FEDERAL RESERVE BANK OF DALLAS

Issue 3 May/June 2005

Southwest Economy

Dallas Fed Introduces Business-Cycle Indexes for Texas Metros

The frequency and severity of cyclical swings in a local economy are important to businesses and consumers because such cycles impact production and inventory decisions, employment and unemployment. Analyzing the overall direction of a local economy, however, can be difficult and confusing. Often the handful of local economic indicators gives mixed signals. For example, if the unemployment rate and job growth both increase, is the local economy picking up or weakening? Often it is not clear.

To more clearly define regional business cycles, the Dallas Fed has developed composite indexes that aggregate the movements of key economic indicators for nine Texas metropolitan areas. The Metro Business-Cycle Indexes use statistically optimal weights so that movements in the indexes best represent the underlying co-movements in the indicators and thus the underlying

(Continued on page 2)

A Fitter, Trimmer Core Inflation Measure

Speaking of the challenge in interpreting monthly inflation numbers during his tenure on the Federal Reserve Board, former Vice Chairman Alan Blinder said, "The name of the game then was distinguishing the signal from the noise, which was often difficult. The key question on my mind was typically: What part of each monthly observation on inflation is durable and what part is fleeting?"[1]

Blinder's conception of a component of monthly inflation that is durable as opposed to fleeting—that represents signal rather than noise—corresponds to what most economists call *core inflation*. Core inflation, understood in this way, represents the underlying trend in inflation once temporary swings have

(Continued on page 4)

INSIDE:

A Fitter, Trimmer Core Inflation Measure

(Continued from front page)

been smoothed out. Because what is temporary and what is lasting can only be known with the benefit of hindsight, the true core inflation rate for any given month cannot be known with certainty until well after the fact. In real time—as the data arrive and policy decisions need to be made—the best that economists can do is estimate the core inflation rate.

Measures of inflation that exclude food and energy prices are probably the best-known core inflation gauges. In fact, the measures excluding food and energy—which government statisticians include in their releases of the Consumer Price Index (CPI), Producer Price Index (PPI) and the price index for Personal Consumption Expenditures (PCE)—are often spoken of as if they were synonymous with core inflation. Properly speaking, though, they represent just one of many potential core measures. To be sure, because of the high short-run volatility of some food and energy prices, there is some rationale for excluding those prices from a measure of core inflation. But as research over the past decade has made clear, much better estimates can be made by taking a more rigorous approach to the problem of which prices to include and which to exclude.

To date, that research has focused primarily on developing better measures of core inflation in the CPI.[2] This article discusses the application of some of the insights and techniques of that line of research to the Federal Reserve Board of Governors' preferred inflation gauge, the PCE price index. (See box titled "The Fed's Favorite Inflation Gauge.") The result is a new measure of core PCE inflation—the trimmed mean PCE—and a somewhat different characterization of the economy's recent inflation experience.

Food and Energy: Signal or Noise?

Consider the following data from March 2005. More than 200 expenditure categories go into the PCE. Table 1 shows the 10 categories with the biggest price increases from February to March 2005.[3] Note that the price changes are not annualized—they are one-month percentage changes. By way of comparison, the change in the overall PCE price index from February to March was +0.46 percent.

Table 1

10 Biggest Price Increases in March 2005

Component	Change from prior month (percent)
Gasoline and other motor fuel	8.0
Purchased fuel oil	5.8
Airline service	4.2
Hotels and motels	4.2
Medical services: labs	3.2
Farm fuel	2.5
Purchased liquid petroleum gas	2.5
Miscellaneous personal services	2.4
Watch, clock and jewelry repair	2.4
Laundry and garment repair	2.4

Table 2

10 Biggest Price Decreases in March 2005

Component	Change from prior month (percent)
Eggs	−4.4
Fresh fruit	−2.6
Women's luggage	−1.8
Men's luggage	−1.8
Intrastate toll calls	−1.8
Photographic equipment	−1.8
Toys, dolls and games	−1.7
Household operation: natural gas	−1.7
Durable house furnishings: textiles	−1.5
Lighting supplies	−1.5

Table 2 lists the 10 components that had the largest price decreases in March 2005. While it's true that food and energy items show up a number of times on both lists, there are many other items as well. Moreover, not all food and energy items had price changes as large as these. Some food components in particular—such as food consumed away from home—are notoriously stable. For example, the price index for "other purchased meals"—which comprises meals purchased at restaurants and bars—rose by just 0.15 percent in March. That small price volatility is typical for food purchased and eaten away from home—making its exclusion from a measure of core inflation questionable.

Clearly, in any given month, excluding only food and energy items still leaves very volatile components in the price index. And, excluding all food and energy items may throw out some useful information.

The Fed's Favorite Inflation Gauge

Since February 2000, the Federal Reserve Board's semiannual monetary policy reports to Congress have described the Board's outlook for inflation in terms of the PCE. Prior to that, the inflation outlook was presented in terms of the CPI. In explaining its preference for the PCE, the Board stated:

> The chain-type price index for PCE draws extensively on data from the consumer price index but, while not entirely free of measurement problems, has several advantages relative to the CPI. The PCE chain-type index is constructed from a formula that reflects the changing composition of spending and thereby avoids some of the upward bias associated with the fixed-weight nature of the CPI. In addition, the weights are based on a more comprehensive measure of expenditures. Finally, historical data used in the PCE price index can be revised to account for newly available information and for improvements in measurement techniques, including those that affect source data from the CPI; the result is a more consistent series over time.
>
> —Monetary Policy Report to the Congress, Federal Reserve Board of Governors, Feb. 17, 2000

The Trimmed Mean Technique: A Little Off the Top (and Bottom)

How, then, do we decide which items to exclude or include more rigorously? In a study focusing on the CPI and PPI, Bryan, Cecchetti and Wiggins make a statistical case for the use of trimmed means as a method for estimating core inflation.[4] In spite of the arcane-sounding name, the concept of a trimmed mean is simple. In fact, trimmed means should be familiar to any follower of international figure skating. In the wake of the controversies surrounding the judging at the 2002 Winter Olympics, the International Skating Union adopted a scoring system in which a skater's highest and lowest marks are discarded before the skater's average score is calculated. Trimmed mean inflation rates are derived by a similar procedure.

In any given month, the rate of inflation in a price index like the CPI or PCE can be thought of as a weighted average, or mean, of the rates of change in the prices of all the goods and services that make up the index.[5] Calculating the trimmed mean PCE inflation rate involves looking at the price changes for each of the individual components of personal consumption expenditures—the sort of data contained in Tables 1 and 2. The individual price changes are sorted in ascending order from "fell the most" to "rose the most," and certain fractions of the most extreme observations at both ends of the spectrum are—like a skater's best and worst marks—thrown out, or trimmed. The inflation rate is then calculated as a weighted average of the remaining components.[6]

How many components should be trimmed from the top and bottom of the monthly price-change distributions? Since our aim is to create a more accurate real-time gauge of core inflation, we want our trimming to yield a measure that comes as close as possible to the core inflation we've observed in historical data. (See box titled "Optimal Trimming: The Nuts and Bolts" for more detail.) Following the approach used by Bryan, Cecchetti and Wiggins in their CPI/PPI study, we will treat true core inflation as a smooth underlying trend in actual inflation (*Chart 1*).[7]

For data that run from 1979 through 2002, the amount of trimming that minimizes the distance between the trimmed mean inflation rate and the proxy for the true core inflation rate turns out to be substantial. The optimal trim drops roughly the top 25 percent of components (as a fraction of expenditures) and the bottom 21 percent. That is, from each month's data, we discard the 25 percent of expenditure components whose prices rose the most and the 21 percent whose prices fell the most (or rose the least). The trimmed mean inflation rate is then calculated as the weighted average of the remaining expenditure components, the middle 54 percent. Note that

In spite of the arcane-sounding name, the concept of a trimmed mean is simple. In fact, trimmed means should be familiar to any follower of international figure skating.

Chart 1

Actual PCE Inflation and Proxy for True Core Inflation

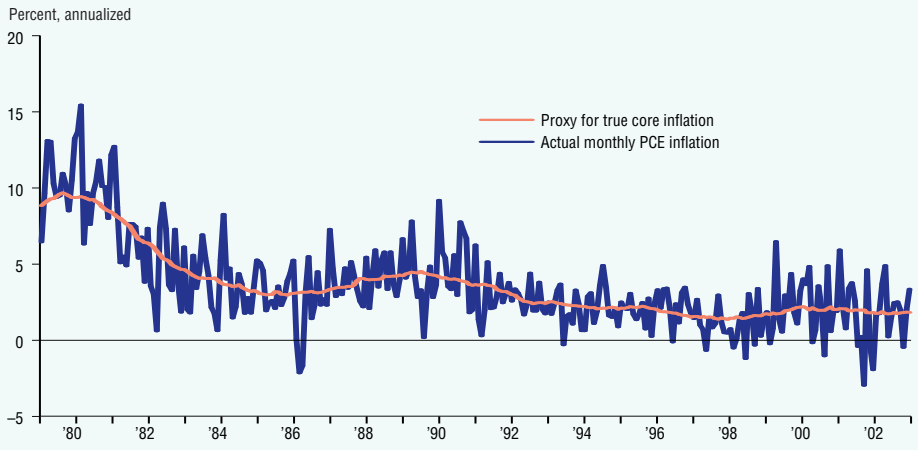

SOURCES: Bureau of Economic Analysis; author's calculations.

The optimally trimmed mean performs much better as an estimator of core PCE inflation than the usual measure excluding food and energy.

the set of goods and services discarded each month—items adding up to roughly 46 percent of expenditures—must include a good deal more than just food and energy, which account for only about 20 percent of total PCE.

So Which Goods Get Trimmed?

As suggested above, some food components, like food purchased and consumed away from home, are rarely excluded when one approaches the trimming problem rigorously. This is a feature of the inflation data that Bryan and Cecchetti (1994) highlighted in their study of the CPI, and it is true of the PCE as well. Chart 2 shows the monthly inflation rate for the PCE component "other purchased meals," together with the upper and lower trim points for the optimally trimmed mean, from 1990 through 2004.

The trim points have the following interpretation. In each month, items whose prices rose by more than the upper trim points in the chart are excluded from the optimally trimmed mean that month, as are items whose prices fell by more (or rose by less) than the lower trim points. There is only a handful of months during this 14-year period in which the purchased meals component was excluded from the optimally trimmed mean.

Food items of this sort are well represented among the components least often excluded from the optimally trimmed mean. Table 3 lists the top 20 least-often-excluded components for the sample period 1977–2004. Food items actually occupy five of the top 10 spots, with "other purchased meals" coming in first. Out of a sample of 335 months, it's excluded only 13 times. The other dominant category in the least-often-excluded list is housing, which shows up in various forms.

Table 4 gives a corresponding list of the top 20 most-often-excluded items. Food items figure prominently here, too, with "fresh vegetables" topping the list. Fuels, financial services and electronics items are also prominent.

How Well Does the Trimmed Mean Perform?

Just as Bryan, Cecchetti and various co-authors found regarding the CPI, the optimally trimmed mean performs much better as an estimator of core PCE inflation than the usual measure excluding food and energy.

In data running from 1979 through 2002, the gain in accuracy from using the optimally trimmed mean rather than the measure excluding food and energy is about 0.77 percentage point annually. That is, compared with the usual measure excluding food and energy, on average the monthly trimmed mean measure would be expected to come closer to true monthly core inflation by just over three-fourths of a percentage point when the inflation rates are expressed in annual terms.

These results compare the performance of one-month inflation rates,

Chart 2

Food Away from Home Is Generally Not Trimmed

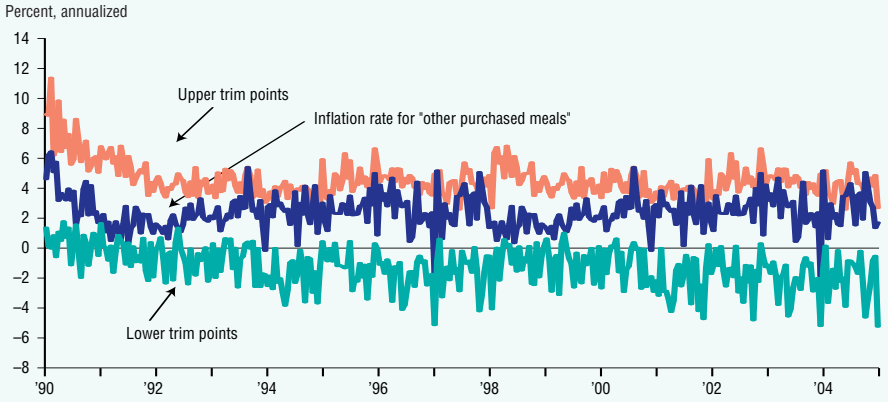

SOURCES: Bureau of Economic Analysis; author's calculations.

Table 3

20 Least-Often-Excluded Components, 1977–2004

Component	Number of months excluded (out of 335)
Other purchased meals	13
Owner-occupied stationary homes	16
Casino gambling	34
Tenant-occupied stationary homes	35
Tenant-occupied mobile homes	40
Purchased meals: elementary and secondary schools	41
Purchased meals: higher education	41
Food furnished to employees: military	41
Food furnished to employees: civilian	42
Club and fraternity housing	50
Tenant group room and board	52
Tenant group employee lodgings	53
Auto repair	54
Owner-occupied mobile homes	57
Military clothing	87
Domestic service paid in cash	88
Household operation, not elsewhere classified	91
Social welfare including child care	94
Medical care: other professional services	95
Dry cleaning	96

Table 4

20 Most-Often-Excluded Components, 1977–2004

Component	Number of months excluded (out of 335)
Fresh vegetables	314
Eggs	314
Computers and peripherals	311
Food produced and eaten on farms	304
Airline services	299
Brokerage charges and investment counseling	298
Software	297
Fresh fruit	296
Purchased fuel oil	294
Gasoline and other motor fuel	286
Farm fuel	285
Poultry	285
Video equipment, excluding TVs	285
Auto insurance net premiums	284
Purchased liquid petroleum gas and other fuel	279
TVs	278
Durable house furnishings: textiles	275
Semidurable house furnishings	274
Commercial bank imputed interest	273
Infants' clothing	273

Optimal Trimming: The Nuts and Bolts

As discussed in the text, we want our trimming to yield a measure that comes as close as possible to a specific proxy for true core inflation, in this case a centered, 36-month moving average of actual monthly PCE inflation. What do we mean by "as close as possible"?

The numbers reported in the article are for the case where the closeness is measured with a root-mean-square-error criterion—that is, the trimmed mean's distance from the proxy for true core inflation is measured by the square root of the average squared monthly deviation between the two series. Each possible amount of trimming—5 percent off the top, 10 percent off the bottom, or 20 percent off the top, nothing off the bottom, and so forth—results in a trimmed mean inflation rate that is some calculable distance from the proxy for true core inflation. The *optimal* trim is the one that minimizes the distance between the trimmed mean and core proxy over our sample period, 1979–2002. This turns out to be the trimming: 25.3 percent off the top, 20.6 percent off the bottom.

Table A shows the value of our measure of fit—the root-mean-square error, or RMSE—for inflation horizons of one, three, six and 12 months, for both the optimally trimmed mean and the measure excluding food and energy. The three-, six- and 12-month inflation rates for the trimmed mean are obtained by cumulating the optimally trimmed series of one-month rates to obtain a price index, then taking three-, six- and 12-month annualized percentage changes of that price index. Smaller numbers are better than larger ones in both Tables A and B.

The optimally trimmed mean also performs better than the measure excluding food and energy in terms of its average error, as can be seen in Table B. The average, or mean, error of an inflation measure is simply the sum of its monthly deviations from the true core proxy divided by the number of months in the sample.

To see the relevance of this last point, suppose that true core inflation is zero in two consecutive months. Imagine that one measure (call it X) estimates core inflation as being +0.25 percent in each of the two months, while a second (Y) estimates it at +1 percent in the first month and –1 percent in the second month. Then Y would have a higher RMSE than X—on average, Y is 1 percentage point away from the truth, versus 0.25 percentage point for X—but it would have a smaller average error than X. Y's average error is zero (the +1 and –1 cancel out) compared with X's average error, which, like X's RMSE, is 0.25 percentage point. If the trimmed mean and excluding food and energy measures followed this pattern—one better in terms of RMSE, the other better in terms of average error—we might be hard-pressed to say which was the better measure. Fortunately, Tables A and B show the trimmed mean is better on both dimensions.

Table A

Root-Mean-Square Errors for Various Inflation Horizons (in percentage points)

	1-month	3-month	6-month	12-month
Trimmed mean	.87	.58	.49	.51
Excluding food and energy	1.63	.94	.72	.76

Table B

Average Errors for Various Inflation Horizons (in percentage points)

	1-month	3-month	6-month	12-month
Trimmed mean	.04	.06	.09	.15
Excluding food and energy	.11	.11	.14	.19

For the average person, however, transitory surges in overall inflation are no less inconvenient simply because they are transitory.

which are quite volatile relative to the slower-moving core series. This is true for both the optimally trimmed mean and the measure excluding food and energy, though less so for the trimmed mean. Looking at the CPI, Cecchetti (1997) emphasized the additional noise reduction that can be achieved by examining longer-horizon inflation rates.[8] Cecchetti's point is equally valid with regard to the PCE. Looking at three-, six- or 12-month inflation rates improves the accuracy of both the trimmed mean and the measure excluding food and energy as gauges of core inflation.

For both measures, six-month changes give the highest accuracy in gauging core inflation. While the longer horizons benefit the measure excluding food and energy more than the trimmed mean, the latter is still the more accurate core inflation gauge. For the three-month inflation horizon, the relative gain in accuracy from using the trimmed mean is almost 0.4 percentage point. For the six- or 12-month horizons, the gain in accuracy is 0.23–0.25 percentage point, a not-insignificant difference. (See Table A in box titled "Optimal Trimming: The Nuts and Bolts.")

Chart 3 gives a visual sense of how the trimmed mean performs relative to the measure excluding food and energy. The chart shows the annualized six-month inflation rates in the two measures, together with the proxy for true core inflation. The series are shown for the full sample period used in the optimal trim calculations, 1979–2002.

What Has the Trimmed Mean PCE Inflation Rate Been Telling Us Lately?

Chart 4 shows the recent behavior of the trimmed mean PCE inflation rate, together with the more common excluding-food-and-energy inflation rate for the three different time intervals. Here are the salient points:

• While both the trimmed mean and excluding-food-and-energy inflation rates decline in 2003, the lows hit by the trimmed mean measure are not nearly as low as those reached by the measure excluding food and energy. For example, the three-month trimmed mean inflation rate falls below 1 percent in only one month of 2003, versus five such months for the inflation rate excluding food and energy. The lows for the six- and 12-month trimmed mean rates are nearer 1.5 percent.

• Both inflation rates began to climb in early 2004. The highs reached in mid-2004, however, are both higher and more sustained in the trimmed mean measure than in the measure excluding food and energy. The three- and six-month trimmed mean inflation rates both spent time in the neighborhood of 2.5 percent.

• Inflation decelerated in the second half of 2004, according to both inflation

Chart 3

Comparison of Core Inflation Measures

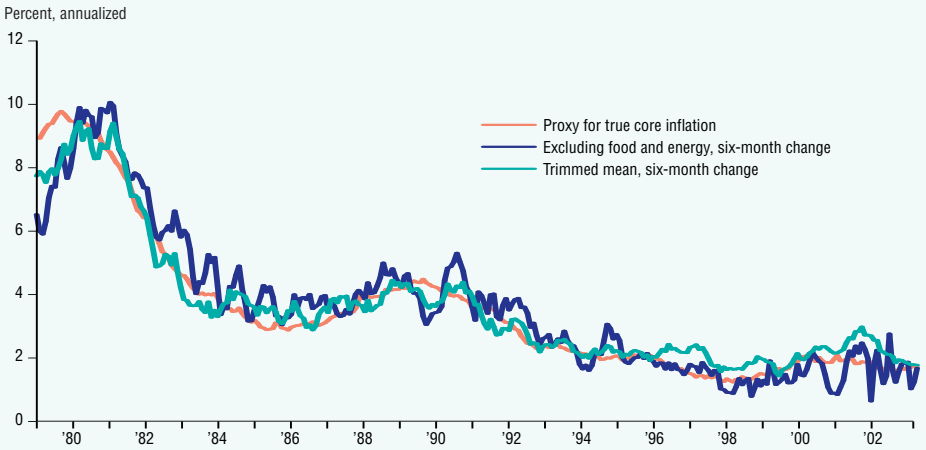

SOURCES: Bureau of Economic Analysis; authors' calculations.

Chart 4

Trimmed Mean and Excluding-Food-and-Energy PCE Inflation

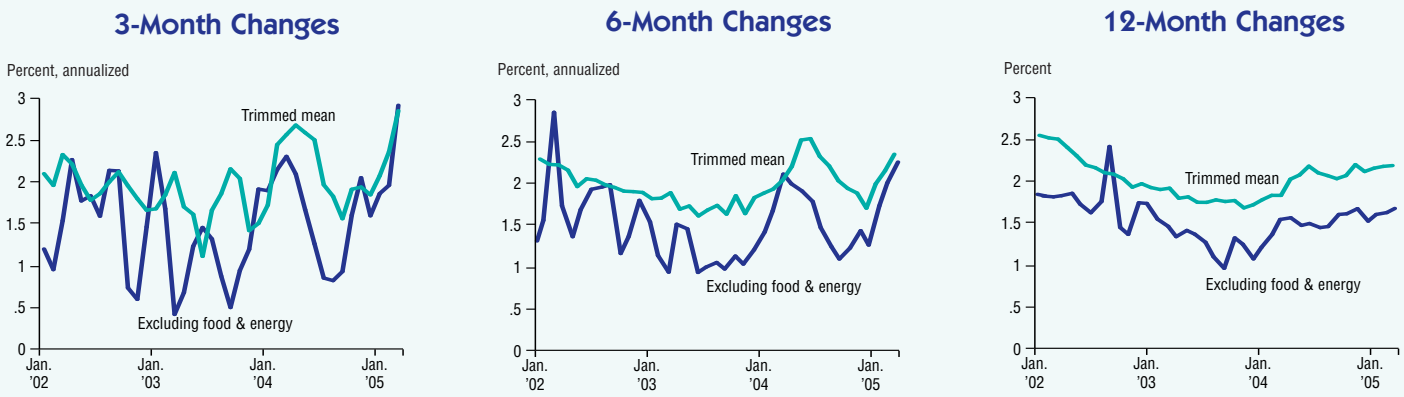

SOURCES: Bureau of Economic Analysis; author's calculations.

measures. This shows up as a decline in the three- and six-month inflation rates and a stabilization in the 12-month rates. The three- and six-month trimmed mean rates bottom out around 1.5 percent, compared with around 1 percent for the three- and six-month troughs in the rate excluding food and energy. Similarly, the 12-month trimmed mean rate stabilizes at about 2 percent, or half a percentage point higher than the 12-month rate excluding food and energy.

• While the 12-month inflation rates in both measures look stable, the three- and six-month rates show that inflation has accelerated since late 2004. Both rates suggest core PCE inflation is currently running above 2 percent.

Why Should We Care?

This article began with a quote from former Fed Vice Chairman Blinder describing a policymaker's difficulties in interpreting monthly movements in the inflation rate. Why the individuals setting monetary policy would care about core inflation—and why, as a result, they continually seek improved estimates of core inflation—is fairly clear. Changes in inflation that are known to be transitory and, thus, soon to be reversed pose less threat to the goal of long-run price stability than more lasting changes.

For the average person, however, transitory surges in overall inflation are no less inconvenient simply because they are transitory. If last month's consumer price inflation was high mainly due to a temporary jump in the prices of food, energy or other items, this does not change the fact that a household's dollars couldn't buy as much food, energy or other items as they otherwise could have.

So why should anyone outside of a central bank care about the latest trimmed mean PCE inflation rate (or any other core measure)? Individuals routinely make decisions that rely, at least implicitly, on forecasts of future inflation—for instance, whether to invest in fixed-income securities or to take on fixed-income obligations. For decisions of this sort, knowledge of whether recent changes in inflation are durable or transitory—signal rather than noise—is likely to be of value.

—Jim Dolmas

Dolmas is a senior economist and policy advisor in the Research Department of the Federal Reserve Bank of Dallas.

Notes

I thank Mark Wynne and Evan Koenig, who provided numerous helpful comments at various stages of this research, and Jennifer Afflerbach, who suggested many improvements in exposition.

1 "Commentary on 'Measuring Short-Run Inflation for Central Bankers,'" by Alan Blinder, Federal Reserve Bank of St. Louis *Review*, May/June 1997.

2 That more rigorous approach was pioneered by Michael Bryan and Stephen Cecchetti. See their article "Measuring Core Inflation," in N. Gregory Mankiw, ed., *Monetary Policy*, Chicago: University of Chicago Press, 1994. For a good survey of these methods, see "Core Inflation: A Review of Some Conceptual Issues," by Mark A. Wynne, European Central Bank Working Papers Series, No. 5, 1999.

3 All data used in this article are from the Bureau of Economic Analysis via Haver Analytics. The data on the detailed components of the PCE index are as reported in Tables 2.4.4U and 2.4.6U in the "Underlying Detail Tables" section of the Bureau of Economic Analysis web site: www.bea.doc.gov/bea/dn/nipaweb/nipa_underlying/Index.asp.

4 "Efficient Inflation Estimation," by Michael Bryan, Stephen Cecchetti and Rodney Wiggins, National Bureau of Economic Research Working Paper Series No. 6183, September 1997.

5 In the CPI, the weight an individual component receives corresponds to its share in consumer spending, on average, over a two-year reference period. CPI weights are thus fixed for two years at a time. Weights in the PCE are slightly more complicated and change from month to month. To a first approximation, the weight a component receives this month is an average of (1) its expenditure share last month and (2) what its expenditure share would be if consumers bought this month's quantities at last month's prices.

6 The weighted median CPI, which is produced by the Federal Reserve Bank of Cleveland and is perhaps familiar to some readers, is an extreme form of trimmed mean. It corresponds to the limiting case where nearly all the price changes in the upper and lower halves of the distribution are trimmed, leaving only the price change of the single component exactly in the middle. Pursuing the skating analogy from the text, imagine that judging panels consist of seven members. The median inflation rate is analogous to a scoring formula that discards a skater's three highest and three lowest marks.

7 In particular, the calculations in this article use a centered, 36-month moving average of monthly inflation rates to proxy for true core inflation—that is, the true core inflation rate in any given month is assumed to be the average of that month's inflation rate together with the inflation rates of the prior 18 months and those of the subsequent 18 months. In a more technical version of this article (forthcoming), I consider other proxies for true core inflation.

8 "Measuring Short-Run Inflation for Central Bankers," by Stephen Cecchetti, Federal Reserve Bank of St. Louis *Review*, May/June 1997.