Federal Reserve Bank of Cleveland Working Paper Series

Improving Inflation Forecasts Using Robust Measures

Randal J. Verbrugge and Saeed Zaman

Working Paper No. 22-23R

May 2023

Suggested citation: Verbrugge, Randal J. and Saeed Zaman. 2023. "Improving Inflation Forecasts Using Robust Measures." Working Paper No. 22-23R. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-202223r.

Federal Reserve Bank of Cleveland Working Paper Series
ISSN: 2573-7953

Improving Inflation Forecasts Using Robust Measures*

Randal Verbrugge
Federal Reserve Bank of Cleveland, US, and
NBER/CRIW

Saeed Zaman
Federal Reserve Bank of Cleveland, US

First version: October 7, 2021
This version: March 31, 2023

Abstract

Theory and extant empirical evidence suggest that the cross-sectional asymmetry across disaggregated price indexes might be useful in forecasting aggregate inflation. Trimmed-mean inflation estimators have been shown to be useful devices for forecasting headline PCE inflation. But is this because they signal the underlying trend or because they implicitly signal asymmetry in the underlying distribution? We address this question by augmenting a "hard to beat" benchmark headline PCE inflation forecasting model with robust trimmed-mean inflation measures and robust measures of the cross-sectional skewness, both computed using the 180+ components of the PCE price index. Our results indicate significant gains in the point and density accuracy of PCE inflation forecasts over medium- and longer-term horizons, up through and including the COVID-19 pandemic. Improvements in accuracy stem mainly from the trend information implicit in trimmed-mean estimators, but skewness information is also useful. An examination of goods and services PCE inflation provides similar inference.

Keywords: median PCE inflation, trimmed-mean PCE, disaggregate inflation, skewness, forecasting

JEL classifications: E31, E37, E52

* We thank Todd Clark, Ed Knotek, Jean-Paul L'Huillier, Pierlauro Lopez, James Mitchell, Rob Rich, and conference participants at the International Institute of Forecasters' 38th International Symposium on Forecasting and the Canadian Economics Association Conference 2022. The views expressed herein are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Cleveland or the Federal Reserve System. Contact information: randal.verbrugge@clev.frb.org and saeed.zaman@clev.frb.org. Research Department, Federal Reserve Bank of Cleveland, P.O. Box 6387, Cleveland, OH 44101.

1. Introduction

Evolution in the value of money – i.e., inflation, or the percentage change in the price level – is a central concern of monetary policy. Accordingly, policymakers at most central banks monitor a range of inflation measures to come to an informed assessment about the underlying inflationary pressures. Over the past decade, increased attention has been paid to trimmed-mean inflation estimators,[1] as these provide signs of any broad-based inflationary pressures or the lack of them (see Mertens, 2016; Verbrugge, 2021).

Recent research has documented the usefulness of trimmed-mean estimators in improving inflation forecasts from a variety of time-series models (e.g., Dolmas, 2005; Mertens, 2016; Meyer and Zaman, 2019; Carroll and Verbrugge, 2019; Ocampo, Schoenle, and Smith, 2022.)[2] The consensus in the literature is that the superior performance of the trimmed-mean estimators in forecasting future inflation results from their ability to signal the trend in inflation. The main rationale behind this consensus is the following: when the underlying distribution is leptokurtic (fat-tailed) and the sample (i.e., the number of components or disaggregates used to compute the aggregate) is not large, as is the case for US inflation,[3] then trimmed-mean estimators are likely to be more accurate estimates of central tendency, compared to the sample mean.

But there is an alternative or complementary explanation for the trimmed-mean estimators' superior predictive performance that has received little attention. In addition to being fat-tailed, as discussed in Section 2, the underlying distribution of inflation components

[1] In this paper, we refer to both median inflation measures (such as the median PCE) and trimmed-mean inflation measures (such as the trimmed-mean PCE) as trimmed-mean measures.
[2] In related and contemporaneous work, Ocampo, Schoenle, and Smith (2022) show that trimmed mean estimators outperform core PCE overall, but during periods when headline inflation is below 2.5 percent, core PCE performs better.
[3] Technically, what matters is not the nominal number of components but rather, given the wide distribution of aggregation weights associated with the components, some notion of an *effective* number.

(disaggregates) is also asymmetric, with the degree of asymmetry evolving slowly over time. Consequently, in forecasting models, when trimmed-mean estimators are added alongside headline inflation measures, as they typically are in practice, the differential between the two provides an implicit signal about the current degree of asymmetry in the underlying distribution of the components. Both theory and extant evidence, reviewed below, suggest that this signal may have notable predictive content. In this paper, we explore this hypothesis and determine the extent to which the superior forecasting performance of trimmed-mean estimators is driven by their implicit signal of asymmetry.

Accordingly, this paper examines both the independent and the joint predictive performance of trimmed-mean estimators and robust asymmetry (skewness) measures to forecast aggregate PCE inflation. Specifically, we make pairwise comparisons of forecast accuracy between univariate, bi-variate, and tri-variate vector autoregressive (VAR) model specifications. In constructing our VAR model specifications, we build upon the "hard-to-beat" Faust and Wright (2013) model, which is a simple univariate AR model in gaps, where the gap is defined as the difference between the inflation measure and long-run inflation expectations of PCE inflation.[4] Our VAR models include additional covariates, a robust skewness statistic, and/or a trimmed-mean inflation measure.

The pairwise comparisons between model specifications allow us to examine both the marginal contribution of skewness measures and trimmed-mean estimators and their joint contribution to potential improvements in the accuracy of PCE inflation forecasts (point and density) above and beyond the univariate AR model in gaps.[5]

[4] Following much of the literature adopting Faust and Wright's univariate AR model, we estimate the slope and intercept parameters of the univariate model. In Section 3, we briefly discuss the advantages.
[5] Our approach should not be confused with more common approaches that posit an asymmetric or nonlinear relationship between slack and inflation (e.g., Ashley and Verbrugge, 2020).

To complete our analysis and provide a broader perspective on forecasting performance, we also assess the accuracy of a model specification embedding the Phillips curve and of a model that has core PCE inflation. Finally, motivated by a growing literature exploring the predictive content of goods and services, we investigate the predictive content of robust measures of goods and services.

Our main finding is that including our robust measures in the AR benchmark forecasting model improves its ability to forecast aggregate PCE inflation. The statistically significant gains in the accuracy of both the point and the density forecasts are achieved for forecast horizons 1.5 years ahead and greater, which are the forecast horizons most relevant to monetary policymakers. Most of the improvements in accuracy are due to the trimmed-mean estimators' ability to signal a trend, with only marginal improvements in their ability to send an implicit signal about the skewness. The statistically significant gains in accuracy are observed over periods when inflation is low, that is, predominantly over the financial crisis and onward sample, including the COVID-19 pandemic period but prior to the inflation surge in mid-2021. We highlight four secondary findings. First, we find slightly stronger support for median PCE over trimmed-mean PCE in forecasting aggregate PCE inflation, and both outperform the exclusion estimator, core PCE. Second, the model specification embedding the Phillips curve is significantly inferior to specifications without the Phillips curve. Third, we generally find support for the Kelly skewness statistic over the Bowley skewness measure. However, the Bowley skewness measure is found to be more useful for estimating stochastic volatility. Last, re-running our analysis separately on goods and services PCE inflation gives results consistent with the findings for headline PCE inflation.

The paper is organized as follows. Section 2 describes the trimmed-mean inflation estimators, the skewness measures, and the data. Section 3 details the model specifications and the design of the forecasting exercise. Section 4 discusses results. Section 5 explores the efficacy of skewness measures for estimating stochastic volatility. Section 6 concludes.

2. Robust Measures and Data

A price index is a stochastic process that is a complicated convolution of thousands of stochastic processes. For example, changes in the personal consumption expenditures price index (PCE price index) are a weighted average of the changes in the indexes of over 180 commodities and services. The weights change over time, reflecting substitution patterns, entry and exit of goods and outlets, and so on.

The evolutions of the underlying stochastic processes are not independent. They reflect a variety of forces such as monetary impulses, changes in transportation costs, transaction technologies and tastes, and productivity growth. They reflect price pressures on *groups* of goods and services. And they reflect idiosyncratic movements as well, themselves driven by changes in information, tastes, technologies, market disruptions, the birth and death of particular outlets, and so on. Any of these influences could be transient or persistent.

One manifestation of the complexity of the evolution of the underlying price process is the cross-sectional distribution of disaggregated component price indexes. Figure 1 depicts a histogram of the monthly inflation rates across 180+ components of the PCE price index for May 2018.

[Figure 1 here]

It is clear that these components experienced significantly different inflation rates in May 2018 and that there are some extreme outliers. The presence of such outliers and the sensitivity of the sample mean to outliers motivate a prominent approach to the estimation of trend inflation: the use of limited-influence inflation estimators, such as a median CPI or trimmed-mean CPI (see Bryan and Cecchetti, 1993 or a median PCE (see Carroll and Verbrugge, 2019) and trimmed-mean PCE (see Dolmas, 2005). Such measures appear to capture trend inflation in as much as they remove noise from inflation, track ex-post measures of its trend,[6] and have been shown to improve inflation forecasting (see, e.g., Smith, 2004; Ball and Mazumder, 2011; Norman and Richards, 2012; and Meyer and Zaman, 2019).

Figure 1 also illustrates that not only is the cross-sectional distribution highly kurtotic, but it is also asymmetric – and typically left-skewed. Indeed, for this reason, the trimmed-mean PCE uses asymmetric trimming. In particular, to ensure that the trimmed-mean PCE price index is unbiased on average over long periods, 24 percent is trimmed from the lower tail, while 31 percent is trimmed from the upper tail (see Dolmas 2005, 2009).

However, the degree of asymmetry is not stable, but changes over time. We illustrate this using two robust asymmetry estimators, Bowley skewness and Kelly skewness statistics, which we define below.

[6] As Dolmas (2005) points out, robust asymmetry estimators are to be preferred, since moment estimators (such as the third centered moment) are all strongly influenced by outliers.

Skewness Statistics: Bowley and Kelly

Following Kim and White (2004) and Dolmas (2005), we define the (weighted) Bowley and Kelly skewness:

$$Skewness\{m\} = \frac{P_a + P_b - 2P_c}{P_a - P_b} \qquad (1)$$

where m refers to Bowley or Kelly skewness, and P_i is the i^{th} percentile of the distribution of component price changes (in a given month), and we have suppressed time subscripts for clarity. When m refers to Bowley, then $a = 75$, $b = 25$, $c = 50$, and when m refers to Kelly, then $a = 90$, $b = 10$, $c = 50$.[7]

For each month, we compute *skewness* statistics over the number of components available.[8] And for each of the skewness statistics, we calculate two measures: one based on disaggregate components' month-to-month (m-o-m) inflation rates and the other based on those components' 12-month trailing inflation rates (y-o-y).

Figure 2 plots Bowley and Kelly skewness measures from 1978 through June 2021 based on disaggregates' 12-month trailing inflation rates.[9] Figure A1, in the online appendix, plots the corresponding skewness measures based on disaggregates' m-o-m rates. Presented are the three-month moving average of these monthly skewness measures. Three observations stand out. First, asymmetry (skewness) displays significant medium-frequency variation. Second, most of the

[7] As implied by the formula, in its construction, the Bowley statistic uses observations in the middle 50 percent of the distribution; that is, it excludes 25 percent of the observations from each tail. Similarly, the Kelly statistic uses observations in the middle 80 percent, excluding 10 percent of the observations from each tail.
[8] Coverage of the PCE has increased over time, particularly in services. For example, in 1960, the Bureau of Economic Analysis (BEA) did not estimate home healthcare consumption, and services such as internet services did not exist. We compute Bowley and Kelly skew statistics using 181 categories of goods and services, which are listed in online appendix A1, Table A2.
[9] Please see Figure A2 in the online appendix for the profile of monthly, 3-month moving average, and 12-month moving average of the Bowley skewness measure, and Figure A3 for the corresponding figures for the Kelly measure.

time, the skew is negative. Third, at times, the two measures of skewness (i.e., Bowley and Kelly) disagree with one another, especially when skewness measures are constructed using disaggregates' 12-month trailing rates. For example, in Figure 2, between 2014 and 2018, the Kelly statistics indicate a strongly negative skew, whereas the Bowley statistics indicate periods in which the skew was positive.

[Figure 2 here]

Why might robust skewness measures have predictive content? There are four reasons. First, leading theories of price-setting behavior (e.g., Ball and Mankiw, 1994) indicate that inflation is linked to asymmetric price adjustment. Second, there is compelling statistical evidence that asymmetry correlates with inflation (e.g., Verbrugge, 1999). Third, as discussed below, a leading approach to estimation of trend inflation involves trimming outliers. To deliver unbiased trend estimates, such trimming must be asymmetric, since asymmetry in the cross-sectional price index distribution would otherwise induce bias (for the same reason that a sample mean departs from a sample median in a skewed sample). However, the degree of this asymmetry is time varying, implying that optimal trimming should similarly be time varying and tied to the current degree of asymmetry. Hence, the time variation in skewness suggests that incorporating information about the degree of asymmetry in empirical models alongside the trimmed-mean estimators may be helpful for forecasting.

Last, the time variation in asymmetry is informative about time variation in the properties of the convolution. Verbrugge (1999) indicates that asymmetry in the cross-sectional distribution is associated with the underlying conditional variance-covariance structure, which is time varying. Accordingly, we hypothesize that a direct estimate of the asymmetry – an estimate that is a nonlinear function of the cross-sectional association or relationship of the underlying stochastic processes –may have beneficial information for inflation forecasting and separately for informing estimates of stochastic volatility in equations defining inflation dynamics.

Median and Skew by Goods and Services

A growing literature has documented the importance of forecasting inflation by separately modeling and forecasting the goods and services sub-categories of aggregate inflation (see Tallman and Zaman, 2017). A recent BIS report (BIS, 2022) advocates looking at a more disaggregated level to better understand the aggregate inflation dynamics. Relatedly, Schoenle and Smith (2022) show that, over time, the US inflationary process has been increasingly driven by idiosyncratic shocks rather than by aggregate shocks; that is, it has become more granular. Motivated by this line of research, we examine whether gains in the accuracy of goods and services inflation forecasts are possible, by computing robust measures (separately) for goods and for services. Furthermore, this decomposition could provide a better understanding of the movements of the aggregate robust measures (e.g., median PCE and the overall skewness). Accordingly, we construct the robust measures (median and skewness) for goods and services PCE. Figure 3 plots median goods PCE inflation and median services PCE inflation alongside median PCE inflation. A quick visual inspection indicates a striking similarity between the median PCE inflation and median services PCE inflation. This suggests that both indexes categorize the median components with similar price changes.[10] To conserve space, Figure A4 in the online appendix plots the skewness measures computed separately by services and goods

[10] Interestingly, in computing the median PCE, over our sample period, about 82 percent of the time (i.e., for 435 out of 533 months), the identified median component belonged to the services category.

categories.

[Figure 3 here]

Data

All of the empirical analysis in the main text uses data at a monthly frequency spanning January 1978 through June 2021.[11] We use data on the personal consumption expenditures price index (PCE), PCE excluding food and energy components (core PCE), and data on both price indexes and nominal expenditure shares of 181 components of PCE.[12] Our target variable of interest is the 12-month PCE inflation rate.[13] Table A1, in the online appendix, provides a complete listing of all the data series, which were retrieved from Haver Analytics.

3. Models and Forecasting Setup

In the inflation forecasting literature, modeling inflation in "gap" form, where the gap is defined as the deviation of inflation from its underlying long-run trend (i.e., long-run inflation), has been shown to be quite helpful in improving the accuracy of inflation forecasts (e.g., Faust and Wright, 2013; Zaman, 2013; Clark and Doh, 2014; and Tallman and Zaman, 2017). In fact, a simple univariate autoregressive (AR) model of inflation in the gap is widely recognized as an "amazingly hard to beat" benchmark (e.g., Faust and Wright, 2013). Accordingly, our design of the forecasting exercise is inspired by modeling inflation in gap form. Specifically, to assess the

[11] In the online appendix, we report selected results based on data spanning July 2021 through December 2022, which was made available after we had completed this paper.
[12] The online appendix (Table A2) lists all of the 181 disaggregated components used to construct the robust asymmetry measure. It is worth mentioning that if we instead use the 153 components that go into constructing core PCE, the resulting estimates of the asymmetry measure are similar to the one obtained with all 181 components.
[13] The Federal Reserve's inflation goal is framed in terms of the 12-month inflation rate in PCE inflation.

marginal contribution of trimmed-mean estimators and skewness measures to improving the accuracy of inflation forecasts, we extend the univariate inflation in the gap model to a multivariate setup.[14] First, we build a bi-variate Bayesian VAR[15] of headline PCE inflation and median PCE inflation, where both inflation measures are specified as deviations from the PTR.[16] We denote this specification *"BVAR: PCE + Median."* We view median PCE inflation as capturing the "medium term" trend in inflation.[17] We model *both* headline and median PCE inflation in deviations from the PTR, to preserve the information implicit in the headline-median gap. If median PCE inflation does convey medium-term trend information, we would expect headline PCE to move toward median PCE inflation over the medium term. We compare the accuracy of the bi-variate BVAR (i.e., *BVAR: PCE + Median)* in forecasting headline PCE inflation to that of the univariate inflation in the gap model. This comparison would give us a sense of the marginal contribution of median PCE inflation above and beyond headline PCE inflation's own history in improving the forecast accuracy of headline PCE inflation. This marginal contribution of median PCE would reflect both the superior measure of the central

[14] Faust and Wright (2013) propose a quarterly AR(1) gap model, and they show that a specification with a fixed slope parameter, rho=0.46 and intercept=0, does slightly better than the unrestricted specification whose slope and intercept are estimated from the data. Since we work with models estimated with data at a monthly frequency, we use a monthly AR(3) gap model. We estimate this AR model, for several reasons. First, since we estimate using Bayesian methods, this conveniently allows us to produce density forecasts. (We note that the point forecasts from the AR model with or without Bayesian estimation are identical.) Second, it is not obvious how a fixed quarterly parameter should be mapped into fixed monthly coefficients. Last, it naturally supports our extension of the univariate AR model to the VAR model by adding the two covariates (trimmed-mean and skewness measures) and allows for a more fair comparison with the univariate AR model.

[15] Bayesian VARs are widely used to forecast macroeconomic variables. We use BVAR models similar to those used in Banbura, Giannone, and Reichlin (2010) and Knotek and Zaman (2019). We set lag length=3 to be consistent with the AR(3) benchmark model. We relegate the BVAR model details to online appendix A2.

[16] PTR is the survey-based long-run (5- to 10-years-ahead) PCE inflation expectations series from the Federal Reserve Board of Governors' FRB/US econometric model.

[17] On its website, the Federal Reserve Bank of Cleveland explicitly states that the median PCE indicator is designed to capture the underlying trend in inflation, where the underlying trend is defined as the "medium-horizon" trend in inflation. Further, when evaluating how well trimmed-mean estimators track the underlying trend inflation, the common practice in the literature is to use the 36-month centered moving average of actual inflation to define trend inflation. These facts support the notion that trimmed-mean estimators best reflect medium-term trend inflation. We recognize that in the literature, researchers often treat trimmed-mean estimators of inflation as reflecting the long-run trend in inflation; however, more recently, there is an increasing recognition that this is not the case.

tendency (signal about the underlying trend) and the implicit signal about the current degree of asymmetry (skewness).

Second, to get a rough approximation of the extent to which skewness contributes to the median PCE's marginal contribution, we construct another bi-variate BVAR of headline PCE and a skewness measure (either Bowley or Kelly). We denote this specification as *"BVAR: PCE + Skew (B),"* when the skewness measure is Bowley, and *"BVAR: PCE + Skew (K),"* when the skewness measure is Kelly. A comparison between this bi-variate BVAR's accuracy in forecasting headline PCE inflation and the one estimated in the previous step, along with the comparison of this bi-variate BVAR with the univariate headline PCE inflation model, would give us a sense of the extent to which skewness is contributing to the marginal contribution of median PCE relative to the signal about the trend to improve the forecast accuracy of headline PCE inflation.

Third, we construct tri-variate BVARs, which incorporate headline PCE, median PCE, and skewness (either Bowley skewness, denoted *"BVAR: PCE + Median + Skew (B),"* or Kelly skewness, denoted *"BVAR: PCE + Median + Skew (K)"*). The comparison of the tri-variate BVAR to the corresponding bi-variate BVAR would give us a sense of the marginal contribution of the "direct" measure of skewness above and beyond that of median PCE and headline PCE, noting that median PCE already embeds an *implicit* signal about the skewness (when added alongside the headline PCE). Similarly, comparing the tri-variate BVAR with the univariate model would give us a sense of the combined usefulness of median PCE and the "direct" measure of skewness in improving the forecast accuracy of headline PCE inflation.

We repeat this exercise by replacing median PCE with the trimmed-mean PCE, which gives us a bi-variate BVAR, *"BVAR: PCE + Trim,"* and tri-variate BVARs *"BVAR: PCE +*

Trim + Skew (B)” and *“BVAR: PCE + Trim + Skew (K).”* Then we repeat replacing the trimmed-mean PCE with core PCE, which gives us bi-variate BVAR, *“BVAR: PCE + Core,”* and tri-variate BVARs *“BVAR: PCE + Core + Skew (B)”* and *“BVAR: PCE + Core + Skew (K).”*

Fourth, we assess the value added of our robust measures in improving the accuracy of the inflation forecasts from the Phillips curve specifications. A long list of papers have documented the inferior accuracy of forecasts from the Phillips curve models compared to forecasts from models with univariate specifications (e.g., Faust and Wright, 2013). More recently, Ball and Mazumder (2020) and Ashley and Verbrugge (2020) show the competitive accuracy of the inflation forecasts from Phillips curve models based on trimmed-mean inflation measures. Accordingly, we examine whether the inclusion of median PCE (or trimmed-mean PCE) and skewness in the Phillips curve specification helps improve accuracy. If it does, are the gains large enough to make the accuracy of the forecast competitive with the univariate benchmark? To preview the result, we find that inclusion of the robust measures helps to improve the forecast accuracy of the Phillips curve model, but the gains are small: the accuracy of the forecasts remains significantly inferior compared to the univariate benchmark. Because of the small gains in accuracy, in the interest of brevity, and to facilitate comparison, we simply report the forecast accuracy from the Phillips curve specification without the robust measures, which we denote as *“BVAR: PCE + UR,”* where UR refers to the unemployment rate gap constructed as the difference between the unemployment rate and the CBO’s estimate of the natural rate of unemployment.

Fifth, to assess the usefulness of robust measures of goods and services inflation in improving the accuracy of goods and services inflation forecasts, we perform two sets of forecasting exercises similar to those described previously. Specifically, in the first exercise, we assess the predictive ability of the robust measures (median and skewness) for goods inflation by estimating three separate BVAR models: "BVAR: G. PCE + Skew (K)," which is a bi-variate VAR of goods PCE inflation and Kelly skewness based on goods inflation; "BVAR: G. PCE + Median," which is a bi-variate VAR of goods PCE inflation and median goods PCE inflation; and "BVAR: G. PCE + Median + Skew (K)," which is a tri-variate VAR of goods PCE inflation, median goods PCE inflation, and Kelly skewness based on goods inflation. In the second exercise, we assess the predictive ability of the robust measures for services inflation by estimating three separate BVAR models: "BVAR: S. PCE + Skew (K)," which is a bi-variate VAR of services PCE inflation and Kelly skewness based on services inflation; "BVAR: S. PCE + Median," which is a bi-variate VAR of services PCE inflation and median services PCE inflation; and "BVAR: S. PCE + Median + Skew (K)," which is a tri-variate VAR of services PCE inflation, median services PCE inflation, and Kelly skewness based on services inflation. For goods and services inflation, we focus on the Kelly skewness measure because, as discussed in the results section, Kelly skewness outperformed Bowley skewness in all the exercises involving aggregate PCE inflation.

Pseudo-Out-of-Sample Forecasting

Even though we have real-time data available for aggregate PCE inflation and the unemployment rate, the availability of real-time data at the disaggregate component level (required to compute the median PCE and skewness) is limited; therefore, we resort to pseudo-out-of-sample forecast evaluation. We perform forecasting evaluation using a recursively expanding window of estimation. All the models, including the univariate AR gap model, are estimated using Bayesian methods, which facilitates computation of the density forecasts. The estimation sample starts in

January 1978 and forecast evaluation is performed over the sample from January 1994 through June 2021. At each recursive run, forecasts are produced up to three years out (i.e., the forecast horizon, h, ranges from h=1 to h=36 months ahead). The models produce forecasts of the PCE inflation "gap," which are then converted into forecasts of the PCE inflation rate by adding to the forecasts of the inflation "gap" the latest estimate of the PTR available at each recursive run. The point forecasts, which are the posterior mean of the density forecasts, are evaluated using the metric of the mean squared forecast error (MSE). To assess the statistical significance of gains in the accuracy of point forecasts between the two models, we use the Diebold and Mariano test (with the Newey-West correction) using the two-sided tests of the standard normal. The density forecasts are evaluated using the widely used metric of the logarithmic predictive score (parametric normal approximation), and the statistical significance is assessed using the likelihood-ratio test of Amisano and Giacomini (2007), where the test statistics use a two-sided t-test.

4. Forecasting Results

Table 1 reports the results of the point forecast evaluation comparing inflation forecast accuracy across several model specifications. The results correspond to model specifications that use Kelly skewness measures constructed based on disaggregates' month-to-month inflation rates;[18] we compute the three-month moving averages as our estimates of the skewness measures that enter the models.[19]

[18] The results based on model specifications in which skewness measures are constructed based on disaggregates' 12-month trailing inflation rates are found to be inferior compared to those obtained using skewness measures constructed from month-to-month inflation rates. Owing to space constraints, we do not report these results in the paper, but they are available upon request from the authors.

[19] The three-month moving average was preferred to other window lengths (e.g., 1, 2, 4, 6, 8, 10, 12).

We find that Kelly skewness contains more predictive content for inflation than does Bowley skewness (see online appendix Tables A4 and A5). The findings that Kelly is preferred to Bowley, that the skewness constructed from the components' month-to-month inflation rates is preferred to the corresponding 12-month trailing rates, and that the three-month window for the moving average of monthly skewness is preferred to other window lengths suggest that for skewness to be useful in forecasting PCE inflation, it matters how the skewness measure is constructed.

To conserve space, the forecast accuracy is reported for select forecast horizons. The top panel of the table reports results corresponding to the full sample (1994-2021), the middle panel corresponds to the pre-Great Recession sample (1994-2007), and the bottom panel corresponds to the financial crisis and onward sample (2008-2021). In each panel, the numbers reported in the first row are the root mean squared error (RMSE) from the benchmark univariate inflation in the gap model, denoted "*AR(3)-PCE*." And the rows below it are ratios that report relative MSEs (relative to MSEs from the AR(3) PCE). Thus, a ratio of more than 1 indicates that the univariate inflation in the gap model is more accurate on average than the model being compared.

The results reported in Table 1 suggest four observations. First, adding trimmed-mean estimators to the model improves the forecast accuracy of the aggregate PCE inflation forecasts for most forecast horizons but worsens forecast accuracy in the near term. The gains in forecast accuracy are greater from including the median PCE than from the trimmed-mean PCE and core PCE. In addition, a larger number of gains in accuracy are classified as statistically significant in the case of median PCE compared to the other two, especially for forecast horizons 18 months ahead and greater.[20]

[20] Our finding that the inclusion of median PCE improves the forecast accuracy of aggregate PCE inflation over the medium- to longer-term horizons is consistent with the findings of Crone et al. (2013), who find similar support for

Second, the inclusion of the Kelly skewness measure with or without the inclusion of trimmed-mean estimators marginally improves the forecast accuracy of the aggregate PCE inflation for most forecast horizons. But for the near-term forecast horizon, Kelly skewness plays a nontrivial role, since its inclusion improves forecast accuracy, primarily by converting statistically significant losses to insignificant losses of smaller magnitude. However, in the sample before the Great Recession, skewness measures did not help improve accuracy.

A deeper examination of the errors, combined with an understanding of the behavior of median PCE and headline PCE inflation, allows us to understand when skewness is useful. It is well recognized that headline PCE inflation moves toward median PCE over time, to close the gap between the two; this explains our robust finding that trimmed-mean indicators (such as median PCE) help improve forecast accuracy at the medium horizon. However, there can be persistent deviations between the two, likely due to persistent relative price shocks. In other words, sometimes it takes a long time for the gap to close, with the period spanning mid-2012 through late 2016 being a prominent example. During this period, forecasts from a model including median PCE are biased upward compared to forecasts from a univariate AR model. Adding the skewness measure to this model, which is negative during this period, generates forecasts that call for less strong inflation. With headline PCE inflation running low, this proved to be more accurate. Similarly, the forecast from an AR model calls for inflation to move up gradually toward the end point estimate implied by the PTR (which during that period is higher than PCE inflation). Adding the skewness measure (which is negative) to this AR model generates a forecast that has inflation moving up more slowly than that coming from the AR model, and with actual inflation continuing to be low, this proves to be a more accurate forecast.

median CPI inflation in forecasting aggregate CPI inflation.

As we will discuss shortly, starting mid-2021, skewness again proved useful (especially early on) because of the persistent nature of the relative price shocks. Prior to the Great Recession period (i.e., from 1994 to 2007), deviations between headline and median PCE were relatively short-lived; hence, skewness played only a minor role.

Third, consistent with the findings in the literature, the bi-variate Phillips curve specification significantly underperforms. The density forecast evaluation results echo the point forecast evaluation results. In the interest of brevity, the results and discussion are relegated to the online appendix A4.

[Table 1 here]

Forecasting Performance during the Great Recession and the COVID Pandemic

We next illustrate the marginal efficacy of our robust inflation measures in forecasting aggregate PCE inflation during crisis periods, which are periods normally associated with heightened uncertainty. We focus on two crisis periods: the great financial crisis (GFC; also known as the Great Recession) and the great pandemic crisis (GPC), which is still ongoing at the time of writing. Specifically, we examine the forecasting performance of our BVAR models for 12-months-ahead forecasts generated during the GFC period spanning October 2007 through June 2009, and the GPC period spanning March 2020 through June 2020. For the latter, i.e., the GPC period, we go only through June 2020 because at the time we compiled results, the available data end in June 2021, which we need to evaluate the 12-months-ahead forecast.

Figure 4, panel (a) plots the forecast errors over the GFC period from three models: the benchmark gap AR(3)-PCE, the BVAR: PCE + Skew (K), and the BVAR: PCE + Median. (The plot for BVAR: PCE + Median + Skew (K) is almost identical to that for the BVAR: PCE +

Median; therefore, we do not show it.) As is evident by big misses, all three models generate forecasts that poorly track the actual PCE inflation during the GFC period. However, the model that includes median PCE inflation experiences relatively smaller errors than the univariate benchmark. During that period, actual PCE inflation came in well below the models' projections, resulting in large errors. Panel (b) in Figure 4 plots the forecast errors from the BVAR: PCE + Trim model alongside the PCE + Median model. Both models performed comparably during this period.

Figure 5, panel (a) plots the forecast errors over the GPC period from the same three models. Again, there is evidence of big misses: all three models generate forecasts during the GPC period that do an inferior job of tracking the actual PCE inflation. However, the model that includes median PCE inflation experiences relatively smaller errors than the univariate benchmark. During this period, actual PCE inflation came in well above the models' projections, resulting in large errors. Panel (b) in Figure 5 plots the forecast errors from the BVAR: PCE + Trim model alongside the PCE + Median model. Both models performed comparably during this period, with the model that includes the trimmed-mean measure performing slightly better than the model with the median measure. It is worth noting that the models' forecast errors during the GPC period are smaller in magnitude than during the GFC period.

Since the writing of this paper, additional data covering the period mid-2021 through December 2022 have become available. In the online appendix A8, we briefly discuss and illustrate the forecasting performance of our selected models over this recent period. This period provides a nice illustration of the forecasting benefits that robust measures can provide under particular circumstances. Early in this period, when inflation began to pick up, it was driven by price spikes in a few sectors (see Almuzara and Sbordone, 2022); thus, trimmed-mean indicators were somewhat late in recognizing the persistent nature of the inflation surge.[21] Later in the period, inflationary pressures became broad-based and elevated. Schoenle and Smith (2022) and Ocampo, Schoenle, and Smith (2022) find that during periods when inflation is broad-based and elevated, trimmed-mean estimators behave similarly to headline PCE; hence, their inclusion does little to improve forecast accuracy. In keeping with these facts, we find that over this period, models with or without trimmed-mean estimators generally perform comparably. However, models with "direct" estimates of skew provided more accurate forecasts. Skewness was unusually positive over this period, and model specifications including skewness projected higher inflation than those that did not.

Overall, the forecast results for the GFC, GPC, and post-pandemic (mid-2021 through December 2022) periods highlight the difficulties in accurately forecasting aggregate PCE inflation. Having said that, one is better off incorporating information from trimmed-mean estimators and Kelly skewness in constructing forecasts of PCE inflation using popular time-series models.

[Figure 4 here]

[Figure 5 here]

Breakdown by Goods and Services

Table 2 reports point forecast evaluation results for goods PCE inflation. Shown are the results for the full sample and two sub-samples. In each panel, the numbers reported in the first row are the RMSE from the univariate model of goods PCE inflation, denoted "*AR(3)-Goods PCE.*"[22] And the three rows below it are ratios that report relative MSEs (relative to the MSE from the

[21] Verbrugge and Zaman (2023) find that, over this period, inclusion of a supply-shock variable greatly improved the forecast from a trimmed-mean model.
[22] Tallman and Zaman (2017), among others, document the superior forecasting properties of the univariate model of goods PCE (and CPI) inflation.

AR(3)-Goods PCE). Thus, a ratio of more than 1 indicates that the univariate model is more accurate on average than the model being compared. The other three models shown are: the BVAR: G. PCE + Skew (K), which is a bi-variate VAR model of goods PCE inflation and a skewness measure computed from the disaggregate components belonging to the goods PCE category; the BVAR: G. PCE + Median, which is a bi-variate VAR model of goods PCE inflation and median goods inflation computed from the disaggregate goods components' inflation rates; and the BVAR: G. PCE + Median + Skew(K), which is a tri-variate VAR model of goods PCE inflation, median goods PCE inflation, and a skewness measure based on goods PCE inflation.

As is evident from Table 2, most entries for the full sample and the financial crisis and onward sample are below one, suggesting the usefulness of the robust measures in improving the point forecast accuracy of goods PCE inflation. However, most of the gains are statistically significant only for the financial crisis and onward period, and that too for models that include the median goods PCE inflation. Similar to the results for headline PCE inflation, the addition of skewness only helps marginally, which suggests that the forecasting prowess of median goods PCE inflation is due to its ability to signal the underlying trend in goods inflation. In contrast to the results for headline PCE inflation, the addition of robust measures worsens the forecast accuracy of goods PCE inflation over the pre-financial crisis period, though the losses are statistically insignificant.

Table 3 reports the corresponding results for services PCE inflation. Similar to the results for goods PCE inflation and headline PCE inflation, the evidence suggests that the addition of robust measures to the univariate gap model of services PCE inflation improves the point forecast accuracy of services PCE inflation. Improvements in accuracy are achieved over the financial crisis and onward period, and the bulk of the gains come from the addition of the median services PCE inflation measures with only marginal improvements from the skewness measure. A comparison between Tables 2 and 3 (bottom panels) suggests that skewness's marginal contribution to improving forecast accuracy is greater for services PCE inflation than for goods PCE inflation.

5. The Usefulness of Skewness for Stochastic Volatility Modeling

As noted in Section 2 above (and in Verbrugge, 1999), asymmetry in the cross-sectional distribution is associated with the underlying (time-varying) conditional variance-covariance structure. This leads to a natural curiosity about whether estimates of skewness could help improve the (quarterly) estimates of stochastic volatility in model equations defining inflation dynamics. To help answer this question, we use the state-of-the-art stochastic volatility in the mean model developed in Chan (2017). Keeping the same notation as in Chan, we list the model equations below:

$$y_t = \tau_t + \alpha_t e^{h_t} + \varepsilon_t^y, \qquad \varepsilon_t^y \sim N(0, e^{h_t}) \tag{7}$$

$$h_t = \mu + \phi(h_{t-1} - \mu) + \beta X + \varepsilon_t^h, \qquad \varepsilon_t^h \sim N(0, \sigma^2) \tag{8}$$

$$\gamma_t = \gamma_{t-1} + \varepsilon_t^\gamma, \qquad \varepsilon_t^\gamma \sim N(0, \Omega) \tag{9}$$

where y_t refers to the observed variable of interest (e.g., inflation), h_t refers to the log-volatility, $\gamma_t = (\alpha_t, \tau_t)'$, and Ω is a 2 x 2 covariance matrix. Because the model allows for time-varying parameters and volatility feedback (that is, estimated volatility could affect the level of inflation (equation 7), the literature refers to the above model as a time-varying parameter stochastic volatility in mean model (TVP-SVM).

In Chan (2017), the variable X in equation (8) is one-period lagged inflation to capture the potential influence of past inflation on current inflation volatility. In our exercise, we estimate the above model by replacing past inflation with the skewness measure in variable X. We estimate the above model separately for aggregate PCE inflation, services PCE inflation, and goods PCE inflation, along with their corresponding skewness measures.[23]

Our objects of interest are the estimates of the parameters β and e^{h_t}. To assess whether skewness provides timely and useful information for estimates of stochastic volatility, we would require the estimate of the parameter β to be significant (when assessed using 68 percent credible intervals), and would expect visual evidence indicating some difference in the estimate of stochastic volatility e^{h_t} relative to the estimate of SV coming from the (default) model specification, which includes past inflation in variable X.

Table 4 reports the estimates of the parameter beta for various model runs. We report the model runs with the Bowley skewness measure, because it was found to be notably more influential compared to the Kelly skewness measure in the estimation of SV. A few observations stand out. First, in all three cases (headline PCE, services PCE, and goods PCE), for the default setting, i.e., where X contains one-quarter lagged inflation, the estimates of beta are trivial and insignificant. Second, for goods PCE inflation, the estimate of beta is of nontrivial magnitude and significant. However, in the case of services PCE inflation, beta is insignificant, though the magnitude of the posterior mean estimate is larger than the estimate based on the default setting. Also, the estimate of beta for headline PCE inflation is significant and nontrivial. Third, whereas

[23] To conserve space, we refer the reader to Chan (2017) for model estimation details. The Matlab code to estimate the TVP-SVM model is available to download from Joshua Chan's website.

in the case of goods PCE inflation, the posterior mean estimate of beta is highly positive, in the case of both services PCE inflation and headline PCE inflation, it is negative. This suggests that an increase in the skewness of goods PCE inflation is associated with increased volatility in goods PCE inflation. In contrast, an increase in the skewness of services PCE inflation is associated with reduced volatility of services PCE inflation, and an increase in the skewness of headline PCE inflation is associated with reduced volatility of headline PCE inflation.

Why is beta estimated to be negative (and statistically significant) for headline inflation (and negative for services inflation), but positive (and statistically significant) for goods inflation? We conjecture that the answer lies in the interaction of these factors: relative price shocks of various types, inflation, inflation expectations and inflation uncertainty, and monetary policy.

In his Nobel lecture, Friedman (1977) postulated that higher inflation generates higher inflation uncertainty. Since then, a large number of studies have examined the empirical relationship between inflation and inflation uncertainty, with mixed results. Ball (1992) found support for Friedman's claim, while Holland (1995) suggested a negative relationship. Holland (1995) argued that higher inflation uncertainty is viewed as costly by the central bank. So, when inflation is high, the central bank will act to reduce inflation uncertainty, resulting in a negative relationship between inflation and inflation uncertainty. More recently, Chan (2017) reconciled the contrasting findings of earlier work by demonstrating that the inflation/uncertainty relationship is time varying; in particular, after about 1990 or so, the correlation between headline inflation and uncertainty switched signs, so that higher inflation was correlated with *lower* uncertainty. We confirm this finding for headline inflation; but looking at the disaggregated level, we show that for goods inflation, its relationship with uncertainty is moderately positive, unlike for services inflation, which exhibits a negative relationship. (The estimated relationships between inflation categories and volatility, i.e., the estimates of time-varying parameter α_t in equation 7, we plot in Figure A5 in the online appendix.)

Next, note that skewness and inflation are robustly positively correlated (unconditionally), for goods, for services, and for headline inflation.[24] Hence, we would expect a negative correlation between skewness and headline inflation (and services inflation) uncertainty over the recent period, and a moderately positive relationship in the case of goods inflation – and this is what we find. Last, we conjecture a role for monetary policy in influencing the signs of these relationships.

There is a consensus among central bankers that movements in goods inflation are typically transitory, so they tend to look past those movements. Conversely, it is well established that households' short-term inflation expectations respond strongly and immediately to relative price shocks in goods (i.e., there is a strong positive contemporaneous correlation between goods inflation and consumer one-year-ahead expectations, whereas there is a weak positive correlation in the case of services inflation, which we confirmed). If monetary policy does not systematically respond to relative price shocks in goods, but the public's inflation expectations do respond strongly to those shocks, it is reasonable to expect the positive correlation between goods inflation and inflation uncertainty (and hence, skewness and inflation uncertainty) that is present in the historical data. Conversely, central bankers closely monitor any developments in headline inflation that are driven by services inflation, given that these are typically persistent and highly correlated with wage inflation.[25] Accordingly, when services inflation (and, in turn, headline

[24] The robust correlation between inflation and skewness was first established in Verbrugge (1999), and we find that this relationship holds true for goods and services inflation as well. We remind the reader that there is not a simple mapping between skewness in goods inflation, skewness in services inflation, and skewness in headline inflation.
[25] We find a much stronger positive correlation between the federal funds rate (the shadow federal funds rate) and

inflation) increases, the public's inflation expectations respond weakly, but the central bank responds with appropriate policy. Hence, the uncertainty associated with headline and services inflation decreases, in turn, displaying a negative relationship between uncertainty and skewness.

Figure 6 plots the (full-sample smoothed) estimates of stochastic volatility for goods PCE inflation (panel a) and headline PCE inflation (panel b). Each panel shows two plots: one labeled "Default," which refers to the model estimation that uses lagged inflation, and the other labeled "Skew," which refers to the model estimation that instead uses skewness measures. A comparison of these two plots within each panel provides us with an assessment of the practical usefulness of the skewness measure for the SV estimation. The plots provide some evidence in support of the skewness measure for goods PCE inflation, as evidenced by the improved precision (defined as the width of the 68 percent credible intervals) of the SV estimates, and the visible differences in the SV estimates from the two approaches during specific periods. Again, in the case of headline PCE inflation, there is some evidence supporting the skewness information, since, during several periods, differences in the estimates of SV are observed. However, there is no evidence of improved precision; if anything, there appears to be a slight worsening in the precision of the SV estimates. Overall, there seems to be some evidence in support of skewness for SV estimation, but economically it does not appear to be meaningful.

[Figure 6 here]

6. Conclusion

This paper explores the usefulness of the trimmed-mean estimators and robust skewness statistics in improving the point and density accuracy of aggregate PCE inflation forecasts. Trimmed-

services inflation compared to goods inflation at the quarterly frequency in the sample spanning 1978 through 2021.

mean estimators have been shown to do well in forecasting aggregate inflation, with the forecast accuracy gains thought to be due to their prowess in tracking the underlying trend. However, we illustrate strong evidence of time variation in the cross-sectional asymmetry computed using the 180+ components of the PCE price index. Such asymmetry is correlated with inflation, suggesting a second reason that trimmed-mean estimators have predictive content: the gap between headline inflation and trimmed-mean inflation provides an implicit signal about skewness. We assess the predictive content of skewness, independent of the information about the future trend embedded within trimmed-mean estimators.

We examine both the joint contribution of these measures and their marginal contributions in possibly improving the point and density forecast accuracy of PCE inflation. Among the trimmed-mean estimators, median PCE inflation's ability to forecast future headline PCE inflation has barely been explored. So, an important secondary contribution of this paper is to examine the usefulness of median PCE in forecasting aggregate PCE inflation. A third important contribution of this paper is to introduce, and examine the usefulness of, median goods PCE and median services PCE – and their respective robust skewness estimates – for forecasting goods PCE and services PCE. Finally, we explore whether robust measures are useful in stochastic volatility modeling.

Based on a forecast evaluation sample covering the period from January 1994 through June 2021, a period that includes large volatility in oil prices, a financial crisis and deep recession, and a severe global pandemic, our results indicate significant gains in the point and density accuracy of PCE inflation forecasts for horizons 18 months ahead and longer. Most of the improvements come from the inclusion of trimmed-mean estimators, with only marginal improvements from the addition of robust skewness estimators. A split sample examination suggests that most of the gains in accuracy are concentrated in the sample spanning the Great Recession and onward, i.e., January 2008 through June 2021, which is a period during which inflation remained low.

We find slightly stronger support for median PCE over trimmed-mean PCE, and both outperform the exclusion estimator, core PCE. We find strong support for Kelly skewness over Bowley skewness; furthermore, it matters whether skewness measures are constructed using the disaggregate components' month-over-month inflation rates or 12-month trailing inflation rates. In our empirical exercises, skewness measures constructed based on components' month-over-month rates proved useful; in contrast, skewness measures based on 12-month rates marginally worsened accuracy, even though aggregate PCE and trimmed-mean estimators enter the models as 12-month trailing rates.

Using a state-of-the-art stochastic volatility in the mean model, we illustrate the modest efficacy of the skewness measure in refining the *contemporaneous* estimates of stochastic volatility in the innovations to the equation defining goods PCE inflation and, in turn, headline PCE inflation.

Over time, the reliance on trimmed-mean inflation estimators as a means of obtaining a signal about both the underlying trend in inflation and future inflation has increased globally. Hence, we view our empirical findings as useful for a broad swath of practitioners interested in forecasting inflation.

Figure 1: Cross-sectional distribution of inflation in PCE price index components, May 2018

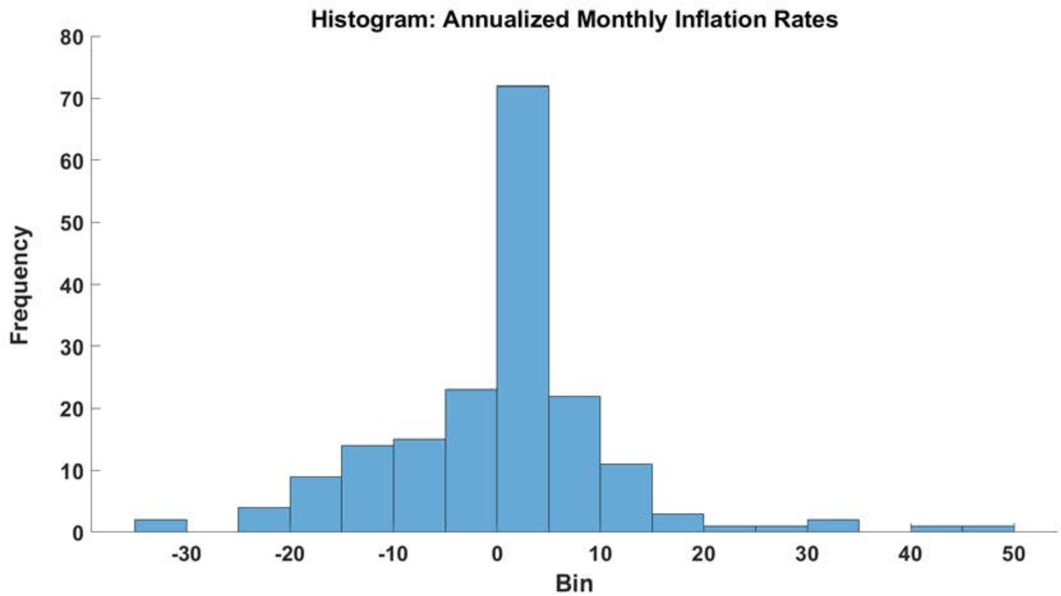

Figure 2: Cross-sectional asymmetry in PCE inflation (12-month %)

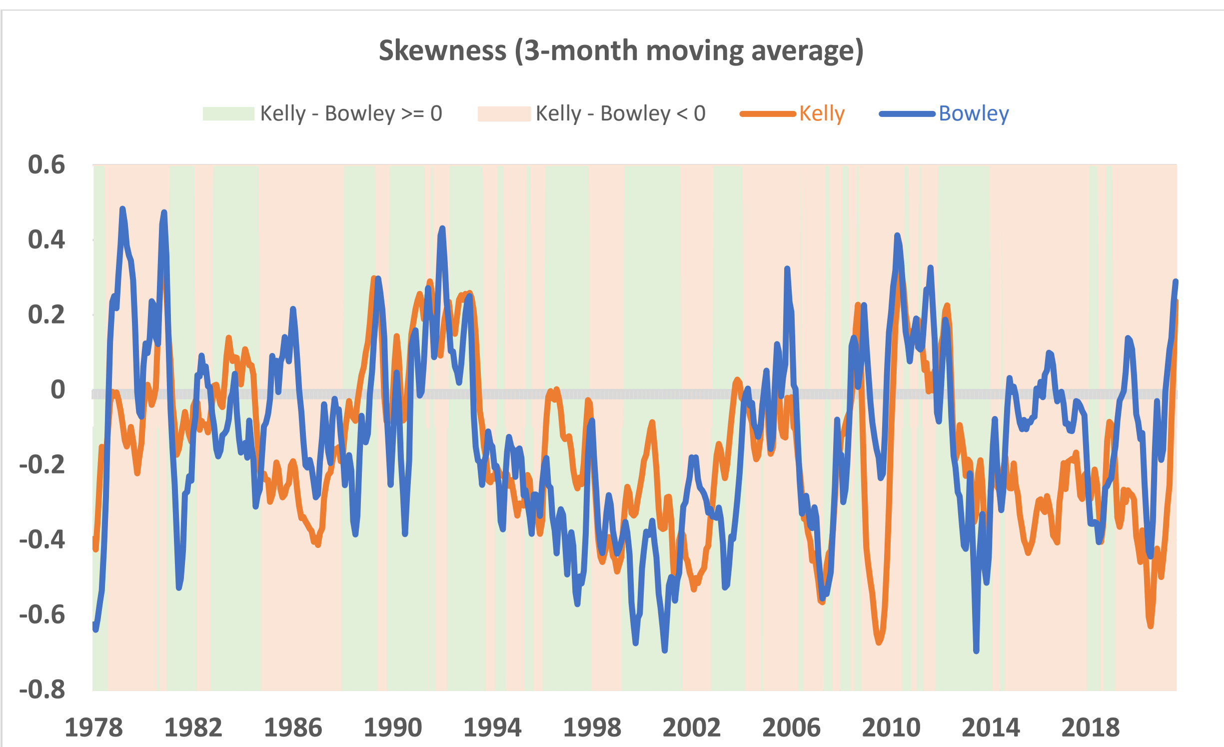

Figure 3: Median by goods and services

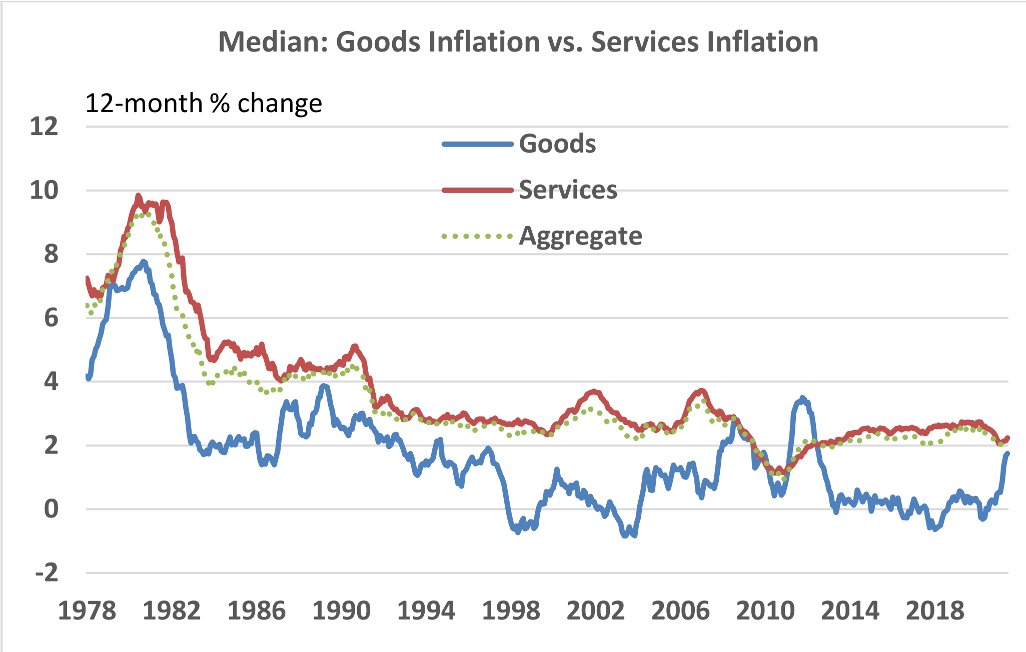

Figure 4: Forecast errors during the Great Recession

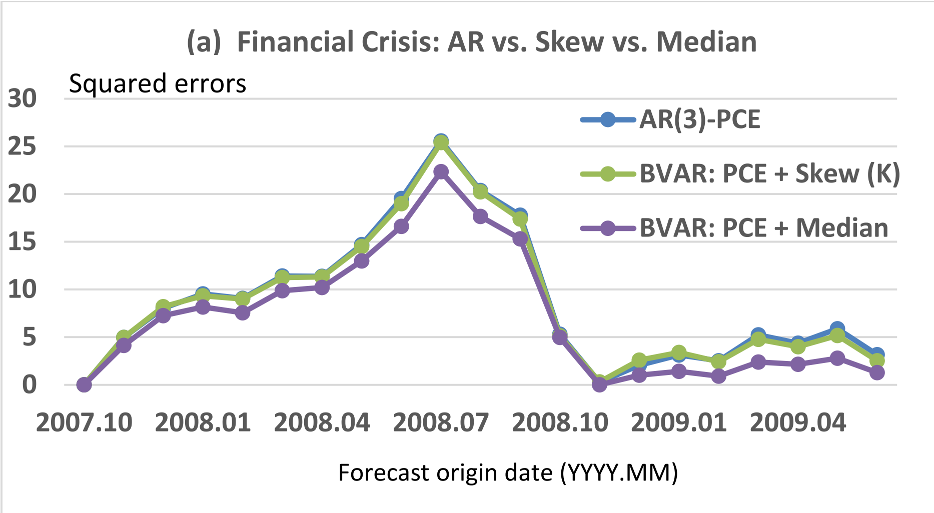

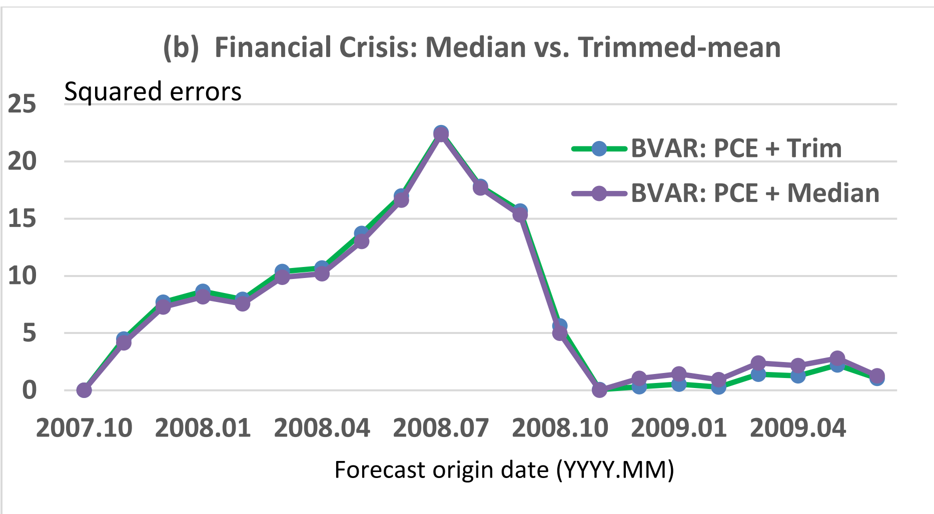

Figure 5: Forecast errors during the great pandemic crisis (COVID-19)

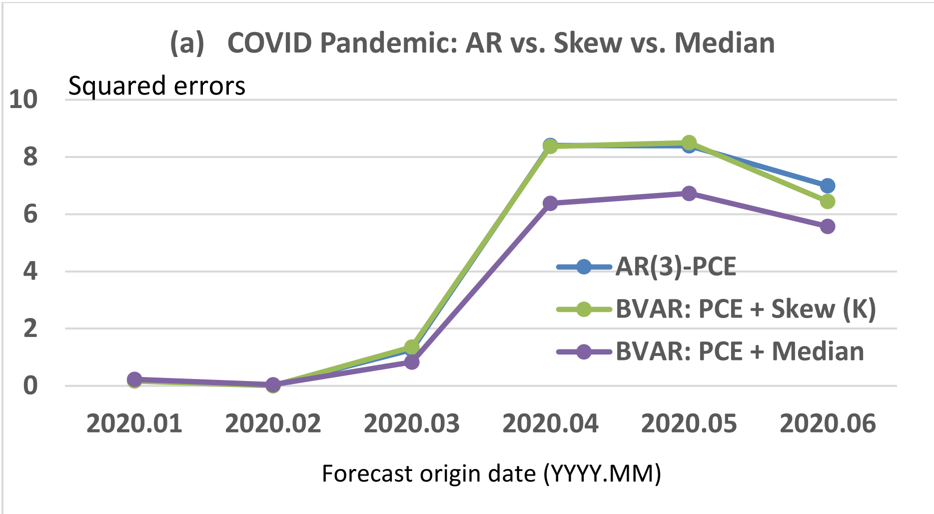

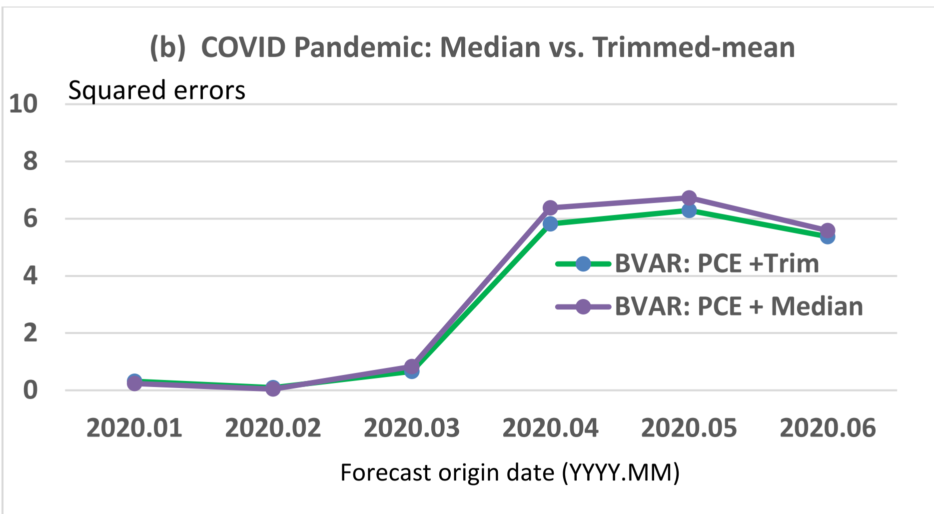

Figure 6: Estimates of SV and precision

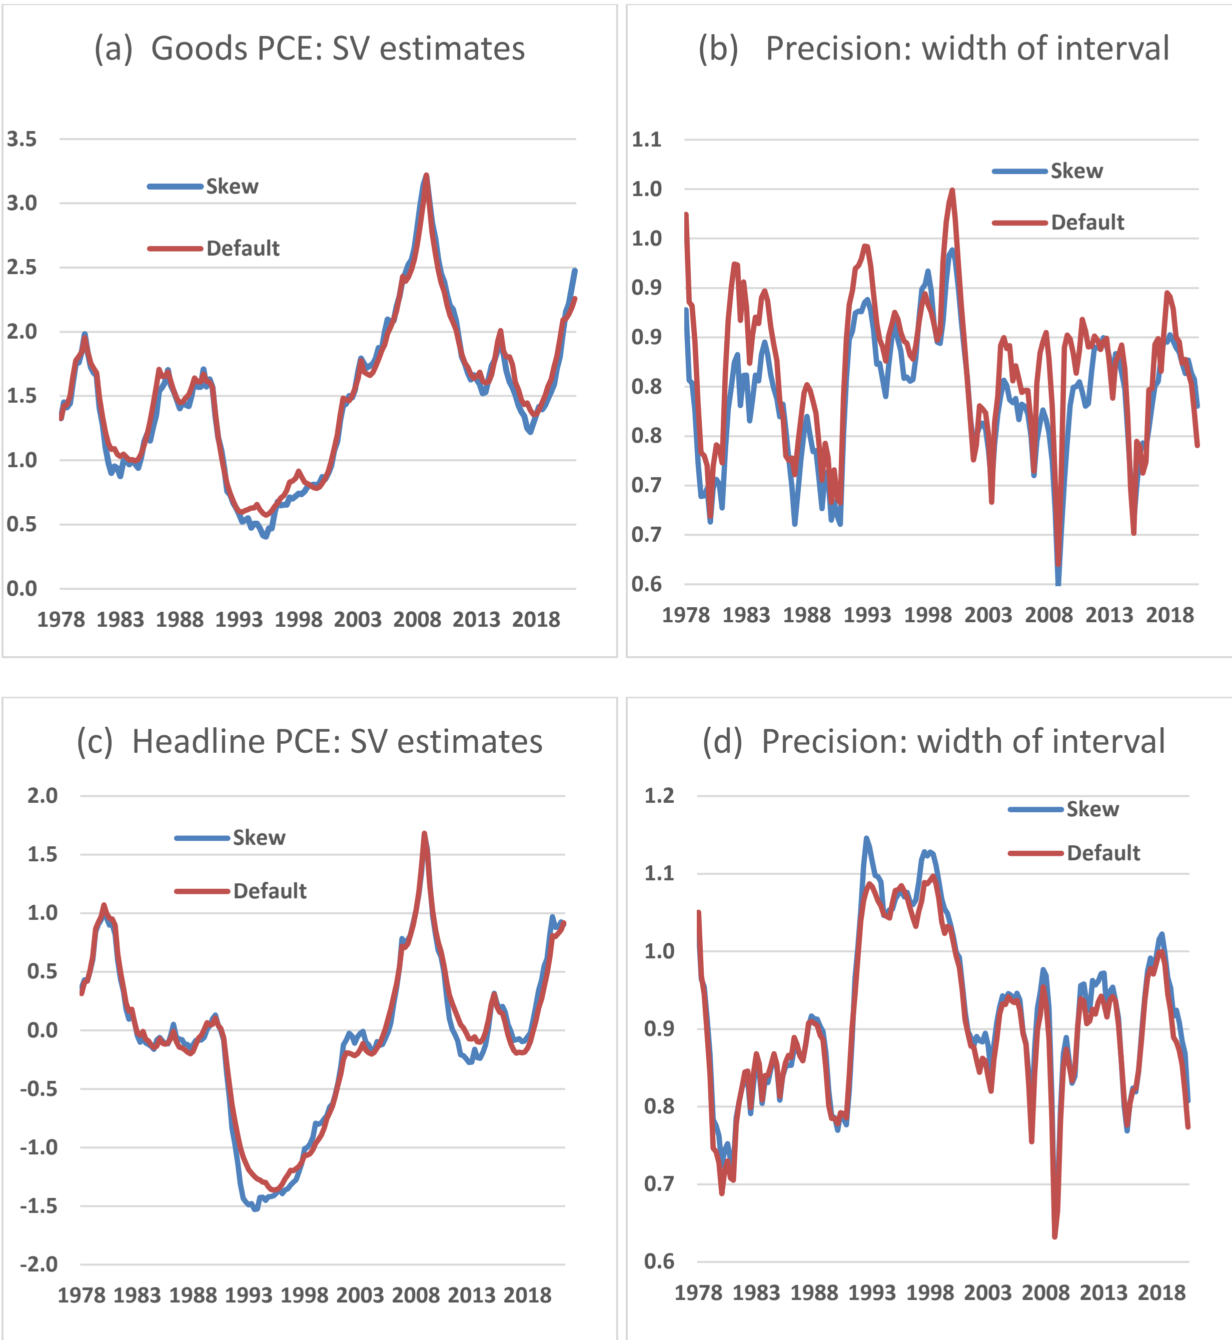

Notes: Panels (a) and (c) plot the posterior mean estimates of the parameter e^{h_t} from the TVP-SVM model specification with lagged inflation (denoted Default), and from the model specification with Bowley skewness (denoted Skew). Panels (b) and (d) plot the corresponding parameters' precision estimates (defined as the width of the 68% credible intervals).

Table 1: PCE inflation out-of-sample point forecasting comparison
[Skew constructed based on month-over-month inflation rates]

Full sample (January 1994 – June 2021)

	h=1M	h=6M	h=1Y	h=18M	h=2Y	h=30M	h=3Y
AR(3)-PCE RMSE	0.265	0.858	1.126	1.064	1.075	1.077	1.044
Relative MSE							
BVAR: PCE + Skew (K)	1.028	0.957*	0.988	0.976	0.959*	0.959*	0.967*
BVAR: PCE + Median	1.046*	0.991	0.893	0.882*	0.879*	0.898*	0.887*
BVAR: PCE + Median + Skew (K)	1.008	0.909	0.889	0.887*	0.876*	0.897*	0.885*
BVAR: PCE + Trim	1.045*	0.997	0.891	0.918	0.913	0.916	0.913*
BVAR: PCE + Trim + Skew (K)	1.011	0.916	0.885	0.922	0.906	0.911	0.911*
BVAR: PCE + Core	1.045*	1.010	1.008	0.997	0.980	0.967*	0.973
BVAR: PCE + Core + Skew (K)	1.045	1.010	1.008	0.997	0.980	0.967*	0.973
BVAR: PCE + UR	1.109*	1.181	1.320*	1.485*	1.628*	1.634*	1.612*

Pre-financial crisis sample (January 1994 – December 2007)

	h=1M	h=6M	h=1Y	h=18M	h=2Y	h=30M	h=3Y
AR(3)-PCE RMSE	0.245	0.553	0.806	0.870	0.941	0.955	0.930
Relative MSE							
BVAR: PCE + Skew (K)	1.009	1.006	0.998	0.989	0.980	0.981	0.995
BVAR: PCE + Median	1.024*	1.053	0.883	0.815	0.787*	0.795*	0.796*
BVAR: PCE + Median + Skew (K)	1.007	1.037	0.888	0.830	0.798*	0.804*	0.802*
BVAR: PCE + Trim	1.019	1.076	0.951	0.910	0.860	0.838*	0.814*
BVAR: PCE + Trim + Skew (K)	0.999	1.052	0.955	0.921	0.866	0.844*	0.818*
BVAR: PCE + Core	1.005	1.030	1.031	1.018	1.004	0.997	1.006
BVAR: PCE + Core + Skew (K)	1.008	1.045	1.046	1.032	1.012	1.000	1.007
BVAR: PCE + UR	1.016	1.220	1.375	1.602*	1.648*	1.769*	1.979*

Financial crisis and onward sample (January 2008 – June 2021)

	h=1M	h=6M	h=1Y	h=18M	h=2Y	h=30M	h=3Y
AR(3)-PCE RMSE	0.285	1.087	1.359	1.097	0.975	0.972	0.833
Relative MSE							
BVAR: PCE + Skew (K)	1.043*	0.943*	0.982	0.953	0.908*	0.924*	0.947*
BVAR: PCE + Median	1.063*	0.976	0.906	0.932	0.793*	0.742*	0.790*
BVAR: PCE + Median + Skew (K)	1.009	0.877	0.901	0.931	0.771*	0.731*	0.781*
BVAR: PCE + Trim	1.065*	0.980	0.883	0.933	0.774*	0.709*	0.808
BVAR: PCE + Trim + Skew (K)	1.021	0.884	0.875	0.929	0.747*	0.697*	0.802*
BVAR: PCE + Core	1.076*	1.004	0.997	0.974*	0.946*	0.942*	0.954*
BVAR: PCE + Core + Skew (K)	1.055	0.955	0.999	0.958*	0.910*	0.927*	0.939*
BVAR: PCE + UR	1.180*	1.179	1.347*	1.603	1.913	1.807	1.894

Notes: The numbers reported in the first row of each panel are the root mean squared error (RMSE) from the univariate AR PCE inflation in gaps (3-lag specification), while the rows below it are ratios that report relative MSEs (relative to the MSE from the AR(3) PCE inflation in gaps). Thus, a ratio of more than 1 indicates that the univariate inflation in gaps model is more accurate on average than the model being compared. The forecast performance is based on an expanding window of estimation spanning the period January 1994 through June 2021 (full sample), and January 1994 through December 2007 (pre-financial crisis sample).
* indicates statistical significance up to the 10% level and is based on the Diebold-Mariano West test

Table 2: Goods PCE inflation out-of-sample point forecasting comparison
[Skew constructed based on month-over-month inflation rates]

Full sample (January 1994 – June 2021)

	h=1M	h=6M	h=1Y	h=18M	h=2Y	h=30M	h=3Y
AR(3)-Goods PCE RMSE	0.644	1.952	2.470	2.334	2.467	2.505	2.407
Relative MSE							
BVAR: G.PCE + Skew (K)	1.025	0.962	0.983	1.022	1.026	1.061	1.136*
BVAR: G.PCE + Median	1.050	0.934	0.888	0.944	0.905	0.925	1.004
BVAR: G.PCE + Median + Skew(K)	1.042	0.911	0.886	0.941	0.902	0.930	1.009

Pre-financial crisis sample (January 1994 – December 2007)

	h=1M	h=6M	h=1Y	h=18M	h=2Y	h=30M	h=3Y
AR(3)-Goods PCE RMSE	0.588	1.278	1.769	1.880	2.051	2.076	1.961
Relative MSE							
BVAR: G.PCE + Skew (K)	1.006	1.017	1.045	1.090	1.103	1.181	1.363
BVAR: G.PCE + Median	1.021	0.923	0.961	1.021	1.023	1.071	1.171
BVAR: G.PCE + Median + Skew(K)	1.017	0.916	0.956	1.015	1.024	1.078	1.185

Financial crisis and onward sample (January 2008 – June 2021)

	h=1M	h=6M	h=1Y	h=18M	h=2Y	h=30M	h=3Y
AR(3)-Goods PCE RMSE	0.697	2.448	2.948	2.478	2.628	2.690	2.485
Relative MSE							
BVAR: G.PCE + Skew (K)	1.039	0.948	0.974	1.002	0.946	0.973	1.059
BVAR: G.PCE + Median	1.072	0.936	0.856	0.899*	0.792*	0.790*	0.904*
BVAR: G.PCE + Median + Skew(K)	1.061	0.909	0.857	0.897*	0.778*	0.790*	0.907*

Notes: The numbers reported in the first row of each panel are the root mean squared error (RMSE) from the univariate AR PCE inflation in gaps (3-lag specification), while the rows below it are ratios that report relative MSEs (relative to the MSE from the AR(3) PCE inflation in gaps). Thus, a ratio of more than 1 indicates that the univariate inflation in gaps model is more accurate on average than the model being compared. The forecast performance is based on an expanding window of estimation spanning the period January 1994 through June 2021 (full sample), and January 1994 through December 2007 (pre-financial crisis sample). * indicates statistical significance up to the 10% level and is based on the Diebold-Mariano West test

Table 3: Services PCE inflation out-of-sample point forecasting comparison
[Skew constructed based on month-over-month inflation rates]

Full sample (January 1994 – June 2021)

	h=1M	h=6M	h=1Y	h=18M	h=2Y	h=30M	h=3Y
AR(3)-Services PCE RMSE	0.174	0.466	0.661	0.683	0.704	0.728	0.752
Relative MSE							
BVAR: S.PCE + Skew (K)	0.968*	0.935	0.968	0.983	0.991	0.983	0.982
BVAR: S.PCE + Median	1.026	1.022	0.979	0.991	0.997	0.997	1.000
BVAR: S.PCE + Median + Skew(K)	0.998	0.954	0.989	1.003	1.005	0.995	0.995

Pre-financial crisis sample (January 1994 – December 2007)

	h=1M	h=6M	h=1Y	h=18M	h=2Y	h=30M	h=3Y
AR(3)-Services PCE RMSE	0.187	0.375	0.555	0.622	0.668	0.692	0.732
Relative MSE							
BVAR: S.PCE + Skew (K)	0.949*	0.982	1.018	1.053	1.084	1.108*	1.107*
BVAR: S.PCE + Median	1.006	1.067	1.016	1.026	1.046	1.072	1.079*
BVAR: S.PCE + Median + Skew(K)	1.005	1.010	1.033	1.056	1.080	1.104	1.103*

Financial crisis and onward sample (January 2008 – June 2021)

	h=1M	h=6M	h=1Y	h=18M	h=2Y	h=30M	h=3Y
AR(3)-Services PCE RMSE	0.159	0.552	0.767	0.672	0.484	0.482	0.466
Relative MSE							
BVAR: S.PCE + Skew (K)	0.994	0.911	0.940*	0.929*	0.898*	0.921	0.972
BVAR: S.PCE + Median	1.054*	1.001	0.958*	0.964*	0.949*	0.945	0.983
BVAR: S.PCE + Median + Skew(K)	0.986	0.927	0.963	0.957*	0.924*	0.946	0.996

Notes: The numbers reported in the first row of each panel are the root mean squared error (RMSE) from the univariate AR PCE inflation in gaps (3-lag specification), while the rows below it are ratios that report relative MSEs (relative to the MSE from the AR(3) PCE inflation in gaps). Thus, a ratio of more than 1 indicates that the univariate inflation in gaps model is more accurate on average than the model being compared. The forecast performance is based on an expanding window of estimation spanning the period January 1994 through June 2021 (full sample), and January 1994 through December 2007 (pre-financial crisis sample). * indicates statistical significance up to the 10% level and is based on the Diebold-Mariano West test

Table 4: Estimates of parameter beta

Model	Posterior Mean	68% Credible Bands
Headline PCE inflation		
Default (past inflation)	0.002	-0.005, 0.009
B. Skew	**-0.199**	**-0.330, -0.068**
Services PCE inflation		
Default (past services inflation)	0.001	-0.006, 0.007
Services B. Skew	-0.121	-0.260, 0.017
Goods PCE inflation		
Default (past inflation)	0.003	-0.005, 0.011
Goods B. Skew	**0.320**	**0.116, 0.525**

Notes: The numbers reported under the column labeled "Posterior Mean" refer to posterior mean estimates of the parameter beta obtained by estimating the TVP-SVM model using quarterly data. "B. Skew" refers to the Bowley skewness measure, "Services B. Skew" refers to the Bowley skewness measure constructed using the components underlying the services PCE category, and "Goods B. Skew" refers to the Bowley skewness measure constructed using the components underlying the goods PCE category. Quarterly values of the skewness measures are computed by averaging the monthly estimates of the skewness. The numbers in bold indicate significant values.

References

Almuzara, Martín, and Argia M. Sbordone. 2022. "Inflation Persistence: How Much Is There and Where Is It Coming From?" *Liberty Street Economics* (blog). Federal Reserve Bank of New York. April 20, 2022. https://ideas.repec.org//p/fip/fednls/94064.html.

Amisano, Gianni, and Raffaella Giacomini. 2007. "Comparing Density Forecasts via Weighted Likelihood Ratio Tests." *Journal of Business & Economic Statistics* 25 (2): 177–90. https://doi.org/10.1198/073500106000000332.

Ashley, Richard, and Randal J. Verbrugge. 2020. "Finding a Stable Phillips Curve Relationship: A Persistence-Dependent Regression Model." Working paper 19-09R. Federal Reserve Bank of Cleveland.

Ball, Laurence. 1992. "Why Does High Inflation Raise Inflation Uncertainty?" *Journal of Monetary Economics* 29 (3): 371–88. https://doi.org/10.1016/0304-3932(92)90032-W.

Ball, Laurence M., and N. Gregory Mankiw. 1994. "Asymmetric Price Adjustment and Economic Fluctuations." *The Economic Journal* 104 (423): 247–61. https://doi.org/10.2307/2234746.

Ball, Laurence M., and Sandeep Mazumder. 2011. "Inflation Dynamics and the Great Recession." *Brookings Papers on Economic Activity* 2011 (1): 337–81. https://doi.org/10.1353/eca.2011.0005.

Ball, Laurence M., and Sandeep Mazumder. 2020. "The Nonpuzzling Behavior of Median Inflation." In *Changing Inflation Dynamics, Evolving Monetary Policy*, edited by Gonzalo Castex, Jordi Gali, and Diego Saravia, 27:49–70. Series on Central Banking Analysis and Economic Policies. Santiago, Chile: Banco Central de Chile. https://hdl.handle.net/20.500.12580/4880.

Bańbura, Marta, Domenico Giannone, and Lucrezia Reichlin. 2010. "Large Bayesian Vector Auto Regressions." *Journal of Applied Econometrics* 25 (1): 71–92. https://doi.org/10.1002/jae.1137.

BIS. 2022. "II. Inflation: A Look under the Hood." In *BIS Annual Economic Report*. https://www.bis.org/publ/arpdf/ar2022e2.htm.

Bowley, A. L. 1920. *Elements of Statistics*. 4th ed. New York: C. Scribner's Sons.

Bryan, Michael F., and Stephen G. Cecchetti. 1993. "Measuring Core Inflation." Working paper 4303. National Bureau of Economic Research. https://doi.org/10.3386/w4303.

Bryan, Michael F., and Stephen G. Cecchetti. 1999. "Inflation and the Distribution of Price Changes." *The Review of Economics and Statistics* 81 (2): 188–96. https://doi.org/10.1162/003465399558148.

Carroll, Daniel R., and Randal J. Verbrugge. 2019. "Behavior of a New Median PCE Measure: A Tale of Tails." *Economic Commentary*, no. 2019–10 (July). https://doi.org/10.26509/frbc-ec-201910.

Chan, Joshua C. C. 2017. "The Stochastic Volatility in Mean Model with Time-Varying Parameters: An Application to Inflation Modeling." *Journal of Business & Economic Statistics* 35 (1): 17–28. https://doi.org/10.1080/07350015.2015.1052459.

Clark, Todd E., and Taeyoung Doh. 2014. "Evaluating Alternative Models of Trend Inflation." *International Journal of Forecasting* 30 (3): 426–48. https://doi.org/10.1016/j.ijforecast.2013.11.005.

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): 505–19. https://doi.org/10.1111/jmcb.12013.

Dolmas, Jim. 2005. "Trimmed Mean PCE Inflation." Working paper 0506. Federal Reserve Bank of Dallas. https://ideas.repec.org/p/fip/feddwp/05-06.html.

Dolmas, Jim. 2009. "The 2009 Revision to the Trimmed Mean PCE Inflation Series." Technical note. Federal Reserve Bank of Dallas.

Faust, Jon, and Jonathan H. Wright. 2013. "Chapter 1 - Forecasting Inflation." In *Handbook of Economic Forecasting*, edited by G. Elliott and A. Timmermann, 2, Part A:2–56. Elsevier. https://doi.org/10.1016/B978-0-444-53683-9.00001-3.

Friedman, Milton. 1977. "Nobel Lecture: Inflation and Unemployment." *Journal of Political Economy* 85 (3): 451–72. https://doi.org/10.1086/260579.

Holland, A. Steven. 1995. "Inflation and Uncertainty: Tests for Temporal Ordering." *Journal of Money, Credit and Banking* 27 (3): 827–37. https://doi.org/10.2307/2077753.

Kim, Tae-Hwan, and Halbert White. 2004. "On More Robust Estimation of Skewness and Kurtosis." *Finance Research Letters* 1 (1): 56–73. https://doi.org/10.1016/S1544-6123(03)00003-5.

Knotek II, Edward S., and Saeed Zaman. 2019. "Financial Nowcasts and Their Usefulness in Macroeconomic Forecasting." *International Journal of Forecasting* 35 (4): 1708–24. https://doi.org/10.1016/j.ijforecast.2018.10.012.

Mertens, Elmar. 2016. "Measuring the Level and Uncertainty of Trend Inflation." *The Review of Economics and Statistics* 98 (5): 950–67. https://doi.org/10.1162/REST_a_00549.

Meyer, Brent, and Saeed Zaman. 2019. "The Usefulness of the Median CPI in Bayesian VARs Used for Macroeconomic Forecasting and Policy." *Empirical Economics* 57 (2): 603–30. https://doi.org/10.1007/s00181-018-1472-1.

Norman, David, and Anthony Richards. 2012. "The Forecasting Performance of Single Equation Models of Inflation." *Economic Record* 88 (280): 64–78. https://doi.org/10.1111/j.1475-4932.2011.00781.x.

Ocampo, Sergio, Raphael S. Schoenle, and Dominic A. Smith. 2022. "How Robust Are Robust Measures of PCE Inflation?" Working paper 552. U.S. Bureau of Labor Statistics. https://www.bls.gov/osmr/research-papers/2022/pdf/ec220070.pdf.

Schoenle, Raphael S., and Dominic A. Smith. 2022. "What Can We Learn from 60 Years of PCE Inflation Data?" In *Inflation: Drivers and Dynamics 2022*. Federal Reserve Bank of Cleveland Center for Inflation Research and the European Central Bank. https://www.clevelandfed.org/-/media/project/clevelandfedtenant/clevelandfedsite/events/2022/ev-20220929-inflation-drivers-and-dynamics/schoenle-paper.pdf.

Smith, Julie K. 2004. “Weighted Median Inflation: Is This Core Inflation?” *Journal of Money, Credit, and Banking* 36 (2): 253–63. https://doi.org/10.1353/mcb.2004.0014.

Tallman, Ellis W., and Saeed Zaman. 2017. “Forecasting Inflation: Phillips Curve Effects on Services Price Measures.” *International Journal of Forecasting* 33 (2): 442–57. https://doi.org/10.1016/j.ijforecast.2016.10.004.

Verbrugge, Randal J. 1999. “Cross-Sectional Inflation Asymmetries and Core Inflation: A Comment on Bryan and Cecchetti.” *The Review of Economics and Statistics* 81 (2): 199–202. https://doi.org/10.1162/003465399558166.

Verbrugge, Randal J. 2021. “Is It Time to Reassess the Focal Role of Core PCE Inflation?” Working paper 21–10. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-202110.

Verbrugge, Randal J., and Saeed Zaman. 2023. “The Hard Road to a Soft Landing: Evidence from a (Modestly) Nonlinear Structural Model.” Working paper 23–03. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-202303.

Zaman, Saeed. 2013. “Improving Inflation Forecasts in the Medium to Long Term.” *Economic Commentary*, no. 2013–16 (November). https://doi.org/10.26509/frbc-ec-201316.