RESEARCH RAP Special Report

Inflation Targeting and Revisions to Inflation Data: A Case Study with PCE Inflation*

Calvin Price

July 2011

Introduction

Central banks around the world have come to recognize the importance of maintaining low and stable inflation. One widely employed tool for helping to do so is known as inflation targeting, whereby a central bank sets a numeric goal for inflation. Once this target is publicly stated, the bank can be held accountable for its actions in regard to meeting, or not meeting, this target. Countries that have adopted such a tool have generally had a favorable experience, and there is evidence that inflation targeting is correlated with increased stability in output growth, lower inflation, and more stable inflation expectations (Dotsey, 2006).

While at a broad level the idea of inflation targeting can appear to be a straightforward concept, carrying it out in practice requires a central bank to make many subtle decisions. For example, a central bank must decide how to specify the target: one value or a range of acceptable values. The bank must also decide over what period of time it should measure inflation. This could mean comparing the target to an average of the past three months or perhaps an average over the past year. Yet another decision must be made as to which measure of inflation will be used. Fundamentally, measuring inflation means measuring the growth of a price level, and the variety of choices stem from the variety of different price indices used to measure the nation's general price level.

* The views expressed here are those of the author and do not necessarily reflect those of the Federal Reserve Bank of Philadelphia or of the Federal Reserve System. Calvin Price is a research associate and can be contacted at Calvin.Price@phil.frb.org.

In the United States, two widely followed key inflation measures are the consumer price index (CPI) and the personal consumption expenditures (PCE) price index. Each index represents a general measure of prices paid by consumers, but they differ in several respects. Some examples of these differences include the scope of goods considered (the PCE is more comprehensive) and how frequently the weights within each index are updated (the PCE weights are updated more frequently). The importance of the PCE is also highlighted by the fact that this index is forecast in the "Summary of Economic Projections" produced by FOMC members four times per year.

Another difference between these two price indices involves the process of data revisions that each index undergoes. Here, the difference between the two indices is quite stark: The CPI is typically subject to revision only because of revised seasonal factors or reporting errors, whereas the PCE is subject to continual, and sometimes large, revisions.[1] These data revisions create a potential problem in using the PCE for inflation targeting purposes: The observed relationship between inflation (for any previous point in time) and the stated target could change over time solely because of revisions to the data. Moreover, by the time more accurate data emerge, monetary policy actions may have already been taken based on the initially observed relationship.

This potential problem is the main focus of this *Research Rap Special Report*. We examine the potential severity of this problem by considering how revisions to past PCE-based inflation data have affected the relationship between inflation and a hypothetical target. We find that initially released data tend to be revised upward and that these revisions can alter the data in several ways that are important with respect to inflation targeting. In particular, revisions can alter the size of deviations from target, even to the extent of changing whether inflation is above or below target, and revisions can alter the record of success seen in hitting the target.

Construction of Real-Time Data

Every observation of both headline and core PCE undergoes a systematic process of revision.[2] As an example, consider the value for core PCE inflation in 2007:Q3, which is currently reported as 2.21 percent. The data for that quarter were first released in 2007:Q4, and the value initially reported at that time was 1.87 percent. This observation continued to be revised in subsequent quarters. In 2008:Q3, year-over-year inflation for 2007:Q3 was revised to 2.01 percent, reflecting the annual revisions that the Bureau of Economic Analysis (BEA) typically makes in July of each year. This observation was revised

[1] The CPI data are in final form when first released. However, revisions can be made in response to a reporting error, which is rare, or due to the updating of seasonal factors. The revisions due to seasonal factors are generally small in magnitude and are made annually in the computation of seasonally adjusted consumer price indices. In contrast, revisions to the PCE are more substantial, since they incorporate new data collected gradually over time.

[2] We use quarterly observations of both headline and core PCE inflation; each observation is computed as year-over-year growth. Note that core inflation differs from headline inflation in that it excludes food and energy prices.

again, to 2.22 percent, in 2009:Q3, reflecting a more extensive benchmark revision, which the BEA makes approximately every five years.

Given this revision process, we choose three perspectives from which to view the data: the initial release, the release following the first annual revision, and the latest known value (the "latest" value here is the value recorded as of 2010:Q1). Thus, for each measure of PCE inflation, core and headline, there are three corresponding sequences of data, and every observation in each sequence (initial, first annual, and latest) represents the data as known after a particular revision. The changes/revisions between any two of these perspectives can be found by subtracting the appropriate sequences. For example, subtracting the initial sequence from the latest sequence measures every observation's revision from initial release to latest release.

Constructing such a data set requires having snapshots of the data as they existed at various points in the past. For example, to obtain the initial release of the 1995:Q1 observation, we would need to be able to view the data exactly as someone standing in 1995:Q2 (when the 1995:Q1 observation was first released) would have viewed them. Similarly, obtaining the 1995:Q1 observation after its first annual revision would require having the data exactly as they existed in 1995:Q3 (since annual revisions are made in July). Data that have been recorded so as to represent exactly what was known at some previous point in time are known as a vintage of data (e.g., we just made reference to the 1995:Q2 and 1995:Q3 vintages of data), and a collection of multiple vintages of data is known as a real-time data set. Constructing our desired sequences of quarterly observations requires having vintages of real-time data, since each observation must be taken from a distinct vintage of data. For this purpose, we use the real-time data set for macroeconomists, compiled by the Philadelphia Fed[3]; we take observations for both core and headline PCE inflation over the period 1995:Q4 – 2009:Q4.

Characterizing the Revisions

Let's first consider some facts about revisions to the inflation data. The revisions to our inflation data are plotted in Figures 1a and 1b. These graphs show that, for both core and headline PCE inflation, the revisions from initial release to first annual release are mostly positive, while the revisions from first annual to latest known have been roughly zero, on average. On net, then, the revisions from initial release to latest known are positive. This means both core and headline PCE inflation have generally been revised upward (see also Croushore, 2008). These results are also evident in the summary statistics shown in Figure 2.

[3] Publicly available on the Philadelphia Fed's website: http://www.philadelphiafed.org/research-and-data/real-time-center/real-time-data/

The mean revision in core PCE inflation from initial to latest is 0.19 percentage point, and this average revision is statistically different from zero. The average revision in core PCE inflation from initial to first annual is 0.16 percentage point, which is significantly different from zero, and the average revision from first annual to latest is 0.04 percentage point, which is not significantly different from zero.[4] Similar results can be seen in the table for headline PCE inflation.

We also test for significant autocorrelation in the revisions (see "Autocorrelation Lag 1" in Figure 2). Significantly positive values here indicate the extent to which a large revision in one period is likely to be followed by a large revision (of the same sign) in the next period. Our tests suggest that there is a positive autocorrelation in our revision sequences.

[4] Significance tests refer to two-tailed tests using significance at the 5 percent level. Similar results are obtained for one-tailed tests and after accounting for autocorrelation in the data.

Figure 1a.

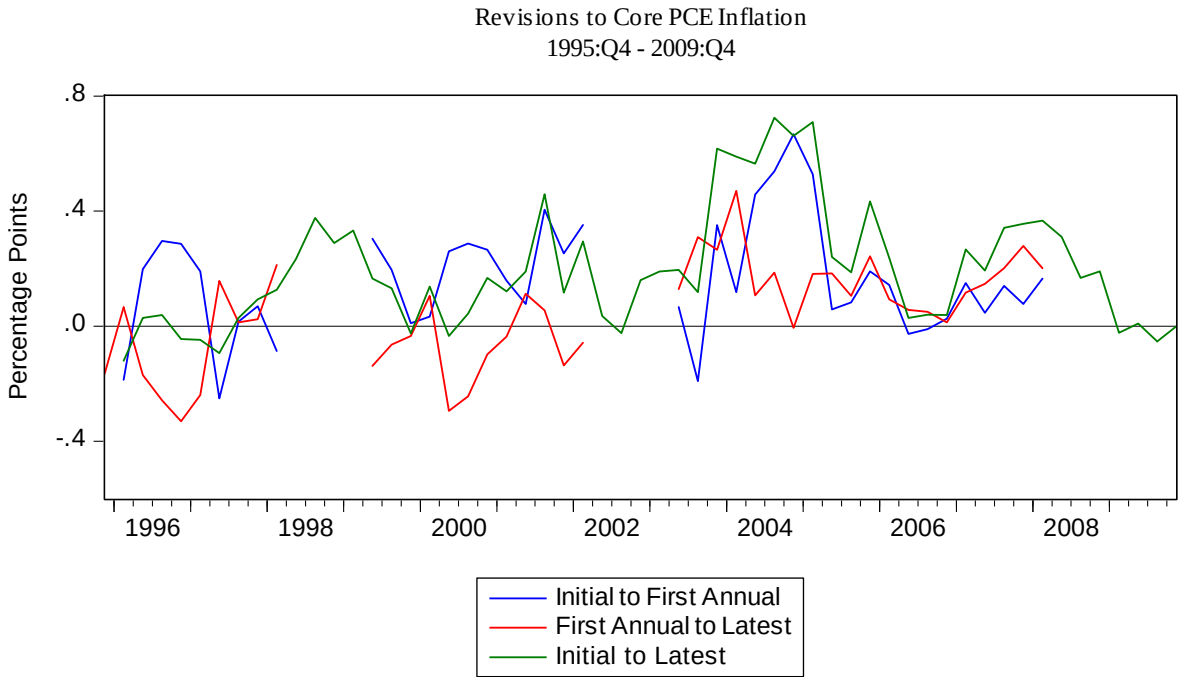

Figure 1b.

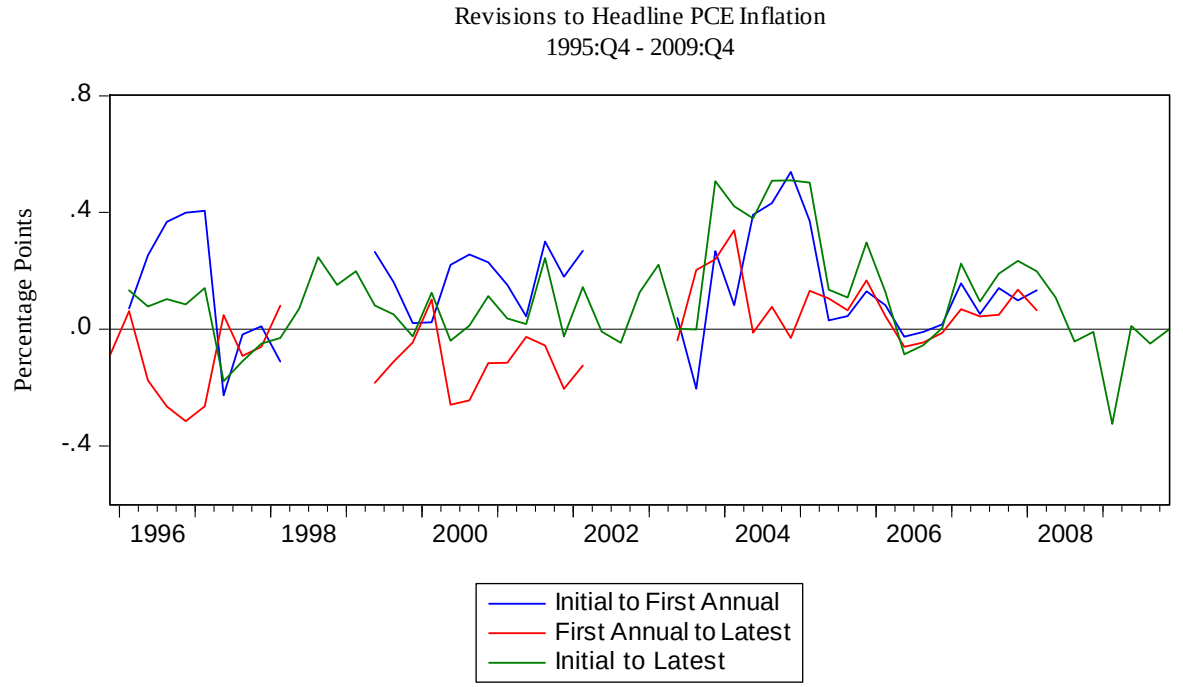

Note: A given observation may be missing in some sequences and present in other sequences. The most common reason for this stems from the BEA's general tendency to not carry out an annual revision in years when a benchmark revision is made. Consequently, when constructing the sequence of data that reflects each observation's value after its first annual revision, in these benchmark revision years we record as missing those observations that would have normally undergone their first annual revision but did not. In addition, these missing values get carried into computations of revisions that involve the first annual sequence, such as the revision from initial release to first annual, and the revision from first annual to latest. In contrast, values will still be recorded for these observations in the initial sequence and latest sequence, and thus the revision from initial to latest will not be missing for these same observations.

Figure 2.

Summary Statistics for Core PCE Inflation

1995:Q4 - 2009:Q4

	Inflation Sequences			Revisions		
	Initial	First Annual	Latest	Initial to First Annual	First Annual to Latest	Initial to Latest
Mean	1.70	1.91	1.90	0.16	0.04	0.19
Median	1.67	1.99	1.87	0.15	0.06	0.17
Maximum	2.45	2.40	2.62	0.67	0.47	0.72
Minimum	0.86	1.11	1.28	-0.25	-0.33	-0.12
Std. Dev.	0.37	0.34	0.35	0.20	0.18	0.21
Observations	56	42	57	41	42	56
Mean Test	34.24*	36.82*	40.35*	5.34*	1.56	6.95*
p-value	0.00	0.00	0.00	0.00	0.13	0.00
Autocorrelation Lag 1	0.78*	0.70*	0.84*	0.40*	0.57*	0.69*
p-value	0.00	0.00	0.00	0.01	0.00	0.00

Summary Statistics for PCE Headline Inflation

1995:Q4 - 2009:Q4

	Inflation Sequences			Revisions		
	Initial	First Annual	Latest	Initial to First Annual	First Annual to Latest	Initial to Latest
Mean	1.97	2.31	2.07	0.15	-0.02	0.10
Median	2.03	2.38	2.04	0.13	-0.03	0.09
Maximum	4.39	3.51	4.35	0.54	0.34	0.51
Minimum	-0.64	0.88	-0.69	-0.23	-0.32	-0.32
Std. Dev.	0.90	0.62	0.94	0.17	0.15	0.17
Observations	56	42	57	41	42	56
Mean Test	16.27*	24.05*	16.71*	5.49*	-0.97	4.56*
p-value	0.00	0.00	0.00	0.00	0.34	0.00
Autocorrelation Lag 1	0.76*	0.59*	0.78*	0.39*	0.54*	0.57*
p-value	0.00	0.00	0.00	0.01	0.00	0.00

NOTE: These tables report summary statistics for core and headline PCE inflation, as measured at three different points in their revision process, as well as for the revisions between these points. "Latest" refers to vintage 2010:Q1. Quarterly observations on inflation are measured on a year over year basis. Test statistics and p-values are reported for the following two tests: 1) a test for zero mean, and 2) a test for zero lag 1 serial correlation. An asterisk denotes significance at the 5 percent significance level.

Of particular concern to policymakers is to what extent an initially known data point may be revised. The correlations between our sequences show to what extent they share a similar pattern of movement; this allows us to judge how similar the general pattern of movement is in the initially released data compared to the revised data (Figure 3).[5] For headline PCE inflation, all of the sequences are very strongly correlated, suggesting very similar movement in the data before and after revision. However, the correlations are noticeably weaker for core PCE inflation, suggesting that the general movement of this series is more likely to be altered after revision.

Figure 3.

Correlations Between Different Sequences of PCE Inflation
1995:Q4 - 2009:Q4

	Initial		First Annual		Latest	
	Core	Headline	Core	Headline	Core	Headline
Initial	1		0.84	0.96	0.78	0.96
First Annual	0.84	0.96	1		0.85	0.97
Latest	0.78	0.96	0.85	0.97	1	

[5] Note that the correlations are, in part, a function of how the observations are computed. Two sequences of rolling year-over-year observations are likely to show a greater correlation than, say, two series of quarter-to-quarter observations.

Revisions and Inflation Targeting

Turning to the importance of revisions to past PCE inflation data in an inflation targeting environment, we consider the data in relation to a hypothetical inflation target of 2.0 percent.[6] Looking at the revisions to inflation data in relation to this hypothetical target, we want to answer three questions: (1) How often do revisions change the material view of inflation relative to target? (2) Do revisions change the general size of past deviations from target? (3) Do revisions alter the historical record of success in hitting the target?

To investigate the first question, we analyze how many times inflation was above or below the 2 percent hypothetical target in each sequence. Because small deviations from the target have little economic significance, we'll say inflation rates of 2.12 percent or higher are above target and inflation rates of 1.88 percent or lower are below target. That is, we use a threshold of 0.12 percentage point; this creates three possible stances for any observation to take relative to target: below target, above target, or roughly on target (i.e., within our tolerance around the target).

If an observation stays within one of these three categories when measured across two different sequences, we say that the two sequences "agree" on that observation. This means the corresponding revision has not materially changed the view of inflation relative to target. In contrast, if the observation does switch between any of these three categories, the two sequences are said to "disagree" about that observation. In this case, the corresponding revision has materially changed the view of inflation, relative to target.

Knowing the extent to which these disagreements occur is important because they can be problematic in regard to policy actions. If policy decisions are based on a known history of data without taking into account future data revisions, then to the extent that these data change over time, the ex post appropriateness of policy decisions may also change. For example, suppose one particular quarter's observation on inflation is 1.75 percent before revision but 2.25 percent after revision. If the central bank ignored the possibility of data revisions, it might have already made a change in its monetary policy stance that seems counter to the later evidence. For this reason, it is clear that such "flips" generated by large-magnitude revisions in the inflation data are potentially problematic, and it is important to know to what extent they may occur in the future by examining the extent to which they have (hypothetically) occurred in the past. This information will also allow policymakers to take data revisions into account.

The results of classifying our inflation observations into agreements and disagreements are shown in Figures 4a and 4b. The figures show a clear contrast between core and headline PCE inflation data:

[6] Some concern may be raised over applying the same target to both core and headline PCE, since these two measures have historically had a gap between them. Applying the same target reflects the expectation that core and overall PCE inflation will converge over time, an expectation stated in the minutes from the FOMC meeting of January 27-28, 2009, "Summary of Economic Projections," p.3.

About 85 to 90 percent of observations are in agreement for headline PCE. For core PCE, however, agreements have occurred for just 60 to 70 percent of observations. For both measures of inflation, the greatest amount of disagreements (lowest proportion of agreements) occurs between the initial release and the first annual release. This suggests that the initial release of the data is subject to meaningful revision and should be viewed as being somewhat uncertain, particularly for core PCE[7].

Figure 4a.

Percentage of Observations in Agreement Between Different Sequences of Inflation, Using Target = 2.0%

Core PCE Inflation (1995:Q4 - 2009:Q4)

	Initial	First Annual	Latest
Initial	NA	61%	70%
First Annual	61%	NA	69%
Latest	70%	69%	NA

Note: Observations within each sequence are put into three categories: below target, above target, or roughly on target (within 0.12 percentage point around target). An "agreement" refers to an observation that has the same category across two different sequences. An observation that changes category across two different sequences is said to be a "disagreement."

[7] Note that the size of the tolerance allowed around zero can affect these results. In general, any change in the size of the tolerance may change some observations from agreements to disagreements as well as some observations from disagreements to agreements. Nevertheless, even when a range of tolerances are inspected, the lowest percentage of agreements generally occurs when the initial sequence is involved, highlighting its propensity to undergo meaningful revision. Also note that when the tolerance is large, headline PCE can have a greater amount of disagreement between sequences than core PCE.

Figure 4b.

Percentage of Observations in Agreement Between Different Sequences of Inflation, Using Target = 2.0%

Headline PCE Inflation (1995:Q4 - 2009:Q4)

	Initial	First Annual	Latest
Initial	NA	85%	88%
First Annual	85%	NA	90%
Latest	88%	90%	NA

NOTE: Observations within each sequence are put into three categories: below target, above target, or roughly on target (within 0.12 percentage point around target). An "agreement" refers to an observation that has the same category across two different sequences. An observation that changes category across two different sequences is said to be a "disagreement."

Comparing the size of the deviations from target between our three sequences, we find that deviations from target are generally larger in revised data, compared with those from initial data. This can be seen by turning our original sequences of observed inflation into sequences of deviations from target and then constructing scatter plots between different pairs of these new sequences (Figures 5a and 5b). A comparison of the deviations from target can be made by looking at the general pattern of how the observations fall in the scatter plot relative to the 45-degree line. If the observations fall tightly along the line, this implies that the deviations are of roughly equal size. In contrast, if the observations generally fall above the line, this implies that the deviations from target using the first annual revision or latest known data are larger than the deviations from target using the initial release. This is, in fact, the general pattern we see in all panels of Figures 5a and 5b. Thus, revisions can alter the historical picture of how large the deviations from target have been.

Figure 5a.

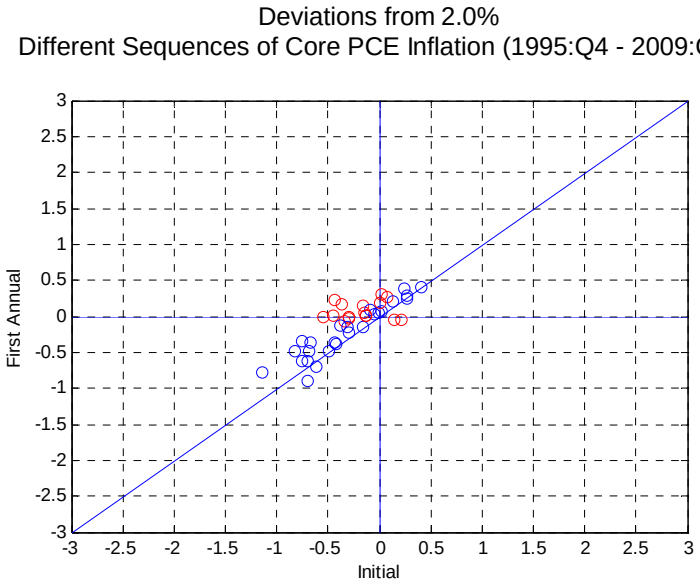

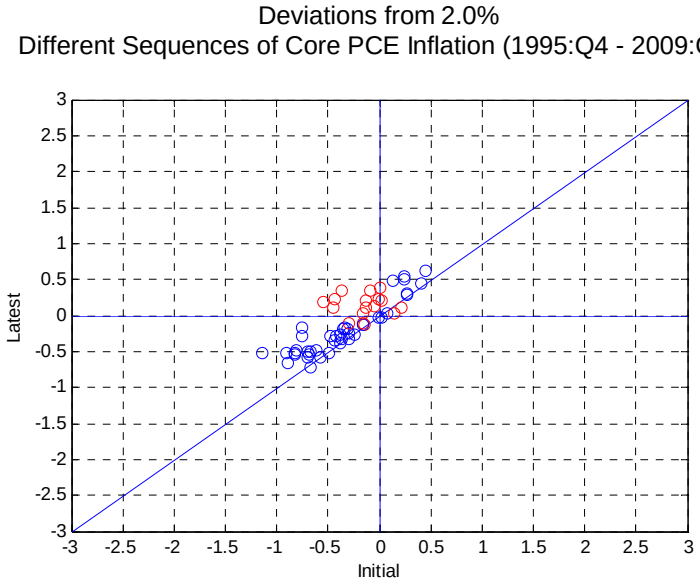

Note: A scatter plot displays a collection of points, with each point having the value of one variable determine its position on the horizontal axis and the value of the other variable determine its position on the vertical axis. An identity line shows points for which the values of the two series are equal (visible in the chart as the 45-degree line). Points in the scatter plot that lie above the identity line represent cases in which the Y variable takes on greater values than the X variable. Points below the identity line show cases in which the value of the X variable is larger than the value of the Y variable. Observations that represent disagreements between the two sequences are shown in red.

Figure 5b.

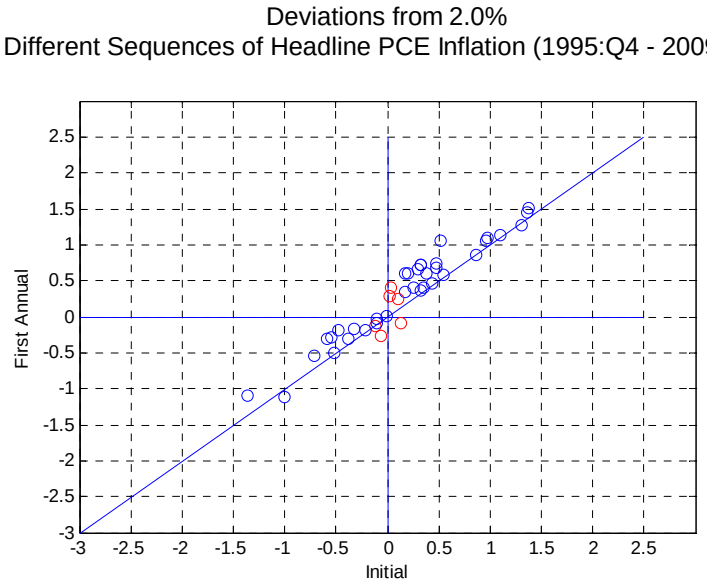

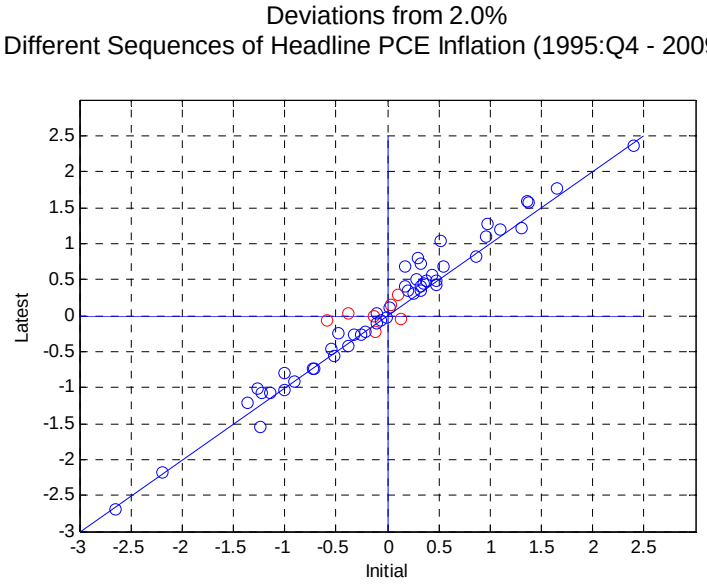

Note: A scatter plot displays a collection of points, with each point having the value of one variable determine its position on the horizontal axis and the value of the other variable determine its position on the vertical axis. An identity line shows points for which the values of the two series are equal (visible in the chart as the 45-degree line). Points in the scatter plot that lie above the identity line represent cases in which the Y variable takes on greater values than the X variable. Points below the identity line show cases in which the value of the X variable is larger than the value of the Y variable. Observations that represent disagreements between the two sequences are shown in red.

We now turn to our third question: Does the general assessment of success in hitting the target change due to revisions? To answer this, we will look at how often inflation can be deemed as hitting our hypothetical target, according to each sequence of data[8]. We will judge an observation to be on target if it is within some tolerance of 2.0 percent; again, we use a tolerance of 0.12 percentage point (Figure 6). We see that the initial sequence of data and the latest known sequence of data show roughly the same proportion of observations as being on target, for both core and headline PCE. However, the first annual sequence of data shows a different picture in both cases: a much larger percentage of on-target observations in the case of core PCE, and a much smaller percentage of on-target observations in the case of headline PCE. This suggests that revisions can change the perceived success of having inflation on target.

Figure 6.

Proportion of PCE Inflation Observations On Target
1995:Q4 - 2009:Q4

	Initial	First Annual	Latest
Core	14%	25%	19%
Headline	14%	9%	18%

NOTE: "On target" means falling within a certain tolerance of the target. We use a tolerance of 0.12 percentage point.

[8] It should be stressed that these data are strictly hypothetical and were computed over a time period in which no official inflation target existed. Therefore, the reader should not extrapolate from these data to infer how much success in hitting the target may occur if an official target did exist. Such a scenario is fundamentally different from the one used to produce these data, and under that scenario, markedly different data could very well be produced. This is only an exercise to give a sense of the potential effect of data revisions, not a comment on policy effectiveness. Note as well that even with an inflation target, the Fed has a dual mandate and so would continue to assess real economic conditions as well as inflation when setting monetary policy.

Conclusions

In this paper we have taken observed PCE-based inflation data from three different points in the revision process and investigated the characteristics of data revisions and potential consequences for inflation targeting. We find that PCE inflation data are typically revised upward between their initial release and their first annual revision, as well as between their initial release and the latest known data.

The consequences of these revisions for inflation targeting are examined in three respects. First, the perceptions of past deviations from target are seen to change after revision. Deviations from target appear to be larger, according to first annual data and latest known data, when compared with the data initially released. Second, we found that revisions can alter the general view of inflation relative to target. These instances are rare for headline PCE but more common for core PCE. Third, we found that the proportion of time that inflation can be considered on target changes depending on from which point in the revision process the data are viewed.

One final note should be considered when interpreting the comparisons made between our three sequences of inflation data. For any fixed quarterly observation, the length of time between when a value is reported for the initial sequence and when that quarter's value is eventually released for the first annual sequence varies, ranging from one quarter to four quarters. Similarly, the length of time between when an observation is reported in the initial sequence and its corresponding value in the latest known sequence also varies from observation to observation, ranging from zero to 56 quarters. Finding a significant revision over a long span of time, while interesting in its own right, may not be of practical concern from a policy perspective. If the conclusion is that policymakers should wait multiple quarters before taking action in order to have significantly different and more accurate data, this may not be practical, since some policy actions must be taken based on the data that are first presented. Thus, perhaps it's unavoidable that there will be greater interest in (though not necessarily greater confidence in) initially released data, relative to revised data, no matter what information one has about future revisions. However, the extent and magnitude of actions taken based on initial data may be tempered if we have a better appreciation of likely future revisions and the uncertainty inherent in initial estimates.

References

Dotsey, Michael. "A Review of Inflation Targeting in Developed Countries," Federal Reserve Bank of Philadelphia *Business Review* (Third Quarter 2006).

Croushore, Dean. "Revisions to PCE Inflation Measures: Implications for Monetary Policy," Federal Reserve Bank of Philadelphia Working Paper 08-8 (2008)