CORE INFLATION AND TREND INFLATION

James H. Stock and Mark W. Watson*

Abstract—This paper examines empirically whether the measurement of trend inflation can be improved by using disaggregated data on sectoral inflation to construct indexes akin to core inflation but with a time-varying distributed lags of weights, where the sectoral weight depends on the time-varying volatility and persistence of the sectoral inflation series and on the comovement among sectors. The modeling framework is a dynamic factor model with time-varying coefficients and stochastic volatility as in Del Negro and Otrok (2008), and is estimated using U.S. data on seventeen components of the personal consumption expenditure inflation index.

I. Introduction

A classic yet still important problem of measuring the rate of price inflation is filtering out the noise in inflation data to provide an estimate of the trend value of inflation. Following Bryan and Cecchetti (1994), we think of trend inflation as the long-term estimate of the inflation rate based on data on prices through the present. Having a good estimate of trend inflation is an important input to monetary policy and to a myriad of private decisions. For example, as this is written, a pressing question in the United States and the Eurozone is how far trend inflation is below the 2% target. Because there are multiple sources of noise in inflation data and because the nature of the noise can change over time, the task of estimating trend inflation is both difficult and of ongoing relevance.

Producing an accurate estimate of trend inflation requires distinguishing persistent variations in inflation from those that are unlikely to persist into the future. Broadly speaking, there are two distinct approaches to this signal extraction problem.

The first approach is to use cross-sectional data on inflation (sectoral-level inflation data), with a scheme that downweights sectors with large nonpersistent variation. The most important example of this approach is the standard measure of core inflation, which excludes food and energy prices (Gordon, 1975; Eckstein, 1981; see Clark, 2001, for a general introduction to core inflation and Wynne, 2008, for a discussion of its history). Other methods that exploit cross-sectional smoothing include trimmed means or medians of sectoral inflation rates (see Bryan & Cecchetti, 1994); these methods impose 0/1 weighting on each component, with weights that vary over time.[1]

(For recent references on core inflation, see Crone et al., 2013.)

The second common approach to the signal extraction problem uses time series smoothing methods. Simple yet effective smoothers include the exponential smoother implied by the IMA(1,1) model of Nelson and Schwert (1977) and the four-quarter average of quarterly inflation (Atkeson & Ohanian, 2001). Stock and Watson (2007) and Cogley and Sargent (2015) provide methods that allow for time variation in the smoother depending on changes in the signal-to-noise ratio of the persistent and nonpersistent components.

We follow this literature on core and trend inflation and consider estimates derived from the price indexes and corresponding expenditure share weights used in the construction of the headline inflation series of interest. A vast literature considers the problem of using other series, such as measures of economic activity, interest rates, and terms of trade to forecast inflation. At an abstract level, the distinction between using only price data, and price data combined with other data, can be thought of as measurement versus forecasting; the focus here is measurement. At a practical level, at least for the United States, some forecasting models using nonprice data can improve on forecasts based solely on prices, but those improvements are small and, in many cases, ephemeral. This underscores the practical relevance of estimates of trend inflation based on constituent sectoral price data.

This paper combines the cross sectional and time series smoothing approaches to examine four questions about the measurement of trend inflation and its relation to core inflation. First, can more precise measures of trend inflation be obtained using disaggregated sectoral inflation measures, relative to time series smoothing of aggregate ("headline") inflation? Second, if there are improvements to be had by using sectoral inflation measures, do the implied sectoral weights evolve over time, or are they stable, and how do they compare to the corresponding sectoral shares in consumption? Third, how do the implied time-varying weights and the resulting multivariate estimate of trend inflation compare to conventional core inflation measures? And fourth, do these trend inflation measures improve on conventional core inflation when it comes to forecasting inflation over the one- to three-year horizon?

We investigate these questions empirically using a univariate and multivariate unobserved-components stochastic volatility outlier-adjusted (UCSVO) model that allows for common persistent and transitory factors, time-varying factor loadings, and stochastic volatility in the common and sectoral components. The time-varying factor loadings allow for changes in the comovements across sectors, such

[1] The Cleveland Fed publishes a median and trimmed mean CPI (https://www.clevelandfed.org/en/Our%20Research/Indicators%20and %20Data/Current%20Median%20CPI.aspx), and the Dallas Fed publishes a monthly trimmed mean PCE inflation index (http://www.dallasfed.org /research/pce/).

as the reduction in energy price pass-through into other prices. Introducing separate sectoral and common stochastic volatility in transitory and permanent innovations allows for changes in the persistence of sectoral inflation and for sector-specific changes in volatility. One source of the changing volatility in the component inflation rates is changes in the methods or underlying data sources used to construct the historical series. A strength of the method used here is that the resulting estimates of historical trends adjust for changes in measurement methods, as well as for fundamental changes in the volatility and persistence of the component series.

At a technical level, the model closest to that used here is Del Negro and Otrok (2008), which has time-varying factor loadings and stochastic volatility (their application is to international business cycles, not inflation). Our model has some differences to fit our application to U.S. sectoral inflation data, including distinct sectoral trends, a common trend, and model-based detection of and adjustment for outliers.

The data we use are seventeen sectors comprising the personal consumption expenditure (PCE) price index for the United States, 1959Q1–2015Q2. Our main findings are: (a) the multivariate trend estimates are more precise than the univariate estimates: posterior intervals for trend inflation using the multivariate model are roughly one-third narrower than intervals based on headline inflation alone; (b) although the implied weights in the multivariate trend on most sectoral components are close to their share weights, the implied weight on some series varies substantially; (c) broadly speaking, the multivariate trend estimate is a temporally smoothed version of core (excluding food and energy) through the 1970s, but starting in the 1980s places more weight on food (both off-premises and food services and accommodation) and less weight on financial services, so that the composition of multivariate trend in the 2000s is roughly similar to inflation for PCE excluding energy; and (d) viewed as forecasts, the multivariate and univariate trend estimates constructed using core inflation improve on forecasts that use headline inflation alone and several other benchmark forecasts, but the forecasting gains are imprecisely estimated.

In addition to the literatures already discussed on core and trend inflation, this work is related to three other large literatures. First, our modeling framework extends work estimating common factors of multiple inflation series, including Bryan and Cecchetti (1993), Cristadoro et al. (2005), Amstad and Potter (2007), Kiley (2008), Altissimo, Mojon, and Zaffaroni (2009), Boivin, Giannoni, and Mihov (2009), Reis and Watson (2010), and Sbrana, Silvestrini, and Venditti (2015). Mumtaz and Surico (2012) introduce stochastic volatility and time-varying factor dynamics into a model of thirteen international inflation rates. Second, the issue of including or excluding energy inflation is related to the literature on changes in the pass-through of energy prices to headline or core inflation, something allowed for in our model (see Hooker, 2002; De Gregorio, Lander-

retche, & Neilson, 2007; van den Noord & André, 2007; Chen, 2009; Blanchard & Galí, 2010; Clark & Terry, 2010; Baumeister & Peersman, 2013). Also related is work that uses variables other than prices to measure trend inflation (e.g., Mertens, 2015; Garnier, Mertens, & Nelson, 2015; Mertens & Nason, 2015).

The next section presents the univariate and multivariate UCSVO models and discusses their estimation. Section III provides the resulting univariate trend estimates for headline, core, and PCE excluding energy. Section IV presents multivariate results, first for the seventeen-sector model, then for a model with only three components: core, food, and energy. Section V compares the forecasting performance of the various trend estimates over the one- to three-year horizon, and section VI concludes.

II. The Unobserved Components Model with Stochastic Volatility, Common Factors, and Outlier Adjustment

A. The Univariate UCSVO Model

The univariate unobserved components/stochastic volatility outlier-adjustment (UCSVO) model used in this paper expresses the rate of inflation as the sum of a permanent and transitory component, where the innovations to both components have variances that evolve over time according to independent stochastic volatility processes and where the innovation to the temporary component can have heavy tails (outliers):

$$\pi_t = \tau_t + \varepsilon_t, \tag{1}$$

$$\tau_t = \tau_{t-1} + \sigma_{\Delta\tau,t} \times \eta_{\tau,t}, \tag{2}$$

$$\varepsilon_t = \sigma_{\varepsilon,t} \times s_t \times \eta_{\varepsilon,t}, \tag{3}$$

$$\Delta\ln\left(\sigma_{\varepsilon,t}^2\right) = \gamma_\varepsilon v_{\varepsilon,t}, \tag{4}$$

$$\Delta\ln\left(\sigma_{\Delta\tau,t}^2\right) = \gamma_{\Delta\tau} v_{\Delta\tau,t}, \tag{5}$$

where $(\eta_\varepsilon, \eta_\tau, v_\varepsilon, v_{\Delta\tau})$ are iidN$(0, I_4)$, and s_t is an i.i.d. random variable that generates outliers in ε_t.

This model expresses the rate of inflation π_t as the sum of a permanent component τ_t (trend) and a transitory component ε_t, equation (1), in which τ_t follows a martingale, equation (2), and the transitory component is serially uncorrelated, equation (3), and in which both innovations follow a logarithmic random walk stochastic volatility processes, equations (4) and (5). Conditional on the stochastic volatility process, the transitory innovation ε_t is modeled in equation (3) as a mixture of normals via the i.i.d. variable s_t, where $s_t = 1$ with probability $(1 - p)$, and $s_t \sim \text{U}[2,10]$ with probability p. This mixture model allows for outliers in inflation—that is, large one-time shifts in the price level—which occur each period with probability p.

The UCSVO model, equations (1) to (5), has only three parameters: γ_ε and $\gamma_{\Delta\tau}$ govern the scale of the innovation in the stochastic volatility process, and p governs the frequency of outliers. At a given point in time, the autocovariance structure of π_t is that of a IMA(1,1) process; however, the outlier distribution of the transitory innovation means that the estimate of τ_t is not always well approximated by the linear exponential smoother associated with a local IMA(1,1) filter.

This difference between equations (1) to (5) and the Stock-Watson (2007) UCSV model is that the USCVO model includes an explicit model-based treatment of outliers. As will be seen below, large infrequent spikes in inflation are observed in the data, especially in the sectoral components.[2] Stock and Watson (2007) made preliminary judgmental adjustments for outliers prior to model estimation; however, that approach is not feasible for real-time trend estimation because it requires knowing whether a large change will mean-revert. Ignoring outliers is not appealing because doing so runs the risk of mistaking a single large outlier for a more systematic increase in the volatility of the transitory component. Because we are interested in real-time trend estimation, equation (3) therefore extends the Stock-Watson (2007) model to make outlier adjustments part of the model by modeling the transitory innovation as a mixture of normals.

B. The Multivariate UCSVO Model

This multivariate UCSVO (MUCSVO) model extends the UCSVO model to include a common latent factor in both the trend and idiosyncratic components of inflation, where the factor loadings are also time varying. Let the subscripts c denote the common latent factor and i denote the sector. The multivariate model is the Del Negro and Otrok (2008) dynamic factor model with time-varying factor loadings and stochastic volatility, extended to have permanent and transitory components and to handle outliers in the transitory disturbance.

The multivariate UCSV model is

$$\pi_{i,t} = \alpha_{i,\tau,t}\tau_{c,t} + \alpha_{i,\varepsilon,t}\varepsilon_{c,t} + \tau_{i,t} + \varepsilon_{i,t}, \tag{6}$$

$$\tau_{c,t} = \tau_{c,t-1} + \sigma_{\Delta\tau,c,t} \times \eta_{\tau,c,t}, \tag{7}$$

$$\varepsilon_{c,t} = \sigma_{\varepsilon,c,t} \times s_{c,t} \times \eta_{\varepsilon,c,t}, \tag{8}$$

$$\tau_{i,t} = \tau_{i,t-1} + \sigma_{\Delta\tau,i,t} \times \eta_{\tau,i,t}, \tag{9}$$

$$\varepsilon_{i,t} = \sigma_{\varepsilon,i,t} \times s_{i,t} \times \eta_{\varepsilon,i,t}, \tag{10}$$

---

[2] An example of such a sectoral outlier is the April 2009 increase in the federal cigarette tax, which resulted in a 22% increase in cigarette prices that month. This tax increase drove a one-time jump in the rate of PCE inflation for other nondurable goods (the category that contains tobacco) in 2009Q2 of 10.4% at an annual rate, well above the 2.7% average rate of inflation for that category in 2008 and 2009 excluding that quarter.

$$\alpha_{i,\tau,t} = \alpha_{i,\tau,t-1} + \lambda_{i,\tau}\zeta_{i,\tau,t} \text{ and}$$

$$\alpha_{i,\varepsilon,t} = \alpha_{i,\varepsilon,t-1} + \lambda_{i,\varepsilon}\zeta_{i,\varepsilon,t}, \tag{11}$$

$$\Delta\ln\left(\sigma^2_{\Delta\tau,c,t}\right) = \gamma_{\Delta\tau,c}\nu_{\Delta\tau,c,t},$$

$$\Delta\ln\left(\sigma^2_{\varepsilon,c,t}\right) = \gamma_{\varepsilon,c}\nu_{\varepsilon,c,t},$$

$$\Delta\ln\left(\sigma^2_{\Delta\tau,i,t}\right) = \gamma_{\Delta\tau,i}\nu_{\Delta\tau,i,t}, \text{ and}$$

$$\Delta\ln\left(\sigma^2_{\varepsilon,i,t}\right) = \gamma_{\varepsilon,i}\nu_{\varepsilon,i,t}, \tag{12}$$

where the disturbances ($\eta_{\tau,c,t}$, $\eta_{\varepsilon,c,t}$, $\eta_{\tau,i,t}$, $\eta_{\varepsilon,i,t}$, $\zeta_{i,\tau,t}$, $\zeta_{i,\varepsilon,t}$, $\nu_{\Delta\tau,c,t}$, $\nu_{\varepsilon,c,t}$, $\nu_{\Delta\tau,i,t}$, $\nu_{\varepsilon,i,t}$) are i.i.d. standard normal.

Equation (6) represents sector i inflation as the sum of a latent common factor for trend inflation, $\tau_{c,t}$, a latent common transient component, $\varepsilon_{c,t}$, and sector-specific trends and transient components, $\tau_{i,t}$ and $\varepsilon_{i,t}$, and where the factor loadings evolve according to a random walk, equation (11). Equations (7) to (10) allow for stochastic volatility in the latent common and sector-specific components, where the stochastic volatility evolves according to the logarithmic random walk, equation (12). Like the univariate model, the multivariate model allows for outliers in the common and sectoral transitory components through the independent random variables $s_{c,t}$ and $s_{i,t}$ in equations (8) and (10), and where the outlier probabilities are p_c and p_i. The trend sectoral inflation is the sum of the contribution of the common latent factor to that sector and the sectoral trend, that is, the sectoral trend is $\alpha_{i,\tau,t}\tau_{c,t} + \tau_{i,t}$. The aggregate trend inflation is the sum of the sectoral trend, weighted by the share weight w_{it} of sector i in total inflation:

$$\text{Aggregate trend} = \tau_t = \sum_{i=1}^{n} w_{it}\left(\alpha_{i,\tau,t}\tau_{c,t} + \tau_{i,t}\right), \tag{13}$$

where n denotes the number of sectors.

The definition, equation (13), of the aggregate trend τ_t nests a range of possibilities, from the common trend providing all the trend movements in sectoral inflation (so that there are $n-1$ cointegrating vectors among the n sectors) to all sectoral inflation being independent with no common trend. In this latter case, the common trend in aggregate inflation would just be the sum of the idiosyncratic trends, weighted by the sectoral share weights.

C. Estimation

The model is estimated using Bayesian methods. The online appendix contains a detailed description of the priors and the numerical methods used to approximate the posteriors. We highlight a few details here.

In the univariate model, priors are needed for the stochastic volatility parameters γ_ε and $\gamma_{\Delta\tau}$, the outlier probability p, and the initial values τ_0, $\ln(\sigma_{\varepsilon,0})$, and $\ln(\sigma_{\Delta\tau,0})$. We use independent uniform priors for γ_ε and $\gamma_{\Delta\tau}$ that are calibrated so that the standard deviations of annual changes in the values of $\ln(\sigma_{\varepsilon,t})$ and $\ln(\sigma_{\Delta\tau,t})$ are distributed U[0,0.2]. The

TABLE 1.—THE SEVENTEEN COMPONENTS OF THE PCE PRICE INDEX USED IN THIS STUDY AND THEIR EXPENDITURE SHARES

Sector	1960–2015	1960–1979	1980–1999	2000–2015
Durable goods				
Motor vehicles and parts	0.053	0.060	0.054	0.042
Furnishings and durable household equipment	0.036	0.044	0.033	0.028
Recreational goods and vehicles	0.029	0.026	0.029	0.032
Other durable goods	0.016	0.015	0.016	0.016
Nondurable goods				
Food and beverages purchased for off-premises consumption[a]	0.117	0.160	0.104	0.077
Clothing and footwear	0.054	0.071	0.051	0.034
Gasoline and other energy goods[a]	0.037	0.044	0.035	0.032
Other nondurable goods	0.078	0.080	0.074	0.081
Services				
Housing and utilities				
Housing excluding gas and electric utilities	0.153	0.146	0.155	0.161
Gas and electric utilities[a]	0.025	0.026	0.028	0.021
Health care	0.114	0.071	0.127	0.155
Transportation services	0.032	0.030	0.034	0.032
Recreation services	0.029	0.021	0.031	0.038
Food services and accommodations	0.064	0.064	0.066	0.061
Financial services and insurance	0.063	0.047	0.068	0.076
Other services	0.081	0.081	0.077	0.087
Final consumption expenditures of nonprofit institutions serving households (NPISHs)	0.020	0.016	0.019	0.026

Each column shows the average expenditure share over the sample period indicated.
[a]Excluded from core PCE.

prior for p is Beta(α,β), where α and β are calibrated to reflect information in a sample of length ten years, with an outlier occurring once every four years. The priors for τ_0, $\ln(\sigma_{\varepsilon,0})$, and $\ln(\sigma_{\Delta\tau,0})$ are independent diffuse normal.

The priors for the multivariate model follow the priors used in the univariate model. Thus, the priors for the various (γ, p) parameters and $\tau_{i,0}$, $\ln(\sigma_{i,\varepsilon,0})$, and $\ln(\sigma_{i,\Delta\tau,0})$ are the ones described in the previous paragraph. The initial values of $\tau_{c,0}$ and $\tau_{i,0}$ are not separately identified, so we set $\tau_{c,0} = 0$. The factor structure of the multivariate model requires a normalization to separately identify the scale of the factor loadings (α_τ, α_ε) and factors (τ_c, ε_c), and this leads us to set $\ln(\sigma_{\Delta\tau,c,0}) = \ln(\sigma_{\varepsilon,c,0}) = 0$. We use an informative prior about the initial values of the factor loadings: letting α_τ be the vector of factor loadings on $\tau_{c,t}$, the prior is $\alpha_\tau \sim N(0, \kappa_1^2 \iota\iota' + \kappa_2^2 I_n)$, where n is the number of sectors and ι is an $n \times 1$ vector of 1's. The parameter κ_1 governs the prior uncertainty about the average value of factor loadings, and the parameter κ_2 governs the variability of each factor loading from the average value. We set $\kappa_1 = 10$ (so the prior is relatively uninformative about the average value of the factor loadings) and $\kappa_2 = 0.4$ (so there is shrinkage toward the average values). The same prior is used for α_ε. The final set of parameters, ($\lambda_{i,\tau}$, $\lambda_{i,\varepsilon}$), governs the time variation in the factor loadings. Following Del Negro and Otrok (2008), we adopt an inverse gamma prior for λ, with scale and shape parameters chosen so that the prior corresponds to T_{Prior} prior observations with $s_{\text{Prior}}^2 = 0.25^2/T_{\text{Prior}}$, where $T_{\text{Prior}} = T/10$ and T is the sample size. Estimation of the posterior proceeds using Markov chain Monte Carlo (MCMC) methods. The stochastic volatility is handled following Kim, Shephard, and Chib (1998), modified to use the Omori et al. (2007) ten-component gaussian mixture approximation for the log chi-squared error. The

MCMC iterations in Stock and Watson (2007) have been corrected for an error pointed out by Del Negro and Primiceri (2015) that applies generally to models with stochastic volatility. Details are presented in the online appendix.

III. Data and Univariate Results

A. The Data

The full data set consists of observations on seventeen components of inflation used to construct the PCE price index. The lowest-level components in NIPA table 2.3.4 consist of sixteen components (four durable goods sectors, four nondurable good sectors, and eight service sectors). Core PCE excludes two of these sixteen components (food for off-premises consumption and gasoline and energy goods), and additionally excludes gas and electric utilities. Because gas and electric utilities does not appear separately in table 2.3.4 but rather is contained in housing and utilities, core PCE cannot be constructed directly from these sixteen components. So that our seventeen-sector treatment nests core, we use addenda data from NIPA tables 2.3.4 and 2.3.5 to further disaggregate housing and utilities into gas and electric utilities and housing excluding gas and electric utilities, for a total of seventeen sectoral components. Expenditure share weights for these components can be computed using the nominal PCE values in NIPA table 2.3.5. The raw data in the sample are monthly observations from 1959M1 to 2015M6. Most of our analysis uses quarterly data constructed by averaging the monthly inflation rates over the three months in the quarter. Throughout, inflation is measured in percentage points at an annual rate. The seventeen components and their expenditure share weights for selected periods are given in table 1.

FIGURE 1.—HEADLINE AND CORE INFLATION
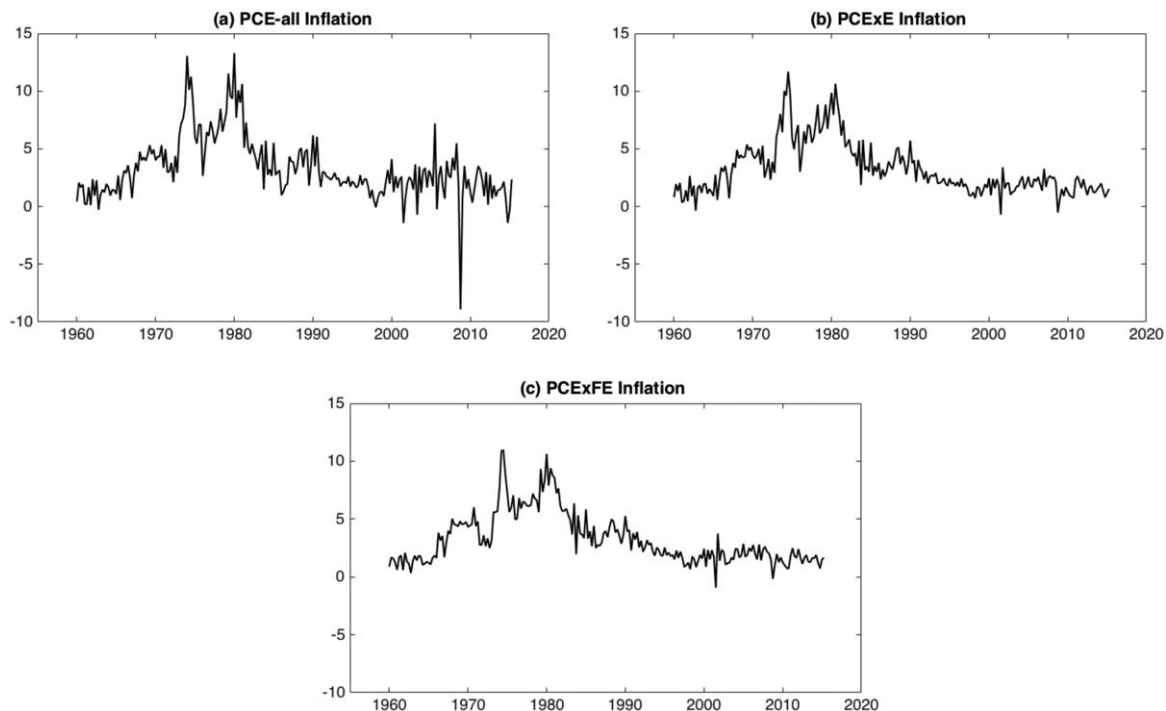
In addition, we consider three aggregate indexes: the headline (all-components) PCE price index (PCE-all), the Bureau of Economic Analysis's PCE price index excluding energy (PCExE), and the BEA core PCE price index excluding food and energy (PCExFE).

The data are all final estimates of these series. Some of the component series have undergone significant methodological changes over the years and have been subject to major historical revisions. For example, in 2013, the price index for financial services was revised, including changing the method for measuring implicitly priced services produced by commercial banks (Hood, 2013). Prior to the revision, the category "financial services furnished without payment" (e.g., checks processed without fees) used imputed prices based on market interest rates, so those prices fluctuated substantially during periods of interest rate volatility. The 2013 revision changed the method for computing the reference interest rate for unpriced financial services, reducing the volatility of this component. Because this revision was implemented retroactively only to 1985, different methods are used to compute this component of the financial services price index pre-1985 and post-1985.

As another example, in the 2009 revision, the category of food and tobacco (which until then had been excluded from core) was distributed across three categories: food and beverages purchased for off-premise consumption, other non-durable goods (which since 2009 includes tobacco), and food services and accommodations; only the first of these is now excluded from core PCE. Because the fully revised series reflects this change, it does not cause a break in the data used in this paper; however, it does mean that previous research on core PCE examined a somewhat different concept from the current definition of core. Changing definitions and measurement methods combined with partial historical adjustment are commonplace, and we return to the implications of these methodological changes below.

B. Univariate Results for PCE-all, PCExE, and PCExFE

Figure 1 plots headline (PCE-all) and the two core inflation series (PCExE and PCExFE). Figure 2 plots the full-sample posterior means for τ_t, $\sigma_{\Delta\tau,t}$, $\sigma_{\varepsilon,t}$, and s_t from the univariate model for each of these inflation measures. The parameter values plotted in figure 2 capture different features of the inflation series plotted in figure 1. Figure 2a plots the posterior means for τ_t. The broadly similar trend estimates reflect the common low-frequency variability in the inflation series (see figure 1); however, there are important differences between the univariate trend estimates, most notably persistently higher trend inflation for PCE-all than for core inflation in the 2000s and large but less persistent deviations of the headline and core trends during the late 1970s and mid-1980s. Over the entire sample period, the mean absolute difference between the estimated trends in PCE-all and PCExFE is 40 basis points; it is 20 basis points for the difference between PCExFE and PCExFE trends. In part, these differences reflect sampling errors associated with estimates, and we present error bands below.

Figure 2b shows estimates of $\sigma_{\Delta\tau,t}$. These too are similar for the three inflation series and reflect the larger trend varia-

FIGURE 2.—FULL-SAMPLE POSTERIOR MEANS FROM THE UNIVARIATE UCSVO MODELS
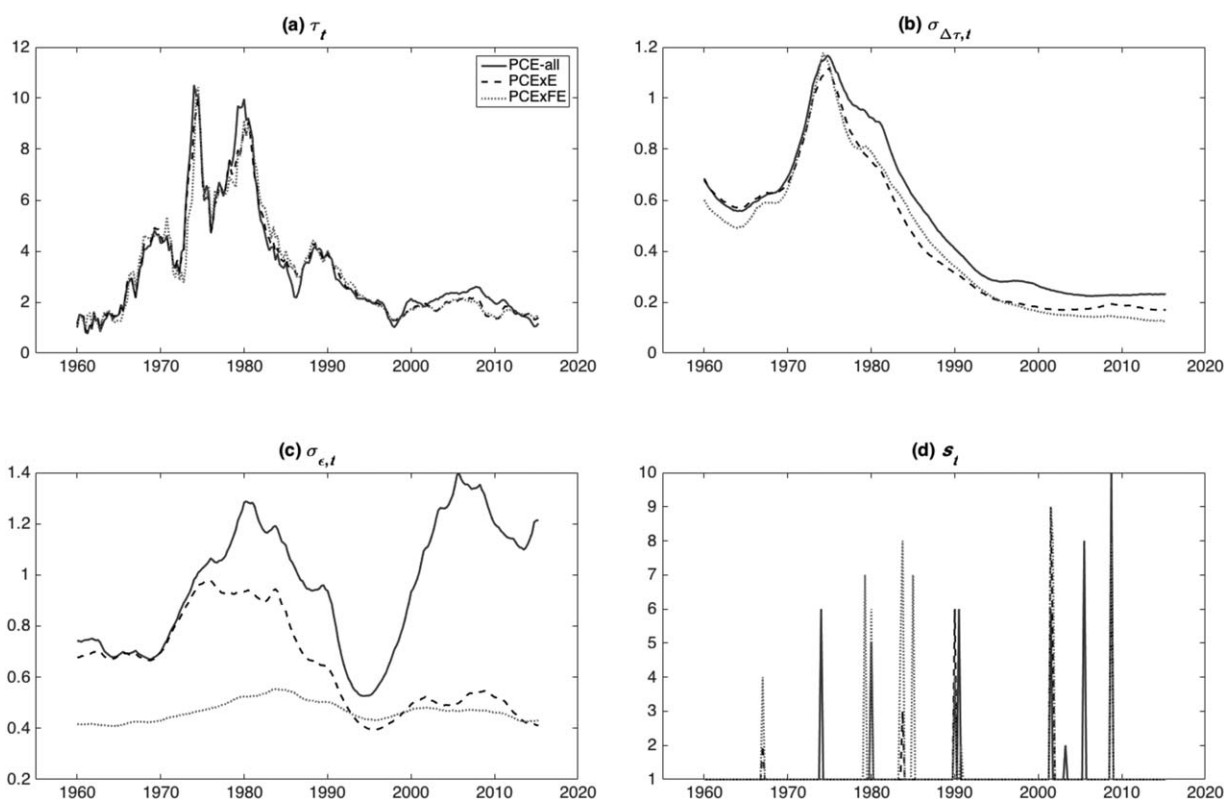
tion in the first half of the sample (when trend inflation increased during the 1970s and fell during the 1980s) than in the second half (when trend inflation was relatively anchored).

Figure 2c shows estimates of $\sigma_{\varepsilon,t}$. These show important variation both over time and between inflation measure. Examination of PCE-all inflation in figure 1 shows relatively little high-frequency volatility during the 1990s followed by a marked increase in volatility in the early 2000s; this is reflected in the estimates for $\sigma_{\varepsilon,t}$ in figure 2. A more subtle feature in figure 1 is the difference between high-frequency variability in the two core inflation measures: their high-frequency volatility is similar in the second half the sample, but PCExFE exhibits much less high frequency than PCExE in the first half of the sample. This too is reflected in the estimates of $\sigma_{\varepsilon,t}$ for the two inflation series.

Finally, figure 2d shows estimates of the outlier scale factors s_t. These factors capture the outliers evident in all the inflation series plotted in figure 1. (Note that s_t measures outliers in standard deviation units, so the absolute size of outliers is larger for headline inflation than the core measures of inflation.)

### IV. Multivariate Results

A. Seventeen-Sector Model

The multivariate model estimates many variables: the common volatilities and trends ($\sigma_{\Delta\tau,c,t}$, $\sigma_{\varepsilon,c,t}$, $\tau_{c,t}$), their

sector-specific counterparts ($\sigma_{\Delta\tau,i,t}$, $\sigma_{\varepsilon,i,t}$, $\tau_{i,t}$), the sector-specific factor loadings ($\alpha_{\tau,i,t}$, $\alpha_{\varepsilon,i,t}$), the common and sector-specific outlier factors ($s_{c,t}$, $s_{i,t}$), and the aggregate inflation trend given in equation (13). The online appendix presents the model's estimates for all of these variables, and we highlight a few of them here.

Figure 3 plots the MUCSVO model's full sample estimates for the aggregate inflation trend, and for comparison it also plots the PCE-all UCSVO estimate. Broadly speaking, the multivariate trend looks more like a time-averaged version of the two core measures (see figure 2) than the univariate trend in PCE-all. The divergence between the univariate PCE-all trend and the multivariate trend is largest in the 1970s, the mid-1980s, and the late 2000s. (Error bands for the estimates are discussed below.) Figure 3 also plots estimates of the volatility for the common factors and common outliers. The time series of volatility for the common trend factor, $\sigma_{\Delta\tau,c,t}$, looks much like the trend volatility estimates from the UCSVO models, and $\sigma_{\varepsilon,c,t}$ evolves much like the corresponding estimates in the UCSVO models for core inflation.

Figure 4 shows estimates for the sector-specific variables for one sector: financial services and insurance. (The online appendix contains the analogous figures for the other sixteen sectors.) As discussed in section III, the price index for the financial services and insurance sector is measured differently before 1985 than after, and this measurement break is evident in the sectoral inflation data plotted in figure 4.

FIGURE 3.—SELECTED RESULTS FROM THE SEVENTEEN-COMPONENT MUCSVO MODEL
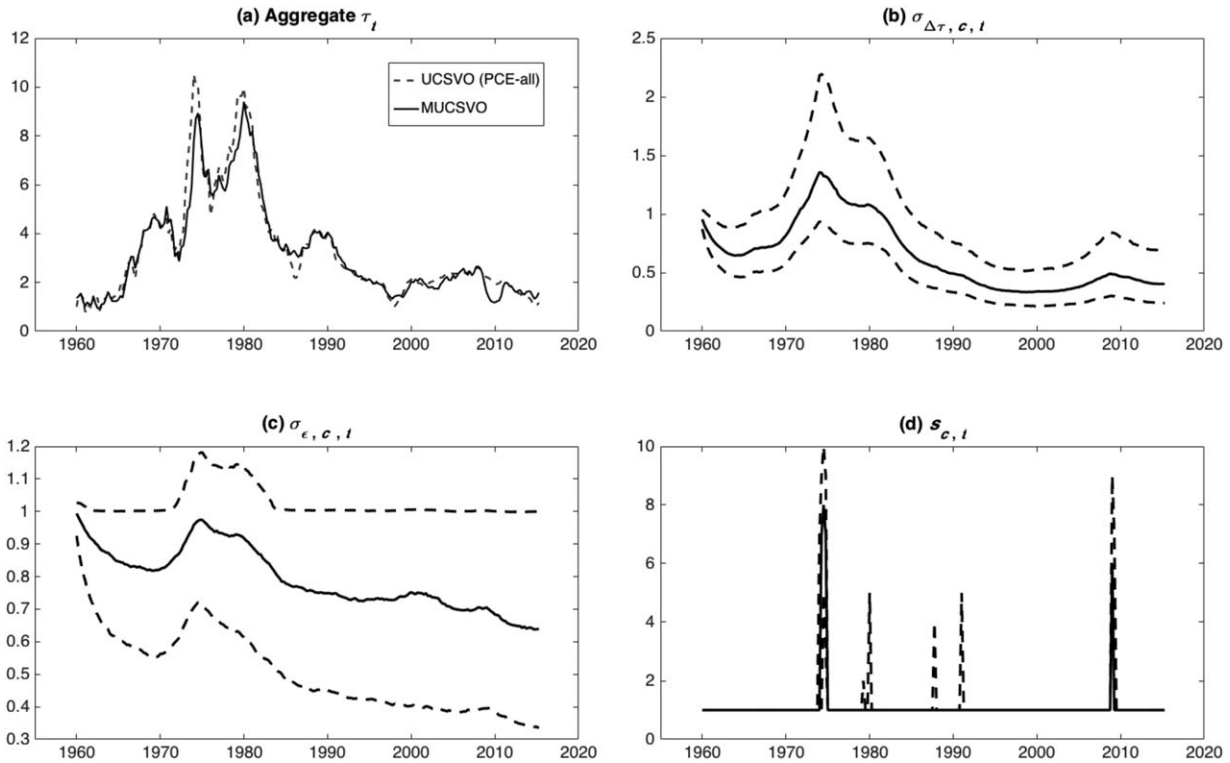
Panel a shows the full sample posterior mean of the aggregate inflation trend computed from the PCE-all UCSVO and MUCSVO models Panels b–d show full sample posterior medians and (point-wise) 67% intervals for $\sigma_{\Delta\tau,c,t}$, $\sigma_{\varepsilon,c,t}$, and $s_{c,t}$.

FIGURE 4.—SELECTED RESULTS FROM THE SEVENTEEN-COMPONENT MUCSVO MODEL: FINANCIAL SERVICES AND INSTITUTIONS
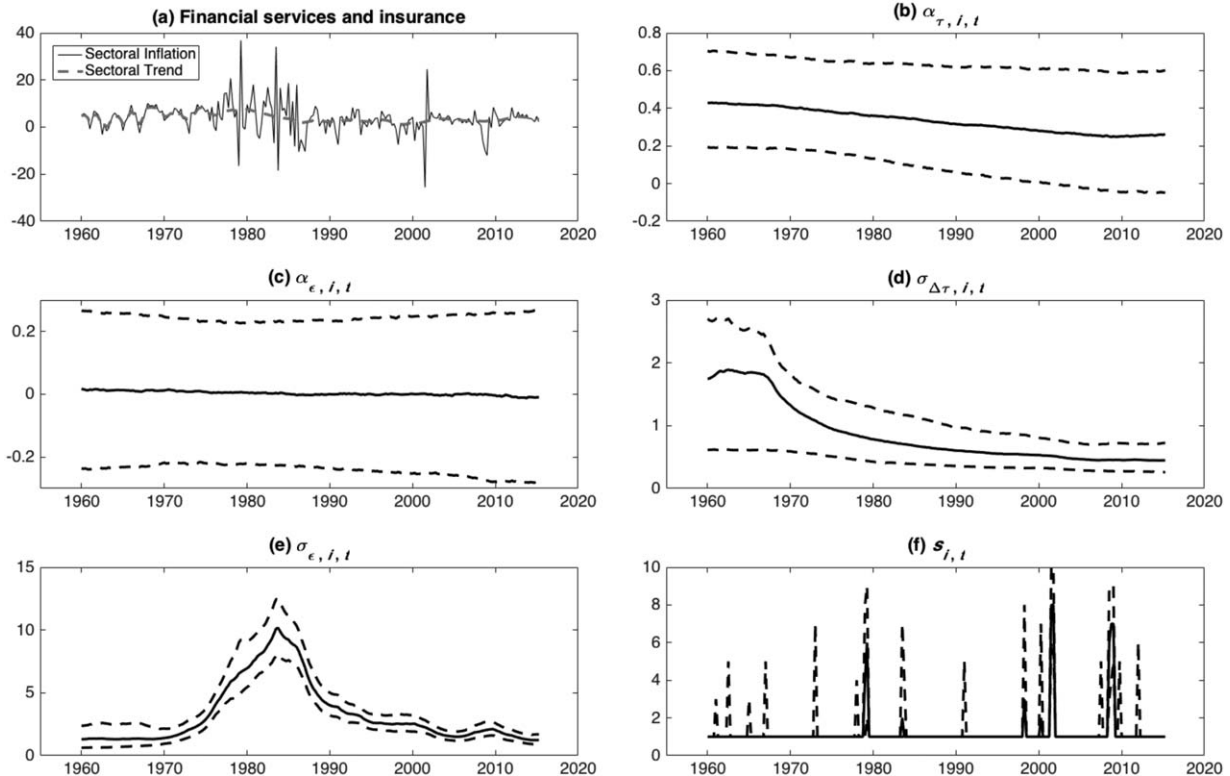
Panel a inflation is the financial services and insurance sector and the full-sample posterior mean of the sectoral trend. The other panels plot the full-sample posterior median and (point-wise) 67% intervals for the sector-specific parameters.
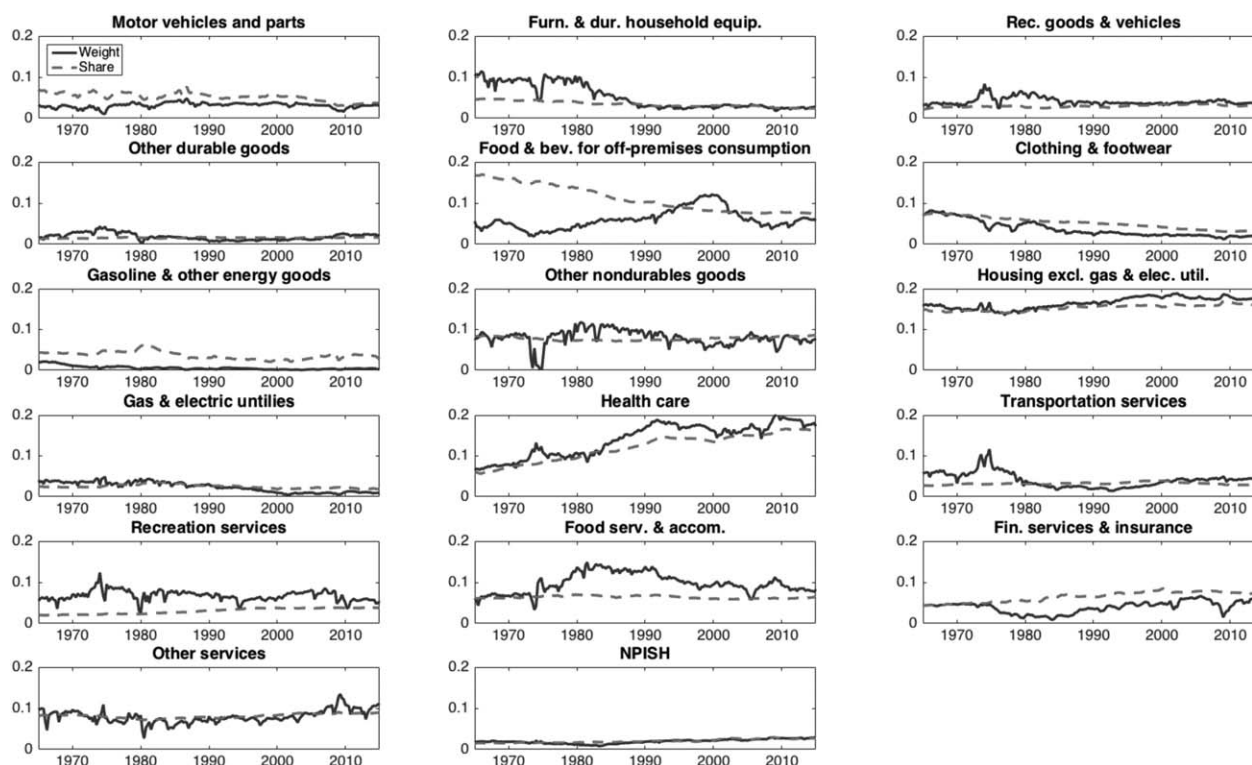
The solid line is the approximate weights on each of the seventeen inflation components (contemporaneous + three lags) in the one-sided MUCSVO trend estimate (solid line), along with the expenditure share (dashed).

The volatility of interest rates in the late 1970s and early 1980s leads to large volatility in this sector's measured inflation, resulting in a large increase $\sigma_{\varepsilon,i,t}$, the volatility of the sector-specific transitory term, ε_{it}. Despite the break in measurement, there is little evidence for a break in the factor loadings, although these are estimated imprecisely, and the appendix shows that this applies to the other sectors as well. There are several sector-specific outliers, both before and after the break in measurement.

The similarities between the estimated trend in the MUCSVO model and the univariate UCSVO estimates using the core inflation measures raise the question of whether the multivariate trend is in effect a temporally smoothed version of core inflation and, more generally, what the time-varying weights implicitly used in the multivariate trend are. For example, consider one-sided estimates of the trend. At any given point in time, the one-sided estimate from the multivariate trend is a nonlinear function of current and past values of the seventeen sectoral inflation rates. Because of the time-varying parameters in the MUCSVO model, these weights evolve over time, and they involve lags because of the time series smoothing implied by the model. The function of current and past values is also nonlinear because of the outlier variable. For these reasons, an exact representation in terms of a time-varying linear weighted average is not feasible. Nevertheless, useful insights into the cross-sectional smoothing can be obtained

by looking at approximate time-varying weights. Specifically, at a given date, a linear approximation to the one-sided trend estimates can be computed using a Kalman filter based on equations (6) to (10), holding fixed the values of the time-varying factor loadings and volatilities ($\alpha_{c,t}$, $\alpha_{i,t}$, $\Delta\ln(\sigma^2_{\Delta\tau,c,t})$, $\Delta\ln(\sigma^2_{\varepsilon,c,t})$. $\Delta\ln(\sigma^2_{\Delta\tau,i,t})$, and $\Delta\ln(\sigma^2_{\varepsilon,i,t})$) at their full-sample posterior mean values at that date and ignoring outliers by setting $s_{c,t} = s_{i,t} = 1$. (The online appendix describes these calculations in more detail.)

Figure 5 plots the approximate linear weights on the seventeen components implicit in the one-sided multivariate estimate of the trend, specifically, the sum of the weights on the current and first three lagged values of the component inflation series. Comparing the approximate MUCSVO weight to the expenditure share shows whether, at a given date, the sector is getting more or less weight in the MUCSVO trend than it does in PCE-all.

As can be seen in figure 5, roughly half of the seventeen components receive weight similar to their expenditure shares. The fact that so many of these weights track expenditure shares is by itself interesting, since the expenditure shares are not used in the MUCSVO model (expenditure shares are used in equation [13] to construct the overall trend based on the seventeen filtered individual trends and the filtered common trend, but not in the calculation of the estimates of the individual and common trends). Components with weights that track expenditure shares include

FIGURE 6.—APPROXIMATE WEIGHTS ON CORE, FOOD, AND ENERGY SECTORS FOR THE SEVENTEEN-COMPONENT MUCSVO ESTIMATED TREND AND EXPENDITURE SHARE
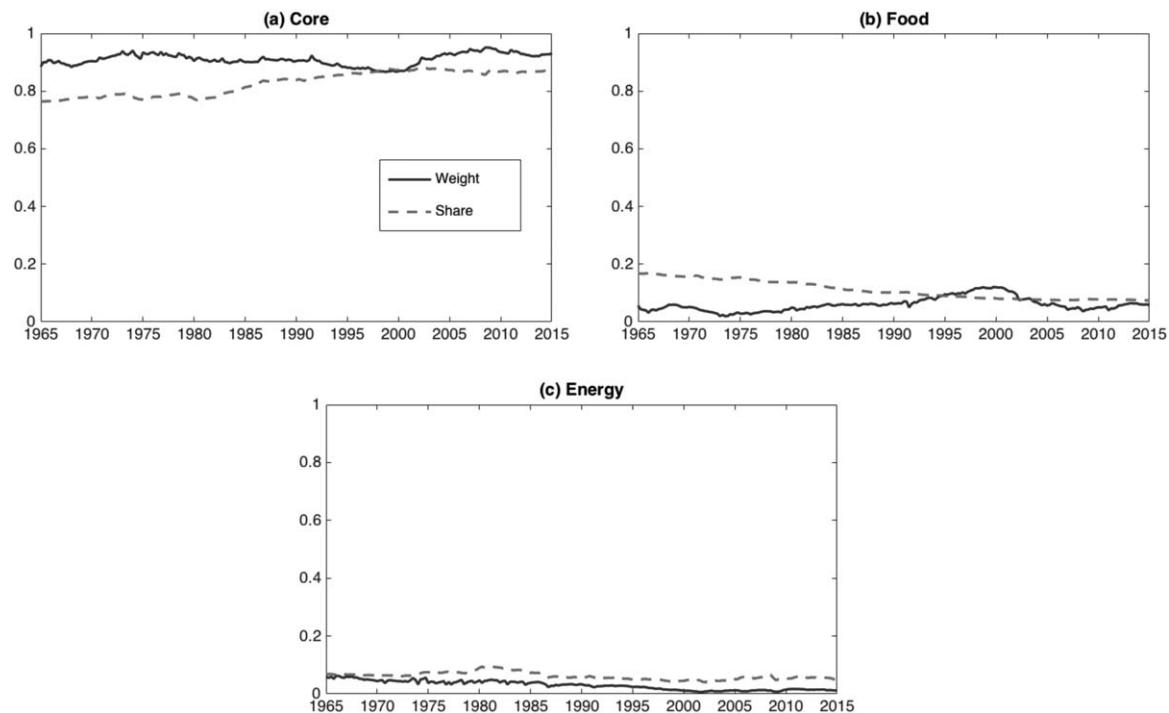
See notes to figure 5.

recreational goods and vehicles, other durable goods, other nondurable goods, housing excluding energy services, health care, transportation services, other services, and NPISHs.

Other series have large swings in their weights. The weight on food and beverages for off-premises consumption ("food at home") increases substantially; since the mid-1990s, it essentially equals its expenditure share. And the weight on food services and accommodations rises from its share in the mid-1970s to nearly double its share in the 1980s and 1990s. Relative to their expenditure shares, the weights fell on financial services and insurance (since the late 1970s), clothing and footwear (since the early 1980s), furnishings and durable household equipment (since the mid-1980s), and gas and electric utilities (since the mid-1990s). Except during the 1960s, gasoline and energy goods receive essentially zero weight.

Figure 6 shows these sectoral weights aggregated to core, food, and energy, where food is food for off-premises consumption, energy is gasoline and other energy goods and gas and electric utilities, and core consists of the remaining fourteen sectors. As can be seen from these weights, the multivariate trend estimate evolved to increase the weight on food and decrease the already low weight on energy, around 1990.

To better understand the reasons for these time-varying weights, we now take a closer look at four of the sectoral inflation rates. Figure 7 plots time series for these inflation rates along with posterior median and (point-wise) 67% pos-

terior intervals for the standard deviation of their idiosyncratic noise components, $\sigma_{\varepsilon,i,t}$. The first inflation series is for food services and accommodations. Inflation in this sector tracks PCE-all inflation for the full sample but has higher idiosyncratic volatility in the 1960s and 1970s. The reduction in the post-1970s idiosyncratic volatility makes this series a better indicator of trend inflation, and thus the series receives more weight in the trend estimate beginning in the late 1970s.

The next inflation series shown in figure 7 is for food and beverages for off-premises consumption. This series is noisy early in the first half of the sample and less so later in the sample. The decrease in volatility of its idiosyncratic transitory innovation makes it a better indicator of trend inflation, so its weight in the estimate of trend inflation increases in the second half of the sample even though its expenditure share is falling.

The third series is furnishing and durable household equipment, which smoothly tracks PCE inflation early in the sample but diverges and exhibits increased volatility from around 1990. While this component receives considerable weight—more than twice its expenditure share—in the MUCSVO trend in the pre-1980 period, its weight drops to its expenditure share since 1990.

The final series is gasoline and other energy goods, which since the early 1970s has exhibited volatility that is an order of magnitude larger than the other sectoral inflation measures. Variations in this series are a poor indicator of trend inflation and the series receives essentially zero-weight in the estimated MUCSVO trend.

FIGURE 7.—SECTORAL INFLATION AND σ_{ε,i,t} FOR FOUR SECTORS
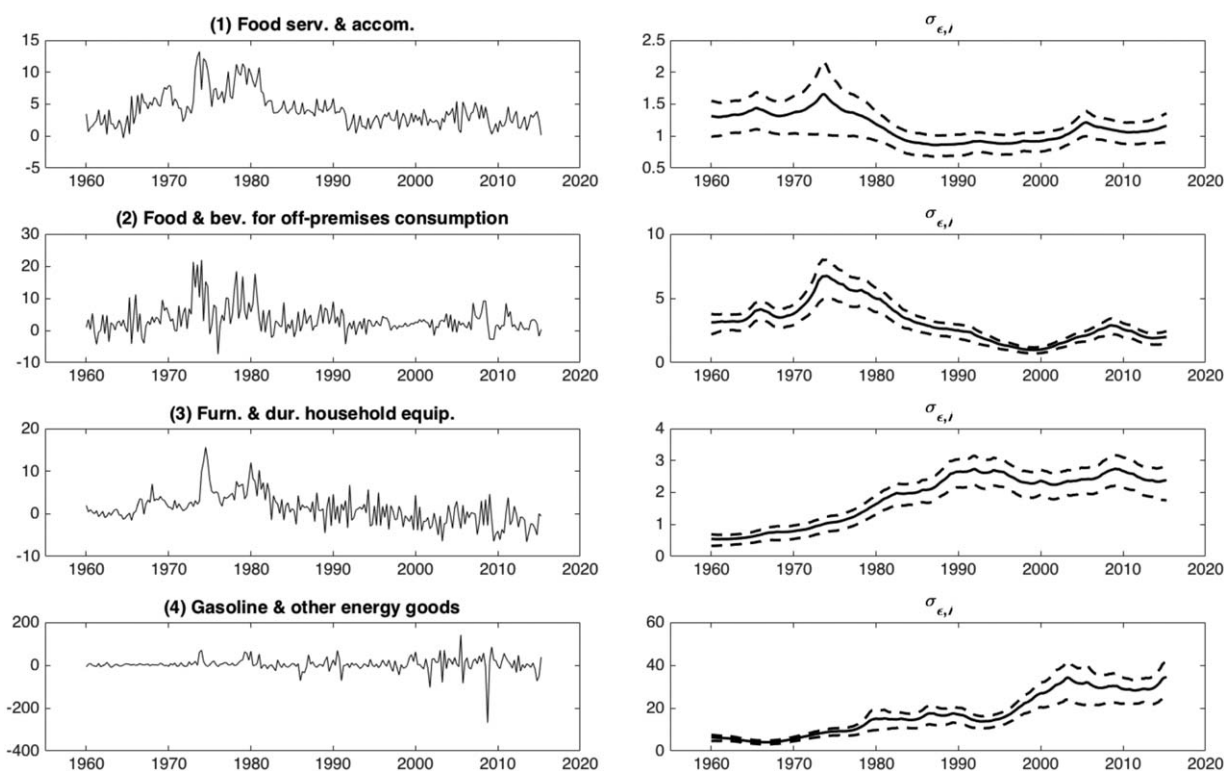
The first column plots sectoral inflation, and the second column plots the median and 67% full sample posterior intervals for $\sigma_{\varepsilon,i,t}$ for the sector.

B. Three-Sector Model

The results for the seventeen-sector model raise the question of whether similar results can be obtained using a simpler three-sector model consisting of core (PCExFE), energy (the two energy components excluded from core, combined with their share weights), and food (off-premises). We therefore estimated this three-component model using the multivariate model of section II. Selected results for this model are presented below, and detailed results are available in the online appendix.

C. Accuracy of the Trend Estimates

One of the motivating questions of this work is whether using sectoral information can improve the precision of the estimator of the trend in headline inflation. Because trend inflation is never observed, the precision of the various estimators cannot be computed directly from the data. In this section, we present model-based accuracy measures based on the width of posterior uncertainty intervals, which are complemented in the next section with a pseudo-out-of-sample forecast experiment.

Figure 8a plots point-wise 90% posterior intervals for the trend in PCE, all computed from the UCSVO and seventeen-component MUCSVO models. The width of these intervals reflects two distinct sources of uncertainty: (a) sig-

nal extraction uncertainty conditional on values of the model's parameters and (b) uncertainty about the model parameters. Because the information set for the multivariate model is strictly larger than univariate model, signal extraction uncertainty is smaller in the MUCSVO model. However, many more parameters are estimated in the MUCSVO model, so parameter uncertainty may be larger, and therefore there is no a priori ranking of the width of posterior intervals in the UCVSO and MUCSVO models.

Examination of figure 8a shows that the MUCSVO intervals are visibly narrower than the UCSVO bands, suggesting a substantial reduction in uncertainty using the information in the multivariate model, even at its cost of additional complexity. Table 2A—summarizes these results by showing the average width of 67% and 90% posterior intervals for the UCSVO and MUCSVO models over the first and second halves of the sample. The 67% and 90% full-sample posterior intervals for the PCE-all trend (labeled $\tau_t^{\text{PCE-all}}$ in the table) are roughly 35% narrower than the corresponding intervals for the univariate model.

Figure 9 shows the corresponding intervals, but for posteriors computed recursively using data from the beginning of the sample through time t. We compute these one-sided posteriors beginning in $t = 1990{:}Q1$ and continuing through the end of the sample (2015Q2). Because these one-sided intervals use less information than the full-sample posteriors, they are necessarily wider, but as the values in figure

FIGURE 8.—NINETY PERCENT FULL SAMPLE POSTERIOR INTERVALS FROM UNIVARIATE AND MULTIVARIATE MODELS
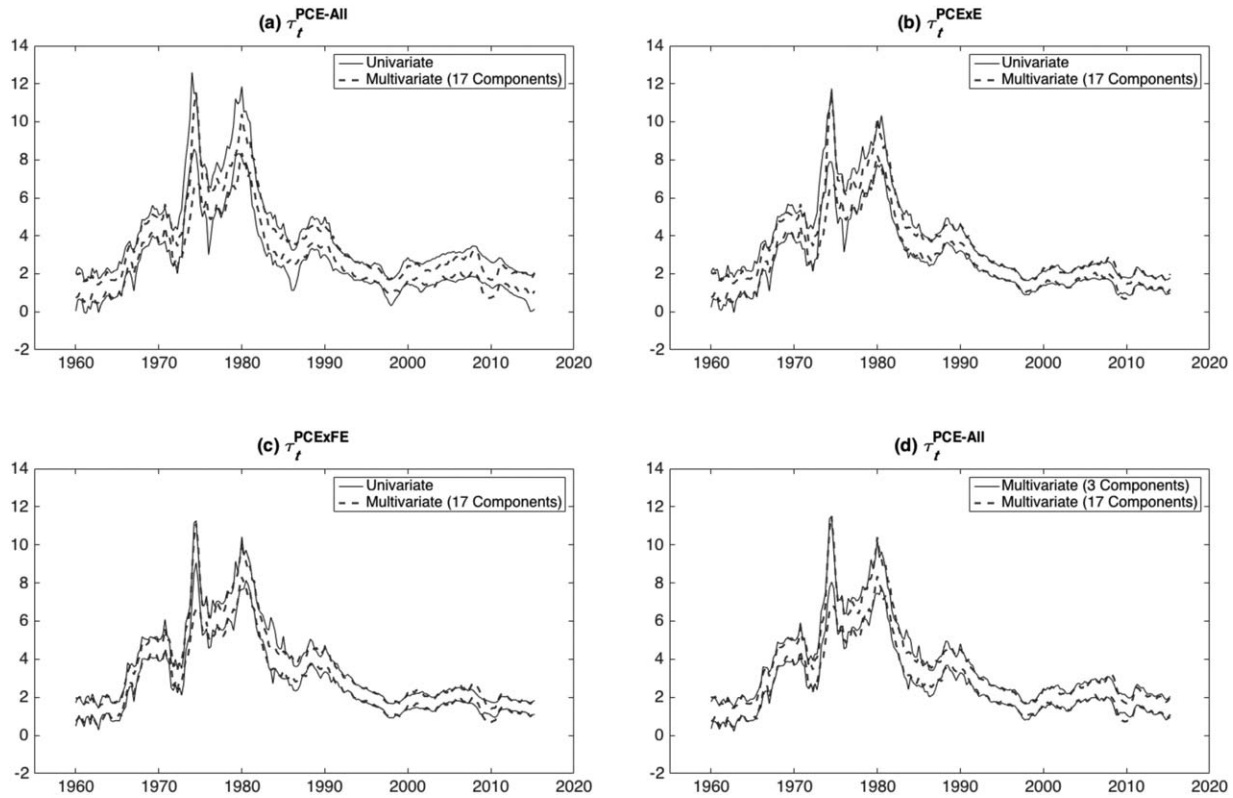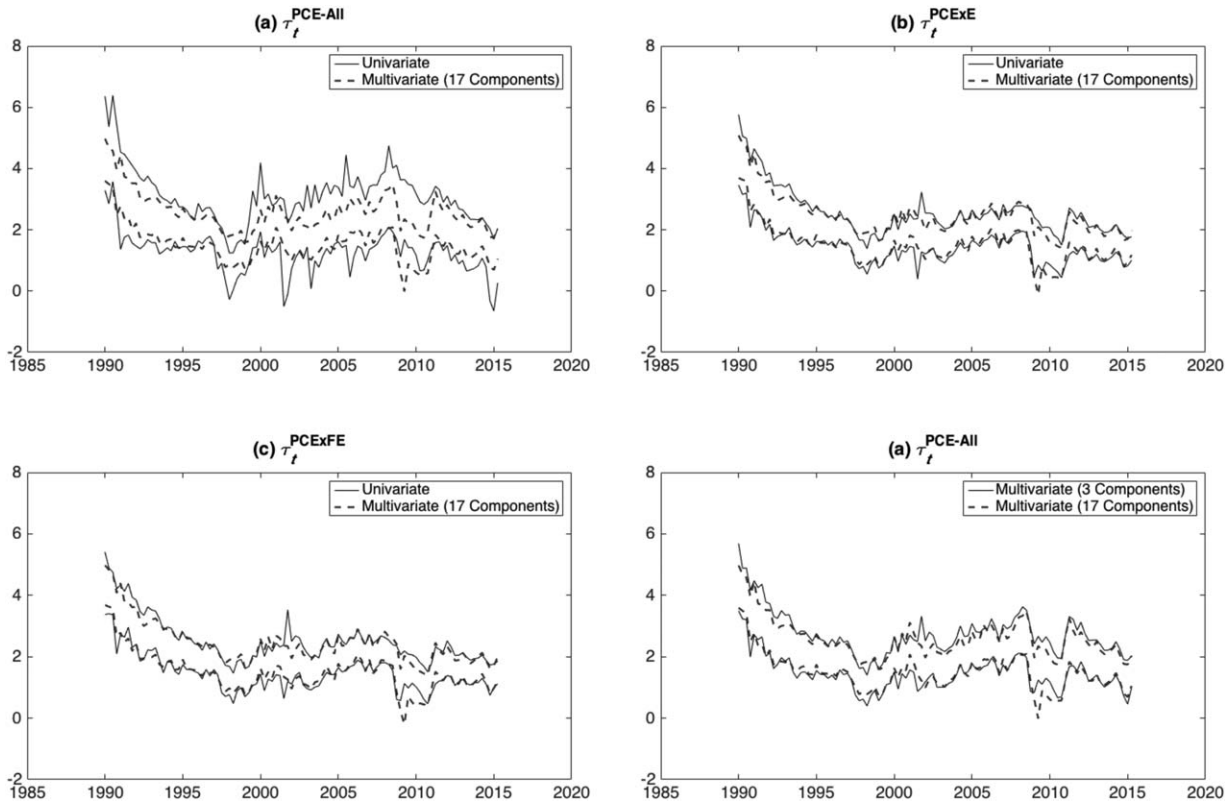

TABLE 2.—AVERAGE WIDTH OF 90% POSTERIOR INTERVALS FOR TREND INFLATION

	Univariate		Multivariate			
Inflation Trend			Three Components		Seventeen Components	
	1960–1989	1990–2015	1960–1989	1990–2015	1960–1989	1990–2015

A. Full-Sample Posterior

	1960–1989	1990–2015	1960–1989	1990–2015	1960–1989	1990–2015
67% intervals						
$\tau^{PCE\text{-}All}$	1.17	0.75	0.88	0.56	0.75	0.49
τ^{PCExE}	1.02	0.45	0.85	0.48	0.73	0.42
τ^{PCExFE}	0.78	0.42	0.81	0.49	0.66	0.41
90% intervals						
$\tau^{PCE\text{-}All}$	2.03	1.32	1.52	0.96	1.27	0.84
τ^{PCExE}	1.77	0.79	1.47	0.84	1.24	0.73
τ^{PCExFE}	1.37	0.73	1.40	0.85	1.14	0.71

B. One-Sided Posterior, 1990–2015

Inflation Trend	Univariate	Multivariate	
		Three Components	Seventeen Components
67% error bands			
$\tau_t^{PCE\text{-}All}$	1.06	0.75	0.64
τ_t^{PCExE}	0.66	0.67	0.58
τ_t^{PCExFE}	0.62	0.68	0.58
90% error bands			
$\tau_t^{PCE\text{-}All}$	1.88	1.31	1.10
τ_t^{PCExE}	1.15	1.17	0.99
τ_t^{PCExFE}	1.09	1.20	0.99

The table shows the average width of (equal-tailed, point-wise) posterior intervals for the inflation trends listed in the first column. Panel A uses the full-sample posterior. Panel B uses a sequence of one-sided posteriors using samples ending in period t, for $t = $ 1990Q1 through 2015Q2.

9b indicate, the seventeen-component MUCSVO model again produces intervals that are roughly 40% narrower than UCSVO model.

The MUCSVO model can also be used to estimate the trend in the core measures of inflation, PCExE and PCExFE, by using equation (13), but with share weights (w_{it}) appropriate for these measures. Comparisons of the posterior intervals for these multivariate estimates and their univariate counterparts are shown in panels b and c of figures 8 and 9, and the average widths of these intervals are shown in table 2. The relative improvements of the accuracy of the multivariate models are much smaller for estimates of the trend in these core measures of inflation. For example, the one-sided multivariate intervals are 10% to 15% percent narrower than their univariate counterparts for core inflation, compared to 40% narrower for headline inflation. And the relative gains for the full sample estimates are even less.

The final panel in figures 8 and 9 compares the three-component and seventeen-component MUCSVO intervals for the trend in headline inflation. The average widths in table 2 suggest that some, but not all, of the accuracy gains for estimating $\tau^{\text{PCE-all}}$ are achieved by the three-component model. Interesting, for estimating the trends in core inflation, the three-component MUCSVO model produces intervals that are wider than the univariate models, indicating

that the increased parameter uncertainty outweighs the signal extraction information.

V. Forecasting Performance

The definition of trend inflation as the forecast of inflation over the long run suggests using forecasting performance to evaluate candidate estimates of trend inflation. Following much of the literature on inflation forecasting using core inflation, we focus on forecasts at the one- to three-year horizon and carry out a pseudo-out-of-sample forecast comparison.

Specifically, we use the one-sided posterior mean estimates of τ_t, denoted by $\tau_{t/t}$ and described in the previous section, from the various models to forecast the average value of inflation over the next four, eight, and twelve quarters, that is, to forecast $\bar{\pi}^{\text{PCE-all}}_{t+1:t+h} = h^{-1} \sum_{i=1}^{h} \pi^{\text{PCE-all}}_{t+i}$ for $h = 4, 8$, and 12 and where $\pi^{\text{PCE-all}}_t$ is the date t value of PCE-all inflation. Forecasts are constructed using $\tau^{\text{PCE-all}}_{t/t}$, constructed from the univariate UCSVO and three- and seventeen-component MUCSVO models and from $\tau^{\text{PCExE}}_{t/t}$ and $\tau^{\text{PCExFE}}_{t/t}$ computed from univariate UCSVO models. The variable being forecast is headline inflation, $\bar{\pi}^{\text{PCE-all}}_{t+1:t+h}$, in all of the experiments even when being forecast by the core trend estimates. We also consider six benchmark forecasts: random

TABLE 3.—MEAN SQUARED FORECAST ERRORS FOR VARIOUS PRICE-BASED INFLATION FORECASTS:
MODEL-BASED ESTIMATED TRENDS AND BENCHMARK FORECASTING MODELS

| | | | A. 1990Q1 through End of Sample | | | |
| | Four-Quarter-Ahead Forecasts | | Eight-Quarter-Ahead Forecasts | | Twelve-Quarter-Ahead Forecasts | |
	MSFE	Difference	MSFE	Difference	MFSE	Difference
Multivariate UCSVO forecasts						
17c (α_t)	0.96 (0.33)	0.00 (0.00)	0.60 (0.14)	0.00 (0.00)	0.50 (0.10)	0.00 (0.00)
3c (α_t)	0.95 (0.34)	−0.01 (0.03)	0.62 (0.15)	0.02 (0.04)	0.52 (0.10)	0.02 (0.04)
17c (α)	0.97 (0.33)	0.01 (0.01)	0.60 (0.14)	0.01 (0.01)	0.51 (0.09)	0.01 (0.01)
3c (α)	0.96 (0.34)	−0.00 (0.04)	0.62 (0.16)	0.03 (0.04)	0.53 (0.11)	0.03 (0.04)
Univariate UCSVO Forecasts						
PCE-all	1.09 (0.42)	0.13 (0.12)	0.81 (0.23)	0.21 (0.11)	0.69 (0.16)	0.19 (0.11)
PCExE	0.87 (0.28)	−0.08 (0.07)	0.58 (0.12)	−0.02 (0.06)	0.49 (0.10)	−0.01 (0.06)
PCExFE	**0.86** (0.25)	−0.10 (0.10)	**0.55** (0.10)	−0.05 (0.09)	**0.48** (0.10)	−0.02 (0.08)
Forecasts using contemporaneous values of inflation						
PCE-all	3.67 (1.67)	2.71 (1.42)	3.32 (1.45)	2.72 (1.33)	3.29 (1.47)	2.78 (1.41)
PCExE	1.23 (0.35)	0.27 (0.10)	0.91 (0.20)	0.32 (0.10)	0.87 (0.17)	0.37 (0.10)
PCExFE	1.18 (0.28)	0.22 (0.13)	0.84 (0.15)	0.24 (0.13)	0.82 (0.15)	0.31 (0.11)
Forecasts using four-quarter averages of inflation						
PCE-all	1.51 (0.61)	0.55 (0.31)	1.23 (0.44)	0.63 (0.32)	1.12 (0.36)	0.62 (0.30)
PCExE	0.94 (0.30)	−0.02 (0.06)	0.66 (0.14)	0.06 (0.05)	0.56 (0.10)	0.06 (0.04)
PCExFE	0.91 (0.26)	−0.05 (0.09)	0.61 (0.11)	0.01 (0.08)	0.53 (0.09)	0.03 (0.07)
			B. 1990Q1 through End-of-Sample, excluding 2008Q4			
	Four-Quarter-Ahead Forecasts		Eight-Quarter-Ahead Forecasts		Twelve-Quarter-Ahead Forecasts	
	MSFE	Difference	MSFE	Difference	MFSE	Difference
Multivariate UCSVO Forecasts						
17c (α_t)	0.61 (0.10)	0.00 (0.00)	**0.48** (0.07)	0.00 (0.00)	**0.41** (0.08)	0.00 (0.00)
3c (α_t)	**0.58** (0.08)	−0.03 (0.03)	**0.48** (0.08)	0.00 (0.04)	0.42 (0.09)	0.01 (0.04)
17c (α)	0.62 (0.10)	0.01 (0.01)	0.49 (0.07)	0.01 (0.01)	0.42 (0.08)	0.01 (0.01)
3c (α)	**0.58** (0.08)	−0.03 (0.03)	0.48 (0.08)	0.01 (0.04)	0.43 (0.09)	0.02 (0.04)
Univariate UCSVO forecasts						
PCE-all	0.66 (0.10)	0.05 (0.08)	0.63 (0.13)	0.15 (0.10)	0.57 (0.14)	0.16 (0.10)
PCExE	0.59 (0.10)	−0.02 (0.04)	0.50 (0.08)	0.02 (0.05)	0.45 (0.10)	0.04 (0.05)
PCExFE	0.61 (0.10)	0.00 (0.05)	0.49 (0.08)	0.02 (0.06)	0.45 (0.11)	0.04 (0.06)
Forecasts using contemporaneous values of inflation						
PCE-all	3.08 (1.28)	2.47 (1.28)	3.04 (1.24)	2.56 (1.22)	3.13 (1.34)	2.72 (1.34)
PCExE	0.92 (0.15)	0.31 (0.11)	0.81 (0.14)	0.33 (0.11)	0.80 (0.15)	0.39 (0.10)
PCExFE	0.95 (0.15)	0.34 (0.11)	0.80 (0.13)	0.32 (0.11)	0.79 (0.15)	0.38 (0.11)
Forecasts using four-quarter averages of inflation						
PCE-all	0.96 (0.20)	0.35 (0.18)	0.98 (0.27)	0.51 (0.24)	0.92 (0.25)	0.51 (0.25)
PCExE	0.63 (0.10)	0.02 (0.04)	0.56 (0.09)	0.08 (0.04)	0.51 (0.09)	0.10 (0.03)
PCExFE	0.65 (0.10)	0.04 (0.05)	0.54 (0.08)	0.07 (0.05)	0.49 (0.10)	0.08 (0.05)

The entries labeled "MSFE" are the mean square forecast errors. The entries labeled "Difference" are the difference between that row's MSFE for and the MSFE for the seventeen-component multivariate UCSVO model. HAC standard errors are in parentheses. Minimum MSFE forecasts for a given horizon are in bold. Units are squared percentage points at an annual rate. The multivariate UCSVO models labeled (α_t) allow for time varying factor loadings as described in the text. The models labeled (α) use time-invariant factor loadings.

walk models using (separately) lagged PCE-all, lagged PCExE, and lagged PCExFE, and the Atkeson-Ohanian (2001) four-quarter random walk model computed using (separately) PCE-all, PCExE, and PCExFE. And because of the weak evidence of time variation in the factor loadings in the MUCSVO models, we also show results from MUCSVO models that impose constant factor loadings by setting $\lambda_{i,\tau} = \lambda_{i,\varepsilon} = 0$ in equation (11). Forecasts are constructed from $t = $ 1990:Q1 through the end of the sample.

Table 3 summarizes the results. Table 3A shows results for the entire 1990Q1–end of sample period. For each forecast, the table reports the sample mean square forecast error (MSFE) together with its estimated standard error and the difference between the forecast's MSFE and the MSFE of the seventeen-component MUCSVO model, together with

its standard error. The values of these MSFEs are greatly affected by the large outlier in $\pi^{\text{PCE-all}}_{2008Q4}$ evident in figure 1. Table 3B shows results from the same forecasting exercise, but with this single observation omitted from the sample values of $\bar{\pi}^{\text{PCE-all}}_{t+1:t+h}$. We concentrate our discussion on the table 3B results.

Five results stand out from this forecasting experiment. First, forecasts that use moving averages of past inflation are more accurate than forecasts that do not. All of the moving-average forecasts (from the simple four-quarter moving averages to the more sophisticated moving averages in the UCSVO and MUCSVO models) produce markedly more accurate forecasts than the corresponding forecasts using only contemporaneous values of inflation. Second, forecast accuracy is improved by downweighting some sectors, most

notably energy. Forecasts that put little or no weight on energy, whether by using the core inflation measures or the MUCSVO models, are more accurate than forecasts based on headline inflation, regardless of the moving-average filter used. Third, the UCSVO forecasts have smaller MSFE than the four-quarter moving-average forecasts, suggesting that there are gains from using forecasts that adapt to the changing persistence in the inflation process. Fourth, the MUCSVO models with and without time variation in the factor loadings perform similarly. And finally, there are only small (perhaps zero) marginal improvements in accuracy for the MUCSVO forecasts relative to the core-inflation UCSVO forecasts.

VI. Discussion and Conclusion

Previous work has shown that the random-walk-plus-white-noise unobserved components model with stochastic volatility provides a simple but flexible univariate framework for describing the persistence and volatility of inflation, estimating its trend, and forecasting future inflation. This paper has investigated a multivariate extension of that model that allows sectoral inflation potentially to improve on the univariate estimates of trend inflation, much like traditional core inflation does for headline inflation.

The analysis leads to two major conclusions. First, there are substantial gains from using sectoral inflation over using headline inflation. The multivariate estimates of the trend in aggregate (headline) inflation are more accurate than the univariate estimates regardless of whether accuracy is measured by model-based uncertainty or pseudo-out-of-sample forecasting accuracy. But second, the analysis suggests that much of this improved accuracy can be achieved from univariate estimates constructed from traditional core measures of inflation. Model-based uncertainty measures suggest that univariate estimates of the trend in core inflation are nearly as accurate as multivariate estimates of these same trends. Moreover, the pseudo-out-of-sample experiments suggest little difference in the accuracy of these estimates of core trend inflation for forecasting future headline inflation.

The results also lead to two other conclusions. The first is that the reduced volatility of food prices, relative to before the mid-1980s, led the multivariate model to include food in the trend estimate post-1990, with a weight close to its expenditure share. This finding suggests paying more attention to PCExE than to PCExFE. Second, the UCSVO models (univariate or multivariate) have the advantage of producing measures of precision of trend estimates (posterior coverage regions). Currently, the width of these 67% regions is approximately 0.6 percentage point using the seventeen-variable or univariate core trend estimates. We see merit to reporting these estimates of the precision of trend inflation along with estimates of that trend.

Finally, we highlight three areas where the analysis might be extended. First, there are a myriad of ways the multivariate model might be changed by, for example, including additional factors or allowing for different dynamics. We experimented with several of these before settling on the specification used here, but our experiments were far from exhaustive. Second, we investigated three- and seventeen-component models, but much finer sectoral disaggregation is possible. Our initial look at more finely disaggregated data suggested substantial challenges associated with instability in measurement, but clever modeling might address those challenges. Finally, and most important, this analysis has used quarterly averages of monthly inflation rates. Real-time analysis would benefit from directly modeling the monthly data. Our experiments applying the UCSVO and MUCSVO models directly to the monthly data yielded forecasts that were less accurate than the forecasts from the quarterly data. (Results are reported in the online appendix.) This suggests that a successful monthly model will require alternative specifications for the transitory and trend innovations.

REFERENCES

Altissimo, Filippo, Benoit Mojon, and Paolo Zaffaroni, "Can Aggregation Explain the Persistence of Inflation?" *Journal of Monetary Economics* 56 (2009), 231–241.

Amstad, Marlene, and Simon M. Potter, "Real Time Underlying Inflation Gauges for Monetary Policy Makers," unpublished manuscript, Federal Reserve Bank of New York (2007).

Atkeson, A., and L. E. Ohanian, "Are Phillips Curves Useful for Forecasting Inflation?" *Federal Reserve Bank of Minneapolis Quarterly Review* 25:1 (2001), 2–11.

Baumeister, Christiane, and Gert Peersman, "Time-Varying Effects of Oil Supply Shocks on the U.S. Economy," *American Economic Journal: Macroeconomics* 5 (2013), 1–28.

Blanchard, Olivier J., and Jordi Galí, "The Macroeconomic Effects of Oil Price Shocks: Why Are the 2000s So Different from the 1970s?" in Jordi Galí and Mark J. Gertler, eds., *International Dimensions of Monetary Policy* (Chicago: University of Chicago Press, 2010).

Boivin, Jean, Marc P. Giannoni, and Ilian Mihov, "Sticky Prices and Monetary Policy: Evidence from Disaggregated Data," *American Economic Review* 99 (2009), 350–384.

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

——— "Measuring Core Inflation" (pp. 195–215), in N. Gregory Mankiw, ed., *Monetary Policy* (Chicago: University of Chicago Press, 1994).

Chen, Shiu-Sheng, "Oil Price Pass-Through into Inflation," *Energy Economics* 31 (2009), 126–133.

Clark, Todd E., "Comparing Measures of Core Inflation," *Federal Reserve Bank of Kansas City Economic Review* 86: Second Quarter (2001), 5–31.

Clark, Todd E., and Stephen J. Terry, "Time Variation in the Inflation Passthrough of Energy Prices," *Journal of Money, Credit and Banking* 42 (2010), 1419–1433.

Cogley, Timothy, "A Simple Adaptive Measure of Core Inflation," *Journal of Money, Credit, and Banking* 34 (2002), 94–113.

Cogley, Timothy, and Thomas J. Sargent, "Measuring Price-Level Uncertainty and Instability in the U.S., 1850–2012," this REVIEW 97 (2015), 827–838.

Cristadoro, Riccardo, Mario Forni, Lucrezia Reichlin, and Giovanni Veronese, "A Core Inflation Indicator for the Euro Area," *Journal of Money, Credit, and Banking* 37 (2005), 539–560.

Crone, Theodore M., N. Neil K. Khettry, Loretta J. Mester, and Jason A. Novak, "Core Measures of Inflation as Predictors of Total Inflation," *Journal of Money, Credit and Banking* 45 (2013), 505–519.

De Gregorio, José, Oscar Landerretche, and Christopher Neilson, "Another Pass-Through Bites the Dust? Oil Prices and Inflation," *Economia* 7 (2007), 155–208.

Del Negro, Marco, and Christopher Otrok, "Dynamic Factor Models with Time-Varying Parameters: Measuring Changes in International

Business Cycles," Federal Reserve Bank of New York staff report 326 (2008).

Del Negro, Marco, and Giorgio E. Primiceri, "Time Varying Structural Vector Autoregressions and Monetary Policy: A Corrigendum," *Review of Economic Studies* 82 (2015), 1342–1345.

Eckstein, Otto, *Core Inflation* (New York: Prentice Hall, 1981).

Garnier, Christine, Elmar Mertens, and Edward Nelson, "Trend Inflation in Advanced Economies," *International Journal of Central Banking* 11 (2015), 65–136.

Gordon, Robert, "Alternative Responses of Policy to External Supply Shocks," *Brookings Papers on Economic Activity* 1 (1975), 183–206.

Hood, Kyle K., "Measuring the Services of Commercial Banks in the National Income and Products Accounts: Changes in Concepts and Methods in the 2013 Comprehensive Revision," *Survey of Current Business* (February 2013), 8–19.

Hooker, Mark A., "Are Oil Shocks Inflationary? Asymmetric and Non-linear Specifications versus Changes in Regime," *Journal of Money, Credit, and Banking* 34 (2002), 540–561.

Kiley, Michael T., "Estimating the Common Trend Rate of Inflation for Consumer Prices and Consumer Prices excluding Food and Energy," FEDS working paper 2008-38 (2008).

Kim, Sanjoon, Neil Shephard, and Siddhartha Chib, "Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models," *Review of Economic Studies* 65 (1998), 361–393.

Mertens, Elmar, "Measuring the Level and Uncertainty of Trend Inflation," this REVIEW (forthcoming).

Mertens, Elmar, and James M. Nason, "Inflation and Professional Forecast Dynamics: An Evaluation of Stickiness, Persistence, and Volatility," manuscript, Board of Governors of the Federal Reserve System (2015).

Mumtaz, Haroon, and Paolo Surico, "Evolving International Inflation Dynamics: World and Country-Specific Factors," *Journal of the European Economics Association* 10 (2012), 716–734.

Nelson, C. R., and G. W. Schwert, "Short-Term Interest Rates as Predictors of Inflation: On Testing the Hypothesis that the Real Rate of Interest Is Constant," *American Economic Review* 67 (1977), 478–486.

Omori, Yasuhiro, Siddhartha Chib, Neil Shephard, and Jouchi Nakajima, "Stochastic Volatility with Leverage: Fast and Efficient Likelihood Inference," *Journal of Econometrics* 140 (2007), 425–449.

Reis, Ricardo, and Mark W. Watson, "Relative Goods' Prices, Pure Inflation, and the Philips Correlation," *American Economic Journal: Macroeconomics* 2 (2010), 128–157.

Rich, Robert, and Charles Steindel, "A Review of Core Inflation and an Evaluation of Its Measures," Federal Reserve Bank of New York staff report 236 (2005).

Sbrana, Giacomo, Andrea Silvestrini, and Fabrizio Venditti, "Short Term Inflation Forecasting: The M.E.T.A. Approach," manuscript, Bank of Italy (2015).

Stock, James H., and Mark W. Watson, "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking* 39 (2007), 3–33.

van den Noord, Paul, and Christophe André, "Why Has Core Inflation Remained So Muted in the Face of the Oil Shock?" OECD Economics Department working paper 551 (2007).

Wynne, Mark A., "Core Inflation: A Review of Some Conceptual Issues," *Federal Reserve Bank of St. Louis Review* 90:3, Part 2 (2008), 205–228.