working paper

15 27

Tracking Trend Inflation: Nonseasonally Adjusted Variants of the Median and Trimmed-Mean CPI

Amy Higgins and Randal Verbrugge

FEDERAL RESERVE BANK OF CLEVELAND

Working Paper 15-27 November 2015

Tracking Trend Inflation:
Nonseasonally Adjusted Variants of the Median and Trimmed-Mean CPI

Amy Higgins and Randal Verbrugge

We make five contributions. We demonstrate that extant trimmed-mean and median CPI construction procedures depart from Bureau of Labor Statistics index construction procedures, and that the departures don't make much of a difference. We produce nonseasonally adjusted variants of the trimmed-mean CPI and median CPI, and demonstrate that these are useful real-time estimates of trend inflation; the NSA median CPI outperforms the median CPI, but both SA and NSA variants of the median and the trimmed-mean CPI easily dominate the so-called "core" CPI. We introduce superior ex post measures of trend inflation. We demonstrate that a small amount of time-series averaging reaps large rewards. Finally, we discuss using model-averaging as a new direction for simple and robust trend inflation indicators.

Keywords: inflation measurement, trimmed-mean inflation estimators, seasonal adjustment, trend inflation, time averaging.

JEL classification: E31, E32, E37.

Suggested citation: Higgins, Amy, and Randal Verbrugge, 2015. "Tracking Trend Inflation: Nonseasonally Adjusted Variants of the Median and Trimmed-Mean CPI," Federal Reserve Bank of Cleveland, working paper no. 15-27.

Amy Higgins is at the Federal Reserve Bank of Boston. Randal Verbrugge is at the Federal Reserve Bank of Cleveland and can be reached at randal.verbrugge@clev.frb.org. The authors thank Mike Bryan, Domenico Giannone, Ed Knotek, and participants at the Cleveland Fed brown bag seminar.

1. Introduction

Since controlling inflation is a central monetary policy goal, monetary policymakers focus intently on inflation signals. But they face a major difficulty: inflation data contains a lot of transitory shocks. Responding to these shocks would be a bad idea, since this would translate into policy swings and reversals, and introduce uncertainty and volatility into the economy. Instead, policymakers attempt to respond to sustained movements in inflation (i.e., to underlying "trend" inflation). But the presence of the transitory "noise" in inflation data makes it difficult to detect early warnings of such sustained movements. Thus policymakers exert much effort trying to discern the underlying trend in the midst of the noisy inflation data. Numerous statistical approaches to this important problem have been proposed. But an approach first proposed in the early 1990s has proven to be one of the most useful, and has withstood the test of time: the Median CPI, introduced in Bryan and Cecchetti (1993)[1], and produced by the Federal Reserve Bank of Cleveland.

Our paper studies the median CPI (and its close cousin, the trimmed mean CPI), and makes five chief contributions. First, we note that the current production procedures for the Cleveland Fed median CPI and trimmed mean CPI depart from Bureau of Labor Statistics index production methods, and we explore whether the departures make much difference. (They don't.) Second, we introduce two new real-time trend inflation indicators, namely non-seasonally-adjusted variants of the median CPI and the trimmed mean CPI, and demonstrate that the new measures are useful and provide reliable signals of trend inflation. Third, we provide a comparison of alternative trend inflation indicators, using both new and established metrics. We hope that our results will, at long last, convince even the most diehard fans of the so-called "core CPI" that it has been definitely superseded.[2] Fourth, we demonstrate that a small amount of smoothing, via time series averaging, reaps large benefits in terms of removing noise from trend inflation indicators. Finally, we discuss using model-averaging as a new direction for simple and robust trend inflation indicators.

[1] See Bryan and Cecchetti (1993, 1994, and 1995), Cecchetti (1997), and Bryan, Cecchetti and Wiggins (1997).

[2] Detmeister (2011) is the paper most closely related to this study, and it reaches many similar conclusions regarding the PCE price index. On the topic of "core" measures, see also Rich and Steindel (2007), who state: "…we cannot identify a compelling analytical reason, on either an ex ante or ex post basis, to concentrate attention on a measure of inflation that excludes food and energy prices." Wynne (2008) provides a discussion about core inflation – a concept that will be termed "trend inflation" in this paper – and its measurement.

2. Median and Trimmed Mean CPIs: Introduction

Since at least the 1970s, the CPI has been recognized to be susceptible to sharp movements as a result of large changes in a few item indexes. The short-term volatility in the inflation measure made it difficult to discern persistent inflation movements. In 1975, Robert Gordon proposed an alternative CPI measure, namely the CPI less food and energy, as a means of extracting transitory noise from the CPI. The rationale was that food and energy components often experience sharp price movements that soon dissipate; therefore, excluding these components from the inflation data will effectively remove some noise that might otherwise mask the signal. This alternative CPI measure became popular, and it remains one of the CPI indexes produced by the Bureau of Labor Statistics (BLS). (While this measure is usually termed the “core CPI”, we regard this as a misnomer because, as we demonstrate below, it can depart substantially from trend inflation. The BLS refers to this measure as CPI-U less food and energy.)

Bryan and Pike (1991) noted some conceptual and statistical problems associated with the core CPI. First, on a conceptual level, this exclusion strategy assumes that essentially all of the large transitory shocks come from food and energy; but this assumption is false. Second, the strategy implicitly assumes that food and energy cannot have persistent relative price trends. However, it is easy to come up with plausible scenarios under which this condition would be violated, particularly for energy prices. Under this condition, leaving these components out will give rise to a signal that is biased. Third, in point of fact, some of the components excluded from this measure are actually quite stable, while some of the components that remain in the index are subject to big transitory shocks fairly frequently.

For these reasons, Bryan and Cecchetti (1993) proposed a different approach, namely using *bounded influence estimators* such as a median or trimmed mean, to measure each month’s central tendency of the cross section of price index changes. A bounded influence estimator is a statistic that, by design, limits the influence of unusual observations. In the CPI context, as we note below, such estimators make a month-to-month judgment about which components to exclude, based upon the realized price index movements that month. In particular, on a month by month basis, all those index changes that are unusually large or small compared to the bulk of the index changes are excluded, regardless of the sector in which they happen.[3]

Why does a trimmed mean or median succeed in filtering out noise? The basic idea is a statistical one, and may be illustrated by example. Suppose that one is interested in getting a reliable estimate of the average rent of a rental unit in Cleveland. Suppose that a random sample of 25 units was taken, and upon reviewing the data, it is discovered that one of the units had a stated monthly rental value of $22,500.[4] Clearly that is an extremely unusual unit in a city where the average rent is closer to $780. That $22,500 rental price is an *outlier*, a rent that is – in this case – far larger than the average rent in Cleveland. In this example, that one huge rent would, by itself, raise the *sample* average by well more than $500, leading to an erroneous conclusion about the average rent in Cleveland. Conversely, one would get a far more accurate reading on the typical rent in Cleveland if one decided to throw out, or *trim*, the highest and lowest rents in the sample, and took as one's estimate the average of the remaining 23 units. For the same reason, the *median* of a sample can sometimes be a more reliable indicator of the mean of the underlying distribution than the average of the sample. The trimmed mean and the median are termed bounded influence estimators because such estimators are designed to keep a small number of unusual data points from having an improperly large influence on the conclusions that are drawn. A large body of research in the statistics literature demonstrates that such estimators are preferable when the underlying distribution tends to generate outliers fairly often.

Outliers are often seen in price index changes, suggesting that limited influence estimators might be useful. In the CPI context, the median and the trimmed mean will remove (or trim) all the monthly price movements in underlying price indexes that are atypically large or atypically small, resulting in an estimate of a CPI change that is based only upon the typically-sized index movements that month. (We provided explanations of the construction of these measures in the next section.) The median CPI and trimmed mean CPI thus provide more accurate readings on the general movement of prices that month.

The theory is sound; but how well have the median CPI and trimmed mean CPI performed in practice? Since the pioneering work of Bryan and Pike (1991) and Bryan and

[3] Below, we provide more detail about how these estimators are computed. See appendix 1 for a descriptive table that shows some volatile components that are often excluded in the Trimmed mean CPI, but which are included in the CPI less food and energy. Bakhshi and Yates (1999) discuss the theoretical case for trimmed mean measures.
[4] Our internet search in May 2015 yielded a few units in Cleveland at this rent or higher.

Cecchetti (1993), both the median CPI and trimmed mean CPI have been the subject of intense scrutiny. The character of the favorable results obtained in the original studies has largely held up over the years: the majority of studies conclude that both the median CPI and trimmed mean CPI are more accurate indicators of overall inflation than is the core CPI. Most of the research studies have focused upon forecasting ability. Bryan and Cecchetti theselves conducted numerous follow-up studies, sometimes with coauthors (see, e.g., Bryan and Cecchetti 1994, and Bryan, Cecchetti and Wiggins 1997), and these researchers, along with others, developed similar measures for other countries as well (see, e.g., Bryan and Cecchetti 1999 and Aucremanne 2000). Results have been overwhelmingly positive. Clark (2001) found that at a two year forecasting horizon, the median and trimmed mean CPI outperformed the core CPI. Cogley (2002) demonstrated that both the median CPI and the trimmed mean CPI outperformed the core CPI in forecasting future CPI-U inflation as several forecast horizons, as did Smith (2004) and Rich and Steindel (2007). Brischetto and Richards (2007) examined trimmed means (and medians) from four economies, and found that these measures outperformed the others in tracking the inflation trend and in forecasting. Meyer and Zaman (2013) investigated the use of the Median CPI in quarterly macroeconomic forecasting models and found that the inclusion of the Median CPI leads to improved forecasting accuracy of Headline and Core CPI inflation as well as the Federal Funds Rate. Meyer and Venkatu (2014) found that the Median CPI and Trimmed mean CPI outperform both the CPI-U and the core CPI in out-of-sample forecasting.

However, one notable study, Crone et al. (2013), found – in seeming contrast to the works cited above – that the traditional core CPI outperformed the Median CPI at forecasting long-run CPI-U in at the six-month horizon (the measures were neck and neck at 12 months out, and the Median was better at longer horizons). This study was distinctive in that it focused upon prediction of inflation in the CPI-U between now and 6, 12, 24, and 36 months from now, using as the predictor the *12-month* change in the respective index.[5] In contrast, the studies above used either the 1-month or the 3-month change in the respective index as the predictor. This choice is consequential; as the figures below indicate, and as our results regarding time-averaging

[5] Meyer and Pasaogullari (2010) also used 12-month changes in core inflation measures as a predictor, and – depending upon the time period examined – obtained some results that are similar in character. They also compared how various core inflation indicators performed relative to some standard forecasting models, and compared regression-based approaches to naïve approaches which simply used the actual past 12-month change in the core inflation measure as the predictor. Taken as a whole, their results were generally quite supportive of the usefulness of the Median CPI and Trimmed mean CPI for forecasting purposes.

demonstrate, averaging over 12 months will smooth out many of the temporary fluctuations in the CPI-U that the Median CPI is designed to remove. Yet this time-averaging also makes detecting early warning signs of changes in inflation more difficult.[6] Still, this study conveyed a novel insight: for forecasting purposes, it is worth considering using time-averaged series as a predictor, and in that context, the median CPI might not outperform all other contenders at shorter horizons.

Thus, the median CPI is a useful forecasting tool, although perhaps not best in all circumstances. But is forecasting the CPI-U the *sina qua non* of a trend inflation indicator? We say, no.[7] After all, if forecasting inflation is the researcher's main purpose, he or she should adopt the best practices from the forecasting literature. For example, forecast averaging – averaging the forecasts from a suite of models – is often superior to using a single model. And there is no reason to restrict attention to forecasts based upon only one predictor, either; for example, Meyer and Zaman (2013) used many variables simultaneously, and a Phillips curve model is often used for inflation forecasting (see, e.g., Ashley and Verbrugge 2015). Further, one might also consider using several trend inflation measures simultaneously, and making use of survey measures such as the Survey of Professional Forecasters to help anchor the long-run prediction. (The importance of adjusting for changes in trend inflation in inflation forecasting is discussed in Clark and Bednar, 2014.) In any case, the CPI-U itself is highly volatile, so much so that in many cases, a forecast miss would be exactly what was wanted.

What, then, should a good indicator of trend inflation do? The primary purpose of a trend inflation indicator is to, in a timely manner, filter out noise and transitory influences in the current inflation signal, and accurately reflect the current actual trend in inflation.[8] How does one assess that? We outline below our somewhat novel means of operationalizing this criteria. In prior work, a centered 36-month moving average has often been used as the measure of the trend. We here propose two trend estimates that are superior. But before we show how to assess the performance of the Median CPI, we discuss the mechanics of constructing it. We provide these details because we will highlight current FRBC practices that depart from BLS index

[6] Technically speaking, a 12-month change alters the *phase* of the underlying series, so that turning points only become evident after a lag of several months or more.
[7] Detmeister (2011) also makes this point.
[8] An analogy, originally due to Michael Bryan, is seasonal adjustment; the goal of seasonal adjustment is to remove movements from an index that are, in a statistical sense, unrelated to trend movements, so as to better expose that trend.

construction practices, and we wish to explore whether these current FRBC practices can be improved.

3. Constructing the Median and Trimmed mean CPI

To understand how a Median or Trimmed mean CPI is constructed, it is important to "look under the hood" and examine how the CPI-U itself is constructed. There are eight major categories used to calculate the CPI-U. These categories are: Food and Beverage, Housing, Apparel, Transportation, Medical Care, Recreation, Education and Communication, and Other Goods and Services. Each of these eight categories can have up to five levels of sub aggregations. At the establishment level, the BLS measures price *changes* of particular items, and then constructs index *changes* using weighted averages of these price changes. For example, the BLS constructs its "Gasoline-Unleaded regular" Index change for Chicago in March by taking a weighted average of all the unleaded gasoline price changes in March that it measured from gas stations in its Chicago sample. Starting from index changes like this, the BLS builds up to more aggregated index changes, like "Energy Services in the Midwest" or even the CPI-U itself, by taking weighted averages of lower-level index changes. (As explained below, the weight of a particular index corresponds to its share in the average consumer's budget.) The core CPI change consists of a weighted average of changes in all the indexes except for food and energy indexes. For this index, the weights of the remaining items have to be increased, so that the sum of the weight over the remaining items remains the same as in the CPI-U.

The median CPI and trimmed mean CPI are in the same family of aggregate indexes, except for the choice of the underlying index changes, and the weights used. Under current production methods, the basic building blocks for both measures are a particular set of disaggregated index changes. Forty-five disaggregated goods and services index changes are used, corresponding to a combination of various sub-aggregation levels. (Some examples include "Food away from home," "Owners' Equivalent Rent in the Midwest," and "Footwear.")[9] The set

[9] With four exceptions, all of the 45 components are national indexes. The four exceptions are the four regional Owners' Equivalent Rent (OER) components. The use of four regional OER components, rather than the national OER index, owes to the research of Brischetto and Richards (2007). The authors noted that disaggregating the national OER component into 4 regional OER components would reduce the influence of OER on the Median CPI and Trimmed mean CPI. The Relative weight for the National OER component is historically around 20-25 percent. With such a large weight on one component, National OER is very likely to be the median component in the Median CPI series.

of indexes is essentially comprehensive: the change in the CPI-U is simply the weighted average of the changes in these 45 indexes, where the weights used in the averaging ultimately derive from expenditure weights in the average U.S. consumption basket. (We say “essentially” comprehensive because there is a tiny “non-sampled” index in the CPI that is ignored by current FRBC production methods, as we explain later.)

The trimmed mean CPI is also a weighted average of index changes. However, the trimmed mean CPI uses a smaller set of indexes, since one excludes (or “trims”) index changes that are far from the typical index change that month – in the case of the Trimmed mean CPI, the top and bottom 8% are excluded – and then adjusts the index aggregation weights accordingly so that they still sum to one. The Median CPI is an extreme version of this, in which only the median index change is retained. Bryan and Cecchetti (1993) explain how both statistics are generated as *weighted* statistics; this implies that the resulting trimmed mean and median are both consistent with the official CPI aggregation weights, yielding statistics that are fully consistent and, in principle, unbiased.

Weighted trimmed means may be constructed using the following procedure. Each month, download the changes in the 45 indexes from the BLS. Then rank the 45 index changes in order, from greatest to smallest. Also, each month, calculate the correct updated aggregation weight for each index, as described in detail in Appendix 2. Then rescale the weights so that they sum to one.[10] Now, in a spreadsheet, place the index changes in order in one column, and place the corresponding weights in the second column. In a third column, compute a cumulative sum of all the weights: thus, for example, in the third row, the third column contains the sum of the weights of the first three entries.

For the median CPI, find the two entries in the third column such that the first number is less than 0.5, while the number immediately below it is greater than 0.5. The index change corresponding to the first number is the weighted median.[11]

Computing a trimmed mean is a little more involved, but only a little. There are some additional steps. Suppose one wanted to compute a 16% trimmed mean. Start with the spreadsheet that was already constructed. For the lower trim, find the two entries in the third column such that the first number is less than 0.08, while the number immediately below it is

[10] BLS methods for constructing special indexes such as the core CPI are slightly different and somewhat involved. See appendix 3 for a technical explanation.
[11] In other contexts, “the” weighted median is estimated as a weighted average of these two numbers.

greater than 0.08. Highlight the second number; label this *l*. Now, for the upper trim, find the two entries in the third column such that the first number is less than 0.92, while the number immediately below it is greater than 0.92. Highlight the first number; label this *h*. Now add up all the weights corresponding to all the indexes in between *l* and *h*, inclusive; call this sum *tms* (for "trimmed mean sum"). In a fourth column, construct new weights: each entry corresponds to the original weight in column 2, divided by (1 – *tms*). The new weights should sum to one. The trimmed mean is simply the weighted sum of all the index changes in between *l* and *h*, constructed using the new weights.

Clearly the correct weights are central to the procedure. We turn next to that issue.

4. Weights and Seasonality

The weights (or, in BLS terminology, "relative importance values") for each component index are crucially important when calculating inflation measures. Particular components make up more or less of overall consumer spending, and the weights associated to each component reflect the typical share of consumer expenditures for that particular goods or services category. The BLS publishes component weights every December. However, monthly weights, necessary to produce the Median CPI and the Trimmed mean CPI, are not published. As adequate documentation from the BLS was not available at the time of the inception of the Median CPI, Bryan and Cecchetti devised a methodology to calculate monthly expenditure weights for each component.

Current FRBC procedures depart from BLS production methods in three ways. First, the FRBC formula used to calculate the monthly weights, while quite close to that of the BLS,[12] is effectively lagged by one month (Appendices 2 and 3 provide details concerning the different methodologies used to construct the monthly weights).

Second, the BLS produces a non-sampled index (and an associated weight) – call it index number 46 – that is ignored by the FRBC method. This item is "Unsampled new and used vehicles." In fact, there are a number of unsampled items in the CPI, but the presence of the others is already incorporated in the 45 published indexes used by the FRBC method. See Appendix 6 for details regarding the weight for this 46th component.

[12] http://www.bls.gov/cpi/cpiriar.htm

We investigated whether either of these first two departures was consequential (for details, see Appendices 4 and 6.). Our investigation suggests that neither makes much difference in practice: if BLS methods for weights and for index 46 had been used instead, the historical differences in the Median CPI and Trimmed mean CPI would have been trivial.

The third difference merits more discussion. The FRBC method exclusively uses seasonally-adjusted data. But to produce the CPI, the BLS constructs weights exclusively using non-seasonally-adjusted (NSA) data. While the BLS produces both NSA indexes and seasonally-adjusted (SA) indexes, its SA indexes are constructed by seasonally adjusting some of the underlying NSA indexes it already produces. (The seasonal adjustment practices of the BLS are somewhat involved; for details, see "Seasonal Adjustment in the CPI" at http://www.bls.gov/cpi/cpisapage.htm.) One might imagine that adhering to BLS practice, and using only NSA data to produce weights (and then indexes), might be the best thing to do. After all, how could one improve matters by using different weights than those actually used to produce the CPI-U?

Constructing a NSA series has another distinct advantage: an NSA median index and an NSA trimmed mean index will rarely need to be revised. Conversely, the SA indexes produced by the BLS are revised at least once a year, necessitating revisions of the median CPI and trimmed mean CPI. In particular, every February, the BLS reports changes and additions made to the CPI data series and to its seasonal adjustment procedures for the year. The size of estimated seasonal factors often changes, and in addition, the BLS makes an annual determination as to which of the underlying indexes are going to be seasonally adjusted. Thus, it is even possible that in a given year, the published data for some components of the CPI may change. For example, a SA data series may be made available for a data series which was previously only reported as NSA, or a previous SA data series may no longer be available and in the future may only be reported as a NSA data series. Revisions to the CPI can be sizeable (see Knotek and Zaman, 2015). These kinds of changes usually result in revisions to the median and trimmed mean CPI series, although the revisions to these series are generally smaller.

On the other hand, it is possible that using NSA data could be harmful in this context, for two reasons. First, since the goal is to filter out noise, it is possible that the use of SA indexes and constructing SA weights might filter out seasonal noise that would otherwise contaminate the signal. (There is some evidence, for example, that using SA data can help forecast accuracy, although this is not always the case; see, e.g., Bagshaw (1985), Plosser (1979), or Bell and Sotiris (2010).) Second, as noted by Brischetto and Richards (2007) and Roberts (2005), it is possible that trimming and seasonality would interact in a manner that would lead to downward bias in a trimmed mean CPI. The argument is that strongly seasonal items might display, only in one month of the year (say), large price changes in one direction or the other. These might well be trimmed. But this trimming would systematically remove the largest price movement in one direction; meanwhile, smaller price movements in the other direction would not be trimmed.

These considerations led us to investigate the benefits of estimating a NSA Median and a NSA Trimmed mean CPI series – then seasonally adjusting the series as necessary – as alternative core inflation measures.

5. Performance Criteria

In this study, we investigate the median and trimmed mean CPI, and their NSA variants, compared to several other prominent trend inflation indicators that have been proposed. While we constructed several alternative NSA median and trimmed mean CPI series starting in January 1983 and ending in March 2015, we will focus on a particular one, which adheres to BLS methods (see Appendix 6 for some additional details and results concerning the various alternative methodologies explored). To evaluate the usefulness of each trend inflation indicator, we use a set of criteria similar to those of Wynne (2008), Clark (2001), and Rich and Steindel (2007).[13]

1. **Transparency of construction**. A core inflation estimate should be straightforward and easily understood by the public and policymakers, and it should be verifiable – agents independent of the producer of the estimate should be able to replicate the series.[14]
2. **Timely**. A core inflation estimate should both be computable with little delay. But also, it should be able to quickly detect turning points, so that it can provide an up-to-date signal of the current inflation trend.

[13] Mick Silver (2007) discussed a wide range of similar criteria. Note that we do *not* include forecasting ability.
[14] Detmeister (2011) notes that there a variety of roles which trend inflation indicators might perform. While one might consider making use of different indicators for different roles, it may be preferable, for clarity of communication, to select a single indicator.

3. **Similarity of means**. The mean of the constructed core inflation estimate should be roughly the same as that of the underlying inflation series over long periods; in other words, the core inflation measure should not have long-run bias.
4. **Not volatile**. A less volatile core inflation series is preferable to one with more volatility.
5. **Historical ability to track the actual underlying inflation trend**.[15] A good trend inflation indicator should, in retrospect, be seen to have closely tracked the actual underlying trend in inflation, and also (see criterion 2) to have quickly detected turning points in that trend. (We discuss how we measure the actual underlying trend below.) If a trend inflation indicator closely tracks the actual trend, then it follows that the indicator will automatically be useful in forecasting. But we view forecasting ability *per se* as less important.

There are two other considerations worth mentioning. First, before a trend inflation indicator is adopted, it should have a track record. This prevents the sort of data-mining problem emphasized by Rich and Steindel (2007) – that a particular indicator could be tweaked so that, at a given point in time, it matches the historical record, but subsequently it underperforms. Second, an indicator which is not revised is preferable to one which is subject to revision. If the estimate of trend inflation changes a lot when new observations become available, this will both limit its usefulness, and will potentially unsettle the public.

In empirical work, occasionally one observes the use of a trend line derived from a two-sided filter, such as the Hodrick-Prescott filter, as an approximation to the (unknown) trend inflation indicator that market participants were actually using, or as a measure of trend inflation for econometric modeling.[16] These uses are misguided. Two-sided filters distort econometric inference (see Ashley and Verbrugge, 2009); and their use as a stand-in trend inflation indicator is conceptually misguided since, by design, information from the future is required to compute them in the first place. Market participants would not use the Hodrick-Prescott filter in a one-sided manner (i.e. in real time), owing to severe end-sample problems.

---

[15] Blinder (1997) criticizes the use of trend estimates which treat the past and future data symmetrically, and cannot be used in real time. In our view, this criticism is misplaced. The only way we can evaluate a trend inflation estimate is by looking at the data we have observed. If a trend inflation indicator has closely tracked trend inflation in the past, we have some confidence that it will continue to do so in the future. And in a given month, the chief goal of a trend inflation indicator is to discern what is trend versus what is noise.

[16] We distinguish these uses of a Hodrick-Prescott trend line from a distinct one, namely using it as an ex-post indicator of trend inflation that could serve as a measurement goal – see Section 6.

6. The Measurement Goal: Actual Trend Inflation (Ex-Post)

The main purpose of a trend inflation indicator is to give an accurate, real-time reading on actual trend movements in inflation. But this is about the same thing as saying that the indicator should, in real time, accurately reflect the *low-frequency* movements in the CPI-U. Implicitly, the actual trend in inflation consists only of those movements in the CPI-U which are not transitory (i.e., which are not *high-frequency*). In mathematics, we might write this as

$$\pi_t = \pi_t^T + \pi_t^N$$

where π_t is actual CPI-U inflation in month *t*, π_t^T is the actual trend in inflation, and π_t^N is transitory noise in inflation. Of course, this decomposition is not knowable in month *t*: only after the fact can one really know if a given movement turned out to be transitory. The chief purpose of a real-time inflation trend indicator is to provide a reliable signal of the currently-unknown trend π_t^T.

But even ex-post, π_t^T is still not *observable* in the data; there are many ways to perform a historical decomposition of π_t into two parts. Thus, there may be disagreement about the best way to obtain a good estimate $\hat{\pi}_t^T$ of π_t^T from the data. But to assess the relative performance of different trend inflation indicators , selecting some estimate $\hat{\pi}_t^T$ is unavoidable: one must choose some $\hat{\pi}_t^T$ and treat this as actual trend inflation, the benchmark against which the contenders are judged. While there have been alternative $\hat{\pi}_t^T$ choices in the literature, most previous research has used a centered 36-month moving average (MA(36)) of π_t (see, e.g., Bryan, Cecchetti, and Wiggins 1997).

But there are at least two drawbacks to using an MA(36) of π_t as an estimate of π_t^T. The first drawback is conceptual, and relates to the purpose of the real-time trend inflation indicator. An MA(36) explicitly includes inflation that is 18 months in the future. But we view 18 months in the future as a little too far away. Why? Because, according to most economists, monetary policy influences inflation at this horizon. As a result, an MA(36) "trend inflation" signal will sometimes implicitly contain "the *change* in trend inflation induced by the monetary policy reaction to the trend inflation signal it received." But what is wanted by the monetary policymaker is a measure of trend inflation that will inform that reaction. Thus, ideally what is wanted is an estimate of what inflation *would have done* absent a monetary policy response. It may be difficult to estimate this; but surely an MA(30), which has a shorter 15-month-ahead window, is at least better.[17] The second drawback is more pragmatic: this particular 36-month measure allows too much high-frequency noise to become part of the estimator.[18] In Appendix 5, we provide some graphical evidence for this excess-noisiness-in-trend-estimate feature, and then prove its presence formally.

In this study, we propose two alternative estimates of $\hat{\pi}_t^T$ which are superior. For the first, "BK30", we use the Baxter-King (1999) frequency-domain filter.[19] We use this filter to remove all the historical movements in the CPI-U which dissipated over the course of 30 months or less, leaving only the more persistent historical movements in the CPI-U. How does frequency filtering work? An analogy comes from the world of audio, and in particular, loudspeakers. High quality speaker units typically consist of three speakers in one enclosure: a woofer, a tweeter, and a midrange speaker. The woofer is a large speaker, and is dedicated to low audio frequency sounds, such as all frequencies that are 150Hz or lower. The tweeter is a small speaker, and is dedicated to high audio frequency sounds, such as all frequencies that are 2,000 Hz or higher. The midrange speaker is responsible for the frequencies in between these two cutoff frequencies. But to dedicate each speaker to a particular frequency range, the speaker system has to break up, or partition, the incoming audio signal into these three parts. This task is accomplished by a "crossover" unit. In like manner, an economic time series can be broken up or partitioned by frequency: partitioned, say, into the low frequencies (i.e., all movements or cycles with periods of 30 months or more) and high frequencies (i.e., all movements or cycles with periods less than 30 months).

An advantage of defining a trend by frequencies is that the filter will ignore even very large movements in the CPI-U, if those movements are transitory. In contrast, a moving average

[17] Some authors select an even shorter window. Giannone and Matheson (2007) characterize their estimation target $\hat{\pi}_t^T$ as a centered 24-month moving average of annual inflation; Bryan and Cecchetti (2001) used a centered 24-month moving average as their target in the Brazilian context, as do Brischetto and Richards (2007) in their multi-country study. Others use an even longer window. Meyer and Venkatu (2014) suggest using an annualized percent change in headline CPI over the future 36 months as an appropriate measure, stating that a 24-month ahead annualized percent change in CPI-U would not be long enough for relative price shocks to level out. (Of course, their horizon partly reflects their focus on forecasting.) The implicit window of Rich and Steindel (2007) is longer than a decade; these authors remove all movements in inflation that are shorter than 8 years (32 quarters).
[18] On these points, see also Brischetto and Richards (2007), Giannone and Matheson (2007), and Wynne (2008).
[19] Using the Ashley-Verbrugge (2009) frequency-domain filter (in a two-sided manner) yields a series that is nearly identical. We are not the first to propose using such methods; earlier, Rich and Steindel (2007) and Giannone and Matheson (2007) both made use of frequency-domain methods to isolate the trend.

must allow those movements to have full weight – although longer moving averages will, admittedly, place relatively little weight on movements which are transitory.

As our second measure of trend inflation, we improve upon a centered 30-month moving average by using two passes of otherwise standard moving average filters. In particular, we first apply a centered 24-month moving average to generate an intermediate series, and then apply a centered 12-month moving average to generate a final series. While this "MA(24,12)" series is similar to a 30-month moving average, it is superior in that it removes almost all high-frequency noise.[20] We think that both of our trend inflation measures strike the right balance between the amount of information incorporated and the amount of transitory noise allowed. Neither includes information from too far ahead—in which case, the signal does not just include the natural dynamics of inflation, but also includes the response of inflation to monetary policy—and neither allows in too much transitory noise, which would lead to an unreliable "trend" estimate.[21] Thus, we believe that our filters well captures the low-frequency part of inflation that past research has attempted to isolate, while avoiding distortion of the signal arising from monetary policy response

We plot both of our measures in Figure 1, along with annualized monthly CPI-U inflation. Both appear to be reasonable estimates of trend inflation. (In Garciga and Verbrugge (2015), we consider more sophisticated trend inflation estimates.)

---

[20] This superior removal of high-frequency noise can be seen graphically: our MA(24,12) series is much smoother than an MA30. More formally, it can also be seen by looking at the frequency response of the filter. See Appendix 5.
[21] While the MA(24,12) filter admittedly uses information from 18 months ahead (rather than the 15 months ahead in MA(30)), this information receives miniscule weight: about 0.0035, or 0.3%. In an ordinary MA(36), the weight is roughly ten times larger. Conversely, relative to an MA(30), the MA(24,12) filter places more weight on information closer to the current month.

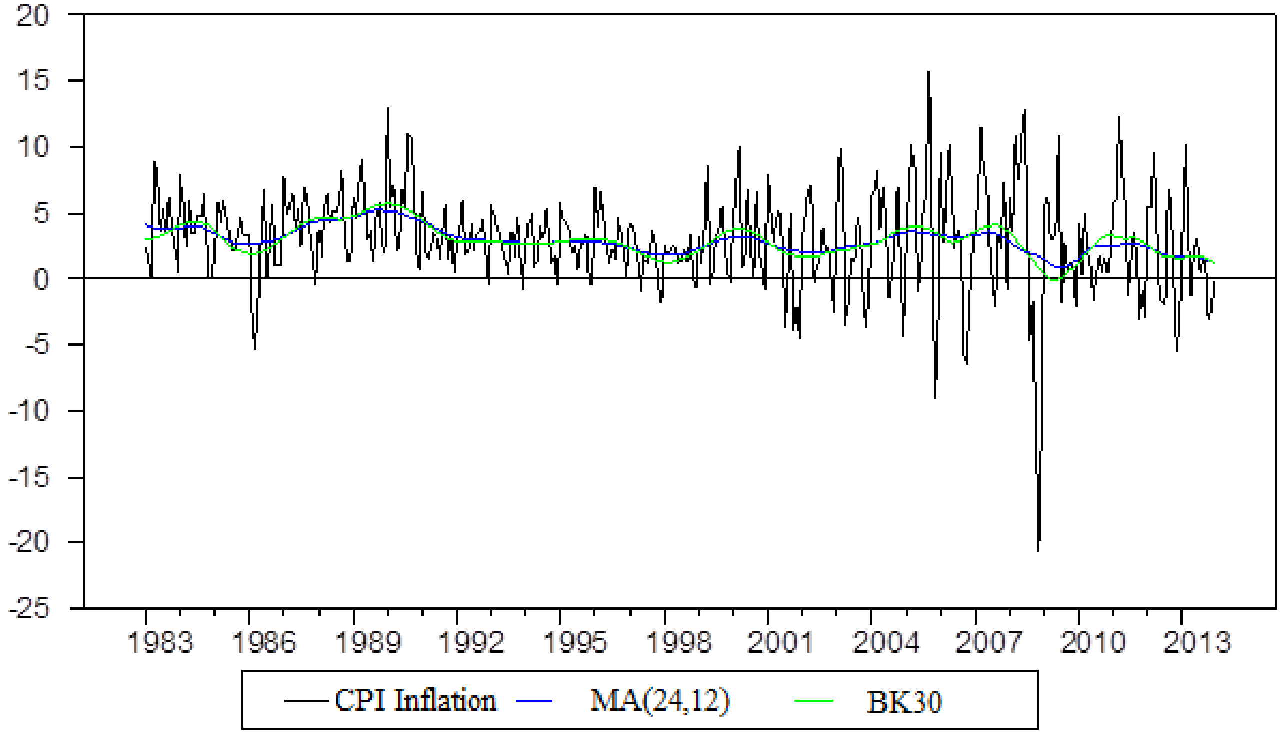
20
15
10
5
0
-5
-10
-15
-20
-25
1983
1986
1989
1992
1995
1998
2001
2004
2007
2010
2013
CPI Inflation
MA(24,12)
BK30

Figure 1

7. Analysis

Criterion 1 is a pre-requisite and rules out the consideration of a number of sophisticated trend inflation indicators that have been proposed in the literature, such as those based upon structural vector autoregressions, factor analysis, or Kalman filtering.[22] We view such sophisticated indicators as useful, but not well-suited for communicating monetary policy decisions to the public. For our analysis, we evaluate the following eight potential trend inflation indicators[23] against our trend inflation measure.

1. 12 Month CPI-U (12-month change in CPI-U)
2. CPI-U less-food-and-energy ("Core CPI")
3. Median CPI
4. Trimmed Mean CPI
5. NSA Median CPI
6. NSA Trimmed Mean CPI
7. SA-NSA Median CPI
8. SA-NSA Trimmed Mean CPI

(In Section 7.2, we also investigate some averaged series.) To construct the "SA-NSA" measures, we simply seasonally adjust our NSA measures.

Each of these eight trend inflation indicators is readily understood, and thus the first criterion is satisfied. Criterion 2 is somewhat qualitative, and we evaluate performance by graphing each series against the inflation trend. Criteria 3 is a simple computation; criterion 4 is measured by examining the variance of monthly changes in each index. Criterion 5 is evaluated by computing the mean squared error of each index, where error is the difference between the trend inflation measure and the trend.

7.1 Empirical Results

We start with criterion 2, and look at the 12-month change in the CPI-U, the only time-aggregated series we are considering at present. (In Figure 1, we also plot the 12-month change in the core CPI, which we will discuss in Section 7.2.) Averaging over 12 months smooths out

[22] See Wynne (1999) for a useful chart outlining various approaches to trend inflation indicators.
[23] Note the median CPI and trimmed mean CPI are based upon the FRBC calculations and methods, while the NSA median and NSA trimmed mean CPI use BLS methods.

some temporary fluctuations in the CPI-U. The 12-month CPI often gives the wrong signal for a half-year or so, and in magnitude can err badly around turning points. But – basically by construction – it eventually picks up any turning point in the underlying CPI-U trend.

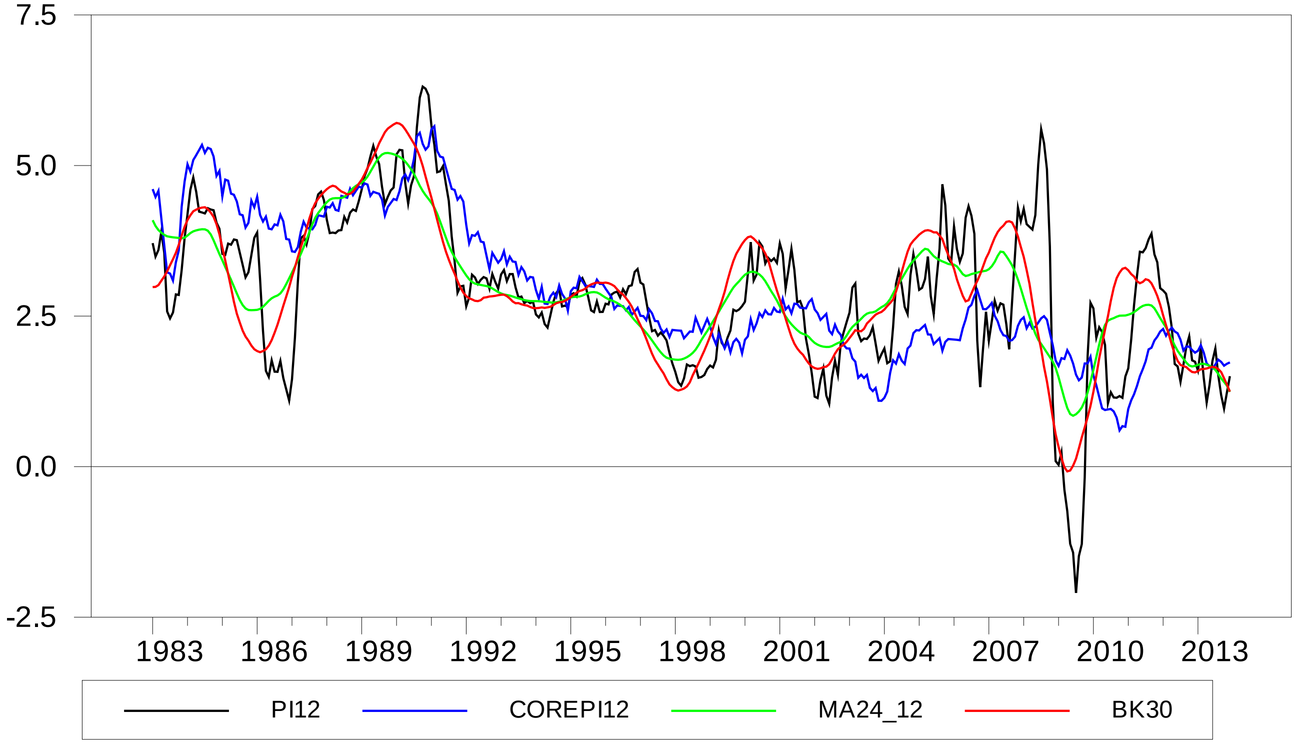

Figure 2

Next, we plot in Figure 3 the monthly movements in the core CPI and the monthly movements in the Median CPI, along with trend inflation as measured by the MA(24,12) trend. Regarding core CPI inflation, it is difficult to draw any conclusions about a turning point from a series that is so volatile. Some degree of time aggregation would be essential before the measure became useful; but Figure 1 provides little support to the idea that such aggregation would be terribly successful. The volatility critique is also true for both the NSA Median and the NSA Trimmed mean (which are not depicted): seasonal volatility must be removed before these series become useful at detecting changes in the trend. Even the SA Trimmed CPI (also not depicted) is too volatile prior to some amount of time aggregation. While SA Median CPI inflation is also somewhat volatile relative to the time-averaged series depicted in previous figures, it provides a better signal and tends to hang closer to the trend, although one might wish that it did a better job at detecting turning points more rapidly. Notice that Figure 3 provides no support whatsoever to the notion that the core CPI dominates the median CPI as a measure of core inflation.

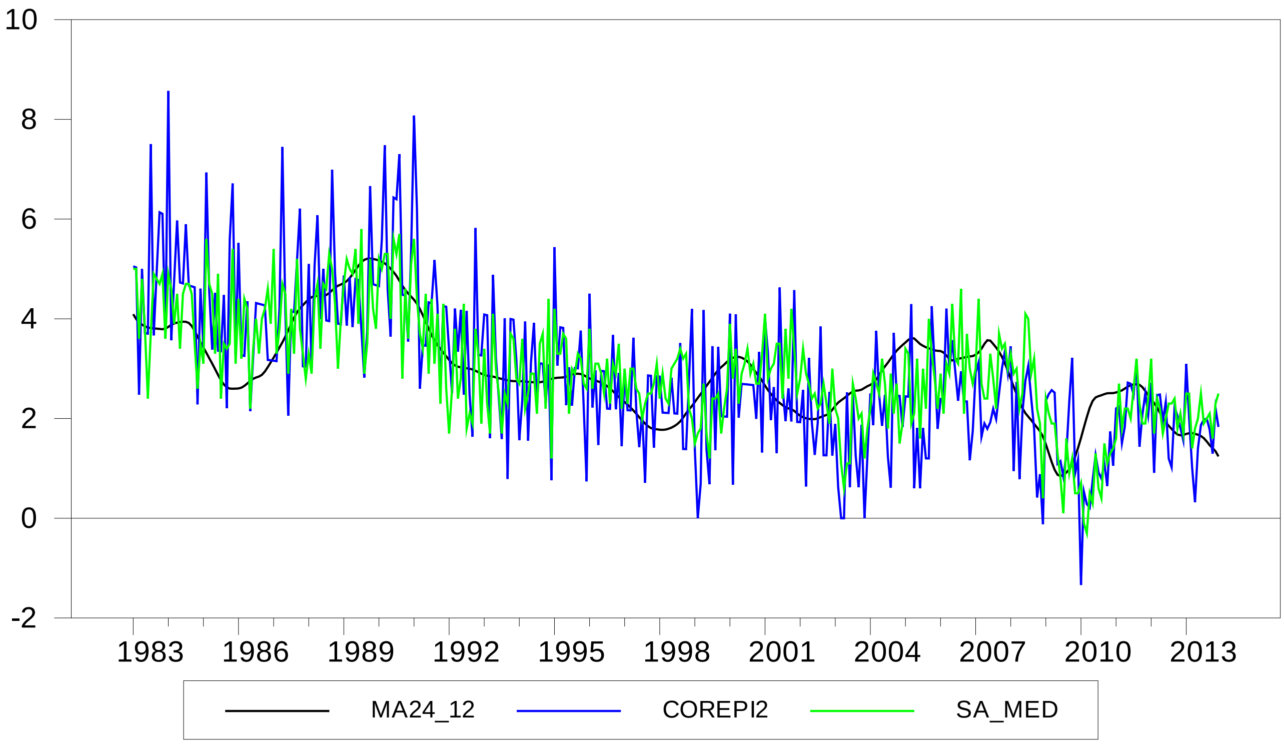

Figure 3

In order to examine criteria 3 through 5 of our performance, we calculate the mean of each contending core inflation series measure, the variance of monthly changes in each series, and the MSE (and maximum squared error) of each series relative to our two trend inflation benchmarks.[24] In reading the MSE statistics on the Tables 1 and 2, numbers which lie within about 0.04 of each other should probably not be considered statistically different from one another.

[24] Note: all calculations below use data ranging from 1/84-9/2013, since adequate data near the endpoints are needed to compute the trend estimates.

Inflation Measure	Mean	Variance of monthly changes	MSE vs. BK30	MSE vs. MA(24,12)
Trend: BK30	2.90	0.013	--	--
Trend: MA(24,12)	2.89	0.004	--	--
Headline CPI-U	2.90	11.60	8.50	8.82
12 Month CPI-U	2.90	0.15	0.73	0.69
Core CPI	2.86	2.36	1.86	1.52
Median CPI	2.94	0.93	1.13	0.79
Trimmed mean CPI	2.77	1.25	1.12	0.91
NSA Median	2.83	1.59	1.44	1.09
NSA Trimmed mean	2.79	3.60	2.47	2.28
SA-NSA Median	2.82	0.76	1.02	0.68
SA-NSA Trimmed mean	2.78	1.47	1.30	1.11

Table 1: Performance of Trend Inflation Indicators

Regarding criterion 3, we observe that the trimmed mean CPIs and the NSA median CPI come in somewhat below the trend. It is possible that this reflects downward bias, although these small differences are minor. [25] Moreover, there is no evidence that the NSA trimmed mean CPI has a bias different from the trimmed mean CPI.

All the trend inflation indicators remove substantial amounts of volatility from the CPI-U. The least volatile indicator is the only one that is time-averaged: the 12 Month CPI-U. The core CPI is over twice as volatile as the median CPI, and over three times as volatile as the SA-NSA median. Both NSA series are quite volatile, and are thus not very useful as such… though it is worth pointing out that the NSA Median is substantially less volatile than the core CPI.

Indeed, the core CPI only manages to outperform one trend inflation indicator: the NSA trimmed mean. As we note below, when using NSA data, it makes sense to use more aggressive

[25] As noted by Bryan and Cechetti (2001), price index change distributions are often skewed. If the price change distribution is skewed, the use of bounded influence estimators such as the median and trimmed mean may be inefficient and produce biased estimates. The authors test for skewness by calculating the median percentile of the price change percentile. Over time, the median percentile should hover around 50 if the price change distribution is symmetric. We also calculate the median percentile of our price change distribution, and present the results in appendix 8. The average over our sample period is slightly above 51.

trimming. More aggressive trimming puts the NSA trimmed mean below the core CPI in terms of volatility.

The biggest surprise here is the performance of the 12 Month CPI. This trend inflation indicator actually stacks up well against the others considered here (although, as we noted when discussing Figure 2, it is somewhat deficient in picking up turning points). First, as noted above, it is the least volatile. Second, on an MSE basis, this measure is in first place when considering the BK30 trend inflation benchmark, and is tied for first place when considering the MA(24,12) trend inflation benchmark. (A priori, one might have expected relatively better performance in the MA(24,12) MSE category, given that the 12-month change in the CPI-U is, mathematically, almost half of the signal in a 30-month centered moving average.)

What are the conclusions so far? First, the core CPI has been decisively rejected as an appropriate measure of core inflation in the CPI (see also Detmeister (2011) and Crone et al. (2013) for similar findings regarding core PCE inflation). Second, two trend inflation indicators do pretty well across the board: the 12-month CPI-U, and the SA-NSA Median CPI. As the former series is time-averaged, we next investigate whether various types of averaging could be used to create an improved trend inflation indicator.

7.2 Improving trend inflation indicators using averaging

7.2.1 Time averaging

The time plots of SA-NSA Median CPI indicate that there is still some transient noise obscuring the signal. Would time averaging help? Practitioners certainly seem to think so: the use of 12-month averages is very common. We have already noted that such averages “bake in” a lagged response to inflation turning points. (Another drawback of longer averages like 12-month averages is that even these rates of inflation are not immune to sharp movements, but can still change rapidly – and not only because of some movement in prices that just occurred, but also because of sharp movements in prices that happened a year earlier.) In addition to annual averaging, it seems worthwhile to investigate more modest amounts of time averaging. Short moving averages, such as ordinary four-month moving averages, or various types of weighted averages, seem promising to investigate. Shorter averages still allow fairly rapid detection of turning points, which longer smoothing might obscure. But we will also examine a much longer moving average, namely the exponential-smoothing average of the CPI suggested in Cogley

(2002), which involves averaging 36 months of CPI-U data. This *exponentially-smoothed* CPI is a one-sided geometric distributed lag of current and past headline inflation. We use the parameter setting suggested by Cogley, adapted to monthly data. (See also Higgins and Verbrugge 2015; that study explores, albeit in an abbreviated manner, averages over a wider span of months.)

First, we explore whether time averaging allows the core CPI to become a reasonable indicator of trend inflation. But the answer is pretty clear: no. Returning to Figure 2, the 12-Month core CPI is decidedly unimpressive as a measure of trend inflation: too high in the mid-1980s, lagging the deceleration in trend inflation in the early 1990s, and providing almost no signal of trend inflation changes after that. It seems mainly to pick up the decrease in inflation after the 1980s, which might explain why it is useful in forecasting longer-run inflation.

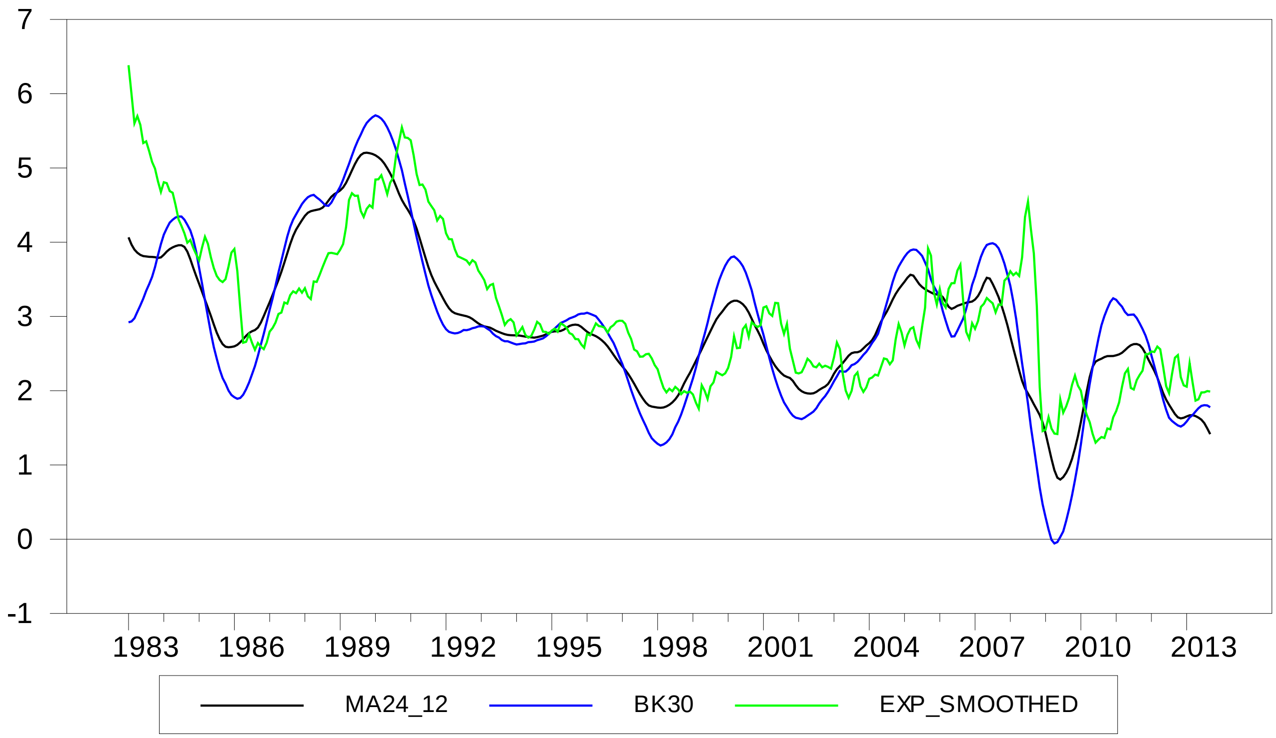

Figure 4

The behavior of the exponentially smoothed CPI is mixed. It is surprisingly choppy, given that it averages 36 months of inflation data. And following directly from its construction, it often lags turning points. In the early 1980s and again in the early 1990s, it was quite delayed in picking up the signal of decelerating inflation. During the Great Recession, it provided a highly erroneous signal for almost a year – trend inflation began to decline in June 2007, but the exponentially smoothed CPI mimicked the 12-month CPI in remaining high and then even accelerating until July 2008. (But after that, it quickly plummeted.) It provided a muddy signal during 2009 and 2010, with the quality of the signal differing depending upon which measure of trend inflation is being used as the metric … although in its defense, that period was a notoriously difficult and puzzling one for inflation dynamics.

We also explore a four-month moving average of four series: the CPI, the core CPI, the median CPI, and the SA-NSA median. In addition, we also consider a 3-month asymmetric moving average of the NSA median, (the "A-MA(3) SA-NSA median"), given by the formula

$$A\text{-}MA(3)\ SA\text{-}NSA\ median = \frac{1}{6}\left\{3\left(SANSAm_t\right) + 2\left(SANSAm_{t-1}\right) + \left(SANSAm_{t-2}\right)\right\}$$

where $SANSAm_t$ is the SA-NSA median at time t. This asymmetric moving average puts half of the weight on the current signal at time t, but allows some signal from previous months too, with declining weights. As we noted, the use of a three-month moving average will slightly reduce its timeliness at detecting turning points, but this may be outweighed by enhanced clarity of the signal. Finally, we also investigate 12-month changes in both the SA-NSA median and the SA-NSA trimmed CPI.

As did Bryan and Cechetti (2001) and Detmeister (2011), we find that there is a notable gain to a modest amount of time-averaging.[26]

[26] Detmeister (2011) found, in keeping with Higgins and Verbrugge (2015), that there appears to be an optimal amount of time-averaging. For example, the trimmed mean PCE appeared to perform best at tracking trend inflation when averaged over about six to nine months. (For forecasting, he found that one might want to use even longer averages.)

Inflation Measure	Mean	Variance of monthly changes	MSE vs. BK30	MSE vs. M(24,12)
12-Month CPI	2.90	0.15	0.73	0.69
12 Month Core CPI	2.90	0.02	1.10	0.66
12 Month Median CPI	2.98	0.01	0.99	0.53
12 Month NSA Median CPI	2.85	0.01	0.95	0.48
12 Month NSA Trimmed CPI	2.81	0.02	0.83	0.41
4 Month CPI-U	3.19	1.20	2.81	3.07
4 Month Core CPI	2.90	0.15	1.06	0.69
4 Month Median CPI	2.95	0.06	0.92	0.52
4 Month SA-NSA Median	2.82	0.05	0.87	0.46
A-MA(3) SA-NSA Median CPI	2.83	0.13	0.82	0.46
Exponentially Smoothed CPI	2.98	0.03	0.75	0.42
Model-Averaging Indicator	2.90	0.06	0.57	0.32

Table 2: Performance of Averaged Trend Inflation Indicators

Over long periods of time, time averaging cannot alter the mean of the series; but it does reduce its variance dramatically. Perhaps surprisingly, given its poor ability to pick up turning points, the exponentially-smoothed CPI does quite well in terms of maintaining a small distance to both alternative trend inflation series. We have already noted the surprising performance of the 12-month CPI. Also noteworthy: In the case of the SA-NSA median CPI, 3-month non-symmetric averaging is arguably superior to 12-month smoothing. While the former leaves more volatility in the measure, it appears to be somewhat better at tracking trend inflation as measured by BK30 (and there is only a small drop in performance vis-à-vis MA(24,12)). And Figure 5 indicates that this measure is among the best of the measures investigated at signaling turning points … though its performance in this regard cannot really be considered stellar. Current theory does not explain why a median would lag trend inflation turning points, as it appears to do. (Garciga and

Verbrugge (2015) investigate this issue.) The core CPI, in either its four-month or 12-month form, performs quite poorly relative to the alternatives.

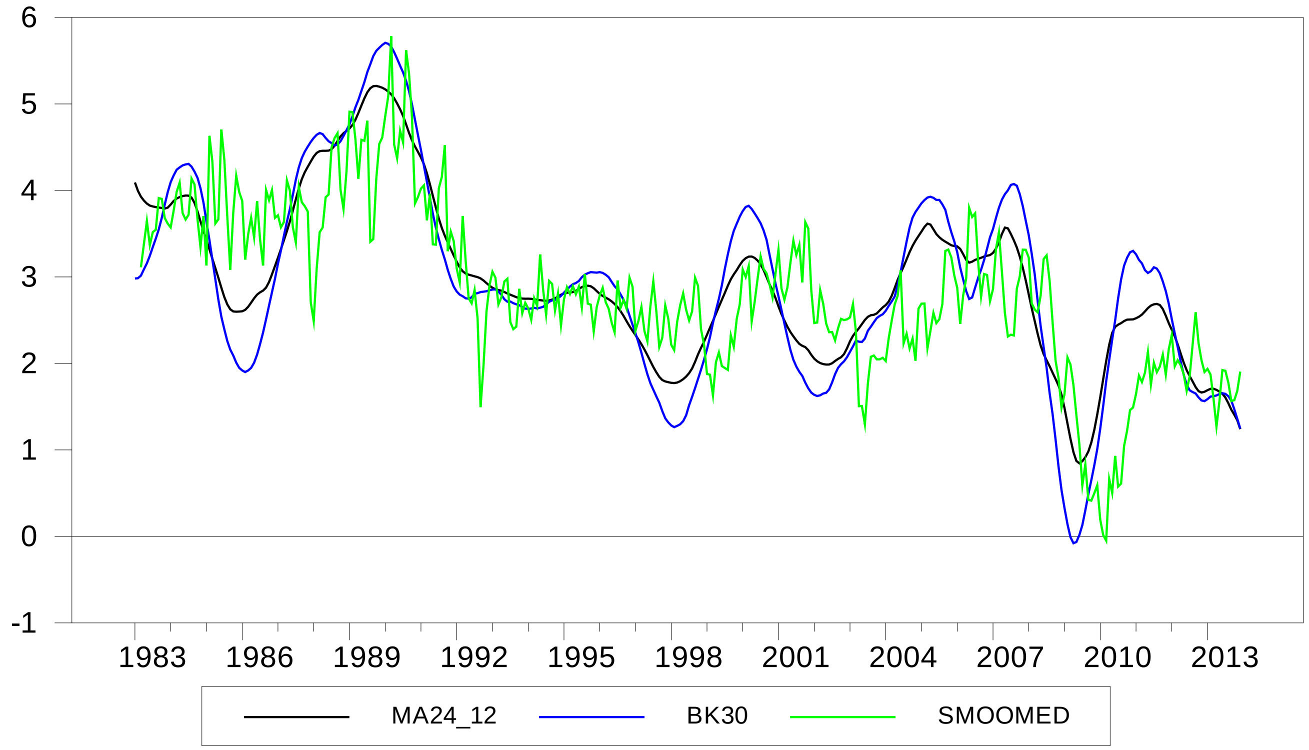

Figure 5

7.2.2 Model averaging

Perhaps we can do better in terms of MSE. We take an idea from the forecasting literature, forecast averaging, and form the simple average of three top performers: the 12-month CPI, the exponentially-smoothed CPI, and the A-MA(3) SA-NSA median. In terms of MSE, this trend inflation indicator is clearly the best of all the measures investigated, as noted in Table 2. But Figure 6 indicates that, despite this performance, it does not quite solve the turning point weaknesses of its components.

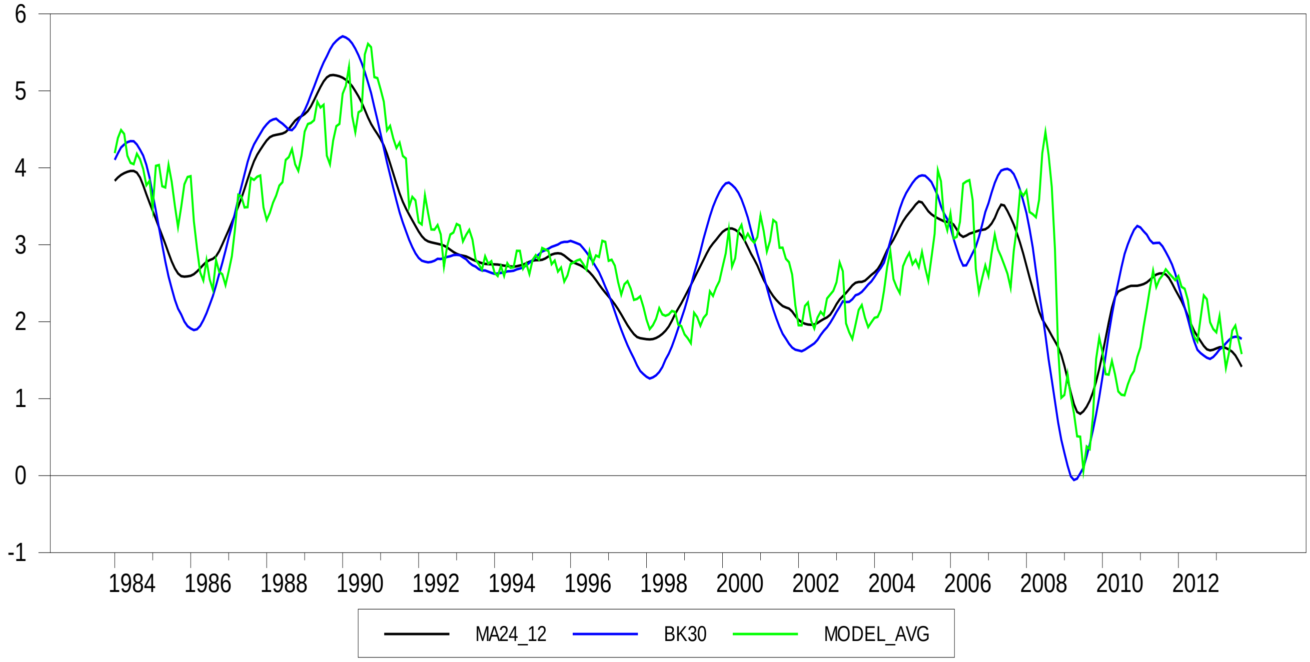

Figure 6

Conclusion

For decades, the median CPI has proven to be an excellent forecasting tool. But in our view, the central purpose of a core inflation measure is not to forecast inflation, but rather to offer an accurate signal of the current level of trend inflation, i.e. to filter out noise from the incoming inflation signals. How well do the Cleveland Fed's core inflation measures do in this regard? To answer this question, one has to come up with an adequate representation of the trend in inflation; in this paper, we put forward two new alternatives.

We also constructed two new core inflation measures that are not seasonally adjusted: an NSA median CPI and an NSA trimmed mean CPI. In their construction, we adhered closely to BLS index production methods. Non-seasonally-adjusted series like these are useful because they will rarely need to be revised – even though, for most uses, they will still have to be seasonally adjusted.

We investigated the performance of these series (in raw form, and in seasonally adjusted form) in comparison to several prominent core inflation measures, including the ever-popular so-called “core” CPI. The core CPI performs poorly. It is completely dominated by trend inflation indicators born in Cleveland, and even by the 12-month version of the headline CPI.

On the basis of our analysis, another trend inflation indicator performs well in terms of its ability to stay close to trend inflation: the exponentially smoothed CPI, developed by Cogley (2002). We also found that an asymmetric three-month moving average of the seasonally-adjusted NSA median CPI performed well in terms of its ability to match turning points, and in terms of its ability to track trend inflation.

Our results led us to propose a new core inflation measure: an model-averaged indicator formed by taking the average of the three-month-averaged SA-NSA median CPI, the 12-month CPI, and the exponentially smoothed CPI. In keeping with typical results in the forecasting literature, this averaging proves to be rather successful, at least in terms of staying close to trend inflation. However, it does not manage to track turning points as good as we would like. And while we are optimistic about trend inflation indicators in this class, and about our new NSA versions of the median CPI and trimmed mean CPI, it is too early to say whether these will stand the test of time. Until then, we recommend the continued use of the median CPI and trimmed mean CPI produced by the Federal Reserve Bank of Cleveland.

Of course, our other recommendation is clear: as a trend inflation indicator, the core CPI should be scrapped.

Appendix 1. Items in the median and trimmed mean

The median CPI is the price index change that lies at the center of the distribution … and that price index always has a name. It's interesting to check out which items often end up at the median. We here display the top 5 most occurring items in the NSA median CPI and the median CPI, between January 1983 through December 1997, and then the more recent period, January 1998 through March 2015.[27] Owing to the interest surrounding OER, we also indicate how frequently each OER index ends up as the median.

Most of the time, both the NSA median and the median select the same index. It is also interesting to note that the index Food Away From Home was, over the entire period, the item *most commonly selected* as the median. This highly-stable and reliable indicator of overall inflation is automatically excluded in the Core CPI![28]

NSA vs SA top 5 Median Components for January 1983 through December 1997

NSA Median Comps	SA Median Comps
Other food at home (5% of the time)	OER Midwest (5.5% of the time)
OER Northeast (7% of the time)	OER South(7% of the time)
OER South (7% of the time)	OER West (7.7% of the time)
OER West (7.7% of the time)	Renters' cost (12% of the time)
Food away home (13% of the time)	Food away from home (14% of the time)
OER Midwest (3.88% of the time)	OER Northeast (3.88% of the time)
All 4 OER regions (26% of the time)	All 4 OER regions (24% of the time)

NSA vs SA top 5 Median Components for January 1998 through March 2015

NSA Median Comps	SA Median Comps
OER Midwest (8% of the time)	OER Northeast (9% of the time)
Rent of primary residence (8% of the time)	Food away from home (10.6% of the time)
OER West (11.5% of the time)	Rent of Primary residence (11% of the time)
Food away from home (12% of the time)	OER Midwest (11.5% of the time)
OER South (13.5 % of the time)	OER South (18% of the time)
OER Northeast (4.8% of the time)	OER West 8.6% of the time)
All 4 OER regions (38% of the time)	All 4 OER regions (47.8% of the time)

We next compare the components which are excluded (or trimmed) from the trimmed mean CPI. The following tables show the components which are excluded from the calculation more than 50 percent of the time. This demonstrates that some components are inherently volatile. It is also interesting to note how frequently the procedure *fails* to trim many food and energy components. This signifies that at the 8 percent trim level, food and energy components are not always the most volatile components in the CPI-U. (Note: Items highlighted in yellow are

[27] In order to gauge frequency, the time periods between Jan. 1983 and Mar. 2015 must be split to reflect the addition and condensing of particular components used the in median and trimmed mean calculations.
[28] We are not the first to note this fact; see, e.g., Bryan and Cecchetti (1994).

components that are inherently volatile despite the use of seasonal adjustment, which tends to dampen volatility.)

NSA vs SA components trimmed more than 50% of the time for January 1983 through December 1997

NSA Trimmed Comps	SA Trimmed Comps
Infants and toddlers apparel (51.6% of the time)	Other apparel commodities (51.6% of the time)
Men's and boy's apparel (58.8% of the time)	Public transportation (52% of the time)
Fuel oil and other household fuel commodities (63% of the time)	Used cars and truck (53% of the time)
Fruits and vegetables (67% of the time)	Infants and toddlers apparel (66.6% of the time)
Motor fuel (78.8% of the time)	Fruits and vegetables (73% of the time)
Women's and girls' apparel (82% of the time)	Fuel oil and other household fuel commodities (73% of the time)
	Motor fuel (76% of the time)

NSA vs SA components trimmed more than 50% of the time for January 1998 through March 2015

NSA Trimmed Comps	SA Trimmed Comps
Infants' & Toddlers' Apparel (52.6% of the time)	Watches & Jewelry (50.7% of the time)
Leased Cars and Trucks (52.6% of the time)	Public Transportation (51% of the time)
Men's & Boys' Apparel (53% of the time)	Energy Services (51% of the time)
Energy Services (53% of the time)	Lodging Away From Home (57% of the time)
Watches & Jewelry (55% of the time)	Fresh Fruits & Vegetables (58.9% of the time)
Fresh Fruits & Vegetables (58% of the time)	Infants' & Toddlers' Apparel (58.9% of the time)
Fuel Oil and Other Fuels (66.6% of the time)	Car and Truck Rental (70% of the time)
Car and Truck Rental (66.6% of the time)	Fuel Oil and Other Fuels (84% of the time)
Lodging Away From Home (84.5% of the time)	Motor Fuel (89% of the time)
Motor Fuel(86% of the time)	
Women's & Girls' Apparel (90% of the time)	

Appendix 2. Calculation of weights

We demonstrate the construction of the relative weight for a given component, using cereals and bakery products as an example. Using SA indexes to construct the relative weight, where time t = March 1999, the following formula is used. I_{Mar99}^{CBake} denotes the index for cereals and bakery products in March 1999, and the relative importance of this component is denoted by R_{Mar99}^{CBake} . The non-normalized weight for R_{Mar99}^{CBake} is given by

$$R_{Mar99}^{CBake} = R_{Dec98}^{CBake} * \left(\frac{I_{Mar99}^{CBake}}{I_{Dec98}^{CBake}}\right)$$

To construct the normalized weight, one must adjust all the relative weights to ensure all weights sum to 100. To obtain the normalized weight, divide the non-normalized weight by the analogous "updated relative importance" for the entire CPI, which has an initial "relative importance" of 100 percent in December 1998.

$$R_{Mar99}^{CPI} = 100 * \left(\frac{I_{Mar99}^{CPI}}{I_{Dec98}^{CPI}}\right)$$

The normalized relative importance formula for cereals and bakery products in March 1999 is as follows,

$$\Phi_{Mar99}^{CBake} = R_{Dec98}^{CBake} * \frac{\left(\frac{I_{Mar99}^{CBake}}{I_{Dec98}^{CBake}}\right)}{100 * \left(\frac{I_{Mar99}^{CPI}}{I_{Dec98}^{CPI}}\right)}$$

This can also be rewritten as a recursive formula

$$R_{Apr99}^{CBake} = R_{Dec98}^{CBake} * \left(\frac{I_{Mar99}^{CBake}}{I_{Dec98}^{CBake}}\right)\left(\frac{I_{Apr99}^{CBake}}{I_{Mar99}^{CBake}}\right) = R_{Mar99}^{CBake}\left(\frac{I_{Apr99}^{CBake}}{I_{Mar99}^{CBake}}\right)$$

And similarly

$$R_{Apr99}^{CPI} = R_{Mar99}^{CPI}\left(\frac{I_{Apr99}^{CPI}}{I_{Mar99}^{CPI}}\right)$$

implying that

$$\Phi_{Mar99}^{CBake} = \Phi_{Feb99}^{CBake}\frac{\left(\frac{I_{Mar99}^{CBake}}{I_{Feb99}^{CBake}}\right)}{\left(\frac{I_{Mar99}^{CPI}}{I_{Feb99}^{CPI}}\right)}$$

Note, that while this is the correct aggregation weight for computing the level of the CPI, in order to compute CPI inflation as a function of item price relatives, as in: $\pi \coloneqq \frac{I_t^{CPI}}{I_{t-1}^{CPI}} =$
$\sum \omega_{i,t}\left(\frac{I_t^i}{I_{t-1}^i}\right)$,
one must transform these weights as in

$$\omega_{i,t}^{\pi} = \frac{\Phi_t^i I_{t-1}^i}{I_{t-1}^{CPI}}$$

Appendix 3. FRBC monthly index weight

The FRBC calculates a monthly weight for each component used in the median and trimmed mean CPI calculation. The methodology is similar to the one suggested by the BLS; however, the weights calculated by the FRBC are lagged by one month. Our research indicates that this time lag has little influence on either the Median CPI or the Trimmed mean CPI. We present below the formula utilized by the FRBC. Since the FRBC method is very similar to the BLS method described in appendix 2, we will skip to the recursive formula.

Using the same example presented in appendix 2, given SA indexes to construction the relative weight for a given component, for example cereals and bakery products, where time t=March 1999 the following recursive formula is used

$$\Phi_{Mar99}^{CBake} = \Phi_{Feb99}^{CBake} \frac{\left(\frac{I_{Feb99}^{CBake}}{I_{Jan99}^{CBake}}\right)}{\left(\frac{I_{Feb99}^{CPI}}{I_{Jan99}^{CPI}}\right)}$$

Notice the one month lag that is created.

Appendix 4. FRBC vs BLS methods: weight-updating

Though the FRBC uses only 45 components, and the FRBC method of calculating weights is lagged by one month, this makes little difference in comparison to the BLS methods. Below we graph the Median CPI using the two alternative methods; BLS-type procedures, applied to seasonally-adjusted data, with 45 components, are termed “BLS-type methods”. Clearly, the difference in the weight-updating formula is of little consequence.

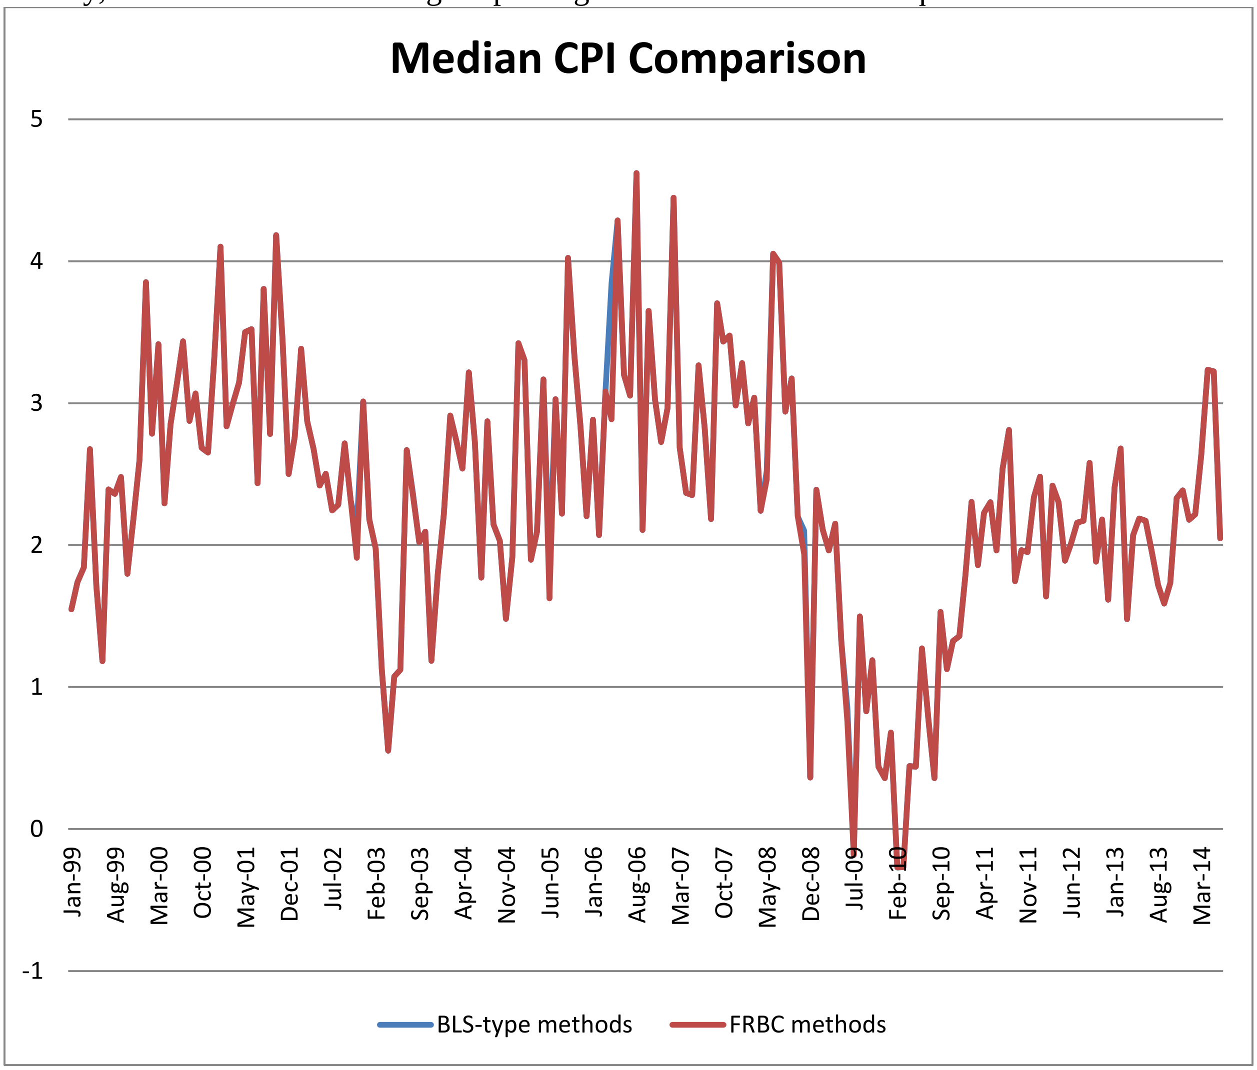

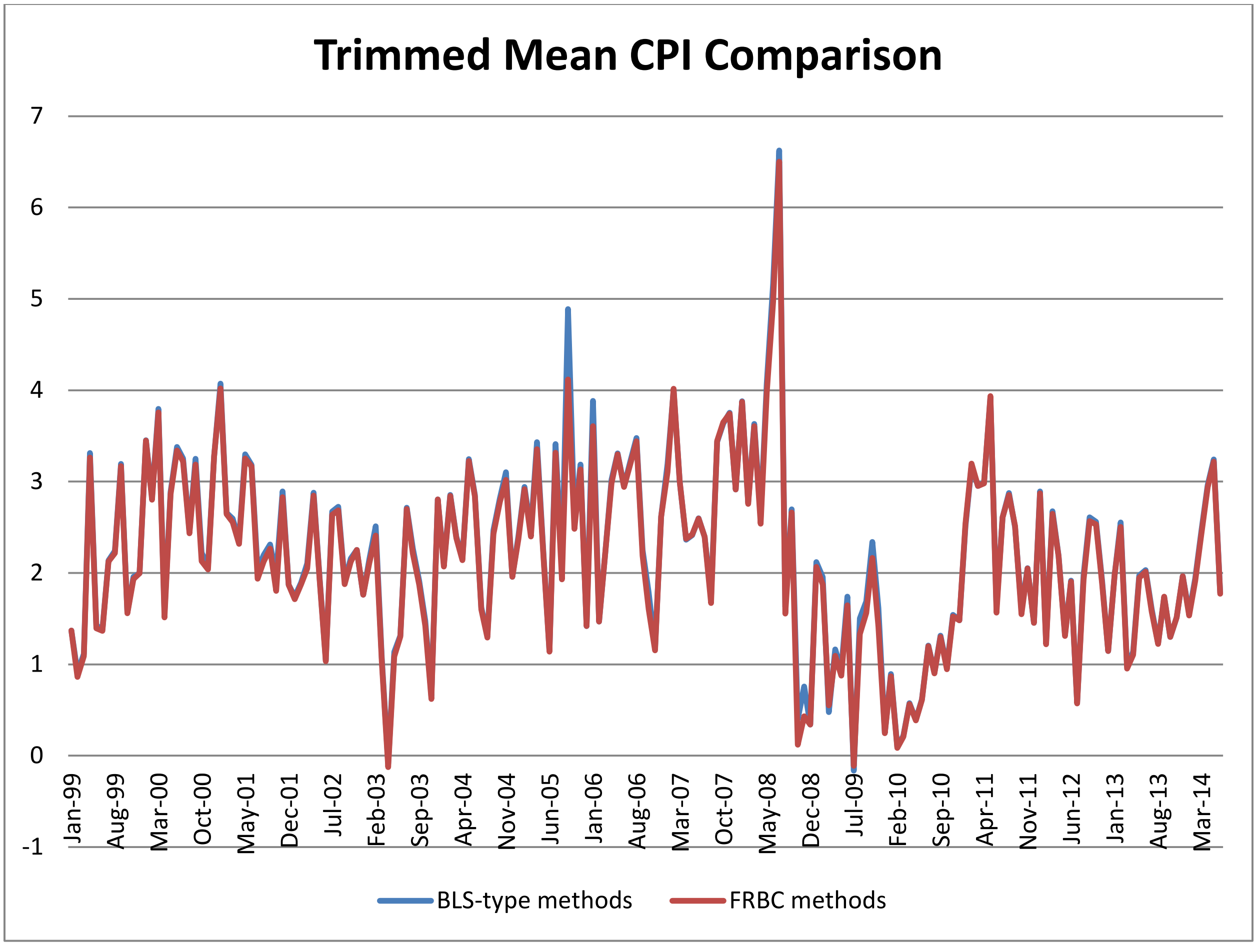

The index construction differences matter only a little more for the Trimmed mean CPI.

Appendix 5. Shortcomings of the 36 or 30 month moving average

The graphs below demonstrate why 36-month or 30-month centered moving average filter is not entirely suitable as measure of trend inflation. First, we simply plot the 30-month moving average against our MA(24,12) series. The latter series is far smoother, and better captures trend movements.

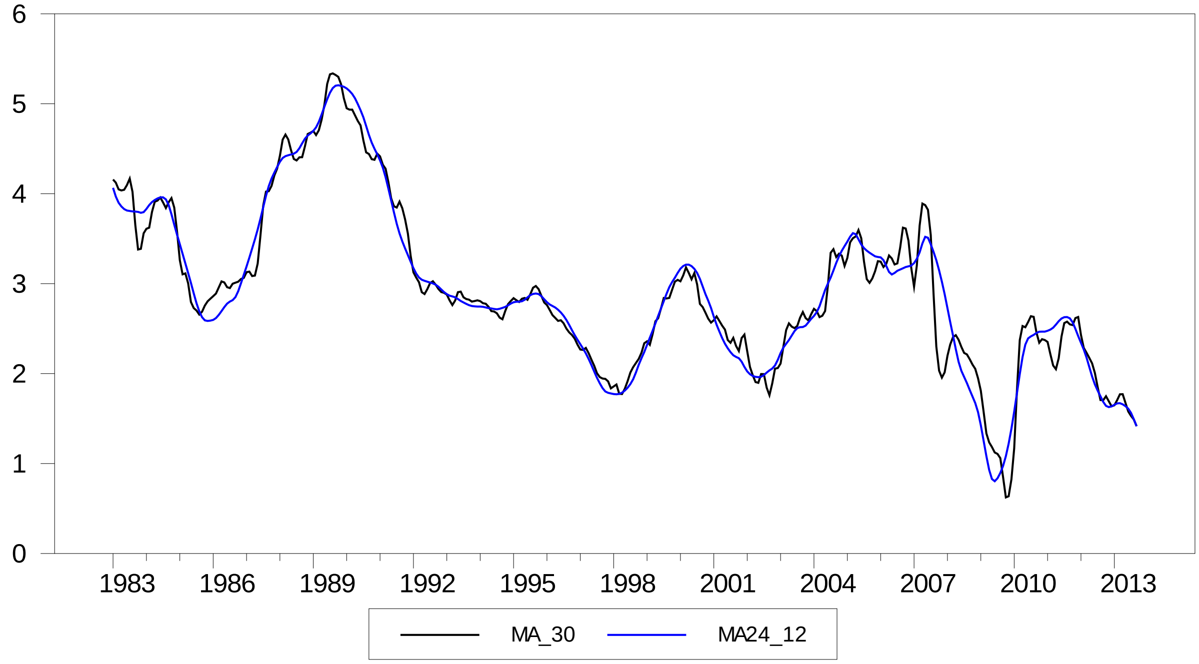

Next we plot the amplitude of the MA(30) filter as a function of period (the amplitude plot of the MA(36) filter is qualitatively similar). To understand this graph, recall that it is possible to decompose any time series into cycles – i.e., sine and cosine waves – with different periods. The *very* long run trend is akin to a "cycle" of more than 10,000 periods, a sine or cosine wave that takes more than 10,000 periods to revert. Business cycle fluctuations are related to cycles with periods between 18 months and 8 years. Seasonal fluctuations will be concentrated in cycles of 12 months, with "echoes" at some higher frequencies, like 24 months. Fluctuations which quickly reverse themselves are comprised of cycles of with periods less than 12 months.

As noted above, the figure is a plot of the amplitude of the 30-month moving average filter. Roughly speaking, the amplitude displays the proportion of each cycle that the filter allows to "leak through" and become part of the trend. Ideally, if one wished to isolate all movements that take 30 months or longer to dissipate, the trend would consist only of those cycles with period 30 months or more; thus the amplitude of the ideal filter would be a vertical rectangle of height 1 on the left side of the graph, and be zero for all periods less than 30 months. As the graph shows, the MA(30) filter allows parts of rapidly-moving cycles to leak through, which accounts for the "jaggedy" appearance of the time series in the first figure. The MA(24,12) filter is clearly superior.

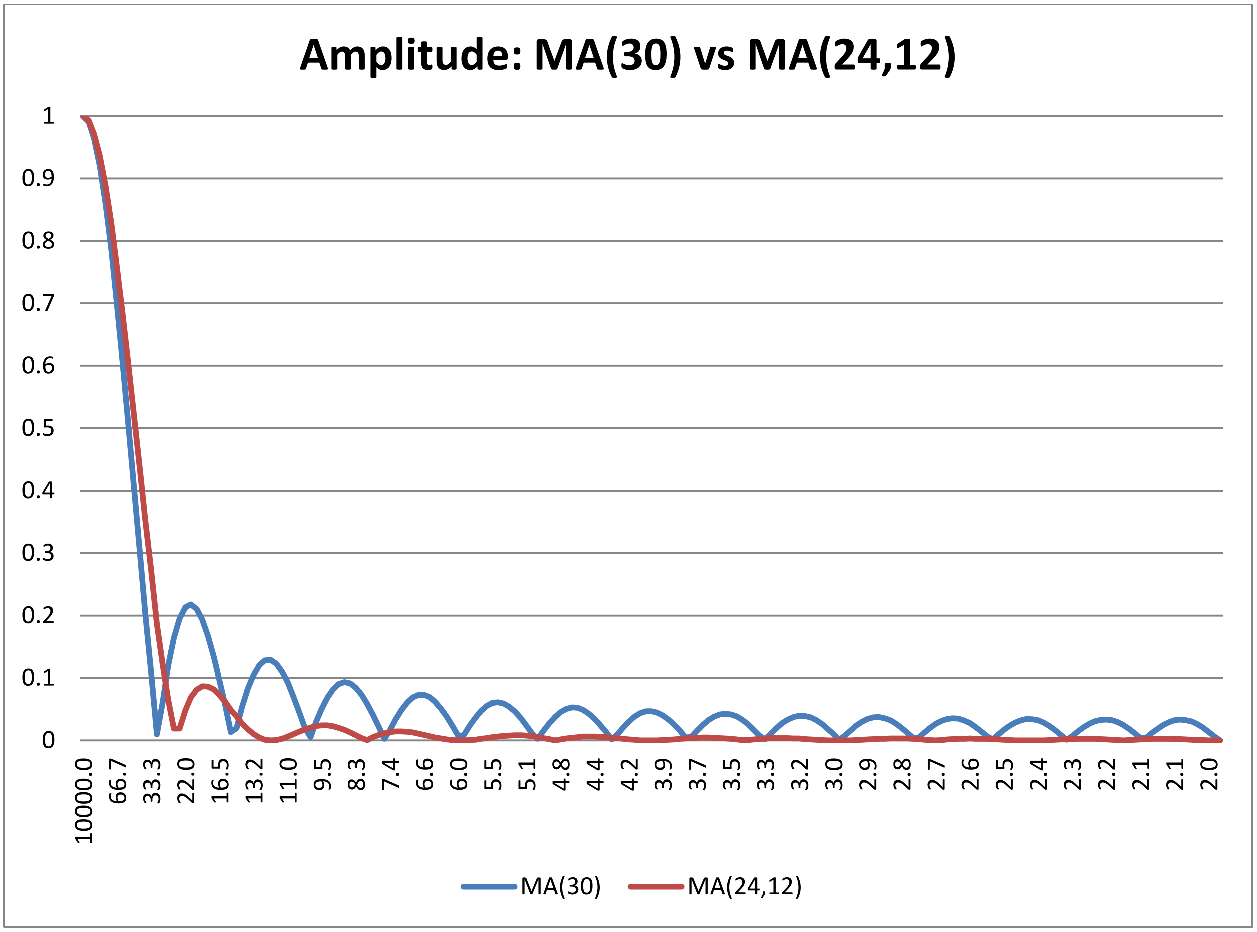
Amplitude: MA(30) vs MA(24,12)
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
10000.0
66.7
33.3
22.0
16.5
13.2
11.0
9.5
8.3
7.4
6.6
6.0
5.5
5.1
4.8
4.4
4.2
3.9
3.7
3.5
3.3
3.2
3.0
2.9
2.8
2.7
2.6
2.5
2.4
2.3
2.2
2.1
2.1
2.0
MA(30)
MA(24,12)

Appendix 6. Other versions of FRBC-type series

The following table displays the mean, variance, and a representative MSE result of alternative Median and Trimmed mean CPI calculations. Our SA-NSA variants are here for comparison, as we used a slightly different trend estimate in this table. The "FRBC" series are the currently published median and trimmed mean series. The "Timeliness" series uses the more timely (BLS) weight updating, but continues to make use of SA data. The "Timeliness + 46" series use the more timely weight, and include the 46th component to the series. As noted above, the BLS produces a relative importance for Unsampled New and Used Vehicles; however, it does not publish an index. In aggregating to the CPI, the price movements for this index are imputed from the price movements of new and used vehicles. Thus, to construct the index of the 46th component, we also use the monthly index for New and Used Vehicles, weighted by the Unsampled New and Used Vehicles relative importance value. The "Timeliness + 46, NSA wgt" method makes use of the BLS methodology of using NSA weights, but SA indexes (in other words, monthly NSA weights calculated NSA indexes, and then these weights are applied to SA indexes).

The evidence here suggests that the departures of the FRBC from BLS production methods are, at worst, innocuous.

Calculation Method	Mean	Variance	MSE (vs. BK-type trend)
Headline CPI-U	2.92	14.40	12.52
SA-NSA Median CPI	2.85	1.17	1.00
SA-NSA Trimmed Mean CPI	2.79	1.79	1.33
FRBC Median CPI	2.95	1.31	1.15
Timeliness Median CPI	2.95	1.32	1.16
Timeliness + 46 Median CPI	2.95	1.32	1.16
Timeliness + 46, NSA wgt, Median CPI	2.95	1.35	1.17
FRBC Trimmed Mean CPI	2.83	1.59	1.11
Timeliness Trimmed Mean CPI	2.80	1.59	1.13
Timeliness + 46 Trimmed Mean CPI	2.83	1.59	1.15
Timeliness + 46, NSA wgt, Trimmed Mean CPI	2.84	1.63	1.19

Appendix 7. Asymmetry

There is asymmetry in the distribution of price index changes underlying the CPI-U in the U.S. (see Bryan and Cecchetti 1999 and Verbrugge 1999, among others), with differing levels of asymmetry in different countries (see, e.g., Bakhshi and Yates 1999 and Figueiredo 2001). Since skewness is present in the data, the median CPI and the trimmed mean CPI may be biased estimators: positive asymmetry leads ordinary trimmed means and medians to be biased downwards, if the goal is to track actual CPI-U movements over time. As we note below, adjusting for such asymmetry is not straightforward, since the degree of asymmetry is unobserved and may change over time.

Following the methodology of Bryan and Cechetti (2001), we calculate the mean percentile in the price index change distribution. Using SA data from January 1983 to March 2015, we calculated the percentile in the weighted distribution of price changes that equals the overall CPI-U movement that month. If the series is not skewed, then the average mean percentile should be 50 percent.

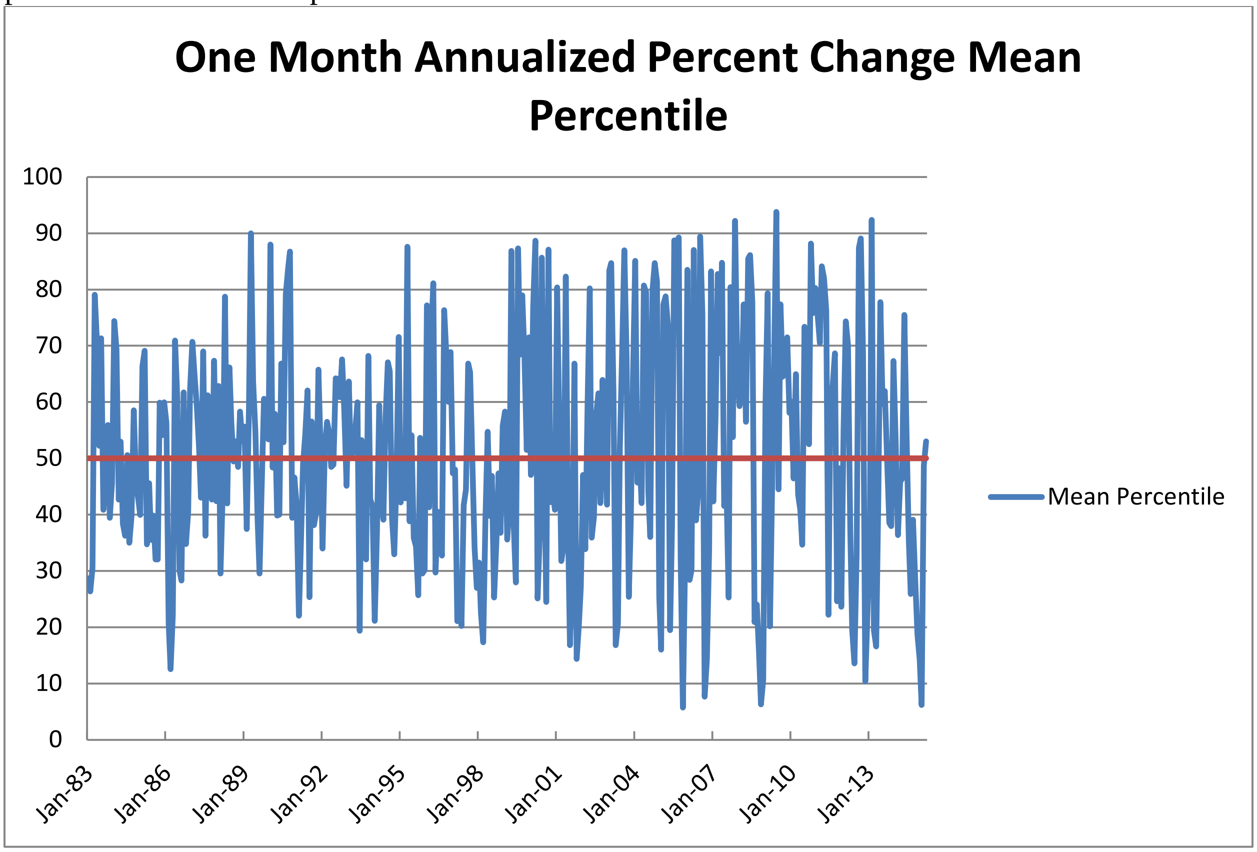

Clearly, the mean percentile varies quite a lot on a month-by-month basis. Skewness is not uniform, nor is it always positive: it can be negative, sometimes substantially negative. Computing the realized mean percentile each month, and then using that percentile as the basis of trimming that month, would defeat the purpose of the core inflation indicator. Nonetheless, there is evidently low-frequency variation in the mean percentile; in other words, if the mean percentile is high this month, it will probably be high next month as well. This raises the possibility that one could estimate those low-frequency movements in some fashion, and then compute trims based upon some function of the currently-predicted mean percentile (rather than the actual realized mean percentile). One version of this procedure would be to compute trims based upon the average level of the mean percentile over some previous horizon, in effect using a one-sided moving average to estimate the trend in the mean percentile. (Experiments along these lines, however, did not prove to be fruitful.) Regarding such averages, the average mean percentile between January 1983 through March 2015 is 51.22, while the average between January 1998 through March 2015 is 52.40.

Bibliography

Ashley, Richard, and Randal Verbrugge (2009) "Frequency Dependence in Regression Model Coefficients: An Alternative Approach for Modeling Nonlinear Dynamic Relationships in Time Series." *Econometric Reviews,* 28(1–3), 4-20.

Aucremanne, Luc (2000) "The Use of Robust Estimators as Measures of Core Inflation," National Bank of Belgium Working Paper No 2.

Bagshaw, Michael L., 1985. "Forecasting and Seasonal Adjustment," Federal Reserve Bank of Cleveland *Working Paper* No. 85-07.

Bakhshi, Hasan and Tony Yates (1999). "To trim or not to trim: an application of a trimmed mean inflation estimator to the United Kingdom." Bank of England Working Papers Series 97.

Baxter, Marianne, and Robert G. King, "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series," *The Review of Economics and Statistics* 81 (1999), 575-593.

Bell, William R., and Ekaterina Sotiris (2010) "Seasonal Adjustment to Facilitate Forecasting: Empirical Results." Manuscript, US Census Bureau.

Blinder, Alan S. May/June 1997. "Commentary." Federal Reserve Bank of St. Louis, Re-view.

Brischetto, Andrea, and Anthony J. Richards (2006). "The Performance of Trimmed mean Measures of Underlying Inflation." Reserve Bank of Australia Research Discussion Paper 2006-10, December.

Bryan, Michael F. and Christopher J. Pike. (1991) "Median Price Changes: An Alternative Approach to Measuring Current Monetary Inflation." Federal Reserve Bank of Cleveland Economic Commentary, December

Bryan, Michael F., and Stephen G. Cecchetti. "The Consumer Price Index as a Measure of Inflation," Economic Review, Federal Reserve Bank of Cleveland (Quarter 4, 1993) pp. 15-24

Bryan, Michael F., and Stephen G. Cecchetti. "Measuring Core Inflation," Monetary Policy, N. Gregory Mankiw, Ed., University of Chicago Press, 1994, pp. 195-215

Bryan, Michael F., and Stephen G. Cecchetti. "The Seasonality of Core Prices," Economic Review, Federal Reserve Bank of Cleveland (Quarter 2, 1995), pp. 12-23.

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

Cecchetti, Stephen G. (1997) "Measuring Short-Run Inflation For Central Bankers," Economic Review of the Federal Reserve Bank of St. Louis, 79 (May/June 1997) 143-156 (NBER Working Paper No. 5786)

Clark, Todd E. (2001). "Comparing Measures of Core Inflation." Federal Reserve Bank of Kansas City Economic Review, v. 86, no. 2 (Second Quarter), pp. 5-31.

Clark, Todd E., and William Bednar (2014) "Methods for Evaluating Recent Trend Inflation." Federal Reserve Bank of Cleveland Economic Trends, 03.28.14.

Cogley, Timothy (2002). "A Simple Adaptive Measure of Core Inflation." Journal of Money, Credit, and Banking, vol. 34, no. 1 (February), pp. 94-113.

Crone, Theodore M., N. Neil K. Khettry, Loretta J. Mester, and Jason A. Novak (2013) "Core Measures of Inflation as Predictors of Total Inflation." Journal of Money, Credit and Banking 45.2-3, 505-519.

Detmeister, Alan K. (2011). "The Usefulness of Core PCE Inflation Measures," Federal Reserve Board: Finance and Economics Discussion Series No. 2011-56.

Figueiredo, Francisco M.R. (2001) "Evaluating Core Inflation Measures for Brazil." Working Paper Series 14.

Giannone, Domenico, and Troy D. Matheson. (2007) "A New Core Inflation Indicator for New Zealand." International Journal of Central Banking, December, 145-180

Garciga, Christian, and Randal Verbrugge (2015) "Real-Time Estimation of Trend Inflation Using Asymmetric Trimming." Manuscript in Preparation, Federal Reserve Bank of Cleveland.

Gordon, Robert J. (1975). "Alternative Responses of Policy to External Supply Shocks." *Brookings Papers on Economic Activity* 1.6,183–206

Higgins, Amy, and Randal Verbrugge (2015) "Is a Nonseasonally Adjusted Median CPI a Useful Signal of Trend Inflation?" Federal Reserve Bank of Cleveland *Economic Commentary*, No. 2015-13.

Knotek, Edward S., and Saeed Zaman. (2015). "Nowcasting US Headline and Core Inflation." Federal Reserve Bank of Cleveland Policy Discussion Paper, (14-03).

Meyer, Brent H., and Mehmet Pasaogullari (2010) "Simple Ways to Forecast Inflation: What Works Best?" Federal Reserve Bank of Cleveland *Economic Commentary*, No. 2010-17.

Meyer, Brent H., and Guhan Venkatu. (2012) "Trimmed-Mean Inflation Statistics: Just Hit the One in the Middle." Federal Reserve Bank of Cleveland Working Paper No. 12-17 (September-13).

Meyer, Brent H., Guhan Venkatu, and Saeed Zaman (2013). "Forecasting inflation? Target the middle." Federal Reserve Bank of Cleveland *Economic Commentary*, April.

Meyer, Brent, and Saeed Zaman (2013). "It's Not Just for Inflation: The Usefulness of the Median CPI in BVAR Forecasting," Federal Reserve Bank of Cleveland, Working Paper no. 13-03.

Plosser, Charles I. (1979) "Short-term Forecasting and Seasonal Adjustment." *Journal of the American Statistical Association* Vol. 74, No. 365, 15-24.

Rich, Robert and Charles Steindel (2007) "A comparison of measures of core inflation," *Economic Policy Review*, Federal Reserve Bank of New York, Dec, 19-38.

Smith, Julie K (2004) "Weighted Median Inflation: Is This Core Inflation?" Journal of Money, Credit and Banking 36.2, 253-263.

Verbrugge, Randal (1999) "Cross-Sectional Inflation Asymmetries and Core Inflation: A Comment on Bryan and Cecchetti." *The Review of Economics and Statistics* 81.2, May 1999, 199-202.

Wynne, Mark (2008). "Core Inflation: A Review of Some Conceptual Issues." Federal Reserve Bank of St. Louis *Review*, May/June 2008, 90(3, Part 2), 205-28.