Comparing Two Measures of Core Inflation: Some Additional Perspective

Evan F. Koenig[1]
Federal Reserve Bank of Dallas
September 14, 2018

Executive Summary: Trimmed-mean PCE inflation does not clearly dominate ex-food-and-energy "core" PCE inflation in real-time forecasting of headline PCE inflation. However, trimmed-mean inflation is a superior communications tool. That is because trimmed-mean inflation more successfully filters out headline inflation's transitory variation, leaving only cyclical and trend components. A corollary is that armed with a trusted measure of longer-run trend inflation, it is relatively easy to use the behavior of trimmed-mean inflation to draw inferences about slack. (Alternatively, armed with a trusted measure of slack, it is relatively easy to use the behavior of trimmed-mean inflation to draw inferences about inflation's longer-run trend.) Finally, over periods where we have relevant data, real-time trimmed-mean inflation has been a less-biased estimator of "true" (*i.e.*, revised) headline inflation.

● Comparing mean values of latest-vintage data across headline and core PCE inflation measures is of limited interest unless one is contemplating switching from a long-run target for headline PCE inflation to a long-run target for core PCE inflation. Taking as given the FOMC's decision to specify a long-run target for headline inflation, of greater interest are the average deviations of *first-release* trimmed-mean and ex-food-and-energy inflation from *latest-vintage* headline inflation. It is, after all, first-release core inflation data that you will be watching and reacting to in real time, while it is the latest-vintage headline numbers that are relevant to an *ex post* assessment of inflation-control performance. We have real-time trimmed-mean data starting in 2005:Q2 and real-time ex-food-and-energy PCE inflation starting in 1996:Q1. Table 1 shows mean and median values of the relevant alternative inflation measures over sample periods that begin on those dates.

Table 1. Mean and Median Inflation Rates

	Mean Inflation		
	Headline*	Ex F&E**	Trimmed Mean**
2005:2 - 2018:Q2	1.77	1.61	1.78
1996:1 - 2018:Q2	1.81	1.60	1.90
	Median Inflation		
	Headline*	Ex F&E**	Trimmed Mean**
2005:2 - 2018:Q2	1.95	1.52	1.79
1996:1 - 2018:Q2	1.96	1.60	1.95

*Latest-available (September 11, 2018) vintage.
** First release or as close as possible to first release. (The earliest-available vintage of trimmed-mean inflation is 2005:Q2.)

[1] This memo benefitted greatly from research assistance provided by Emil Mihaylov, and from the comments of Jim Dolmas.

Conclusion: Early-release trimmed-mean PCE inflation is, on average, a more accurate gauge of headline inflation than is early-release ex-food-and-energy PCE inflation. If you want a sense of whether trend headline inflation is at, above, or below the FOMC's 2-percent longer-run target, you will likely want to pay more attention to trimmed-mean inflation releases than to ex-food-and-energy inflation releases.

● Analysts sometimes use inflation data to draw inferences about slack. (For example, the Congressional Budget Office historically has made revisions to its estimates of the natural rate of unemployment based on the behavior of inflation.) The strength and robustness of the relationship between early-release estimates of an inflation measure and latest-vintage slack is important for assessing the usefulness of that measure as a rule-of-thumb indicator of resource utilization.[2] We regressed early-release estimates of inflation (de-trended using SPF long-run inflation expectations) on a constant and latest-vintage CBO estimates of the unemployment gap. Results suggest that trimmed-mean inflation is more strongly and more reliably related to labor-market slack than is either headline inflation or ex-food-and-energy inflation. See Table 2 and Figures 1, 2, and 3.

Table 2. Which Inflation Measure is Most Closely, and Reliably, Related to Slack?

	Inflation Measure	Constant (S.E)	U – U* (S.E.)	Adj. R^2
2005:2 – 2018:2	Headline – SPF	- 0.146 (0.315)	- 0.124 (0.104)	0.033
	Ex F&E – SPF	**- 0.241** (0.070)	**- 0.136** (0.031)	0.487
	Trimmed Mean – SPF	0.035 (0.079)	**- 0.197** (0.039)	0.651
1996:1 – 2018:2	Headline – SPF	- 0.367 (0.170)	- 0.048 (0.067)	- 0.001
	Ex F&E – SPF	**- 0.531** (0.073)	- 0.049 (0.031)	0.051
	Trimmed Mean – SPF	- 0.120 (0.065)	**- 0.152** (0.033)	0.432
1981:4 – 2018:2	Headline – SPF	- 0.173 (0.170)	- 0.047 (0.077)	- 0.001
	Ex F&E – SPF	- 0.139 (0.145)	- 0.069 (0.065)	0.017
	Trimmed Mean – SPF	- 0.124 (0.061)	**- 0.198** (0.037)	0.442

Notes:
• First-release inflation data were used whenever possible. When first-release inflation was unavailable (before 2005:Q2 for trimmed-mean inflation and before 1996:Q1 for conventional-core inflation) the earliest-available vintage was used instead.
• "SPF" denotes 10-year inflation expectations, from the Survey of Professional Forecasters. The unemployment gap is lagged 4 quarters, but results were qualitatively similar with a 1-quarter lag.
• Coefficients that are statistically significant at the 1-percent level are bolded.

[2] Alternatively, the analyst armed with a real-time slack measure in which she has confidence might hope to use the behavior of inflation to draw inferences about longer-run inflation expectations. Then it will be useful for early releases of the inflation measure to be strongly and robustly related to longer-run expectations, without having to control for a wide range of influences other than slack.

Conclusion: Deviations of trimmed-mean inflation from SPF long-run inflation expectations are a better indicator of whether slack remains in the labor market (as gauged, *ex post*, by the CBO) than are deviations of either headline or ex-food-and-energy inflation. Put another way, trimmed-mean inflation is more successful at filtering out transitory inflation variation than is ex-food-and-energy inflation: The deviation of trimmed-mean inflation from trend inflation (as captured by SPF long-run expectations) better approximates inflation's cyclical component.

● Which core inflation measure is most useful for predicting future headline inflation? You can't accurately answer that question without thinking carefully about how best to go about estimating forecasting equations in real time. Koenig, Dolmas, and Piger ("The Use and Abuse of 'Real-Time' Data in Economic Forecasting," *Review of Economics and Statistics*, 85, 2003) show that if you are going to be forecasting using first-release data, then you should estimate your forecasting equation with first-release data on its right-hand side. Analysts often, instead, use end-of-sample-vintage real-time data (the most up-to-date vintage available in real time) on the right-hand side of their real-time forecasting equations. For the left-hand-side variable, the obvious choice is end-of-sample data. However, there are potential gains in coefficient precision (hence, forecast accuracy) from stripping unforecastable noise from the left-hand-side variable before estimation. Gains will be most evident in smaller samples. When forecasting inflation, depending on the forecast horizon, stripping out unforecastable noise could mean using trimmed-mean or ex-food-and-energy inflation as the dependent variable, even if it is headline inflation that you are ultimately interested in forecasting (Koenig and Atkinson, Federal Reserve Bank of Dallas *Staff Papers*, Issue 16, 2012).[3] Lacking time for a complete analysis of real-time inflation forecasting, we undertook two very simple exercises, described below.

Rule-of-Thumb Forecasting

A simple rule of thumb is to set your forecast of inflation over the next four quarters equal to observed inflation over the most recent 4-quarter period. There's no estimation, here. The relevant "observed inflation" is first release. The variable being forecasted is latest-vintage headline inflation. Results are shown in Table 3A.

Table 3A. Rule-of-Thumb Inflation Forecasting

	RMSE		
Rule-of-Thumb Forecast	2006:2–2018:2	1996:1–2018:2	1981:4–2018:2
Lagged Headline	1.472	1.229	1.411
Lagged Ex Food & Energy	1.060	0.998	1.218
Lagged Trimmed Mean	1.134	0.985	0.923

Note: When first-release inflation was unavailable (before 2005:Q2 for trimmed-mean inflation and before 1996:Q1 for conventional-core inflation) the earliest-available vintage was used instead.

[3] Depending on the data-revision process, it might also mean using early-release data on the left-hand-side of the forecasting equation. See, again, Koenig, Dolmas, and Piger ("The Use and Abuse of 'Real-Time' Data in Economic Forecasting," *Review of Economics and Statistics*, 85, 2003).

The above results may be distorted by the aftermath of the financial crisis, which brought about a sharp decline in inflation. It is highly unlikely that anyone would have relied on a rule-of-thumb inflation forecast during that period. Excluding 2009:Q1 – 2009:Q4 from the calculations yields the results displayed in Table 3B.

Table 3B. Rule-of-Thumb Inflation Forecasting (excluding 2009)

	RMSE		
Rule-of-Thumb Forecast	2006:2–2018:2	1996:1–2018:2	1981:4–2018:2
Lagged Headline	1.021	0.945	1.274
Lagged Ex Food & Energy	0.823	0.870	1.162
Lagged Trimmed Mean	0.840	0.807	0.807

Note: When first-release inflation was unavailable (before 2005:Q2 for trimmed-mean inflation and before 1996:Q1 for conventional-core inflation) the earliest-available vintage was used instead.

Conclusion: Over the periods for which we have real-time trimmed-mean and ex-food-and-energy inflation, these two series perform about equally well as rule-of-thumb predictors of headline inflation. Both core inflation series perform notably better than lagged headline inflation.

Recursive Real-Time Forecasts of Headline Inflation

We also recursively estimate an inflation-forecasting equation of the form

$$\pi(t) = \alpha + \beta_1\pi(t-4) + \beta_2\pi^c(t-4) + \beta_3\pi^e(t-4) + \gamma u(t-4),$$

where u is the unemployment rate, π is 4-quarter headline inflation, π^c is either ex-food-and-energy or trimmed-mean inflation, and π^e is SPF long-run inflation expectations. End-of-sample-vintage data are used on the equation's left-hand side and first-release (or as close to first release as possible) data are used on the right-hand side. The first forecast is for 2006:Q2, using first-release data for 2005:Q2. The final forecast is for 2018:Q2, using first-release data for 2017:Q2. We use two different sample starting points: 1996:Q1 and 1982:Q4. The real-time forecasts are compared with latest-vintage headline PCE inflation. The 4-quarter period immediately following the financial crisis (2009:Q1 – 2009:Q4) is excluded from both estimation and forecast evaluation. Results are shown in Table 4A.

Table 4A. Real-time Forecasts of Headline Inflation

	RMSE: 2006:Q2 – 2018:Q2 (ex 2009)	
Start of Sample	π^c = ex-food-and-energy	π^c = trimmed-mean
1996:Q1	1.21	1.12
1982:Q4	0.86	0.89

When the sample period used for estimation is short, forecast performance is slightly better using the trimmed mean to measure core inflation. That advantage disappears (and forecast performance improves) when the sample period is extended back into the 1980s.

As noted above, more-precise coefficient estimates (and, so, better forecasts) can sometimes be obtained by stripping noise from the dependent variable. That fact suggests there could be an advantage to replacing $\pi(t)$ on the left-hand side of the above equation with $\pi^c(t)$. Again, left-hand-side data are end-of-sample vintage, and *forecasts are compared with latest-vintage headline PCE inflation data*. Results are displayed in Table 4B.

Table 4B. Real-time Core Inflation Forecasts as Predictors of Headline Inflation

	RMSE: 2006:Q2 – 2018:Q2 (ex 2009)	
Start of Sample	π^c = ex-food-and-energy	π^c = trimmed-mean
1996:Q1	0.87	0.83
1982:Q4	0.88	0.88

Forecast performance with a short sample is much improved, regardless of which core inflation measure is used, but there is no change in forecast performance when the estimation period is extended back to the 1980s. Interestingly, the root-mean-square forecast errors reported in Table 4B are no better than those obtained from simple rule-of-thumb forecasting based on lagged core inflation. (Compare the RMSEs reported in Table 4B with the left-column entries in Table 3B.)

Conclusion: At a 4-quarter horizon, evidently, real-time forecasts of core inflation usefully serve, also, as forecasts of headline inflation. It makes relatively little difference whether the core measure is ex-food-and-energy inflation or trimmed-mean inflation.

Figure 1. De-trended headline inflation is only loosely related to labor-market slack

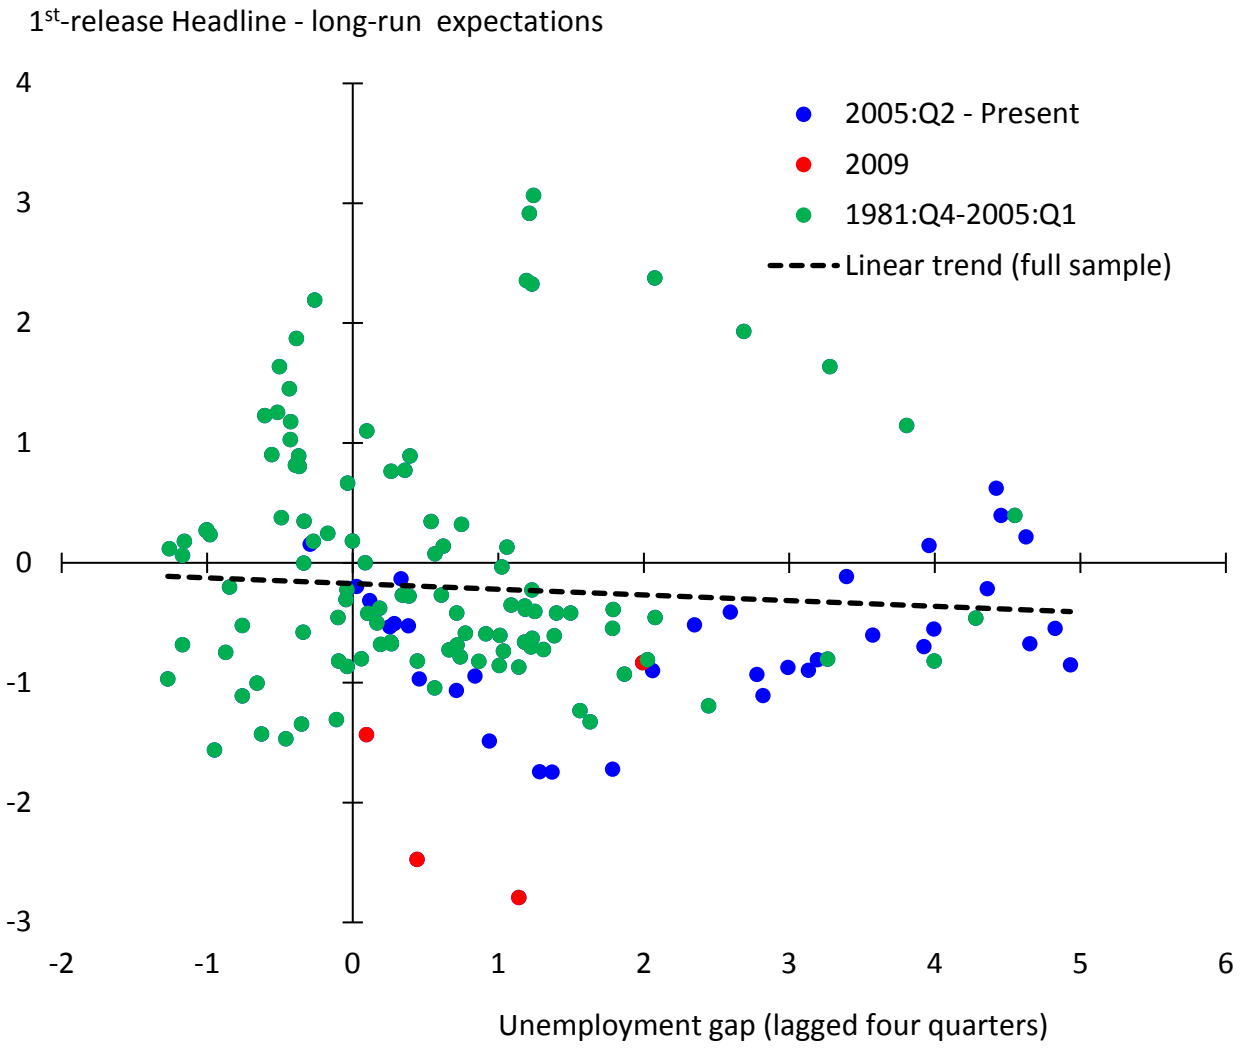

Figure 2. De-trended ex-food-and-energy inflation is only loosely related to labor-market slack

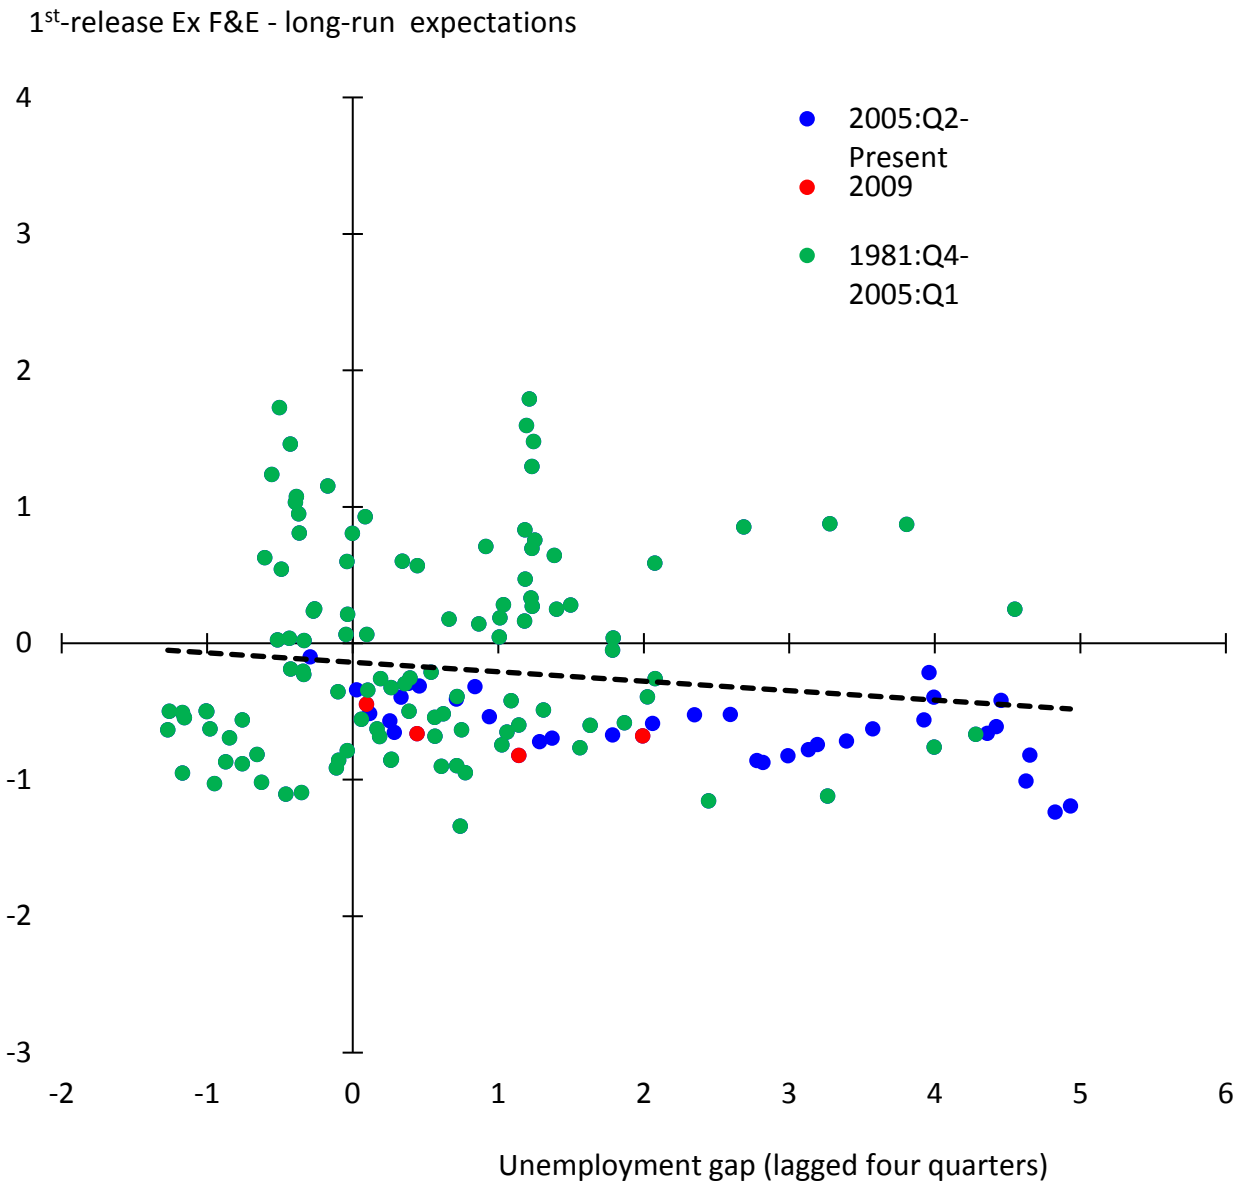

Figure 3. De-trended trimmed-mean inflation shows a fairly strong connection to labor-market slack