4.29 (1.19)
	M2		
Consumer price index excluding food and energy	3.80 (1.25)	3.59 (1.25)	3.36 (1.18)
Weighted median	3.76 (1.22)	3.02 (1.25)	2.68 (1.19)
15% trimmed mean	3.80 (1.24)	3.36 (1.26)	3.11 (1.17)

Note: The table reports the forecasts using a regression of average inflation K months into the future on twelve monthly lags of the nominal interest and the growth rates of either the monetary base, M1, or M2, industrial production, the CPI, and the candidate measure of inflation. Included are the fitted value for the forecast using data through December 1992 and standard errors that incorporate parameter uncertainty with the covariance matrix of the coefficient estimates computed using the Newey and West (1987) procedure with $K + 1$ lags.

6.5 Conclusion

This paper examines the use of limited-influence estimators as measures of core inflation. Specifically, we study the CPI excluding food and energy, and several estimates based on trimming the outlying observations of the cross-sectional distribution of inflation in each month, including the weighted median. Our use of these estimators is motivated by the observation that non-monetary economic shocks can, at least temporarily, produce noise in reported inflation statistics. As an example, we show how, when the distribution of sector-specific supply shocks is asymmetric, costly price adjustment can result in transitory movements of average inflation away from its long-run trend.

We are encouraged by the finding that the limited-influence estimators are superior to the CPI in several respects. They have a higher correlation with past

money growth and provide improved forecasts of future inflation. Furthermore, unlike the all items CPI, the limited-influence estimates appear to be unrelated to future money growth.

Within the class of inflation measures we consider, the weighted median CPI fares best in virtually all of the statistical criteria we examine. Such a finding is not particularly surprising, given the nature of the problem we have outlined. A disproportionate share of the noise in the price data comes from the extreme tails of the distribution of price changes, and so systematically eliminating the tails of the distribution should give us a more robust measure of the persistent component of inflation.

What is missing from our analysis is a fully satisfactory model of the money growth–inflation relationship. This prevents us from addressing a number of interesting propositions, such as the degree to which monetary policy reacts to temporary aggregate supply and aggregate demand shocks. Also absent from consideration is the related issue of long-run bias in inflation measurement that results from permanent changes in the expenditure weights. From the perspective of a policymaker interested in short-run indicators of monetary inflation, we suspect that such biases are of secondary importance. Nevertheless, we believe that the long-run properties of limited-influence estimators of inflation remain an important area for future research.

References

- Ball, Laurence, and Stephen G. Cecchetti. 1990. Inflation and uncertainty at short and long horizons. *Brookings Papers on Economic Activity* 21 (1): 215–45.
- Ball, Laurence, and N. Gregory Mankiw. 1992. Relative-price changes as aggregate supply shocks. NBER Working Paper no. 4168. Cambridge, Mass.: National Bureau of Economic Research, September.
- Blinder, Alan S. 1982. The anatomy of double-digit inflation. In *Inflation: Causes and effects*, ed. Robert E. Hall, 261–82. Chicago: University of Chicago Press.
- Bryan, Michael F., and Christopher J. Pike. 1991. Median price changes: An alternative approach to measuring current monetary inflation. *Economic Commentary*, Federal Reserve Bank of Cleveland, 1 December.
- Dickey, David A., and Wayne A. Fuller. 1981. Likelihood ratio statistics for autoregressive time series with a unit root. *Econometrica* 49 (July): 1057–72.
- Eckstein, Otto. 1981. *Core inflation*. Englewood Cliffs, N.J.: Prentice-Hall.
- Hoover, Kevin. 1991. The causal direction between money and prices. *Journal of Monetary Economics* 27 (June): 381–423.
- Newey, Whitney K., and Kenneth D. West. 1987. A simple, positive definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55 (May): 703–8.
- Vining, Daniel R., and Thomas C. Elwertowski. 1976. The relationship between relative prices and the general price level. *American Economic Review* 66 (September): 699–708.

Comment Stephen P. Zeldes

Bryan and Cecchetti provide an alternative method of summarizing the vast amount of information in the many components of consumer prices. Rather than looking at reported aggregate consumer price index (CPI) inflation, which is approximately a weighted average of the inflation rates of all of the individual components, they calculate the median inflation and the 15% trimmed mean inflation. Bryan and Cecchetti (BC) refer to these measures as estimates of "core inflation." These measures turn out to have intriguing properties: they have a lower variance and higher persistence than the CPI inflation, they can be used to improve on forecasts of future inflation, and they are related to past money growth yet do not forecast future money growth.

BC document these intriguing properties, but they are not very clear as to what these measures are supposed to represent; that is, they never really define core inflation. My main comment at the conference has been partially addressed here but remains partly relevant. Reading the paper is like watching the television game show "Jeopardy." The category is inflation. BC have provided us with the answer: the trimmed mean. Now, for two hundred dollars, what is the question to which this is the answer? Unfortunately, the buzzer rang before either they or I could completely figure it out.

In the previous literature, there seemed to be two different notions of core inflation. First, Otto Eckstein defines core inflation as "the trend increase in the cost of the factors of production. It originates in the long-term expectations of inflation in the minds of households and businesses, in the contractual arrangements which sustain the wage-price momentum, and in the tax system." He essentially considers a Phillips curve augmented by inflation expectations and supply shocks. The core inflation rate corresponds to the inflation expectations term: what inflation would be if there were no supply shocks, and unemployment were at the natural rate. He was trying to get at the persistence of inflation in order to say something about future inflation.

The other popular usage of the term *core inflation* is in the popular press: the rate of change of CPI excluding food and energy. The intuition is that the persistence of the food and energy component is less than the rest, so this measure can provide a better indicator of future inflation than looking at total inflation.

BC consider an alternative measure. They compute inflation rates for each of thirty-six components of the CPI, and calculate two types of trimmed means, in each case eliminating the outliers on each side of the distribution. These measures are the 15 percent trimmed mean and the median (or 50 percent trimmed mean).

The first question one probably should ask is of what use is a summary

statistic such as core inflation? It is unlikely that any one statistic is going to be a sufficient statistic, that is, that it will completely summarize all of the relevant information in relative price movements. But if one statistic can capture, even imperfectly, the information in a large volume of data, it can clearly be of use to economists, policymakers, and the general public.

One of the goals of the paper is to generate an inflation measure that is useful for formulating monetary policy. Exactly what information would monetary policymakers like to have, and how does this relate to core inflation? Presumably, they would like the answer to the following very broad question: For each possible path of monetary policy, what is the expected path of real output and inflation? For example, holding the growth of the monetary base constant, what is the expected rate of inflation and real output growth in each of the next ten years? Or consider another example. The inflation rate has been constant at an annualized rate of 4 percent for a number of quarters. In the current quarter inflation turns out to be 6 percent. Does this merit a change in monetary policy? Presumably, the answer depends on the persistence of the inflation innovation in the absence of any change in monetary policy, that is, on some conditional forecast of the path of future inflation.

All of the proposed measures of core inflation relate in some way to expectations about future inflation, although none is strictly formulated as a rational forecast conditional on an information set. One way of formulating a measure of core inflation might be: what is the expected path of future inflation given a future path of monetary policy and the past behavior of all individual price indices? This is not exactly what BC do, and doing so would require taking a stand on the underlying economic model.

BC argue that their measures should yield a better estimate of the underlying inflation in the economy, that is, that their measures provide more information about future inflation than simply looking at the current inflation rate. The theoretical argument requires a few steps, involving the link between outliers and persistence. These steps are as follows: (a) skewness in the distribution of the shocks to relative prices causes a change in aggregate inflation, and (b) these changes, and only these changes, are transitory.

Consider each step in turn. The first link clearly needn't be true. In a competitive model with no sticky prices and where relative prices adjust fully to changes in demand and supply, the aggregate price level is tied down by the demand and supply of money, and money is neutral. In this model it makes no sense to look at the individual components: if, given money growth, food prices have a large increase because of a drought, then the price of other commodities will grow sufficiently less slowly to keep the overall inflation rate at the same rate it would have been in the absence of a drought. In this model, relative price changes are independent of aggregate price changes, so it would be misleading to eliminate or downweight the relative price changes that are outliers.

BC rely instead on the Ball-Romer model with sticky prices, in which

skewed changes in relative prices can change aggregate inflation. BC show that in a simple version of this type of model, the trimmed mean provides an estimate of what inflation would be in the absence of these supply shocks. An important extension of this paper would be to derive the properties of the trimmed mean in a fully specified dynamic model of this type. It is unclear whether the results presented here would continue to hold in such a model.

Even if the trimmed mean eliminates supply shocks, there is still the second link: are the aggregate shocks to inflation resulting from skewed relative price shocks transitory while other sources are not? BC argue that in a version of the Ball and Romer model with the simplest dynamics included, the supply shocks would be transitory. However, in a different model, one could imagine permanent shocks to inflation causing skewness in the underlying distribution of relative price changes. In this case it would be misleading to eliminate these shocks from our measure of inflation. It is also clearly possible to have transitory money supply and money demand shocks, as well as other types of transitory aggregate demand shocks.

The overall message here is that the appropriate measure of core inflation must be model based. Additional work needs to be done deriving the properties of the trimmed mean in more general models of inflation.

Next, turn to the empirical properties of the trimmed means. First, the univariate properties. Relative to the standard CPI inflation, BC show that their measures have lower variance and higher persistence (first order serial correlation). They are also slightly lower during the 1979 to 1981 period, and do not fall during the transitory 1986–87 slowdown in CPI inflation. Next, BC find that for M1 and M2 growth, money growth Granger-causes trimmed mean inflation, but trimmed mean inflation does not Granger-cause money growth. Although I'm not quite sure why this occurs, the results suggest that the trimmed means might be able to avoid some of endogeneity problems of previous work. Finally, BC regress long horizon (one to five year) inflation against different inflation measures over the previous twelve months. They find that using the median or trimmed mean provides a moderately better forecast of future inflation than does standard CPI inflation or the CPI excluding food and energy. Also, in some cases the median or trimmed mean helps predict future inflation even after including past money growth and other past variables. These results are very intriguing.

Although I am still not sure what is meant by the term core inflation, I think BC have taken an important first step in documenting that the median and trimmed mean are measures that can provide some important summary information about future inflation. It is impressive that we learn about the time-series properties of aggregate inflation using only the cross-sectional distribution in relative price changes. We can therefore add this to the list of possible sources of information about inflation expectations that were mentioned during

the conference: forecasts of professional forecasters, commodity prices, bond prices (comparing nominal bonds to hopefully-soon-to-be-issued indexed bonds), and univariate aggregate time-series models. The results of this paper should interest policymakers and will likely promote future theoretical and empirical research on this topic.