Monitoring Inflation in a Low-Inflation Environment

Michael F. Bryan*

Research Department

Federal Reserve Bank of Cleveland

May 15, 2007

ABSTRACT:

Central bank economists (and others) are expected to spot changes in inflation with expedience, a responsibility that is made all the more critical given the (presumed) long and variable lags that exist between monetary policy and inflation, and the relatively large confidence intervals around the typical inflation forecast.

This working paper discusses the use of trimmed-means as high-frequency measures of the inflation trend. Given the noise that characterizes retail price data, I argue that these estimators help policymakers identify small shifts in the inflation trend earlier than the headline or traditional core inflation measures—a substantial advantage when conducting policy in a low-inflation environment. However, while this methodology significantly improves the signal-to-noise ratio in high-frequency data, the usefulness of these estimators declines with the frequency of the data, becoming almost negligible when the data extend to horizons of a year or beyond.

* *Vice President and Economist, Federal Reserve Bank of Cleveland.* *Linsey Molloy provided the research assistance for this paper. Pat Higgins was responsible for programming the bootstrap experiments. John Carlson provided the programs and guidance on the use of the change-point statistical tests. Any mistakes, of course, are solely the responsibility of the author.*

I. INTRODUCTION

Over the past twenty years, central banks around the world have successfully reduced inflation to relatively low levels. Today, the achievement of price stability is regarded as one the most important, if not *the* primary goal of their policies. In an environment where the central bank has fixed expectations on low inflation, small changes in the inflation trend can potentially have large and adverse effects on economic behavior, especially if the central bank loses some of its credibility. Indeed, the ability to quickly identify and react to undesired deviations in the inflation path can be crucial to the successful implementation of the central bank's policies.

Unfortunately, central banks set their policies with imperfect knowledge. The interest rate or monetary levers they maneuver are only loosely correlated with their inflation objective, and then only after a potentially long, unknown lag. This causes central banks to rely heavily on inflation forecasts, for which standard errors are understandably large. Since 1983, for example, the RMSE of the median economist's year-ahead CPI forecast was about 1 percentage point, producing a confidence interval around the inflation forecast that is uncomfortably large.[1]

In light of these uncertainties, central banks carefully monitor incoming price data and continually evaluate whether the inflation trend appears to be developing in an undesirable way, and if it is, react before the new trend becomes embedded into expectations. The more expeditiously they can spot a change in the inflation trend, the sooner they can react.

---

[1] Bryan and Molloy (2007) document the recent track record of economists' growth and inflation forecasts.

In this paper, I argue that trimmed means—averages computed after systematically eliminating a certain percentage of the extreme high and low price changes—are an effective way to filter noise that obscures the inflation trend in the price data. Filtering this noise substantially reduces the time it takes for the central bank to identify small changes in the inflation trend relative to either time-series averaging or the CPI excluding food and energy (core CPI).

In section II, I outline the types of noise that make changes in the inflation trend difficult to identify and I describe commonly used procedures to reduce that noise, including trimmed means. Section III reports on experiments that test the effectiveness of these estimators in helping to reveal changes in the inflation trend using bootstrap techniques and historical, ex-post analysis. Section IV concludes.

II. TRIMMED-MEANS: A SIMPLE NOISE REDUCTION PROCEDURE

Data "noise" is arguably the single most important problem confounding the ability of policymakers and others to gauge the variables they wish to act upon. For example, seasonal variations obscure the more persistent patterns in the vast majority of spending, production, and price data that are of primary interest for most business cycle analysis. Since 1998, seasonal fluctuations account for about 40 percent of the monthly variation in the CPI and a little more than 80 percent of the monthly variation in the core CPI.

One way to reduce seasonal noise in the monthly price-change data is to compute changes on a year-over-year basis. While this is effective in reducing seasonal patterns in the data, it comes at the cost of excluding all of the information in the data at frequencies of less than a year. Those in need of higher-frequency information—monthly or quarterly data—

generally apply any one of a number of seasonal adjustment techniques, which reduce the seasonal influence on the data while (hopefully) preserving the nonseasonal information.[2]

Seasonality in the price data has a significant idiosyncratic character. For example, fruits and vegetable prices tend to rise sharply every December and January, and fall nearly as much in July and August. Lodging away from home tends to show the opposite pattern, rising disproportionately every summer, then declining sharply between November and February. Other goods experience other seasonal patterns—apparel prices are higher in the spring and fall, lower in the winter and summer, while prices for education tend to get adjusted every September, and recreation costs get adjusted every April. For this reason, price data are usually adjusted for seasonal influences at the detailed, component level.

Seasonality influences the cross-sectional distribution of retail price changes (table 1). The individual seasonal patterns across the CPI components result in a cross-sectional distribution of price changes that has exceptionally long tails. The average cross-section variance in the unadjusted monthly data was 34.2 percent since 1998, 64% greater than after the components are seasonally adjusted (average cross-section variance of 20.9 percent).[3] This type of distribution suggests a particular solution—reduce the influence of the tails in the computation of the aggregate inflation rate by way of a trimmed-mean.

The use of trimmed-means as a noise-reduction technique has been suggested by others, notably Bryan and Cecchetti (1994), Bryan, Cecchetti, and Wiggins (1997), Roger (1997), Smith (2004), and most recently Dolmas (2005.) A (symmetric) trimmed-mean is

[2] A reference to seasonal adjustment as a data filter can be found in Ghysels (1997).
[3] These cross-section statistics, and the trimmed-mean estimates evaluated in this paper are produced from 41 components for the years 1998-2007, and 36 components from 1967-1997. These components are readily available from the U.S. Department of Labor's Bureau of Labor Statistics, and are listed on the Federal Reserve Bank of Cleveland's website at: http://www.clevelandfed.org/research/Inflation/US-Inflation/mcpitable.cfm.

computed as the weighted average of the data ordered in magnitude over goods *x*(*i*), with their associated weights *w*(*i*),as in equation (1),

$$\bar{x}_{\alpha} = \frac{1}{1-2\frac{\alpha}{100}} \sum_{i \in I_{\alpha}} w_i x_i \tag{1}$$

where I_{α} is the set of observations symmetrically excluding the most extreme $\alpha/100$ and $1-\alpha/100$ percentiles of the distribution.[4] So $\bar{x}_o$ is the weighted mean of the data, $\bar{x}_{50}$ is the weighted median, and the $\bar{x}_{\alpha}$'s for $0 < \alpha < 50$ are the set of trimmed-mean estimators between the mean and median.

A simple inspection of the aggregate seasonality that exists in the data reveals a clear reduction in seasonal noise as a result of aggregating retail price data using trimmed-means (table 2). Seasonally adjusting the component price data reduces the monthly variance in the aggregate CPI from 16.95 percent to 10.40 percent—a reduction of 6.55 percentage points. Similarly, while excluding food and energy from the aggregate inflation statistic reduces its monthly variance, it does little to eliminate seasonal noise. Seasonal adjustment of the core components reduces the monthly variance in the core CPI by 4.95 percentage points. However, trimming the extreme observations of the data clearly reduces seasonal noise. The 16 percent trimmed mean (i.e., $\bar{x}_8$) eliminates all but 1.13 percentage points of seasonal variance in the data, while the median CPI eliminates all but 0.13 percentage point of seasonality.

[4] When the distribution of the component price changes is persistently skewed, asymmetric trimmed-means have been employed, such that the resulting estimator is centered on the mean percentile of the distribution, rather than the 50th percentile. Asymmetric trimmed-means have been used for New Zealand by Roger (1997), for Brazil by Bryan and Cecchetti (2001), and for U.S PCE price data by Dolmas (2005).

There are other, more effective ways to reduce seasonality in the data than cross-sectional trimming. The characteristics of certain kinds of seasonal noise are easily defined and therefore allow a more direct elimination by time-series methods that measure repeated, periodic patterns in data. But not all seasonality can be eliminated in this way. For example, irregular seasonal movements, such as patterns that occur at frequencies not delineated by calendar time (e.g., holidays that fall during different weeks of the year), or are not annually consistent (e.g., winters and summers of varying intensity) are not well captured by these rather precise techniques. And of course, seasonality is but one type of noise. Sampling error, for example, can affect data in a way that obscures the longer-term cyclical patterns of interest to the policymaker. Indeed, there are an uncountable number of "one-off" changes that users of high-frequency price data discount in order to better gauge the cyclical pattern they wish to influence. These commonly include changes related to taxes and any number of temporary supply disruptions like strikes, or natural calamities.

In the price data, there are two commonly used approaches for reducing high-frequency, nonperiodic noise. One way is to selectively exclude the offending data prior to aggregation. This is the so-called "excluding food and energy" approach, where there is specific information regarding the transitory nature of some component that justifies its elimination from the set of component data.[5] A general criticism of this approach is the uncertainty over what constitutes "specific" information. Every price is subject to idiosyncratic, transitory forces of varying degrees and knowing which deserve exclusion at any point in time subjects the data to the particular tastes and biases of the data user.

[5] Support for this approach and a caution against "filtering" approaches like trimmed means can be found in Poole (2004).

A popular alternative is to evaluate incoming price data on the basis of its year-over-year change, a frequency long enough, presumably, for high-frequency noise to be eliminated from the data. But as in the case of seasonality, year-over-year changes eliminate all information in the data at frequencies of less than a year and thus, greatly reduce the timely application of the data.

Several other papers have documented the use of trimmed-mean estimators as a method for reducing high-frequency noise in the inflation statistic. For example, Bryan and Cecchetti (1994) show that trimmed-means of the CPI produce better year-ahead CPI forecasts and are more strongly correlated with past money growth—two indications that this filtering technique does not eliminate information useful to monitoring the inflation process. This approach has been used by others for alternative market baskets (Smith 2004, and Dolmas 2005), and for different countries (such as Roger [1997], Mio and Higo [1999], and LaFleche [1997]).

Viewed as a noise-reduction technique, two qualities of trimmed-means should be emphasized. First, the power of this approach increases as the inflation signal falls relative to the noise. That is, the usefulness of a trimmed-mean inflation measure increases as the inflation signal declines. Second, a key advantage to a trimmed-mean price statistic is that it preserves the high-frequency quality of the data. When it is combined with time series averaging (i.e. 12-month percent changes), which is essentially just another noise-reduction method, the usefulness of trimmed-means is substantially reduced.

These qualities are made clear in Figure 1a, which shows the accuracy of the 16% trimmed-mean CPI ($\bar{x}_8$) relative to the all-items CPI in predicting the all-items CPI changes over the next 12-months. RMSE's are shown over moving-ten year horizons beginning in 1982. For the monthly data, aggregating after symmetrically trimming 8 percent from each tail of the price-change distribution consistently produces a superior estimate of the future CPI trimmed than the CPI does for itself. As the inflation trend falls over time, however, the relative predictive accuracy of the trimmed-mean increases relative to the all-items CPI. That is, as the strength of the inflation "signal" falls, the relative influence of high-frequency noise in the data rises.

Also shown in the figure is the relative predictive power of the two inflation measures on a 12-month basis. Note that when the data are time-series averaged, the relative forecast accuracy of this trimmed-mean is greatly reduced, at times even favoring the all-items CPI. As time-series averaging of the data is just another way to reduce high-frequency noise, the relative signal of the all-items CPI is more similar to the trimmed-mean at this frequency. (The RMSE's for similar experiments between the median CPI and the all-items CPI are shown in figure 1b, the results of which are essentially identical.)

I next evaluate whether preserving the high-frequency of the data by the use of trimmed means reduces the time it takes to identify permanent changes in the inflation trend—a valuable asset for a central bank trying to achieve, and maintain, price stability in a low-inflation environment.

III. AN EVALUTION OF BREAKS IN THE INFLATION TREND

In this section, I report the results of tests designed to answer the following question: Given a change in the inflation trend, does filtering the data by way trimmed-means provide a more timely signal of the trend shift than the Core CPI or time-series averaging? In all tests, the methodology used for identifying trend breaks is the multiple-structural break identification algorithm of Bai and Perron (1998). Details of the model are described in the Appendix, (which has been reproduced from Carlson et al. [2000].) The Bai-Perron approach is computationally efficient without being too restrictive in terms of the number of breaks, and it can easily accommodate a single, autoregressive process like inflation. The number of breaks, their timing, and the constant (mean of the inflation trend) are all estimated from a series of sequential Wald tests. All of the tests reported in this section employ a univariate form of the Bai and Perron methodology, where breaks are assumed to be in the constant of a linear autoregressive inflation equation.

III A. Identification of Breakpoints—Controlled Bootstrap Experiments

I first conducted a controlled experiment, identifying breaks in a constructed set of inflation data. Specifically, I created 100 time series from bootstraps of each of the following four inflation measures; the aggregate CPI ($\bar{x}_0$), the median CPI ($\bar{x}_{50}$), the 16 percent trimmed-mean CPI ($\bar{x}_8$), and the CPI excluding food and energy (core CPI).[6] The median and 16 percent trimmed mean CPI were chosen because earlier research has identified them as reasonable representatives among the potential set of trimmed mean estimators.[7] For example, Bryan and Cecchetti (1994) and Smith (2004) show that the median CPI is a superior signal of the future inflation trend, while the 16% trimmed-mean CPI is the minimum variance estimator of CPI data.[8]

---

[6] The bootstraps were constructed from data observed between 1991 and 2007, a relatively low and (statistically) stable inflation trend.
[7] Computational costs prevented running the full set of Monte Carlo simulations in this version of the working paper.
[8] See Bryan and Cecchetti (1994).

For each data series, mean breaks in the inflation trend of various magnitudes are introduced into the data at some time *t*, and the Bai-Perron algorithm is run on each series, beginning with the time-series spanning 1 to *t*+1, successively adding new (lower mean) observations until that algorithm identifies the break in the series with a 95 percent probability. Figure 2a shows the results of this bootstrap experiment using a break in the trend of 1.5 percentage points.

Breaks in the data were identified in half of the constructed series for the core CPI, the median CPI, and the 16% trimmed-mean CPI, all within a few months of each other—four to six months after the break was introduced into the data. There was only a tiny gain in break identification for the 16 percent trimmed-mean and the median CPI relative to the core CPI, although all measures substantially outperformed the overall CPI, which took about two years before the Bai-Perron procedure was capable of identifying a break in half of its series. The relative failure of the CPI to pick up the break is somewhat misleading. The sample over which these bootstraps were drawn, 1991 to 2007, was a period of rather extreme volatility in oil prices, which clearly obscures the ability of the change-point test to spot the shift in the inflation trend.

Figures 2b and 2c show the results of identical experiments, but for successively smaller trend breaks. In figure 2b, the break in the inflation trend is only 1 percentage point, and the improvement in the trimmed-mean measures relative to the traditional core CPI becomes more obvious. Here, break points were identified in half of the median CPI and 16% trimmed-mean CPI series within eight and seven months, respectively, but it took eleven months for the core CPI. In figure 2c, the break is reduced further to ½ percentage point. In this case, a break point was revealed in half of the bootstrapped 16% trimmed-mean data within 18 months of the shift in the inflation trend, 22 months for the median CPI, and not until 37 months for the core CPI.

These bootstrap experiments are suggestive of a gain in signal-to-noise above and beyond the mere extraction of food and energy from the inflation data, but only suggestive. Given that the trimmed-mean measures enjoy a reduced time-series variance relative to the core and the all-items CPI, the fact that mean breaks in trend were found earlier for these measures is not remarkable. A more compelling test is whether, based on historical data, the trimmed-mean measures would have alerted policymakers to a new inflation trend earlier than the all-items CPI or the core CPI.

III B. Identification of Breakpoints—Historical Monthly Data

Bai-Perron break-point analysis identifies three inflation regimes (two clear break points) for the inflation trend over the 1973-2007 period. A high inflation period runs from January 1973 to August 1981, over which CPI inflation averaged 9.4 percent; a moderate inflation period runs from September 1981 to December 1990, with a CPI average of 4.1 percent; and a low-inflation period runs from January 1991 to the present, with the CPI averaging 2.7 percent. Using these ex-post breaks, I repeating the procedure used in the earlier section, sequentially adding monthly data to each of our inflation measures until the Bai-Perron test reveals the break in the inflation trend. The results of these experiments are shown in figures 3a and 3b.[9]

---

[9] The Bai-Perron test requires we set a "break-fraction", the minimum proportion of the sample before the end (beginning) of the data after (before) which a break date will not be indicated. For the bootstrap experiments, this break was set a 5%, but due to the autocorrelation in the historical data, the break fraction was set at 15% for these tests. If the break signal is strong enough, data breaks can still be signaled as having occurred, but in such cases the date specified for the break will be in error. While this could be a limitation for these tests in a

For the inflation trend break from high to moderate inflation beginning in September 1981 (when the CPI trend dropped by 5.3 percentage points), the CPI signaled a break in trend with the March 1983 data—18 months after we now estimate the break to have occurred (figure 3a and column 1 of table 3A.). The 16% trimmed mean did not pick up the new trend until May 1984 (after 32 months), and the median CPI five months later (October 1984). The traditional core CPI identified the break two months after that (December 1984). The reduction in high-frequency noise achieved by the trimmed means or the traditional core CPI appears to have offered the policymaker no better indication of the shift in the inflation trend than the all-items CPI did when the size of the break was relatively large.

However, the monthly CPI failed to detect a more subtle shift in the inflation trend, such as occurred when inflation fell from an average of 4.1 percent to 2.7 percent in January 1991. It took the monthly CPI six and one-half years to signal this trend break (figure 3b and first column in table 3B). The core CPI performed better and identified the break within 26 months of the realized break date; the 16% trimmed mean identified the break six months before that (September 1992), while the median CPI signaled a break in July 1992, 18 months after the break is now known to have occurred.

As implied by the bootstrapped data, spotting a relatively modest shift in the trend is inhibited by high-frequency noise in the retail price data. While the core CPI strips away one particular source of that noise, namely the volatile food and energy prices, the trimmed-mean data appears to detect a change in the inflation trend more efficiently.

---

policy setting, especially for relatively modest changes in the inflation trend, the restriction does not appear to have had an influence on any of the conclusions drawn from these ex-post historical experiments.

III C. Identification of Breakpoints—Historical Data with Time-Series Smoothing

Data analysts commonly use time-series averages of the data, notably 12-month percent changes, as a method of filtering noise from the data. As a final experiment, I check to see if there are any gains from employing time-series averages of the inflation series. Time-series averaging may conceivably offset its inherent loss of timeliness by substantially improving the signal-to-noise ratio and, on balance, could speed up the identification of a shift in the inflation trend. To investigate this possibility, Bai-Perron tests on historical data are again conducted using two-month, three-month, six-month, and twelve-month time-series averages of each of the four inflation measures (the CPI, the core CPI, the 16% trimmed mean, and the median CPI.) The results of these experiments for both the high-to-moderate and moderate-to-low inflation breaks are report in table 3.

For the high-inflation break of 1981, the monthly CPI provided best signal of all of the indicators (and identical to the core CPI). As the frequency of the data is reduced—that is, as we move from monthly to two-month, three-month, six-month, and twelve-month time-series averages of the data, we actually observe a reduction in the clarity of the break identification for all of the inflation measures—at least as judged by the Bai-Perron methodology.

Even in the low-inflation period, time-series averaging of the data did not reduce the amount of time it took for the measure to identify the break point (except for a small improvement in the signal of the median CPI at the two-month frequency.) This result is intriguing since it suggests that while time-series averaging produces a more stable inflation series, it eliminates much of the signal that the inflation trend has changed.

IV. CONCLUSION

This paper has taken another look at the usefulness of trimmed-mean inflation measures in monitoring inflation. It argues that these measures reduce high-frequency noise in the inflation data effectively and that they are superior to both the standard core CPI and the more commonly used time-series averages of inflation data. The ability of trimmed-mean inflation measures to signal a shift in the inflation trend improves as the shift becomes more subtle. Presumably, large breaks in the inflation series are obvious in the aggregate statistic since the signal is very strong relative to the underlying noise in the data. But as the signal decreases in intensity, such as we might expect in an environment where the central bank is trying to manage inflation around some specific, but relatively low average, the problem of high-frequency noise becomes more troublesome for the policymaker. In this environment, trimmed means appear to offer one way to reduce this noise, while preserving the timeliness of the signal. This could be an important advantage to a central bank which is expected to quickly spot changes in the inflation trend and respond before the higher inflation trend becomes more entrenched in the inflation process.

Table 1: Cross-sectional Distribution Characteristics of Monthly CPI Data
Seasonally and nonseasonally adjusted components (annualized percent, 1998-2007)

	Mean	Variance	Skewness
Unadjusted Component Data	**2.74**	**34.2**	**0.45**
Seasonally Adjusted Data	**2.66**	**20.9**	**0.85**

Table 2: Monthly Time-Series Variance of Alternative Inflation Measures
Seasonally and nonseasonally adjusted component data (percent, 1998-2007)

	Unadjusted	Seasonally adjusted	Variance reduction from seasonal adjustment
CPI	16.95	10.40	6.55
Core CPI	6.06	1.11	4.95
16% Trim	1.82	0.69	1.13
Median	0.79	0.66	0.13

Table 3: Identification of CPI Trend Breaks by Trim and Frequency
(Months until the break in the inflation trend is identified)

A: Break from High (9.4%) to Moderate (4.1%) Inflation (Sept. 1981, break=5.3 percentage points)

	Data Frequency				
	One-month	**Two-months**	**3-months**	**6-months**	**12-months**
CPI	18 months	33 months	39 months	38 months	29 months
Core CPI	39 months	54 months	57 months	52 months	71 months
16% Trim	32 months	42 months	43 months	40 months	35 months
Median	37 months	45 months	46 months	54 months	50 months

B: Break from Moderate (4.1%) to Low (2.7%) Inflation (Jan. 1991, break = 1.4 percentage points)

	Data Frequency				
	One-month	**Two-months**	**3-months**	**6-months**	**12-months**
CPI	78 months	94 months	135 months	159 months	132 months
Core CPI	26 months	26 months	31 months	36 months	42 months
16% Trim	20 months	30 months	38 months	46 months	47 months
Median	18 months	17 months	18 months	22 months	29 months

Figure 1a: Relative Predictive Accuracy of the 16% Trimmed-mean CPI
(Estimated relative to the all-items CPI over moving, 10-year samples)

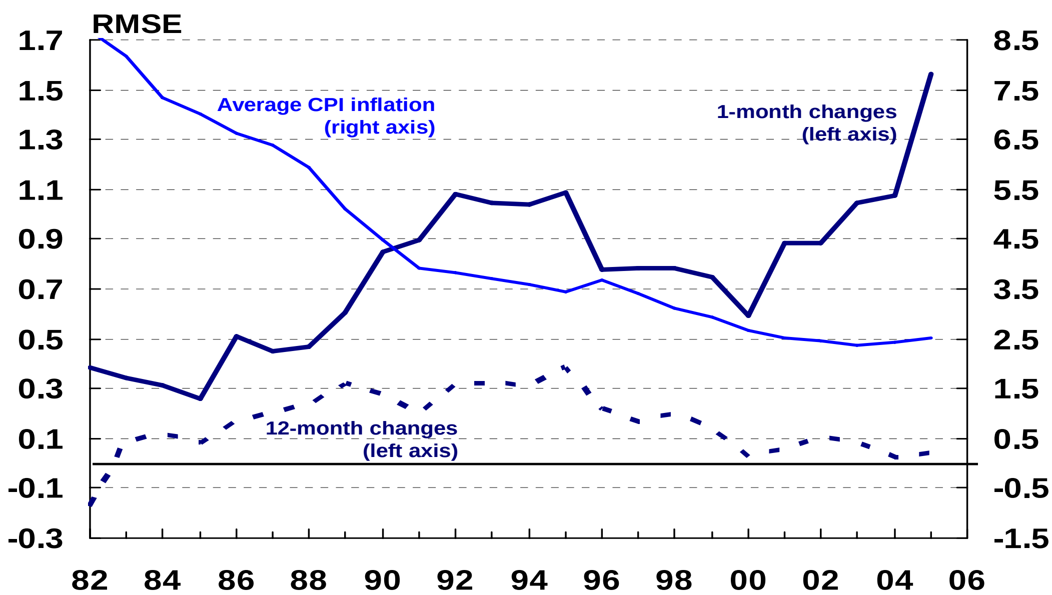

Figure 1b: Relative Predictive Accuracy of the Median CPI
(Estimated relative to the all-items CPI over moving, 10-year samples)

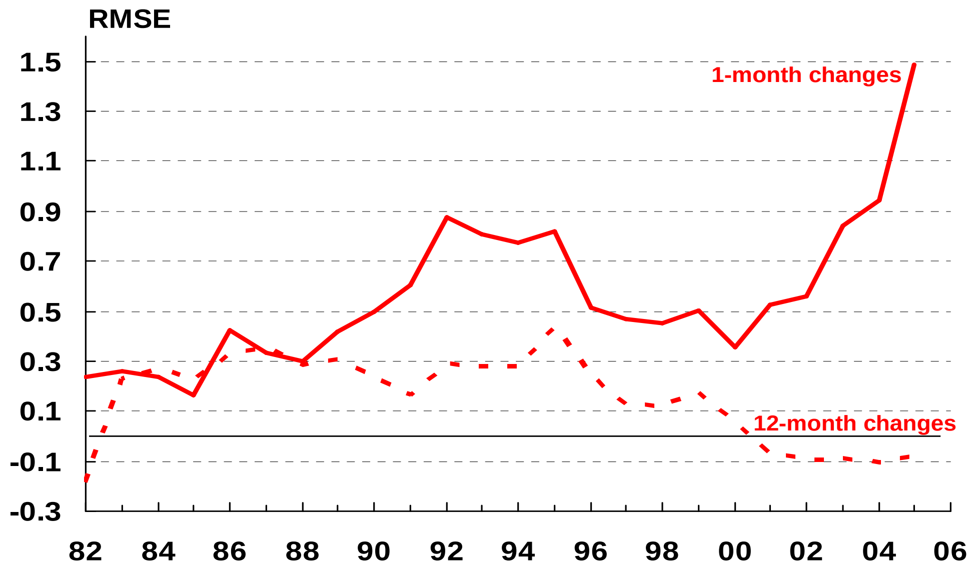

Figure 2a: Identifying Breaks in the Inflation Trend
(monthly data, break = 1.5 percentage points)

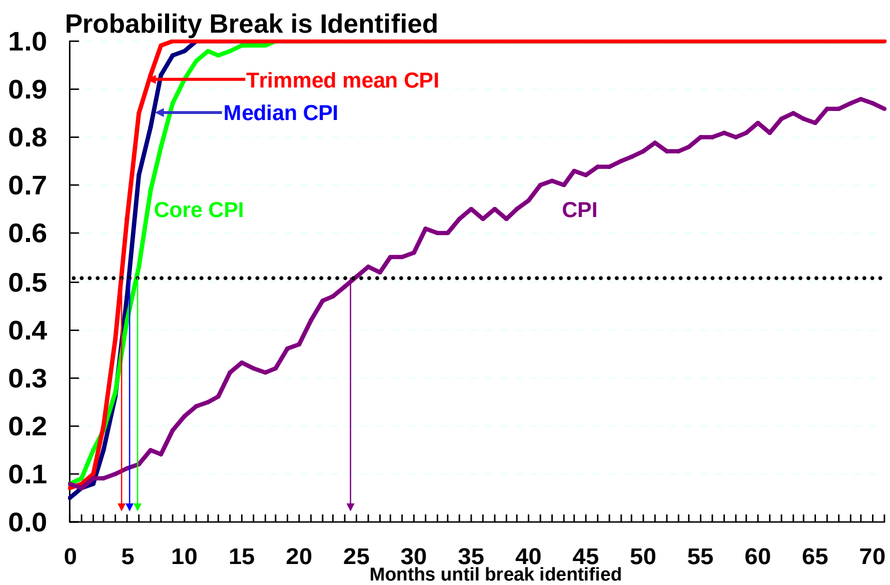

Figure 2b: Identifying Breaks in the Inflation Trend
(monthly data, break = 1.0 percentage point)

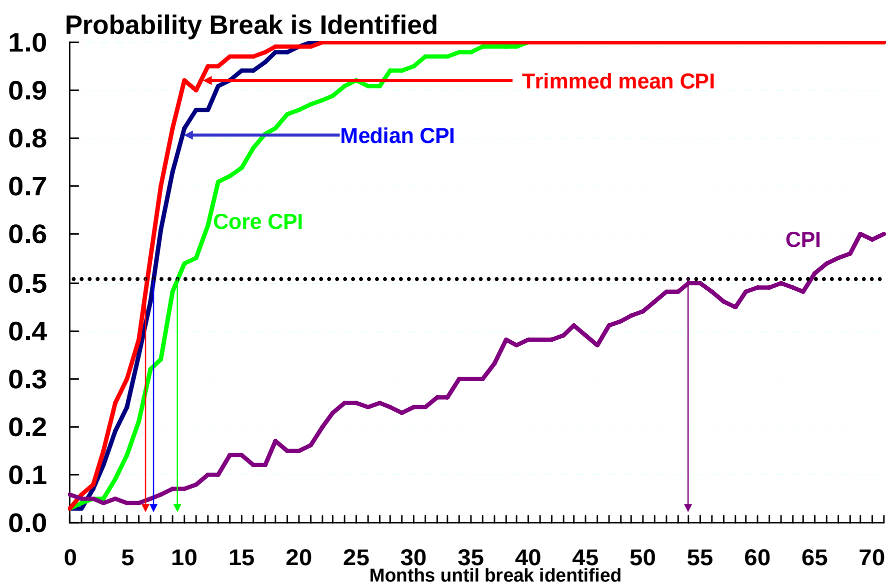

Figure 2c: Identifying Breaks in the Inflation Trend

(monthly data, break = 0.5 percentage point)

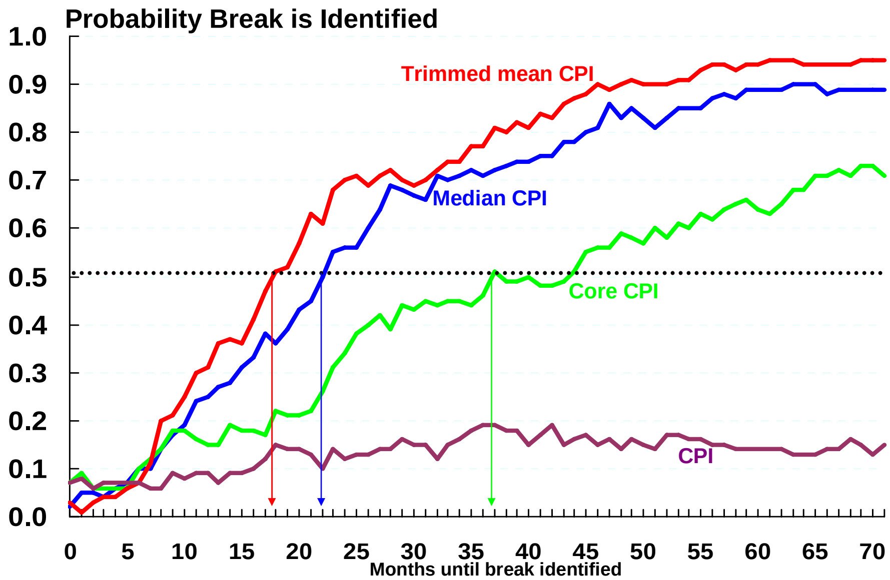

Figure 3a: HISTORICAL CPI BREAK IDENTIFICATION

(monthly data, break = 5.4 percentage points)

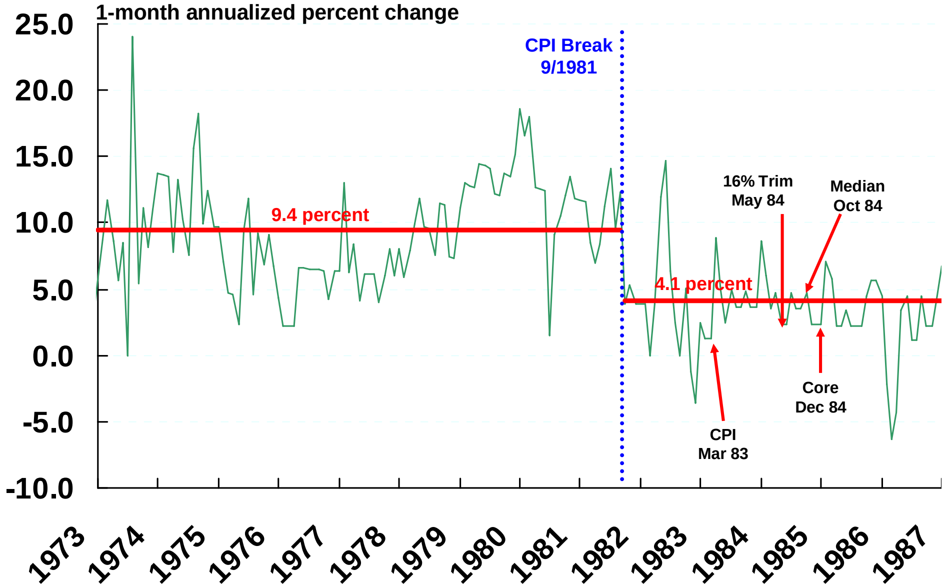

Figure 3b: HISTORICAL CPI BREAK IDENTIFICATION

(monthly data, break = 1.4 percentage points)

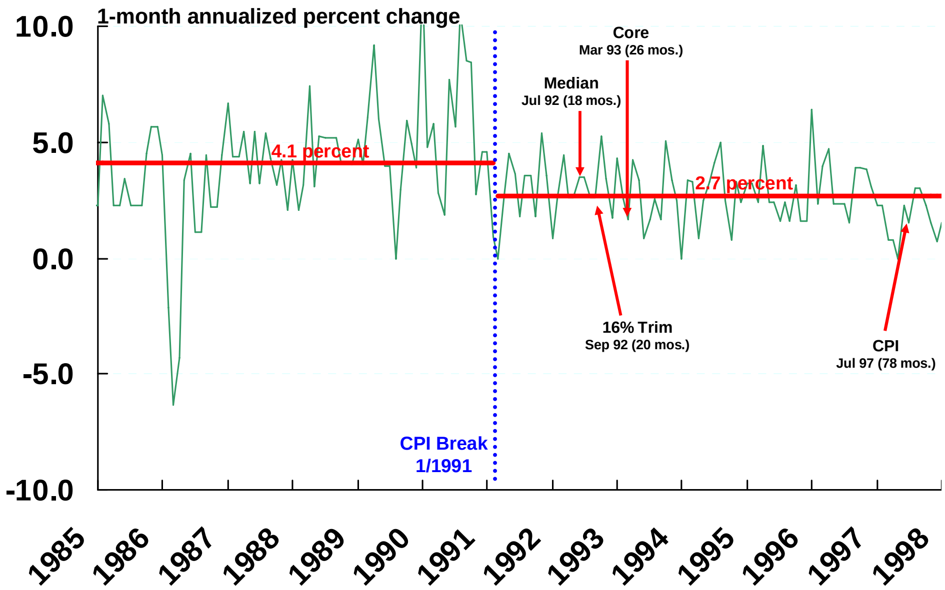

References

1. Bai, Jushan and Pierre Perron, "Estimating and Testing Linear Model with Multiple Structural Changes," Econometrica, Vol. 66, no. 1 (Jan. 1998), pp. 47-78.

2. Bryan, Michael F., and Stephen G. Cecchetti, "Measuring Core Inflation," in Monetary Policy, N. Gregory Mankiw (ed.), National Bureau of Economic Research Studies in Business Cycles, Vol. 29, University of Chicago Press (1994), pg. 195-215.

3. Bryan, Michael F., and Stephen G. Cecchetti, "The Seasonality of Consumer Prices," NBER Working Paper No. 5173, July 1995.

4. Bryan, Michael F., and Stephen G. Cecchetti, "Inflation and the Distribution of Price Changes," The Review of Economics and Statistics, Vol. 81, Issue 2 (May 1999), pg. 188-196.

5. Bryan, Michael F., and Stephen G. Cecchetti, "A Note on the Efficient Estimation of Inflation in Brazil," Working Paper no. 11, Bank of Brazil, March 2001.

6. Bryan, Michael F., Stephen G. Cecchetti, and Rodney L. Wiggins, "Efficient Inflation Estimation," NBER Working Paper No. 6183, September 1997.

7. Carlson, John B., Ben Craig, and Jeffrey C. Schwarz, "Structural Uncertainty and Breakpoint Tests: An Application to Equilibrium Velocity," Journal of Economics and Business, 2000; 52: 101-115.

8. Dolmas, Jim, "Trimmed Mean PCE Inflation," Federal Reserve Bank of Dallas Working Paper 0506, July 2005.

9. Ghysels, Eric, "Seasonal Adjustment and Other Data Transformations," Journal of Business and Economic Statistics, vol. 15, no. 4 (October 1997), pp. 4410-418.

10. LaFleche, Therese, "Statistical Measures of the Trend Inflation Rate," Bank of Canada Review, Autumn, 1997, pp. 29-46.

11. Mio, Hitoshi, and Masahiro Higo, "Underlying Inflation and the Distribution of Price Change: Evidence from the Japanese Trimmed-mean CPI," Monetary and Economic Studies, Bank of Japan, Vol. 17, no. 1 (May 1999).

12. Poole, William, "Inflation Signals and Inflation Noise," Speech at the University of Arkansas at Little Rock," April 6, 2004.

13. Roger, Scott, "A Robust Measure of Core Inflation in New Zealand 1949-1996", Reserve Bank of New Zealand Discussion Papers, July 1997.

14. Smith, Julie K, "Weighted Median Inflation: Is This Core Inflation?", Journal of Money, Credit, and Banking, 36 (2004), pg. 253-263.

15. Smith, Julie K., "Inflation Targeting and Core Inflation," Canadian Journal of Economics, Vol. 38, Issue 3, pg. 1018 (August 2005).

APPENDIX (Reproduced from Carlson, Craig, and Schwarz, 2000)

The Bai-Perron (1998) test is based upon an information criterion in the context of a sequential procedure, and allows one to find the number of breaks implied by the data, as well as estimating the timing of the breaks and the parameters of the processes between breaks. Among the many important advantages offered by this test is that it is not computationally excessive, in spite of what, on the surface, would appear to be an extremely computer intensive and difficult numerical problem. To illustrate, consider the following *m*-break model:

$$y_t = x_t'\beta + z_t'\delta_1 + u_t,\ \ t = 1, 2, \ldots, T_1$$
$$y_t = x_t'\beta + z_t'\delta_2 + u_t,\ \ t = T_1 + 1, \ldots, T_2$$
$$\vdots$$
$$y_t = x_t'\beta + z_t'\delta_{m+1} + u_t,\ \ t = T_m + 1, \ldots, T$$

where y_t = observed value of the dependent variable at time *t*; x_t = p x 1 vector of covariates with corresponding vector of coefficients, β; z_t = q x 1 vector of covariates with corresponding vector of coefficients, δ_j (j = 1, 2, …, m+1); u_t = disturbance (error) at time *t*; and $T_1, \ldots, T_m$ = breakpoint dates, treated as unknowns. Note that in this general framework, not all coefficients need to be subject to structural change. A proposed set of breakpoint dates implies a partition of the data into segments of time, each of which has a separate model associated with it.

The Bai-Perron test looks at all the possible partitions for a given number of such breakpoints, *m*. The partition which minimizes the sum of the squared residuals when one iteratively estimates the parameters β and δ_j associated with it is chosen as the most likely partition. The computational burden rapidly becomes intolerable for $m > 2$ (there are $T!/[T-m]!$ different possible combinations of breakpoints, each of which would imply a different estimate). An important innovation of Bai and Perron is an efficient computational method based on dynamic programming which reduces the number of computations to the order of T^2 for any number of breaks greater than one.

The Bai-Perron procedure applies the least squares principle to estimate β, δ_j (j = 1, …, m+1), and T_i (i = 1, …, m) using T observations on (y_t, x_t, z_t). That is, for each m-partition β and δ_j (j = 1, …, m+1) are estimated by…

$$\text{Min } (Y - X\beta - \overline{Z}\delta)'(Y - X\beta - \overline{Z}\delta) = \sum_{i=1}^{m+1} \sum_{t=T_{i-1}+1}^{T_i} [y_t - x_t'\beta - z_t'\delta_i]$$

In practice, however, the number of breaks is unknown. Bai and Perron propose the following sequential approach: First, estimate the model with a small number of breaks that are thought to be necessary (or start with no break). Second, perform parameter constancy tests for every subsample, adding a break to a subsample associated with a rejection of the null hypothesis of no break using the test $F_T(l + 1 \mid l)$ given by

$$F_T(l+1 \mid l) = \left\{ S_T\left(\hat{T}_1, \ldots \hat{T}_l\right) - \min_{1 \le i \le l+1} \inf_{\tau \in l_{i,\eta}} S_T\left(\hat{T}_1, \ldots, \hat{T}_{i-1}, \tau, \hat{T}_i, \ldots \hat{T}_l\right) \right\} / \hat{\sigma}^2$$

where

$$\Lambda_{i,\eta} = \left\{\tau; \widehat{T}_{i-1} + \left(\widehat{T}_i - \widehat{T}_{i-1}\right)\eta \le \tau \le \widehat{T}_i + \left(\widehat{T}_i - \widehat{T}_{i-1}\right)\eta\right\}$$

and $\widehat{\sigma}^2$ is a consistent estimate of σ^2 under the null hypothesis. Note that for $i = 1$,

$S_T\left(\widehat{T}_1,..,\widehat{T}_{i-1},\tau,\widehat{T}_i,...\widehat{T}_l\right)$ is understood as $S_T\left(\tau,\widehat{T}_1,...\widehat{T}_l\right)$ and for $i = 1+1,$ as $S_T\left(\widehat{T}_1,...\widehat{T}_l,\tau\right)$.

This process is repeated by increasing *l* sequentially until the test fails to reject the null hypothesis.

There are two variants of this sequential procedure proposed by Bai and Perron. One uses a null hypothesis based on a global minimization of the sum of squared residuals to estimate a given number of breakpoints, *m*. The other approach is more strictly sequential, using a null hypothesis of *m* breaks determined sequentially, not globally. Bai and Perron show that an estimation strategy need not simultaneously estimate the location of breaks in order to consistently determine the number of breaks. They extend their procedure to accommodate serial correlation in the disturbance term (giving their method an advantage over alternative approaches suggested by Yoa 1988 and Liu et al. 1997) or, alternatively, to accommodate the inclusion of a lagged dependent variable in the list of regressors.