

Trimmed-Mean Inflation Statistics: Just Hit the One in the Middle

Brent Meyer and Guhan Venkatu

Working Paper 2014-3

March 2014

Abstract: This paper reinvestigates the performance of trimmed-mean inflation measures some 20 years since their inception, asking whether there is a particular trimmed-mean measure that dominates the median consumer price index (CPI). Unlike previous research, we evaluate the performance of symmetric and asymmetric trimmed means using a well known equality of prediction test. We find that there is a large swath of trimmed means that have statistically indistinguishable performance. Also, although the swath of statistically similar trims changes slightly over different sample periods, it always includes the median CPI—an extreme trim that holds conceptual and computational advantages. We conclude with a simple forecasting exercise that highlights the advantage of the median CPI (and trimmed-mean estimators in general) relative to other standard measures in forecasting headline inflation.

JEL classification: E31, E37

Key words: inflation, inflation forecasting, trimmed-mean estimators

The authors thank Julio Blanco, Todd Clark, Alan Detmeister, and Tim Dunne for their helpful comments, criticisms, and guidance. This paper is a revised version of an earlier Federal Reserve Bank of Cleveland working paper (WP12-17) first posted in September 2012. The views expressed here are the authors' and not necessarily those of the Federal Reserve Banks of Cleveland or Atlanta or the Federal Reserve System. Any remaining errors are the authors' responsibility.

Please address questions regarding content to Brent Meyer, Research Department, Federal Reserve Bank of Atlanta, 1000 Peachtree Street, N.E., Atlanta, GA 30309-4470, 404-498-8852, brent.meyer@atl.frb.org; or Guhan Venkatu, Research Department, Federal Reserve Bank of Cleveland, P.O. Box 6387, Cleveland, OH 44101-1387, guhan.venkatu@clev.frb.org.

Federal Reserve Bank of Atlanta working papers, including revised versions, are available on the Atlanta Fed's website at frbatlanta.org/pubs/WP/. Use the WebScriber Service at frbatlanta.org to receive e-mail notifications about new papers.

1 Introduction

Trimmed-mean inflation statistics diagnose the most volatile monthly price changes as noise and “trim” them from the price-change distribution, leaving a clearer inflation signal behind. These measures systematically remove sources of noise on a monthly basis, rather than ad hoc exclusionary measures such as the ex food and energy (“core”) CPI— which implicitly suggests that relative price changes in all other retail price components are inflation signal, even though food prices are no more volatile than the core CPI itself (Clark 2001).

It’s been roughly 20 years or so since trimmed-mean inflation statistics were first examined by Bryan and Pike (1991), and then more rigorously by Bryan and Cecchetti et al (1994, 1997). Bryan, Cecchetti, and Wiggins (1997) documents that the retail price-change distribution is leptokurtic (fat-tailed), arguing that in the presence of excess kurtosis the mean is an inefficient measure of location compared to a trimmed-mean approach. They find that the “optimal” symmetric trimmed-mean CPI measure is an 18% percent trim using the benchmark of the 36-month centered moving average in the headline CPI. In similar work, Bryan and Cecchetti (1994) also focus on the more extreme trim, the median CPI, and tout its use as a “core” inflation measure because it is more strongly correlated with changes in the money supply and a better forecaster than the ex food and energy CPI.¹

In this paper we investigate whether the median CPI is still the appropriate trimmed-mean inflation statistic to use as a measure of underlying inflation. Rather than just focusing on symmetric trims like Bryan and Cecchetti et al, we open our investigation up to the full set of symmetric and asymmetric trims. In somewhat of a departure from other research on underlying inflation, we use gauge the usefulness of a particular trimmed-mean inflation measure by its ability to forecast future inflation. Importantly, we test whether modest differences in forecasting ability are statistically different using the Diebold-Mariano (DM) equality of prediction test. Our sample includes data from 1967 through 2013, and we test for changes in the “optimal” trim by splitting the sample and allowing for rolling-windows. We close with a simple forecasting exercise that highlights the advantage of the median CPI (and trimmed-means in general) relative to other standard inflation measures.

¹In that paper, the optimal trim is the 15 percent trimmed-mean, which is chosen because it was the minimum variance estimator during their sample period.

Others have investigated trimmed-mean inflation statistics since the work of Bryan and Cecchetti et al in the mid-1990s. Smith (2004), using both conditional and unconditional forecasting models, finds that the weighted median CPI outperforms the core CPI. Clark (2001) evaluates a handful of core inflation measures’ ability to track the current inflation trend and forecast future inflation, and finds the 16% trimmed-mean CPI and CPI ex energy to be superior “core” measures. Meyer and Pasaogullari (2010) find the median and the 16% trimmed-mean CPI forecast year-ahead headline inflation about as well as inflation expectations do, and outperform simple forecasting models. Crone, Khettry, Mester, and Novak (2013) found that over longer-horizons (i.e. 24-months), the median CPI yields a forecast significantly superior to that of the headline or ex food and energy CPI index.

Dolmas (2005) applies the trimmed-mean procedure to Personal Consumption Expenditures Price Index (PCE) data. He allows for asymmetric trims and ties their use to the shape of the PCE price change distribution over his sample period. Interestingly, through a series of forecasting exercises, Dolmas points to a “cost” of imposing symmetry in terms of a higher root-mean-squared error (RMSE). However, he does not test to see if these losses in forecast accuracy are significant. Interestingly, Detmeister (2011) created an “core” inflation statistic that takes an “ex post” average of all possible symmetric trims (from essentially a headline measure to the median) and finds it performs “on par” with the Dolmas’ asymmetric trim at tracking trend inflation or predicting future inflation.

Many studies in this area of research—such as Clark (2001) or Cogley (2002)—start with a candidate’s ability to track an in-sample trend, and then evaluate its ability to forecast future inflation in an out-of-sample setting. This paper proceeds in a similar fashion, except that instead of a centered or backward looking trend, our benchmark is annualized inflation over the next 3 years.

As a preview of our results, we find that since 1983, the trimmed-mean CPI measure with the lowest RMSE is the 31-35 percent trimmed-mean—which trims 31 percent off the lower tail and 35 percent off the upper tail. However, the forecast stemming from this trim is not statistically different from that of a wide selection of trimmed-means over a variety of time-periods. Interestingly, this wide-swath of trims with statistically equal forecasting power tends to include aggressive and roughly symmetric trimming points, such as the median CPI.

In general, we find aggressive trimming (close to the median) that is not too asymmetric appears to deliver the best forecasts over the time periods we examine. However, these “optimal” trims vary across periods and are never statistically superior to the median CPI. Given that the median CPI is conceptually easy for the public to understand and is easier to reproduce, we conclude that it is arguably a more useful measure of underlying inflation for forecasters and policymakers alike. We close with evidence that the median CPI and other trimmed-mean estimators generally outperform the headline CPI (and core CPI) in an out-of-sample forecasting test.

2 Data

The Consumer Price Index (CPI) is one of the two major retail prices indexes constructed for the United States. Its data are collected and assembled by the Bureau of Labor Statistics (BLS). In its broadest form, the CPI tracks the inflation experienced by urban consumers, or roughly 87 percent of the US population.²³ The index is currently divided into 211 categories called item strata, which are associated with 8 major groups: Food and Beverages, Housing, Apparel, Transportation, Medical Care, Recreation, Education and Communication, and Other Goods and Services. Also included in the index are taxes and government-charged user fees, such as auto-registration fees.

To collect price information on item strata, BLS employees call or visit thousands of retailers, rental units, and doctors’ offices in 38 urban areas across the US every month. In particular, these employees collect data for 305 items, called entry-level items, associated with the item strata. If an item in this set is no longer available, or has changed in quantity or quality, the field worker selects a substitute, and notes the change. Analysts at the BLS’s national offices then make any adjustments necessary to preserve consistency and comparability in the price data on an item across time.

The weights for the item strata in the CPI are based on information from the Consumer

²There are 38 separate areas that form the full urban sample.

³Strictly speaking, this measure is called the Consumer Price Index for All Urban Consumers, or the CPI-U. A related measure, called the Consumer Price Index for Urban Wage Earners and Clerical Workers, or the CPI-W, tracks the inflation experienced by a subset of all urban consumers, who represent about 32 percent of the US population.

Expenditure Survey (CEX). In this survey, some households provide information on their quarterly purchases, while others provide more detailed bi-weekly diaries, to track more frequently purchased items, such as food and personal care products. Over a 2 year period, the BLS estimates that the CEX provides it with approximately 28,000 weekly diaries and 60,000 quarterly interviews. Since 2002, the BLS has updated CPI expenditure weights every two years. For example, as of January 2012, expenditure weights in the CPI are based on data from the 2008-2010 CEX. In the two years prior to that, beginning in January 2010, expenditure weights were based on data from the 2006-2008 CEX. Before 2002, weights were updated approximately every 10 years.

In between biennial updates of expenditure weights, the BLS assumes that the quantity of items in the CPI market basket does not vary. Expenditure shares, however, will change because of changes in item prices. As such, the biennial weights originally drawn from the CE are updated monthly as follows:

$$W_{i,t+1} = \frac{w_{i,t} \times \left(\frac{p_{i,t+1}}{p_{i,t}} \right)}{\left(\frac{P_{i,t+1}}{P_{i,t}} \right)}$$

where $w_{i,t}$ is the weight of item i at time t , and $\{w_{i,t} | 0 \leq w_{i,t} < 1\}$, $p_{i,t}$ is the index value of item i at time t , and P_t is the index value of the CPI at time t , where $P_t = \sum w_{i,t} p_{i,t}$.

In this analysis, we use data from January 1967 to December 2013, which are divided into 3 periods: January 1967 to December 1982, January 1983 to December 1997, and January 1998 to December 2013. Though the entire CPI is used in each of these three periods, each employs different degrees of aggregation of the item strata. The first period, 1/1967-12/1982, uses 36 components. The second period, 1/1983-12/1997, uses 42 components, and incorporates a change in the way that owner-occupied housing is treated in the CPI. The last period 1/1998-12/2013, uses 45 components, and incorporates a major revision to the CPI's item structure, as well as regional indexes for owner-occupied housing.⁴⁵ In all cases, we use seasonally adjusted indexes for components where they exist.

⁴There are actually 44 indexes from January 1998 to December 2001. The 45th index, leased cars and trucks, begins in January 2002.

⁵For information on the construction of regional indexes for owner-occupied housing, see [<http://www.clevelandfed.org/research/data/us-inflation/revmcpfi.cfm>].

The largest change between the 1967-1982, and 1983-forward time periods is a change in the BLS' methodology for measuring owner-occupied housing costs. In December 1982, the BLS switched to a rental equivalency approach. This change in methodology created a new component—Owners' Equivalent Rent (OER).⁶ OER is the largest single component in the CPI, accounting for roughly 25 percent of the overall index by expenditure weight.

Having one oversized component relative to the rest of the market basket creates a potential issue for trimmed-mean price statistics. As the aggressiveness of the trim increases (toward the median) the influence of this component could grow disproportionately. For example, from January 1998 to July 2007 OER was the median component roughly two-thirds of the time. Brischetto and Richards (2007) recognized this and proposed a solution to reduce the influence of OER on the median CPI, by breaking it up into smaller regional components. The Federal Reserve Bank of Cleveland has adopted this procedure and it has lessened the frequency that an OER component is the median component.⁷⁸

3 Choice of Benchmark

The choice of inflation benchmark is a key aspect of research on underlying inflation measures. Some attempts focus on a candidate measure's ability to track or forecast inflation horizons of a year or less. In our view, this isn't a long enough time period to allow relative price shocks to unwind, and much of what is expressed in the near-term growth rate of headline inflation is noise. We prefer to measure inflation over longer, monetary policy relevant, time horizons that have a greater propensity to exude true inflation—the effect on prices in general that is due to directly to the impact of monetary policy. In this sense, a more appropriate benchmark would be to measure inflation over a 2 or 3 year time horizon.

⁶In 2010, the BLS shifted the relative importance value (weight) of an unpriced component—'secondary residences'—from 'lodging away from home' to OER, and changed the new component name to 'owners' equivalent rent of residences.' The BLS stated the reason for the shift was to lessen the volatility of the shelter. This move did not change the price index for OER, it just added a little less than 1.0 percentage point to its weight (potentially contributing to an entirely different issue for extreme trimmed-mean measures like the median CPI).

⁷See <http://www.clevelandfed.org/research/data/us-inflation/revmcpicfm> for more details.

⁸In some sense, the results we present later in the paper may ease some discomfort with this potential issue, as the median CPI is in the set of statistically similar forecasts in every time horizon we investigate.

Many studies of core inflation treat a longer-run centered moving average as their trend proxy. In fact, much of Bryan and Cecchetti’s work uses the 36-month centered moving average in the headline CPI as its benchmark. However, using a centered moving average would leave this paper open to Blinder’s critique of Cecchetti (1997):

Regarding the criterion, notice that a 36-month centered moving average treats the past and the future symmetrically. That is natural from the viewpoint of a scholar analyzing historical data. But it is very unnatural from the viewpoint of a central banker living in real time. Historically, $t - 1$ and $t + 1$ look more or less the same. But if you must make a decision at time t , there is a world of difference between $t - 1$ and $t + 1$. The past is known; the future is not. More important, the past is the dead hand of history, but the central bank must worry about the future.

We take Blinder’s advice and focus on trimmed-means’ ability to track the trend in future inflation, specifically the annualized percent change in the headline CPI over the next 36 months. This choice of benchmark allows enough time for most relative price changes to work themselves out, leaving just the monetary impulse that is inflation behind.

We could have also chosen the 24-month ahead annualized percent change in the CPI as our inflation benchmark. The results of our empirical tests would largely be unaffected by switching to this benchmark. However, as the gasoline price shock of 2008 unwound, the 24-month annualized percent change in the CPI actually turned negative in early 2010. Given that this “deflation” was largely driven by falling energy prices, it suggests that a 24-month growth rate isn’t quite long enough to dampen the effect of relative price shocks.

4 Evaluating the full set of trimmed-means

As in Bryan, Cecchetti, and Wiggins (1997), we calculate weighted $\alpha, \beta\%$ -trimmed means in the following fashion. We first begin by sorting the monthly price-change data, $\{x_1, x_2, \dots, x_i\}$, and the associated weights, $\{w_1, w_2, \dots, w_i\}$. The percent trimmed from the lower tail is $\alpha = jn$,

while the percent trimmed from the upper tail is $\beta = kn$; α and $\beta = \{0, 1, 2, \dots, 50\}$. The α, β %-trimmed mean is given by:

$$\bar{x}_{\alpha,\beta} = \frac{1}{n - j - k} \sum_{i=j+1}^{n-k} w_i x_i$$

We trim the weighted distribution in 1-percent increments, to create 2500 unique trimmed-mean measures. To make comparisons to the forecast target, the subsequent 36-month percent change in the CPI, we annualize all figures. For example,

$$\bar{x}_{\alpha,\beta}^{ann} = \left(\left(\frac{\bar{x}_{\alpha,\beta} + 100}{100} \right)^{12} - 1 \right) * 100$$

We then compare each trimmed mean, $\bar{x}_{\alpha,\beta}^{ann}$, at time t to the forecast target at time $t + 36$, and generate a summary measure for forecast accuracy, root mean square error (RMSE), as follows:

$$RMSE_{\alpha,\beta} = \sqrt{\frac{1}{t} \sum (\pi_{(t+36)}^{ann} - \bar{x}_{\alpha,\beta}^{ann})^2},$$

where $\pi_{(t+36)}^{ann}$ is the annualized, 36-month percent change in the CPI from time t to $t + 36$.

In attempting to determine which trimmed-mean measure will be the best forecaster of future inflation, we take two approaches:

1) Non-rolling sample: Uses the entire time series, or a specific subsample, to find the trimmed-mean, $\bar{x}_{\alpha,\beta}^{ann}$, with the lowest RMSE.

2) Rolling samples: Uses overlapping, 5-year and 10-year windows. For example, we calculate each $RMSE_{\alpha,\beta}$ for data from the initial 5-year window, t to $t + 60$. The trimmed-mean, $\bar{x}_{\alpha,\beta}^{ann}$, with the lowest RMSE in this initial period is saved as the best trimmed-mean forecaster for time t , along with its trimming percentages α and β . This process is then repeated for the next 5-year window, $t + 1$ to $t + 61$, with the relevant information assigned to time $t + 1$, and so on. In this way, we construct a time-varying, best trimmed-mean forecaster. The process is carried out in the same way for the 10-year rolling windows.

We then take these lowest-RMSE trimmed-mean measures, and compare their squared fore-

cast errors to those from trimmed-means with higher RMSEs, to determine whether the differences are statistically significant. To do this, we use a Diebold-Mariano (DM) test with the pre-whitened quadratic spectral kernel introduced by Andrews and Monahan (1992) to correct the standard errors for heteroskedasticity and autocorrelation.⁹ A common approach is to use a DM test with Newey-West standard errors that employs the Bartlett kernel. However, Clark and McCracken (2011) show that as the prediction horizon increases (and in small samples), this variance estimator becomes biased, overstating the rejection region.

Using the CPI component data described in the above section, we now compute the entire set of symmetric and asymmetric trimmed-mean price statistics. The least aggressive trimmed-mean (or mean in this case) is the headline CPI; and the most aggressive trimmed-mean measure is the median (which trims 49.99 percent from each tail). Unlike Bryan and Cecchetti, we evaluate our candidates for “optimal” trim on the basis of their ability to forecast our benchmark for inflation—that is the annualized percent change in the headline CPI over the next 3 years.

It should be mentioned that we are only investigating trimmed-mean measures over 1-month horizons. Others, such as Detmeister (2011), have suggested that it may be useful to investigate longer component sampling horizons (say, 3-month or 6-month annualized percent changes) and then perform the trimming procedure. While this method would result in a less disperse price change distribution, it would also muddy the inflation signal by allowing relative price changes to becoming embedded in less aggressive trims. For example, a sharp one-month price spike in a single component that undoes itself the next month may not be trimmed out of a procedure that takes the trimmed-means on 3-month growth rates of the components. Yet, those large, opposing price swings are likely indicative of mismeasurement or seasonal adjustment noise and are precisely what trimmed-means were developed to eliminate.

⁹We also tried the Harvey-Leybourne-Newbold (HLN) approach and found qualitatively similar results. However, as noted in Clark and McCracken (2011), the HLN test isn’t guaranteed to yield a positive variance, which lead, in this application, to incomplete results over a few time horizons.

4.1 Root-Mean Squared Errors

Figure 1 plots the RMSEs for all the various trimmed-mean measures over our full sample—from January 1967 to December 2013.¹⁰ The lowest RMSE belongs to $\bar{x}_{43,49}$, which is the trimmed-mean CPI that excludes 43 percent of the lower tail and nearly all (49 percent) of the upper tail. The contour plot reveals a small area (in royal blue), just shy of the symmetric trim (the black line) where all of these trimmed-mean measures have a RMSE between 2.0 and 2.5.¹¹ That next swath (lighter blue) ranges from a RMSE of 2.5 and 3.0. This grouping is roughly symmetric and is fairly large—encapsulating the median CPI all the way to about a 6 percent symmetric trim. Going further out on the contour plot, the deterioration in forecast accuracy tends to follow the same shape; venturing too far away from a symmetric trim leads to a poorer RMSE. In fact, the worst performing trim, the $\bar{x}_{49,0}$ excludes the lower-half of the distribution leaving the just the upper-half remaining, and carrying a RMSE of 7.45 (roughly three times worse than the best performing trim).

Next, we’d like to pay particular attention to the period following what’s commonly referred to as the “Volker Disinflation.”¹² Figure 2 shows the post-1982 contour plot. Post-1982, the $\bar{x}_{31,35}$ owns the best forecast accuracy, with a RMSE of 1.08. That said, there are roughly 100 or so trims (of the 2500 total) that have a RMSE within 0.1 percentage point of the $\bar{x}_{31,35}$. This contour plot differs somewhat from figure 1. While it is still the case that, in general terms, asymmetric trims are associated with worse forecasting performance; the lowest RMSE swath is much larger than in figure 1.¹³

¹⁰The trimmed-mean measures are tracking 3-year ahead inflation, therefore the evaluation period ends in December 2010.

¹¹The lowest RMSE over the whole time period, belonging to $\bar{x}_{43,49}$ is 2.44, so the scale may be overstating the relative forecast accuracy a little.

¹²If the monetary regime, inflation process, or some other structural factor has changed since then, the results in Figure 1 may not hold. Also, the item strata changed in level of aggregation and component definition starting in January 1983. Bryan, Cecchetti, Wiggins (1997) note that the level of kurtosis positively correlated with the number underlying components. The addition of 6 additional components may have increased the kurtosis in the underlying price change distribution, changing the “optimal” trim.

¹³Figures A1 and A2 in the appendix further illustrate the difference between the two time periods. The blue dots signify the minimum RMSE over these sample periods. Over both periods, the RMSE falls as the trimming point increases, say, to about a 10 percent symmetric trim. After that, the gains are much smaller.

4.2 Statistical Significance

To determine the statistical significance of these modest differences in forecasting performance, we employ the Diebold-Mariano (1995) equality of prediction test that utilizes the pre-whitened quadratic-spectral kernel to ensure HAC standard errors. Figure 4 shows the results from our equality of prediction tests for the full sample period (1967-2013). This procedure tests for the mean difference in squared forecast errors between the trim with the lowest RMSE during that time period and all others. The orange areas show (with 90% and 95% confidence) where we fail to reject the null hypothesis that the mean difference forecast errors between the “optimal” trim and the candidate trim are equal to zero. The 45 degree line (in black) represent the symmetric trims.

Figure 3 reveals a large swath of statically indistinguishable forecasts that include a majority of the symmetric trims. Interestingly, over this time period—one that includes the high inflation episode of the late 1970s—some extremely asymmetric cases could not be rejected. In fact, one could exclude the entire lower tail of the price-change distribution, as long as they trimmed at least 12 percent from the upper tail, and remain in the non-rejection area.

However, once we split the sample and focus on the post-1982 period, that extreme asymmetry disappears and the area of statistically indistinguishable predictions shrinks precipitously (illustrated in Figure 4). There appears to be a penalty in terms of forecasting accuracy as the trimming points become too asymmetric. Importantly, more aggressive trimming is needed to reach the edge of the non-rejection area. Focusing on just the symmetric cases for example, a 12% symmetric (6% from each tail) trim is statistically indistinguishable from the lowest RMSE trim in the full sample. Yet, once the 1967-1982 period is excluded from the sample, the amount of trimming necessary to reach the minimum threshold doubles to a 24% symmetric trim (or 12% from each tail). In either sample period, the median CPI appears in the non-rejection region.

Given the modest differences between figures 3 and 4, we were curious about how much the “optimal” trim changed across sample periods. Figure 5 illustrates the changing nature of “optimal” trim in a 5-year rolling-windows framework. Interestingly, it appears that the upper and lower trimming points move in concert, which hints at the symmetric nature of the

process. This modest instability taken together with the wide area of non-rejection in figures 3 and 4, suggests that more aggressive (and roughly symmetric) trimming reduces the likelihood of producing poor forecasts.

Figure 6, which doubles the size of the overlapping windows, also highlights the instability of the optimal trim. Perhaps of more interest though, is that during the disinflationary period of the 1980s and early 1990s, the optimal trim (while aggressive) wasn't symmetric. It appears that the optimal trimming points do not return to moving in concert until the mid 1990s. We find the modest, though insignificant (as highlighted in figure 4), widening of the optimal trimming points through much of the 2000s consistent with the so-called "Great Moderation" and "anchoring" of inflation expectations. This is a relatively sanguine period for price changes, and it isn't a surprise that the inflation signal is (usually) less distorted by relative price changes which would require a more aggressive trim. However, there are still windows that suggest that a more aggressive trim is optimal.¹⁴

In general our results suggest:

- 1) There is a wide range of trimmed-mean measures that deliver roughly "equal" forecasting accuracy.
- 2) That range of trims almost always includes the symmetric trims (after trimming roughly 20 percent or so of the most dramatic price changes from the monthly distribution).
- 3) No single trimmed-mean measure can be declared unequivocally "the best." However, over nearly every time period we examined, the median CPI is included in the set of trims that has statistically equal forecasting accuracy as the lowest RMSE trim in that time period. Given that the median CPI is conceptually and computationally simple, we'd advocate for its continued use as an underlying inflation indicator.

5 An Illustrative Forecasting Test

Given the results above—that there is no single trimmed-mean measure that strictly dominates the median CPI—we still feel it necessary to illustrate its usefulness of the median CPI as

¹⁴We remain keenly interested in whether it is possible to tease out the "optimal" trim in real-time, perhaps through some sort of regime-switching model, but leave this for further research.

an inflation forecaster relative to the headline and “core” CPI.¹⁵ To perform this illustrative forecasting test, we follow the procedure set forth by Crone, Khettry, Mester, and Novak (2013).

In arguing against the use of “core” inflation measures as useful predictors of future inflation (and therefore appropriate guideposts for monetary policymakers to follow), they use a fixed-window rolling regression technique to test whether the year-over-year growth rate in the CPI ex food and energy, CPI ex energy, or the median CPI (produced by FRBC) outperforms forecasts based on the trend in headline inflation. We proceed using their forecasting equation and a similar framework.

The forecasting equation is:

$$\pi_{t,t+h} = \alpha + \beta x_{t-l,t} + \varepsilon_t ,$$

where $\pi_{t,t+h}$ is the annualized growth rate in headline inflation over the horizon t to $t+h$. We allow $t+h = 6$ -, 12 -, 24 -, and 36 -months ahead. Unlike Crone, Khettry, Mester, and Novak (2013) that restrict their independent “core” measure to its respective 12-month growth rate, in $x_{t-l,t}$, we allow $l = 1, 3, 6, 9$, or 12 months.

Allowing for shorter growth rates in the various core measures is motivated largely by the work of Bryan and Cecchetti and is also suggested by Bryan and Meyer (2011). This research shows that the trimmed-means, by efficiently excluding noisy relative prices changes can more accurately forecast future inflation over shorter time horizons (i.e. 3-months). Bryan and Meyer even point out that the 3-month annualized growth rate in the median CPI can more accurately forecast 3-year ahead inflation than the 12-month growth rate in the headline CPI. They also note that as the length of the trend increases—say, looking at the 24-month annualized growth rate in the candidate forecasters—gains from using core inflation dissipate. To us, this actually makes intuitive sense. If, in fact, trimmed-means (and other core measures) are eliminating relative price noise appropriately, then as these noisy price movements unwind over longer horizons, the trend in core measures and headline inflation should converge.

We estimate our forecasting equation using a recursive (expanding-window) strategy that

¹⁵For context, we also include the results for the 16% symmetric trimmed-mean CPI which is released alongside the median CPI by the Federal Reserve Bank of Cleveland.

begins in January 1968—the first date at which we have all the data necessary for estimation—and runs through the last available data point.¹⁶¹⁷ We start with an initial 15-year (180 observation) window (January 1968–December 1982) and then use those coefficients to forecast inflation over the next 6-, 12-, 24-, and 36-months ahead. The first forecast error for each forecasting horizon and specification is then computed. We then iterate through all the available data, expanding the estimation window by 1 month and gather up forecast errors at each step.

We evaluate the forecasting performance of the various inflation measures by calculating the out-of-sample RMSEs and evaluating their statistical significance (using the headline CPI-based forecasts as the benchmark) with the Diebold-Mariano test. The null hypothesis under this test is that the two competing models have indistinguishable finite-sample prediction errors.

We calculate the DM statistic as follows:

Let $d_t = e_{1t}^2 - e_{2t}^2$, where e_{1t}^2 is the squared prediction error from the baseline model at time t . In our case, it is the model that uses lagged growth rates of headline inflation as regressors. e_{2t}^2 is the squared prediction error for the competing models that use underlying inflation measures as regressors. The associated test statistic is $S_{DM} = \frac{\bar{d}}{\sqrt{\sigma^2/n}}$, where \bar{d} is the mean of d and σ^2 is the variance of the loss differential. The DM statistic is a two-sided t statistic estimated using heteroskedasticity and autocorrelation consistent (HAC) standard errors with $h - 1$ lags and a truncated (rectangular) kernel to correct for autocorrelation. We also employ the Harvey, Leybourne, and Newbold (1997) small-sample adjustment to the t -statistic in an attempt to ensure the test is appropriately sized.

Tables 1 through 4 report the results of the out-of-sample forecasting exercise that compares how well lags of headline CPI forecast future inflation relative to the core CPI, 16% trimmed-mean CPI, and median CPI.

Given our prior that the appropriate CPI inflation benchmark is 36-months ahead, we report those results first, in Table 1. We find that the models with the near-term trends in the trimmed-mean measures as independent variables outperform the model with headline inflation

¹⁶In an earlier version of the paper we mimicked the 101 period rolling-window estimation scheme of Crone et al (2013). However, we found those results to be sensitive to both the size of the estimation window and the starting date of the sample.

¹⁷The last available data point was December 2013, so to evaluate 6-month ahead CPI inflation our sample period ends in June 2013. Just as the 36-month ahead inflation evaluation window ends in December 2010.

as the independent variable in terms of RMSE. The DM test confirms that those differences are statistically significant for the 1-month annualized percent change in the 16% symmetric trimmed-mean CPI and the median CPI at the 5 percent and 10 percent significance levels, respectively. Interestingly, models that leverage trends in the core CPI actually carry a higher RMSE than the CPI-based regressions for growth rates longer than 3 months. These results suggests that movements in food and energy prices, like any other relative price change, exhibit a component of inflation. Ignoring these prices, while helpful when removing volatility in the over the near-term, becomes a hinderence over longer trends (12-month and out).

Tables 2, 3, and 4 repeat this forecasting exercise in forecasting headline inflation over the next 24-, 12-, and 6-month horizons. Over every forecasting horizon and for all the growth rates in the independent variables we test, the trimmed-mean measures carry a lower RMSE than the headline CPI . It is also the case the usefulness of trimmed-mean inflation indicators is primarily in their ability to disentangle signal from noise over shorter time horizons (1-month and 3-month growth rates). The 1-month growth rates in the median CPI and 16% symmetric trimmed-mean CPI significantly outperform the headline CPI is forecasting inflation over all the forecast horizons we examine. This is usually the case for the 3-month growth rates as well. It is also the case that the 3-month growth rates in the trimmed-mean measures carry a lower RMSE than the 12-month growth rate in the headline (and core) CPI, a result that is consistent with earlier work by Bryan and Meyer.

Our findings appear to find support for using trimmed-mean inflation estimators, running contrary to the thrust of Crone et al (2013) argument. In our tests, it is never the case that models based on the headline CPI significantly outperform any of the core measures, nor did the CPI-based models ever carry a lower RMSE than any of the trimmed-mean inflation-based models. We attribute some of the inconsistencies between these two papers to differences in sample selection and estimation scheme. We would argue our results are more robust because they less sensitive to arbitrary start dates, window size selection, and we are leveraging all available data. That said, we would also point out that the gains in forecasting accuracy relative to the headline CPI are concentrated over shorter frequencies, a finding that is consistent with their results. Gains in forecasting accuracy on the part of trimmed-mean measures tend to diminish relative to the headline CPI when evaluating inflation indicators over longer-term (12-

month) trends. This is not surprising to us given the nature of noisy relative price swings. We'd expect that over longer horizons these price movements would unwind, alternative measures of underlying inflation would tend to converge to the trend in headline inflation.

6 Conclusion

While we originally set out to find a single superior trimmed-mean measure, we could not conclude as such. In fact, it appears that a large swath of candidate trims hold statistically indistinguishable forecasting ability. That said, in general, the best performing trims over a variety of time periods appear to be somewhat aggressive and almost always include symmetric trims. Of this set, the median CPI stands out, not for any superior forecasting performance, but because of its conceptual and computational simplicity—when in doubt, hit the one in the middle.

Interestingly, and contrary to Dolmas (2005) we were unable to find any convincing evidence that would lead us to choose an asymmetric trim. While his results are based on components of the PCE chain-price index, a large part (roughly 75% of the initial release) of the components comprising the PCE price index are directly imported from the CPI. It could be the case that the imputed PCE components are creating the discrepancy. The trimmed-mean PCE series currently produced by the Federal Reserve Bank of Dallas trims 24 percent from the lower tail and 31 percent from the upper tail of the PCE price-change distribution. This particular trim is relatively aggressive and is not overly asymmetric—two features consistent with the best performing trims in our tests.

Finally, even though we failed to best the median CPI in our first set of tests, it remains the case that the median CPI, among other trimmed-mean statistics, are useful forecasters of future inflation.

References

- [1] Andrews, Donald W.K. and J. Christopher Monahan. 1992. "An Improved Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimator." *Econometrica* 60: 953-966.
- [2] Blinder, Alan S. May/June 1997. "Commentary." *Federal Reserve Bank of St. Louis, Review*.
- [3] Brischetto, Andrea and Anthony Richards. 2006. "The Performance of Trimmed Mean Measures of Underlying Inflation." *Reserve Bank of Australia, Research Discussion Paper*.
- [4] Bryan, Michael F. and Stephen G. Cecchetti. 1994. "Measuring Core Inflation." In e. N. Gregory Mankiw, *Monetary Policy*, 195-215. Chicago: University of Chicago Press.
- [5] Bryan, Michael F., Stephen G. Cecchetti, and Ronald Wiggins. 1997. "Efficient Inflation Estimation." *National Bureau of Economic Research Working Paper*, No. 6183
- [6] Bryan, Michael F. and Brent H. Meyer. 2011, June 1. "Should we even read the monthly inflation report? Maybe not. Then again..." Retrieved from Macroblog: <http://macroblog.typepad.com/macroblog/2011/06/should-we-even-read-the-monthly-inflation-report-maybe-not-then-again.html>
- [7] Bryan, Michael F. and Christopher Pike. 1991. "Median Price Changes: An Alternative Approach to Measuring Current Monetary Inflation." *Federal Reserve Bank of Cleveland Economic Commentary*.
- [8] Clark, Todd E. 2001, Second Quarter. "Comparing Measures of Core Inflation." *Federal Reserve Bank of Kansas City Economic Review*.
- [9] Clark, Todd E. and McCracken, Michael W. 2011. "Advances in Forecast Evaluation." *Federal Reserve Bank of Cleveland Working Paper*, No. 11-20.
- [10] Cogley, Timothy. 2002. "A Simple Adaptive Measure of Core Inflation." *Journal of Money, Credit, and Banking*, 34(1): 94-113.
- [11] Crone, Theodore M., N. Neil K Khettry, Loretta Mester, and Jason Novak. 2013. "Core Measures of Inflation as Predictors of Total Inflation." *Journal of Money, Credit, and Banking*, 45(2-3): 505-519.
- [12] Detmeister, Alan K. 2011. "The Usefulness of Core PCE Inflation Measures." *Federal Reserve Board Finance and Economics Discussion Series*, working paper, No. 2011-56.
- [13] Diebold, Francis X., and Roberto S. Mariano. 1995. "Comparing Predictive Accuracy." *Journal of Business and Economic Statistics*, 13: 253-263
- [14] Dolmas, Jim. 2005. "Trimmed Mean PCE Inflation." *Federal Reserve Bank of Dallas, Research Department Working Paper*, No. 0506.
- [15] Giacomini, R., and White, H. 2006. "Tests of Conditional Predictive Ability." *Econometrica*, 74: 1545-1578.
- [16] Harvey, D. I., Leybourne, S. J., & Newbold, P. 1998. "Tests for Forecast Encompassing." *Journal of Business and Economic Statistics*, 16(2): 254-259.

- [17] 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.
- [18] Smith, Julie K. 2004. "Weighted Median Inflation: Is This Core Inflation." *Journal of Money, Credit, and Banking*, 36: 253-63.

7 Figures and Tables

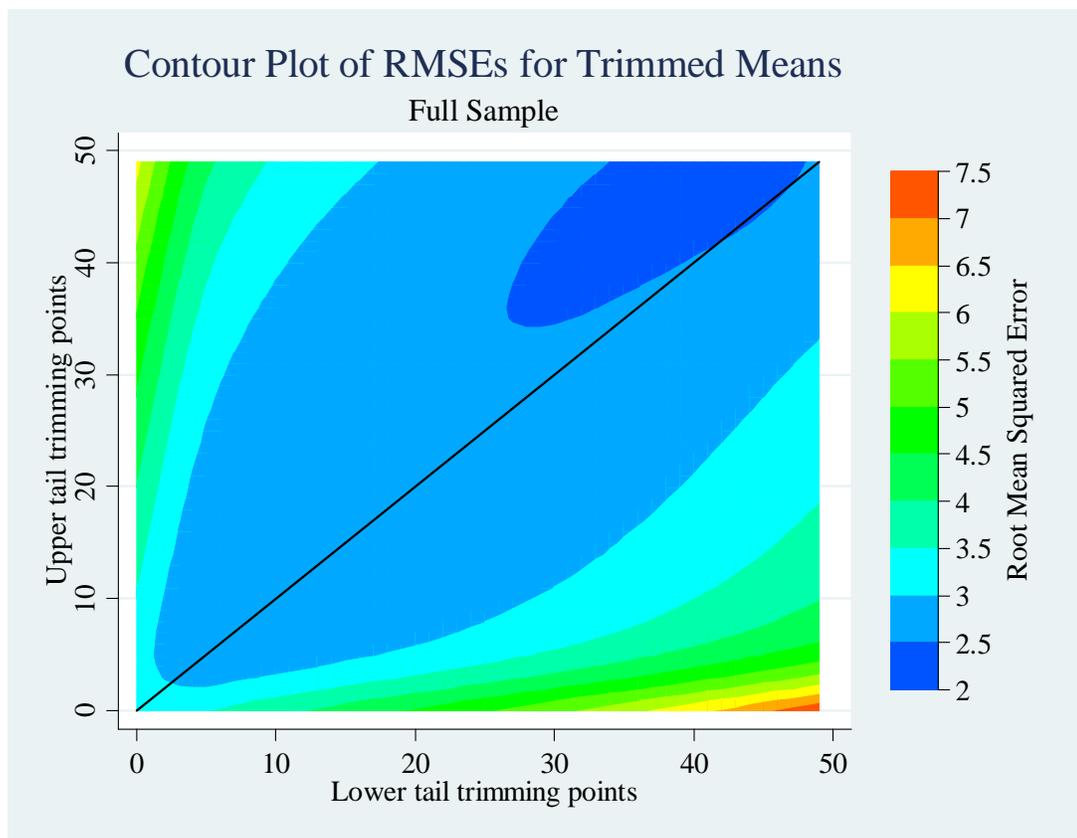


Figure 1: The above plots RMSEs for monthly CPI trimmed-means from forecasts of the 36-month ahead, 36-month annualized percent change in the CPI for the period from January 1967 to December 2013. The lower-tail trim, α , is shown on the horizontal axis, while the upper-tail trim, β , is shown on the vertical axis. The symmetric trimmed-means, where $\alpha = \beta$, are shown by the black line. The lowest RMSE area of the plot extends from the upper-right corner (shown in blue), with an RMSE range of 2.0 to 2.5. Each successive contiguous area increases the RMSE range by 0.5.

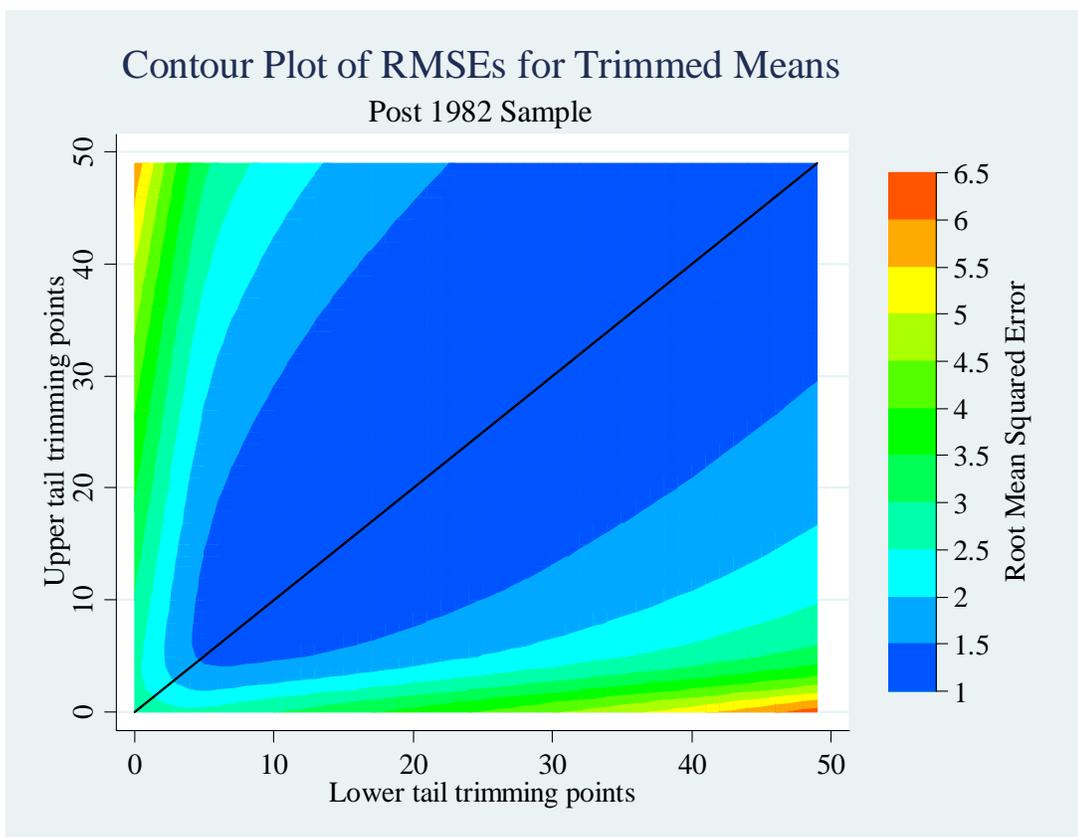


Figure 2: The above plots RMSEs for monthly CPI trimmed-means from forecasts of the 36-month ahead, 36-month annualized percent change in the CPI for the period from January 1983 to December 2013. The lower-tail trim, α , is shown on the horizontal axis, while the upper-tail trim, β , is shown on the vertical axis. The symmetric trimmed-means, where $\alpha = \beta$, are shown by the black line. The lowest RMSE area of the plot extends from the upper-right corner (shown in blue), with an RMSE range of 1.0 to 1.5. Each successive contiguous area increases the RMSE range by 0.5.

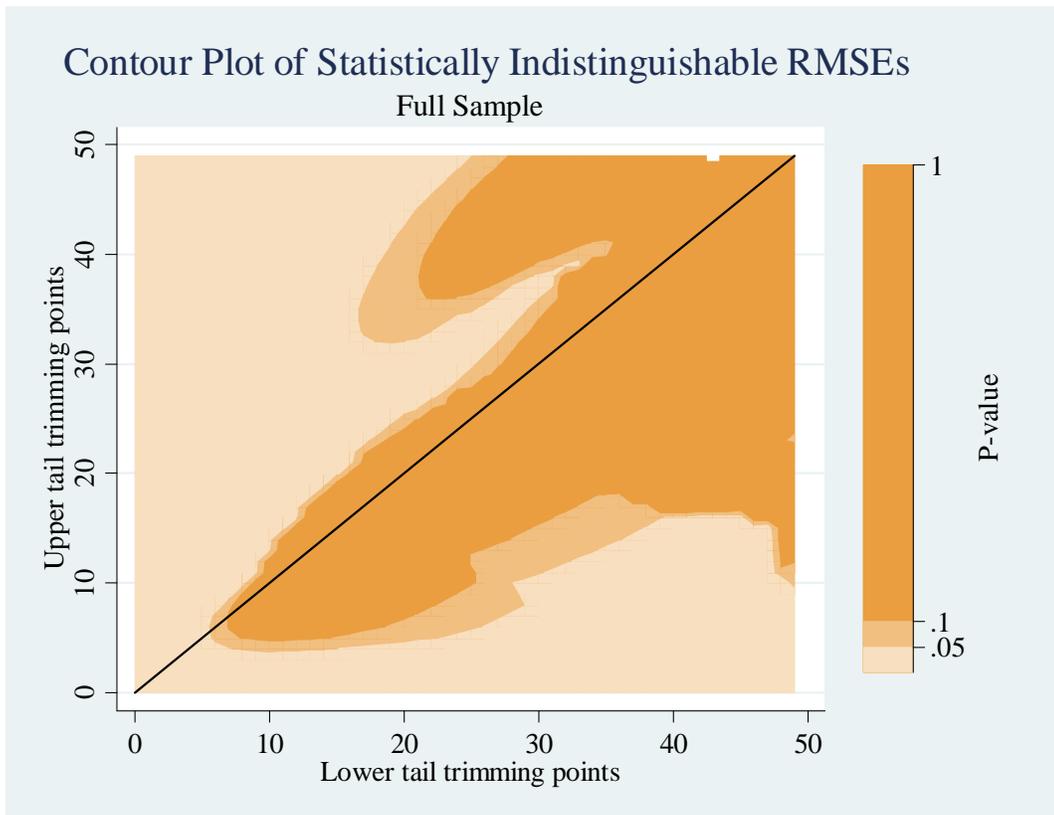


Figure 3: The above figure shows the trimmed-means with RMSEs that are statistically indistinguishable from the lowest-RMSE trimmed-mean of $\bar{x}_{43,49}$ (shown in white) for the evaluation period from January 1967 to December 2010, using the Diebold-Mariano (DM) test with the quadratic spectral kernel and Andrews (1991) optimal bandwidth. The darker orange area shows the trimmed-means that are statistically indistinguishable from the $\bar{x}_{43,49}$ with 90% confidence interval. The 95% confidence interval includes lighter oranges areas. The symmetric trimmed-means, where $\alpha = \beta$, are highlighted with the black line.

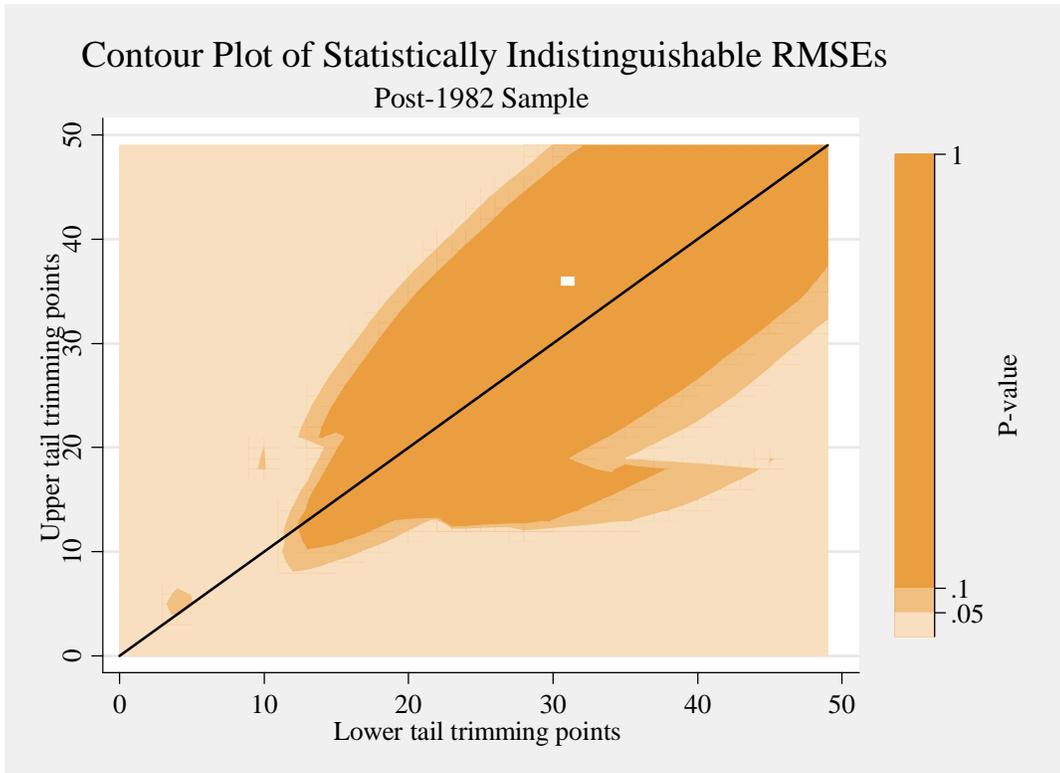


Figure 4: The above figure shows the trimmed-means with RMSEs that are statistically indistinguishable from the lowest-RMSE trimmed-mean of $\bar{x}_{31,36}$ (shown in white) for the evaluation period from January 1967 to December 2010, using the Diebold-Mariano (DM) test with the quadratic spectral kernel and Andrews (1991) optimal bandwidth. The darker orange area shows the trimmed-means that are statistically indistinguishable from the $\bar{x}_{31,36}$ with 90% confidence interval. The 95% confidence interval includes lighter oranges areas. The symmetric trimmed-means, where $\alpha = \beta$, are highlighted with the black line.

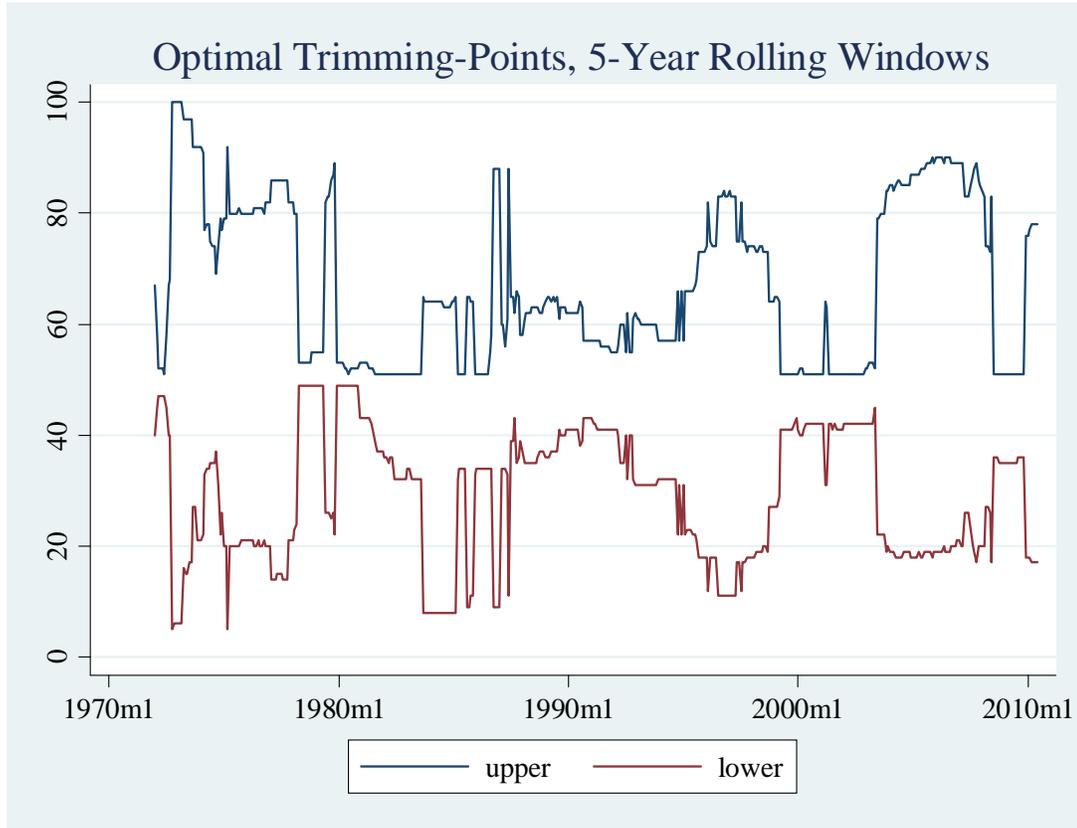


Figure 5: The following shows the trimming points, α and β , of the trimmed-mean, $\bar{x}_{\alpha,\beta}$, that minimize the RMSE for a forecast of the 36-month percent change in the CPI over overlapping 5-year intervals. For example, the first set of observations show the trimming points for the RMSE-minimizing trim over the first 60 months ending February 1972. The second observations, for the 60 month window ending in March 1972, show the trimming points of the RMSE-minimizing trimmed-mean, and so on. The dotted line shows the lower-tail trim, α , while the red line shows 100 minus the upper-tail trim, β .

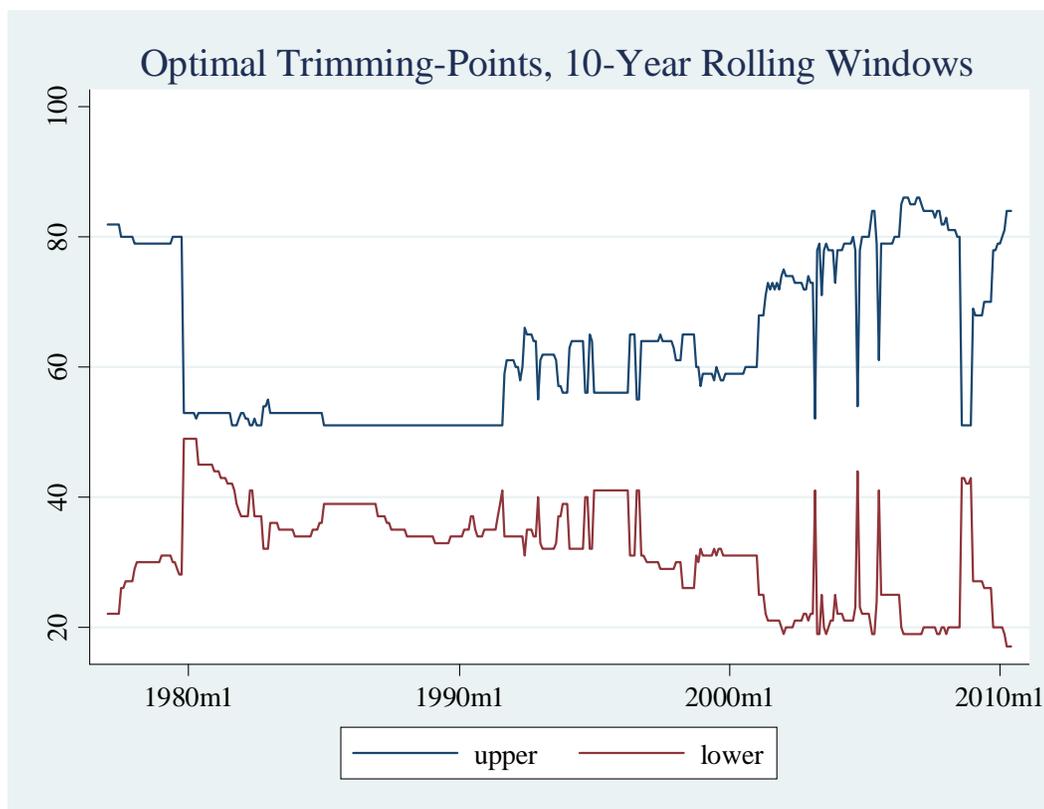


Figure 6: The following shows the trimming points, α and β , of the trimmed-mean, $\bar{x}_{\alpha,\beta}$, that minimize the RMSE for a forecast of the 36-month percent change in the CPI over overlapping 10-year intervals. For example, the first set of observations show the trimming points for the RMSE-minimizing trim over the first 120 months ending January 1977. The second observations, for the 120 month window ending in February 1977, show the trimming points of the RMSE-minimizing trimmed-mean, and so on. The dotted line shows the lower-tail trim, α , while the red line shows 100 minus the upper-tail trim, β .

Forecast Accuracy: 36-months ahead RMSEs						
Annualized percent change over the last:	1 month	3 months	6 months	9 months	12 months	
CPI	2.204	2.023	1.900	1.861	1.887	RMSE
Core CPI	2.084	1.964	1.937	1.953	1.987	RMSE
	<i>0.560</i>	<i>0.240</i>	<i>-0.166</i>	<i>-0.471</i>	<i>-0.582</i>	<i>DM statistic</i>
16% trimmed-mean CPI	1.880	1.792	1.766	1.787	1.838	RMSE
	2.201 **	<i>1.442</i>	<i>1.061</i>	<i>0.811</i>	<i>0.680</i>	<i>DM statistic</i>
Median CPI	1.898	1.821	1.775	1.777	1.823	RMSE
	1.925 *	<i>1.155</i>	<i>0.938</i>	<i>0.924</i>	<i>0.981</i>	<i>DM statistic</i>

'*', '**', '***' denotes significance at the 10%, 5%, and 1% level, respectively

TABLE 1: *Out-of-Sample Test of Forecast Accuracy, 36-months ahead*

Note: The estimated equation is of the form: $\pi_{t,t+h} = \alpha + \beta x_{t-l,t} + \varepsilon_t$; where $\pi_{t,t+h}$ is the annualized growth rate in headline inflation over the horizon t to $t+h$. $x_{t-l,t}$ is the annualized growth rate in the independent variable, where $l = 1, 3, 6, 9,$ or 12 months. The out-of-sample forecast evaluation period is January 1983 through December 2010. The forecast errors are generated recursively, starting with a 180-month base period (January 1968-December 1982). The Diebold-Mariano (DM) equality of prediction test is estimated with HAC standard errors (rectangular kernel) using $h - 1$ lags to control for autocorrelation.

Forecast Accuracy: 24-months ahead RMSEs						
Annualized percent change over the last...	1 month	3 months	6 months	9 months	12 months	
CPI	2.184	1.953	1.782	1.722	1.737	RMSE
Core CPI	1.981	1.818	1.782	1.810	1.864	RMSE
	<i>0.840</i>	<i>0.477</i>	<i>-0.001</i>	<i>-0.369</i>	<i>-0.617</i>	<i>DM statistic</i>
16% trimmed-mean CPI	1.743	1.623	1.585	1.604	1.662	RMSE
	2.825 ***	1.819 *	<i>1.288</i>	<i>1.045</i>	<i>0.890</i>	<i>DM statistic</i>
Median CPI	1.769	1.663	1.598	1.599	1.652	RMSE
	2.300 **	<i>1.375</i>	<i>1.035</i>	<i>0.950</i>	<i>0.914</i>	<i>DM statistic</i>

'*', '**', '***' denotes significance at the 10%, 5%, and 1% level, respectively

TABLE 2: *Out-of-Sample Test of Forecast Accuracy, 24-months ahead*

Note: The estimated equation is of the form: $\pi_{t,t+h} = \alpha + \beta x_{t-l,t} + \varepsilon_t$; where $\pi_{t,t+h}$ is the annualized growth rate in headline inflation over the horizon t to $t+h$. $x_{t-l,t}$ is the annualized growth rate in the independent variable, where $l = 1, 3, 6, 9,$ or 12 months. The out-of-sample forecast evaluation period is January 1983 through December 2011. The forecast errors are generated recursively, starting with a 180-month base (January 1968-December 1982). The Diebold-Mariano (DM) equality of prediction test is estimated with HAC standard errors (rectangular kernel) using $h - 1$ lags to control for autocorrelation.

Forecast Accuracy: 12-months ahead RMSEs						
Annualized percent change over the last;	1 month	3 months	6 months	9 months	12 months	
CPI	2.280	2.092	1.858	1.711	1.657	RMSE
Core CPI	1.914	1.710	1.646	1.655	1.707	RMSE
	<i>1.527</i>	<i>1.219</i>	<i>0.724</i>	<i>0.237</i>	<i>-0.240</i>	<i>DM statistic</i>
16% trimmed-mean CPI	1.691	1.574	1.506	1.489	1.525	RMSE
	3.354 ***	2.163 **	1.653 *	<i>1.513</i>	<i>1.181</i>	<i>DM statistic</i>
Median CPI	1.718	1.607	1.533	1.510	1.550	RMSE
	2.693 ***	1.719 *	<i>1.277</i>	<i>1.089</i>	<i>0.729</i>	<i>DM statistic</i>
* , ** , ***' denotes significance at the 10%, 5%, and 1% level, respectively						

TABLE 3: Out-of-Sample Test of Forecast Accuracy, 12-months ahead

Note: The estimated equation is of the form: $\pi_{t,t+h} = \alpha + \beta x_{t-l,t} + \varepsilon_t$; where $\pi_{t,t+h}$ is the annualized growth rate in headline inflation over the horizon t to $t+h$. $x_{t-l,t}$ is the annualized growth rate in the independent variable, where $l = 1, 3, 6, 9,$ or 12 months. The out-of-sample forecast evaluation period is January 1983 through December 2012. The forecast errors are generated recursively, starting with a 180-month base (January 1968-December 1983). The Diebold-Mariano (DM) equality of prediction test is estimated with HAC standard errors (rectangular kernel) using $h - 1$ lags to control for autocorrelation.

Forecast Accuracy: 6-months ahead RMSEs						
Annualized percent change over the last;	1 month	3 months	6 months	9 months	12 months	
CPI	2.493	2.315	2.094	2.003	1.950	RMSE
Core CPI	2.089	1.897	1.824	1.837	1.877	RMSE
	2.511 **	1.720 *	<i>1.067</i>	<i>0.804</i>	<i>0.428</i>	<i>DM statistic</i>
16% trimmed-mean CPI	1.929	1.836	1.766	1.756	1.774	RMSE
	4.951 ***	2.830 ***	1.925 *	1.839 *	1.737 *	<i>DM statistic</i>
Median CPI	1.971	1.869	1.781	1.762	1.791	RMSE
	4.028 ***	2.191 **	<i>1.468</i>	<i>1.377</i>	<i>1.165</i>	<i>DM statistic</i>
* , ** , ***' denotes significance at the 10%, 5%, and 1% level, respectively						

TABLE 4: Out-of-Sample Test of Forecast Accuracy, 6-months ahead

Note: The estimated equation is of the form: $\pi_{t,t+h} = \alpha + \beta x_{t-l,t} + \varepsilon_t$; where $\pi_{t,t+h}$ is the annualized growth rate in headline inflation over the horizon t to $t+h$. $x_{t-l,t}$ is the annualized growth rate in the independent variable, where $l = 1, 3, 6, 9,$ or 12 months. The out-of-sample forecast evaluation period is January 1983 through June 2013. The forecast errors are generated recursively, starting with a 180-month base (January 1968-December 1982). The Diebold-Mariano (DM) equality of prediction test is estimated with HAC standard errors (rectangular kernel) using $h - 1$ lags to control for autocorrelation.

8 Appendix

Figures A1 and A2 show the information presented in Figures 1 and 2, but only for symmetric trimmed-means (corresponding to the black lines above), i.e., where $\alpha = \beta$. The blue dots identify the lowest-RMSE symmetric trimmed-means.

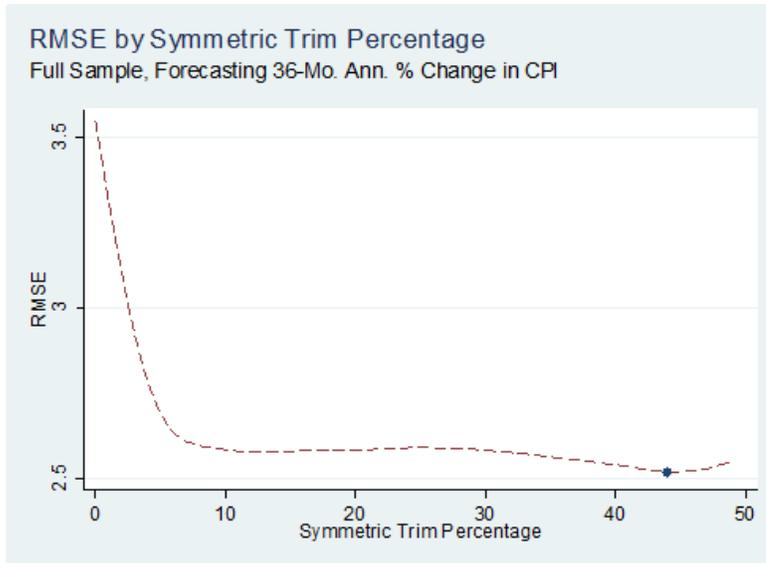


Figure: A1

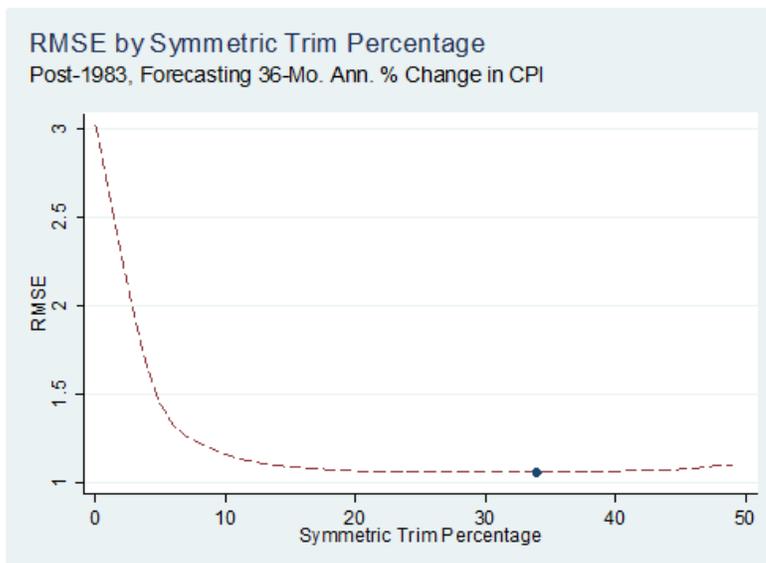


Figure: A2