

TREND, SEASONAL, AND SECTORIAL INFLATION IN THE EURO AREA

James H. Stock
Harvard University

Mark W. Watson
Princeton University

Pages 317-344 | Chapter printout

Changing Inflation Dynamics, Evolving Monetary Policy

Gonzalo Castex, Jordi Gali,
and Diego Saravia
editors



27

Series on Central
Banking, Analysis,
and Economic Policies

Banco Central de Chile
Central Bank of Chile

The views and conclusions presented in the book are exclusively those of the authors and do not necessarily reflect the position of the Central Bank of Chile or its Board Members.

Copyright © Banco Central de Chile 2020
Agustinas 1180
Santiago, Chile

All rights reserved

Published in Santiago, Chile by the Central Bank of Chile

Manufactured in Chile

This book series is protected under Chilean Law 17336 on intellectual property. Hence, its contents may not be copied or distributed by any means without the express permission of the Central Bank of Chile. However, fragments may be reproduced, provided that a mention is made of the source, title, and author(s).

ISBN (print) 978-956-7421-67-1

ISBN (digital) 978-956-7421-68-8

Intellectual Property Registration 2020-A-2881

ISSN 0717-6686 (Series on Central Banking, Analysis, and Economic Policies)

TREND, SEASONAL, AND SECTORIAL INFLATION IN THE EURO AREA

James H. Stock
Harvard University

Mark W. Watson
Princeton University

A central focus of monetary policy is the underlying rate of inflation that might be expected to prevail over a horizon of one or two years. Because inflation is estimated from noisy data, the estimation of this underlying rate of inflation, which we refer to as trend inflation, requires statistical methods to extract the inflation “signal” from the noise. The task of measuring trend inflation is further complicated by the large seasonal fluctuations in many prices, so that attempts to estimate core or trend inflation at a frequency higher than annual must additionally either use seasonally adjusted data or undertake seasonal adjustment as part of the effort to measure trend inflation.

The challenge of estimating trend inflation is particularly acute for the euro-area Harmonized Index of Consumer Prices (HICP) inflation, official values of which are only reported seasonally unadjusted. Figure 1 plots quarterly values of euro-area HICP inflation (in percentage points at an annual rate) from 2001 to 2018. The quarter-to-quarter variation in inflation is large: the standard deviation of quarterly changes in inflation is 2.5 percentage points. HICP inflation is also highly seasonal: over the entire sample period, inflation averaged 1.6 percent, but averaged 4.8 and 2.2 percent in the second and fourth quarters respectively, and 0.1 and 0.3 over the first and third quarters. While some long-run, low-frequency variation in HICP inflation is evident, that variation —the “signal”— is small compared

to the seasonal variation and what appears to be transient, one-off movements in the rate of inflation. The question, “What is the value of trend inflation today?” is an important one for monetary policy, but the answer to it arguably requires more than just staring at figure 1.

One approach to estimating trend inflation is to exploit variation across the components of inflation (across sectors) to reduce noise. The most prominent such estimates are “core” measures (e.g., Gordon, 1975 Eckstein, 1981) that exclude inflation from the volatile food and energy sectors. Alternative core measures include trimmed mean or median of sectorial inflation rates; for example, see the early work by Bryan and Cecchetti (1994) or the paper by Ball and Mazumder in this volume. Ehrmann and others (2018) provide an up-to-date summary of work at the European Central Bank (ECB) involving underlying and sectorial inflation.

The HICP has 12 second-tier components, which we modify to create 13 components by pooling the energy components of housing and transportation into a separate “energy” component. These 13 inflation components are plotted in figure 2. The heterogeneity of the time-series properties of these components is striking. Some sectors exhibit large seasonal variation (for example, clothing), others exhibit large non-seasonal quarterly variation (energy) or outliers (healthcare), and relative price movements impart different lower-frequency trends in each sector. Almost as striking is the apparent variation over time in those time-series properties, for example, the seasonal components of furnishing, clothing, and transportation have increased markedly over this period. The heterogeneity of these components suggests that there could be considerable gains from using a multivariate approach that allows the components to have distinct time-series properties and uses both time-series smoothing and cross-sectional weighting to estimate aggregate HICP trend inflation.

This paper makes three contributions towards measuring trend HICP inflation. First, we estimate an unobserved components (UC) model with stochastic volatility (UCSV), which extends the UCSV model in Stock and Watson (2007) to include a seasonal component. This univariate model is an extension of the textbook unobserved components model¹ to incorporate stochastic volatility to capture

1. Chapter 1 of Nerlove and others (1979) offers a historical survey of UC models in economics. The textbook by Harvey (1989) is a classic reference on analyzing UC models by using Kalman filter methods.


the time-varying importance of the trend, seasonal, and irregular components.²

Second, we extend the multivariate unobserved components/stochastic volatility model of Stock and Watson (2016) to allow each component to have separate seasonals, also with stochastic volatility. We apply this extended model to the 13 HICP components in figure 2 to obtain multivariate estimates of the trend. We find that doing so produces trend estimates that are more precise than those based on the univariate model of aggregate HICP. We also find that this measure of core inflation moves cyclically with real economic activity.

Third, as a byproduct, we also obtain quarterly estimates of seasonally adjusted HICP. Another approach to handling seasonals is simply to use the four-quarter average of quarterly inflation; however, that measure tends to respond sluggishly. Compared with four-quarter rolling inflation, the new seasonally adjusted HICP series has the potential to provide more timely insights into movements of inflation.

Section 1 presents the univariate and multivariate model that we use for aggregate and sectorial inflation. Section 2 uses these models to estimate trend and seasonal factors for euro-area HICP inflation. Section 3 examines the relation between seasonally adjusted inflation and real activity.

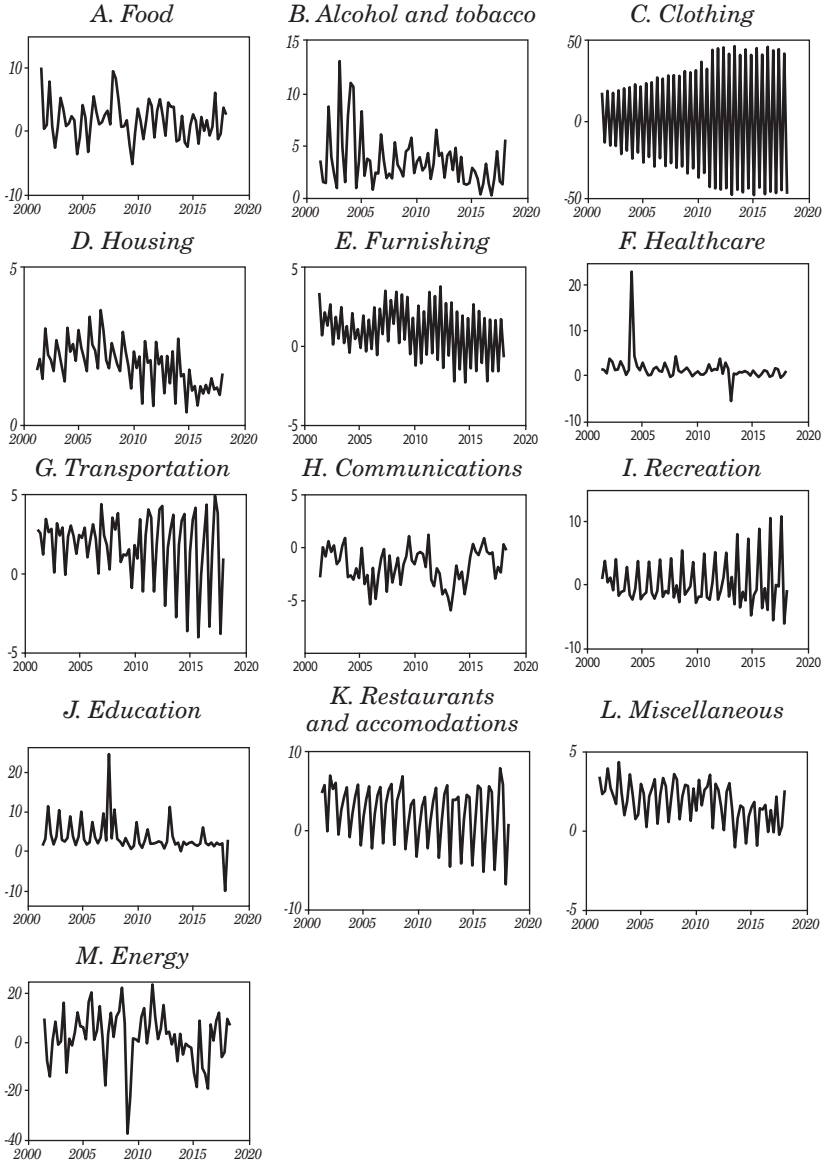
Figure 1. HICP Inflation for the Euro Area



Source: Authors' calculations.

2. Several papers have used related univariate UC models to study the evolution of prices and inflation. Examples include Ball and Cecchetti (1990); Cecchetti and others (2007); Cogley and Sargent (2015); Cogley and others (2015); and Kang and others (2009).

Figure 2. 13 HICP Sectors



Source: Authors' calculations.

Notes: These are the 12 HICP tier-two sectors, with energy excluded from the housing and transportation sectors, and shown separately as the 13th sector.

1. SEASONAL UCSV MODELS

Unobserved components models have a long history in economic time series and have been used for, among other things, data description, forecasting, structural analysis, and seasonal adjustment. Here we present versions of the UC model that can be used to seasonally adjust aggregate inflation and to estimate its trend value. One version of the model is univariate and uses only aggregate inflation; the other is multivariate and models the joint dynamics of sectorial inflation. Both models incorporate stochastic volatility and are known by their acronym UCSV.

1.1 Univariate Seasonal UCSV Model

Inflation is observed quarterly and is denoted by π_t . The UC model decomposes π_t into three unobserved components: trend (τ_t), seasonal (s_t), and irregular (ε_t).

$$\pi_t = \tau_t + s_t + \varepsilon_t. \quad (1)$$

The components are separately identified because they follow distinct stochastic processes. Let $\eta_{\tau,t}$, $\eta_{s,t}$, and $\eta_{\varepsilon,t}$ denote three martingale-difference processes; the trend component follows a martingale:

$$(1 - L)\tau_t = \eta_{\tau,t} \quad (2)$$

so it is dominated by low-frequency, or “trend”, variation; s_t follows the quarterly seasonal process:

$$(1 + L + L^2 + L^3) s_t = \eta_{s,t} \quad (3)$$

so is dominated by variation at the seasonal frequencies with periods 2 and 4 quarters; and the irregular component is unforecastable:

$$\varepsilon_t = \eta_{\varepsilon,t} \quad (4)$$

The unobserved components model (1)–(4) is a version of Harvey’s (1989) “local-level” model, augmented by the seasonal component s_t . Versions of the model (often with more flexible models for the components) are the backbone of model-based seasonal adjustment

methods—e.g., Hillmer and Tiao (1982), Hausman and Watson (1985), and Maravall (1995).

In the non-seasonal version of the local-level model, the estimate of τ_t based on observations of π through date t is the forecast of the future rate of inflation:

$$E(\pi_{t+h} | \{\pi_i\}_{i=1}^t) = E(\tau_{t+h} + \varepsilon_{t+h} | \{\pi_i\}_{i=1}^t) = E(\tau_t | \{\pi_i\}_{i=1}^t) = \tau_{t|t}, \quad (5)$$

where the final equality follows from the martingale assumption for τ_t and the martingale-difference assumption for ε_t .

The seasonal model (3) is specified so that this definition of the trend as the long-run forecast continues to hold for annual averages. Specifically, Harvey (1989), in subsection 6.2, defines a seasonal process to be any time-series process with predicted values that (i) repeat seasonally and (ii) sum to zero over a one-year period. The seasonal process (3) satisfies these two conditions, specifically (i) $s_{T'+jT} = s_{T'+j+4}T$ and (ii) $\sum_{j=1}^4 s_{T'+jT} = 0$, where $s_{r|T}$ is the predicted value of s_r made by using data through time T , for any $T' \geq T$. The seasonal model (3) yields a similar interpretation of $\tau_{t|t}$, but now for annual averages of future values of π : letting $\bar{x}_{i:j}$ denote the sample average of an arbitrary variable x between time i and j ,

$$E\left(\bar{\pi}_{t+j:t+j+3} | \{\pi_k\}_{k=1}^t\right) = E\left(\bar{\tau}_{t+j:t+j+3} + \bar{s}_{t+j:t+j+3} + \bar{\varepsilon}_{t+j:t+j+3} | \{\pi_k\}_{k=1}^t\right) \quad (6)$$

$$= E\left(\tau_t | \{\pi_k\}_{k=1}^t\right) = \tau_{t|t}$$

for $j > 0$, where the penultimate equality follows from the random walk model for τ , $\sum_{j=1}^4 s_{T'+jT} = 0$, and the unpredictability of future ε 's. Thus, as in the model without seasonality, $\tau_{t|t}$ measures the (non-seasonal) forecastable level of inflation.

Examination of the inflation series in figure 1 and figure 2 highlights the need for two modifications of the basic UCSV model. The first modification allows for time variation in the variances of the unobserved components, and the second allows for outliers. We discuss these in turn.

Time-varying variances are added to the model by allowing the shocks in (2), (3), and (4) to follow stochastic volatility processes, say $\eta_t = \sigma_t e_t$, where $e_t \sim$ i.i.d. $N(0,1)$ and σ_t^2 evolves through time as a logarithmic random walk: $(1-L)\ln(\sigma_t^2) = v_t$ with $v_t \sim$ i.i.d. $N(0, \sigma_v^2)$. Kim and others (1998) show how this stochastic volatility model

can be estimated using Gibbs sampling methods by using a mixture of normal densities to approximate the $\log\text{-}\chi_t^2$ density together with standard Kalman smoothing recursions; Omori and others (2007) provide improved approximations. Stock and Watson (2007) incorporate these methods together with ideas in Carter and Kohn (1994), and Kim and Nelson (1999) to estimate a non-seasonal version of the UCSV model.

Outliers are incorporated in the model through additional random multiplicative factors linking the η_t innovations to the i.i.d. $N(0,1)$ shocks e_t . As in Stock and Watson (2016), we use a formulation with $\eta_t = o_t \sigma_t e_t$ where o_t is an i.i.d. outlier term with $o_t = 1$ with probability $1-p$ and $o_t \sim U(2,10)$ with probability p . When $o_t = 1$, there is no outlier, and when $o_t \sim U(2,10)$ there is an outlier with a standard deviation that is between 2 and 10 times larger than in the no-outlier case. In the model for euro-area inflation, we allow outliers only in the irregular component ε_t , as this seems consistent with outliers evident in figure 2; in other applications, outliers might also be appropriate for τ_t and/or s_t .

In summary, the complete UCSV model is (1)–(4) and

$$\eta_{\tau,t} = \sigma_{\tau,t} e_{\tau,t} ; \eta_{s,t} = \sigma_{s,t} e_{s,t} ; \eta_{\varepsilon,t} = o_t \sigma_{\varepsilon,t} e_{\varepsilon,t} \quad (7)$$

$$(1-L)\ln(\sigma_{x,t}) = v_{x,t} \quad \text{for } x = \tau, s, \varepsilon, \quad (8)$$

where $(e_{\tau,t}, e_{s,t}, e_{\varepsilon,t}, v_{\tau,t}, v_{s,t}, v_{\varepsilon,t})$ are mutually independent i.i.d. normal random variables with mean zero, the e terms have unit variance, and each of the v terms has a component-specific variance, say $\sigma_{v(\tau)}$, $\sigma_{v(s)}$, and $\sigma_{v(\varepsilon)}$.

1.2 Multivariate Seasonal UCSV Model

The multivariate model is a generalization of the univariate that includes common and sector-specific versions of the three unobserved components. For each of the $i = 1, \dots, n$ sectors, the rate of price inflation in sector i , $\pi_{i,t}$ follows:

$$\pi_{i,t} = \alpha_{i,\tau} \tau_{c,t} + \alpha_{i,s} s_{c,t} + \alpha_{i,\varepsilon} \varepsilon_{c,t} + \tau_{i,t} + s_{i,t} + \varepsilon_{i,t} \quad (9)$$

where $(\tau_{c,t}, s_{c,t}, \varepsilon_{c,t})$ are common to all sectors, $(\tau_{i,t}, s_{i,t}, \varepsilon_{i,t})$ are sector specific, and $(\alpha_{i,\tau}, \alpha_{i,s}, \alpha_{i,\varepsilon})$ are time-invariant coefficients (factor loadings). The τ, s, ε components follow processes as in the univariate

model, with component/sector-specific parameters. The components are mutually independent, so that dependence across sectors comes from the common components τ_c , s_c , and ε_c . Outliers are allowed in each of the sector-specific $\varepsilon_{i,t}$ components and in the common $\varepsilon_{c,t}$ component.

The multivariate sectorial model is designed so that it (approximately) aggregates to univariate UCSV model. Because of its symmetric structure, aggregation in the multivariate model is straightforward: letting $w_{i,t}$ denote the share weight for sector i at time t

$$\pi_t = \sum_{i=1}^n w_{i,t} \pi_{i,t} = \tau_t^a + s_t^a + \varepsilon_t^a \quad (10)$$

where

$$\tau_t^a = \sum_{c,t} w_{i,t} \alpha_{i,\tau} + \sum w_{i,t} \tau_{i,t} \quad (11)$$

and similarly for the other components. When the share weights are time-invariant, τ_t^a evolves as a martingale, s_t^a follows the seasonal process in (3), and ε_t^a is a martingale difference. And, as in the univariate model, filtered values of $(\tau_t^a, \tau_{c,t}, \tau_{i,t})$ constructed from the multivariate model summarize the forecastable levels in both sectorial and aggregate inflation:

$$E\left(\bar{\pi}_{i,t+j:i,t+j+3} \mid \{\pi_{l,k}\}_{k=1}^t, l=1, \dots, n\right) = \alpha_{i,\tau} \tau_{c,t|t} + \tau_{i,t|t} \quad (12)$$

and

$$E\left(\bar{\pi}_{t+j:t+j+3} \mid \{\pi_{l,k}\}_{k=1}^t, l=1, \dots, n\right) = \tau_{t|t}^a. \quad (13)$$

1.3 Estimation and Inference

We estimate the univariate and multivariate UCSV models by using Bayes' methods that are generalizations of the methods outlined in online appendix to Stock and Watson (2016). We provide an overview here.

The univariate UCSV model is characterized by four sets of parameters: (i) the stochastic volatility innovation standard deviations,

$\sigma_{v(\tau)}$, $\sigma_{v(s)}$, and $\sigma_{v(\varepsilon)}$; (2) the outlier probability parameter p ; (3) the initial values for the standard deviations $\sigma_{\tau,0}$, $\sigma_{s,0}$, and $\sigma_{\varepsilon,0}$; and (4) the initial values of the components τ_0 and $(s_0, s_{-1}, s_{-2}, s_{-3})$. We used independent priors for the parameters:

- $\sigma_v \sim U(0,0.10)$. (A value of $\sigma_{v(\tau)} = 0.10$ implies that the standard deviation of $\ln(\sigma_{\tau+t+40}/\sigma_{\tau,t})$ is approximately 0.3, that is a standard deviation of 30 percent over 40 quarters).
- $p \sim \text{Beta}(a,b)$ with $a = 2.5$ and $b = 37.5$. (This implies that an outlier is expected to occur every four years).
- $\ln(\sigma_{\tau,0})$, $\ln(\sigma_{s,0})$, $\ln(\sigma_{\varepsilon,0})$, and τ_0 follow independent diffuse Gaussian priors.
- $(s_0, s_{-1}, s_{-2}, s_{-3})$ follow a diffuse singular Gaussian distribution, where the singularity enforces $s_0 + s_{-1} + s_{-2} + s_{-3} = 0$.

The multivariate model requires two normalizations. First, the factor structure requires a normalization to separately identify the scales of the factor loadings (α_τ , α_s , α_ε) and the common factors (τ_c , s_c , and ε_c). We normalize the standard deviations of the common factors to be unity for $t=0$. The second normalization is needed because the initial values of the common and idiosyncratic factors (e.g., $\tau_{c,0}$ and $\tau_{i,0}$) are not separately identified. To identify the model, we normalize the common factors to be zero for $t=0$; that is $\tau_{c,0} = 0$ and $(s_{c,0}, s_{c,-1}, s_{c,-2}, s_{c,-3}) = 0$.

The multivariate model also requires a prior distribution for the factor loadings. Let α_τ denote the $n \times 1$ vector of factor loadings for $\tau_{c,t}$; we use the prior $\alpha_\tau \sim N(0, 10^2 \mathbf{1}\mathbf{1}' + 0.4^2 I_n)$, where $\mathbf{1}$ is an $n \times 1$ vectors of ones. This prior is essentially uninformative about the average value of $\alpha_{i,\tau}$ (the first term in the variance), but shrinks the factor loadings toward a common value (the second term in the factor variance). Independent priors of the same form were used for α_s and α_ε .

The empirical results in the next section are based on 60,000 Markov chain Monte Carlo (MCMC) draws from the posterior (discarding the first 10,000 draws) by using the algorithm outlined in Stock and Watson (2016), modified to incorporate the seasonal factor. Error bands are from 68-percent equal-tail credible sets. The 95-percent error bands, which are unreported, are approximately twice as wide as the reported 68-percent bands.

2. THE DATA AND ESTIMATION RESULTS

2.1 Data

There are twelve tier-two components for the euro-area HICP. These consumer spending components are organized by purpose (transportation, housing, recreation, etc.) rather than by type of product (motor vehicles, gasoline, recreational goods, etc.), which is the organizing principle used in the U.S. PCE and CPI data. Because the euro-area sectors are organized by purpose, they contain a mix of both goods and services. For example, the transportation component contains both motor vehicles (a good) and airline transport (a service). Energy is not a separate sector in the HICP tier-two categorization. Because energy prices historically behave differently from other prices, including large outliers and different seasonal patterns, we extracted the major energy components from housing (electricity, gas, liquid fuels, solid fuels, heat energy) and transportation (fuels and lubricants for personal transportation equipment) to form a separate energy component. Thus the 13 components we analyze are energy, housing excluding energy, transportation excluding fuels and lubricants for personal transportation, and the ten remaining unaffected components of the HICP. These are the thirteen sectors shown in figure 2.

The data are available monthly. We temporally aggregated the monthly price indices to quarterly averages and computed sectoral inflation rates as $\pi_{i,t} = 400 \times \ln(p_{i,t} / p_{i,t-1})$, where $p_{i,t}$ is the quarterly price index for sector i in quarter t . Data are available for all sectors as from 2001:Q1, and the first quarterly inflation value is for 2001:Q2. Our sample ends in 2018:Q1.

Spending shares for each sector are available annually. We interpolated the annual average shares to construct quarterly shares by using a random walk interpolator.³ Table 1 lists the 13 sectors, shows the average share weights over the entire sample period and over the first and second subsamples. Shares vary little over the sample period; the largest sector is food (16%) and smallest is education (1%); the energy share is 10 percent.

3. That is, we modeled the unobserved quarterly shares as a random walk, the observed annual shares as the annual average of the quarterly shares, and estimated the quarterly shares by using the Kalman smoother.

Table 1. The 12 Tier-two Sectors of the Euro-Area HICP Plus the Energy Sector

	<i>Average expenditures shares</i>		
	<i>2001-2018</i>	<i>2001-2009</i>	<i>2010-2018</i>
Food	0.16	0.16	0.15
Alcohol and tobacco	0.04	0.04	0.04
Clothing	0.07	0.07	0.06
Housing (excl. energy)	0.10	0.10	0.10
Furnishing	0.07	0.08	0.07
Healthcare	0.04	0.04	0.04
Transportation (excl. energy)	0.11	0.11	0.11
Communications	0.03	0.03	0.03
Recreation	0.10	0.10	0.09
Education	0.01	0.01	0.01
Restaurants and accommodations	0.09	0.09	0.09
Miscellaneous	0.08	0.08	0.09
Energy	0.10	0.09	0.10

Source: Authors' calculations.

Notes: Energy components of housing (electricity, gas, liquid fuels, solid fuels, heat energy) and transportation (fuels and lubricants for personal transportation equipment) were removed from those components and collected into the separate "Energy" category, given in the final row.

2.2 Results

Univariate HICP. The univariate model produces estimates of the volatilities $\sigma_{\tau,t}$, $\sigma_{s,t}$, $\sigma_{\varepsilon,t}$ and the components τ_t , s_t and ε_t . Table 2 shows the estimated values (posterior medians) and 68-percent credible sets for these variables at the beginning, middle, and end of the sample.

The estimated standard deviations of the innovations in τ , s , and ε are relatively constant over the sample period. The level of trend inflation is estimated to have fallen from 2.5 percent in 2001 to 1.5 percent in 2018. The estimated seasonal component shows that aggregate HICP inflation tends to be low in the first and third quarters and high in the second; the seasonal amplitude increased over the sample period.

Table 2. Parameter Estimates for the Univariate UCSV Model for Aggregate Inflation

Posterior medians and 68-percent equal-tail posterior credible intervals

(a) Estimated volatilities and trends from the univariate model

	2001:Q2	2009:Q4	2018:Q1
<i>Standard deviations of shocks to components</i>			
σ_τ	0.44 (0.25, 0.70)	0.55 (0.34, 0.85)	0.52 (0.32, 0.81)
σ_s	0.29 (0.18, 0.45)	0.26 (0.17, 0.40)	0.26 (0.15, 0.43)
σ_ε	0.61 (0.32, 0.90)	0.67 (0.36, 0.99)	0.65 (0.35, 0.98)
<i>Estimates of trend component</i>			
τ_t	2.54 (2.01, 3.12)	1.39 (0.96, 1.84)	1.45 (0.94, 2.03)

(b) Estimates of seasonal factors

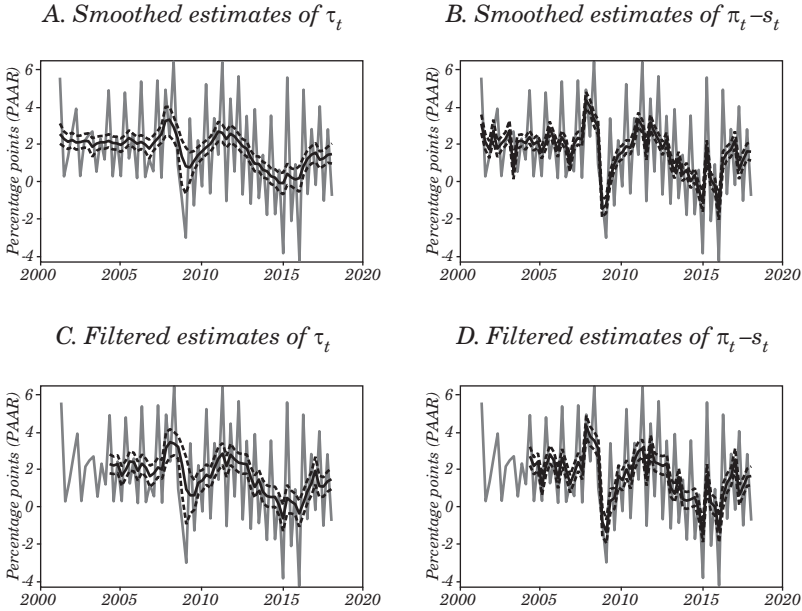
	Q1	Q2	Q3	Q4
2002	-0.06 (-0.67, 0.48)	2.16 (1.74, 2.62)	-1.76 (-2.12, -1.39)	-0.14 (-0.50, 0.22)
2009	-1.78 (-2.20, -1.34)	3.24 (2.85, 3.63)	-2.23 (-2.59, -1.86)	0.81 (0.43, 1.19)
2017	-2.42 (-2.91, -1.97)	3.45 (2.95, 3.95)	-2.22 (-2.64, -1.78)	1.19 (0.77, 1.64)

Source: Authors' calculations.

Figure 3 shows estimated values of τ_t and seasonally adjusted inflation, $\pi_t - s_t$. The upper panels show the posterior estimates based on the full sample (the smoothed estimates) and the lower panel shows estimates based on data through date t (the filtered estimates).⁴ As desired, the estimates of seasonally adjusted inflation evidently eliminate the largest seasonal swings. The 68-percent error bands for seasonally adjusted inflation are wide (1.0 percentage points at the end of the sample). The time path of trend inflation is also uncertain, but, as shown below, the estimates closely track real activity in the euro area.

4. For computational simplicity, the filtered estimates are based on the full-sample estimates of the variance parameters, and are therefore approximations the true one-sided estimates. The filtered estimates are plotted beginning in 2004 because of the diffuse prior for the $t=0$ values.

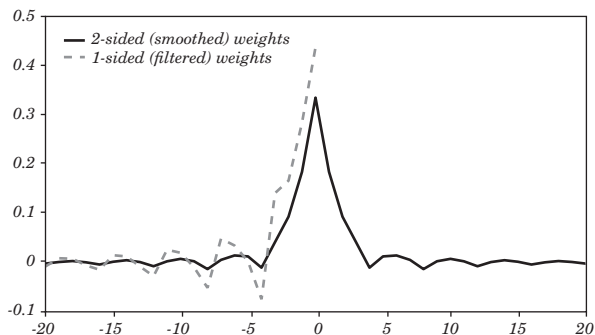
Figure 3. Smoothed and Filtered Estimates from Univariate UCSV Model for Trend (τ_t) and Seasonally Adjusted ($\pi_t - s_t$) HICP Inflation



Source: Authors' calculations.

Notes: The values shown are the posterior median and 68-percent equal-tail posterior credible intervals for the dates shown.

The estimates of τ_t and s_t are weighted averages of the π_{t+j} . For example, the full-sample posterior estimates of τ_t are given by $\tau_t|T = \sum_{j=-t+1}^{T-t} a_{t,j} \pi_{t+j}$, where the weights $a_{t,j}$ depend on the parameters $\{\sigma_{\tau,i}, \sigma_{s,i}, \sigma_{\varepsilon,i}\}_{i=1}^T$. When these parameters are time-invariant and t is not close to the beginning or end of the sample, the weights are time-invariant, that is, $a_{t,j} \approx a_j$. Figure 4 plots these weights constructed by using the sample average of $\{\sigma_{\tau,i}, \sigma_{s,i}, \sigma_{\varepsilon,i}\}_{i=1}^T$ for both the one-sided (filtered) and two-sided (smoothed) estimates of τ_t . By construction, these weights sum to unity (because the zero-frequency pseudo-spectrum of π is determined solely by variation in τ) and the figure indicates that nearly all of the weight is placed on values of $\pi_{t,j}$ for $|j| \leq 4$. These short moving-average weights are optimal because of the relatively high signal-to-noise ratio for the trend ($\sigma_{\tau} / \sigma_{\varepsilon} \approx 0.80$).

Figure 4. Weight Placed on π_{t+j} for Estimating τ_t .

Source: Authors' calculations.

Notes: The weights are computed from the Kalman filter and smoother for a univariate trend + seasonal + irregular model with constant variances computed as the average values of the UCSV model variances.

Multivariate. The univariate model implicitly applies the same time-series filter to each of the 13 sectors making up the aggregate, with the component-wise results aggregated by using share weights. Yet it is clear from figure 2 that the components follow highly heterogeneous time-series processes. For example, the clothing sector appears to be dominated by seasonality, healthcare by a few large outliers but little seasonality, energy by large irregular variation, and the housing sector by components with roughly equal variation. Thus, there plausibly is considerable variation in the UCSV parameters across the 13 components.

These visual impressions are confirmed by the posterior estimates for 13-sector model. Table 3 summarizes some key results. Consider the standard deviations of the innovations in the idiosyncratic components: the estimated values of the $\sigma_\tau/\sigma_\varepsilon$ signal-to-noise ratios range from a high of 1.8 (furnishing) to a low of 0.2 (food and energy). Seasonal signal-to-noise ratios ($\sigma_s/\sigma_\varepsilon$) vary from nearly 4 (clothing) to 0.05 (energy). Most of these standard deviations are reasonably stable over the 2001–2018 sample, but there are exceptions: for example, seasonal fluctuations have become larger in recreation, and irregular fluctuations have become smaller in alcohol and tobacco.

The multivariate model captures the covariance across sectors through the common factors τ_c , s_c , and ε_c . The estimated standard deviation of the innovations in these factors fell by roughly 40 percent from 2001 to 2018; this implies a reduction in the co-variability across the sectors. The estimated factor loadings suggest that much of the comovement arises from the common trend component, less from common seasonals, and very little from common irregular variation.

Table 3. Parameter Estimates from the 13-sector Multivariate UCSV Model

(a) Standard deviation of shocks to common components (τ_ϵ , s_ϵ , ϵ_ϵ)

	2001	2018
σ_τ	0.99 (0.91, 1.04)	0.57 (0.30, 1.00)
σ_s	0.98 (0.89, 1.03)	0.62 (0.32, 1.00)
σ_ϵ	0.99 (0.91, 1.03)	0.67 (0.36, 1.00)

(b) Sector-specific parameters

Sector	Factor loadings		
	α_τ	α_s	α_ϵ
Food	0.72 (0.41, 1.04)	0.29 (0.17, 0.42)	0.05 (-0.32, 0.42)
Alcohol and tobacco	0.06 (-0.15, 0.34)	0.03 (-0.06, 0.14)	0.05 (-0.29, 0.41)
Clothing	0.29 (0.12, 0.47)	0.12 (0.05, 0.19)	0.03 (-0.28, 0.36)
Housing (xE)	0.03 (-0.03, 0.12)	0.01 (-0.01, 0.05)	-0.01 (-0.12, 0.09)
Furnishing	0.24 (0.11, 0.45)	0.10 (0.04, 0.18)	0.02 (-0.15, 0.19)
Healthcare	0.28 (0.16, 0.45)	0.12 (0.06, 0.18)	-0.02 (-0.22, 0.21)
Transportation (xE)	0.36 (0.24, 0.53)	0.15 (0.10, 0.22)	-0.04 (-0.40, 0.33)
Communications	-0.13 (-0.38, 0.23)	-0.05 (-0.16, 0.09)	0.02 (-0.37, 0.39)
Recreation	0.35 (0.18, 0.54)	0.14 (0.07, 0.22)	0.03 (-0.51, 0.53)
Education	0.26 (0.11, 0.46)	0.11 (0.04, 0.19)	-0.01 (-0.23, 0.21)
Restaurants and accommodations	0.47 (0.32, 0.69)	0.19 (0.13, 0.28)	-0.04 (-0.41, 0.40)
Miscellaneous	0.18 (0.08, 0.30)	0.07 (0.03, 0.12)	0.01 (-0.20, 0.23)
Energy	0.39 (-0.02, 0.78)	0.16 (-0.01, 0.32)	-0.01 (-0.45, 0.42)

Source: Authors' calculations.

Notes: The values shown are the posterior median and 68-percent equal-tail posterior credible intervals for the dates shown.

Table 3. Parameter Estimates from the 13-sector Multivariate UCSV Model*(b) Sector-specific parameters (continued)*

<i>Sector</i>	<i>Standard deviation of shocks to sector-specific components</i> ($\tau_{i,t}$, $s_{i,t}$, $\varepsilon_{i,t}$)					
	σ_{τ}		σ_s		σ_{ε}	
	<i>2001</i>	<i>2018</i>	<i>2001</i>	<i>2018</i>	<i>2001</i>	<i>2018</i>
Food	0.30 (0.12, 0.94)	0.28 (0.12, 0.74)	0.17 (0.10, 0.31)	0.17 (0.09, 0.31)	1.60 (1.15, 2.16)	1.54 (1.19, 1.92)
Alcohol and tobacco	0.28 (0.14, 0.53)	0.29 (0.15, 0.50)	0.50 (0.17, 0.89)	0.40 (0.16, 0.69)	1.41 (0.69, 2.60)	0.92 (0.53, 1.34)
Clothing	0.13 (0.07, 0.23)	0.12 (0.07, 0.20)	1.56 (1.11, 2.13)	1.14 (0.79, 1.60)	0.31 (0.15, 0.51)	0.31 (0.15, 0.49)
Housing (xE)	0.13 (0.09, 0.18)	0.13 (0.09, 0.18)	0.10 (0.07, 0.14)	0.11 (0.08, 0.15)	0.13 (0.08, 0.19)	0.13 (0.08, 0.19)
Furnishing	0.21 (0.14, 0.27)	0.22 (0.17, 0.30)	0.15 (0.11, 0.20)	0.15 (0.11, 0.20)	0.12 (0.07, 0.20)	0.12 (0.07, 0.20)
Healthcare	0.12 (0.07, 0.20)	0.11 (0.07, 0.19)	0.13 (0.08, 0.21)	0.13 (0.08, 0.21)	0.77 (0.60, 1.02)	0.47 (0.32, 0.66)
Transportation (xE)	0.11 (0.07, 0.19)	0.11 (0.07, 0.20)	0.26 (0.15, 0.39)	0.28 (0.17, 0.40)	0.37 (0.21, 0.52)	0.39 (0.22, 0.55)
Communications	0.69 (0.46, 1.01)	0.69 (0.48, 1.01)	0.14 (0.08, 0.26)	0.14 (0.08, 0.25)	0.93 (0.59, 1.21)	0.87 (0.53, 1.15)
Recreation	0.17 (0.10, 0.27)	0.17 (0.10, 0.26)	0.35 (0.21, 0.54)	0.68 (0.46, 1.00)	0.28 (0.14, 0.50)	0.32 (0.15, 0.65)
Education	0.15 (0.09, 0.24)	0.15 (0.09, 0.24)	0.23 (0.11, 0.38)	0.22 (0.11, 0.39)	0.71 (0.51, 0.89)	0.77 (0.61, 0.97)
Restaurants and accommodations	0.15 (0.09, 0.23)	0.14 (0.09, 0.22)	0.23 (0.14, 0.37)	0.38 (0.22, 0.57)	0.24 (0.16, 0.35)	0.22 (0.13, 0.36)
Miscellaneous	0.16 (0.11, 0.22)	0.17 (0.12, 0.25)	0.17 (0.11, 0.25)	0.16 (0.10, 0.26)	0.22 (0.14, 0.34)	0.27 (0.18, 0.43)
Energy	1.38 (0.43, 2.63)	1.47 (0.45, 2.77)	0.40 (0.14, 1.03)	0.41 (0.15, 1.08)	7.15 (4.78, 9.28)	8.12 (6.25, 10.32)

Source: Authors' calculations.

Notes: The values shown are the posterior median and 68-percent equal-tail posterior credible intervals for the dates shown.

The multivariate model produces a rich set of results. Figures 5 and 6 illustrate a few of these results. The first four panels of figure 5 show selected results for the transportation sector: the raw data and seasonally adjusted values ($\tau_{i,t} - s_{i,t}$) are plotted in panel (a), the trend and seasonally adjusted values are plotted in panel (b), the seasonals are shown panel (c), and the estimated seasonal standard deviations, $\sigma_{i,s,t}$, are shown in panel (d). Evidently, the multivariate UCSV model


accommodates the increased dispersion in the seasonal evident in panel (c) with increases in $\sigma_{i,s,t}$ in panel (d), and provides a reasonably sharp decomposition into trend, seasonal, and irregular components in panel (b). Panels (e)–(h) show the same results for the clothing sector. From panel (e), seasonal variation in clothing price inflation is so large that it is difficult to discern any variation in the seasonally adjusted series. A change of scale in panel (f) makes the variation in the seasonally adjusted series visible and shows an outlier in 2011. Panel (g) shows that the variance of the seasonal component increases in the first half of the sample, but remains large and approximately constant, in the second half of the sample. The estimates of $\sigma_{i,s,t}$ shown in panel (h) are consistent with this changing seasonal variability. Panel (i) plots healthcare inflation and shows two large outliers. Panel (j) shows the posterior mean estimates of the outlier factor $o_{i,t}$ for healthcare, which successfully pinpoints the outliers in panel (i). Panels (k) and (l) show the analogous results for the energy sector, where outliers are also an important source of variability.

Figure 6 shows the trend estimates for each of the 13 sectors. The sectorial trends differ, but comovement is apparent, most notably during the cyclical downturns in 2008–10 and 2014–15.


As discussed above, the estimates of τ_t from the univariate model are constructed by using weighted averages of aggregate inflation, where the weights sum to unity; the one- and two-sided weights were plotted in figure 4. In the multivariate model, estimates of τ_t are also weighted averages of leads and lags of inflation for each of the sectors. When share weights and variances are time-invariant, lead-lags weights on each sector sum to that sector's share weight. For sectors with low signal-to-ratios, substantial weight is placed on distant leads and lags, but for sectors with high signal-to-ratios, most of the weight is concentrated near the contemporaneous value of $\pi_{i,t}$. figure 7 plots the sector-specific optimal weights from the 13-sector model, and compares these to the weights for the 1-sector model (which are identical for all sectors). Relative to the 13-sector weights, the 1-sector model puts too much weight on contemporaneous values of food, alcohol, and energy inflation (which have a low signal-to-noise ratio) and too little weight on sectors like furnishing and restaurants (which have relatively high signal-to-noise ratios). An implication is that the estimates of the aggregate seasonal and trend components constructed from the sectorial model and data are more precise than the estimates by using only the aggregate data.

Figure 5. Selected Results from the 13-sector UCSV Model


A. Transportation




B. Transportation




C. Transportation



D. Transportation



E. Clothing



F. Clothing






Figure 5. (continued)


G. Clothing




H. Clothing




I. Healthcare




J. Healthcare



K. Energy



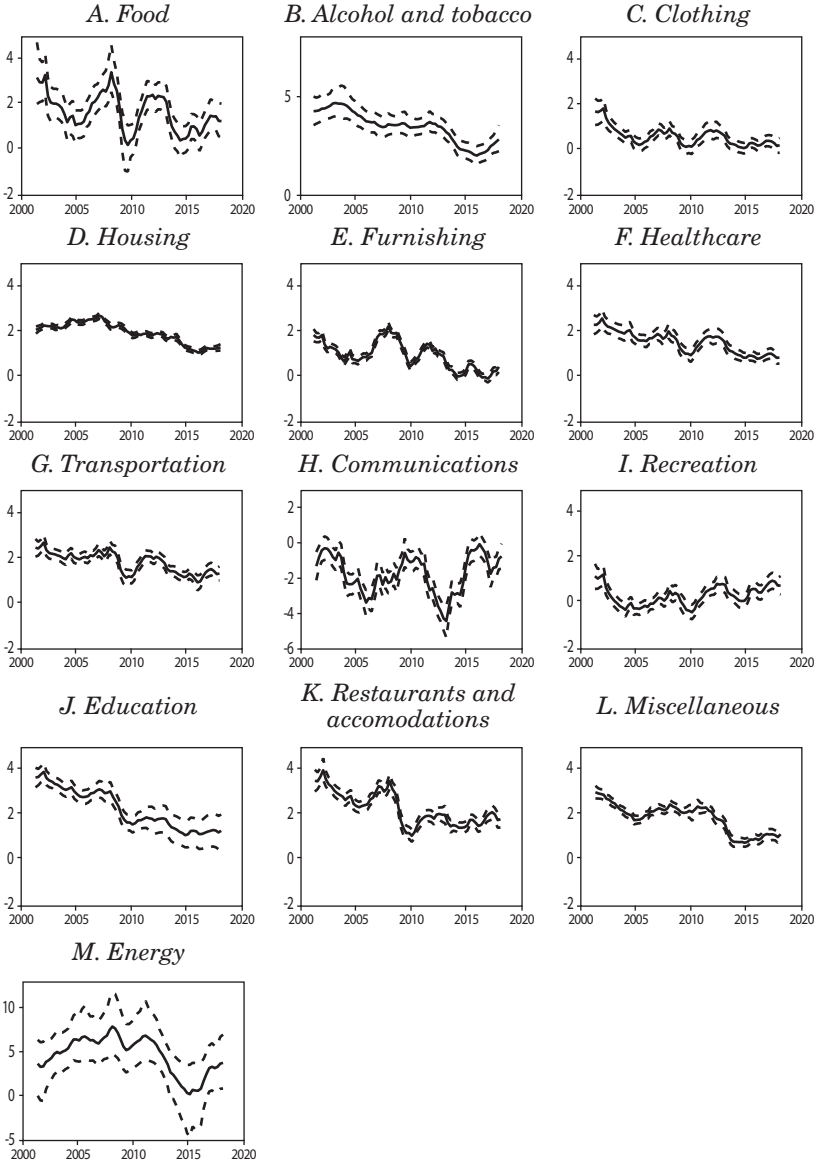
L. Energy



Source: Authors' calculations.

Notes: See text for description of the panels. Error bands are 68-percent posterior credible intervals.

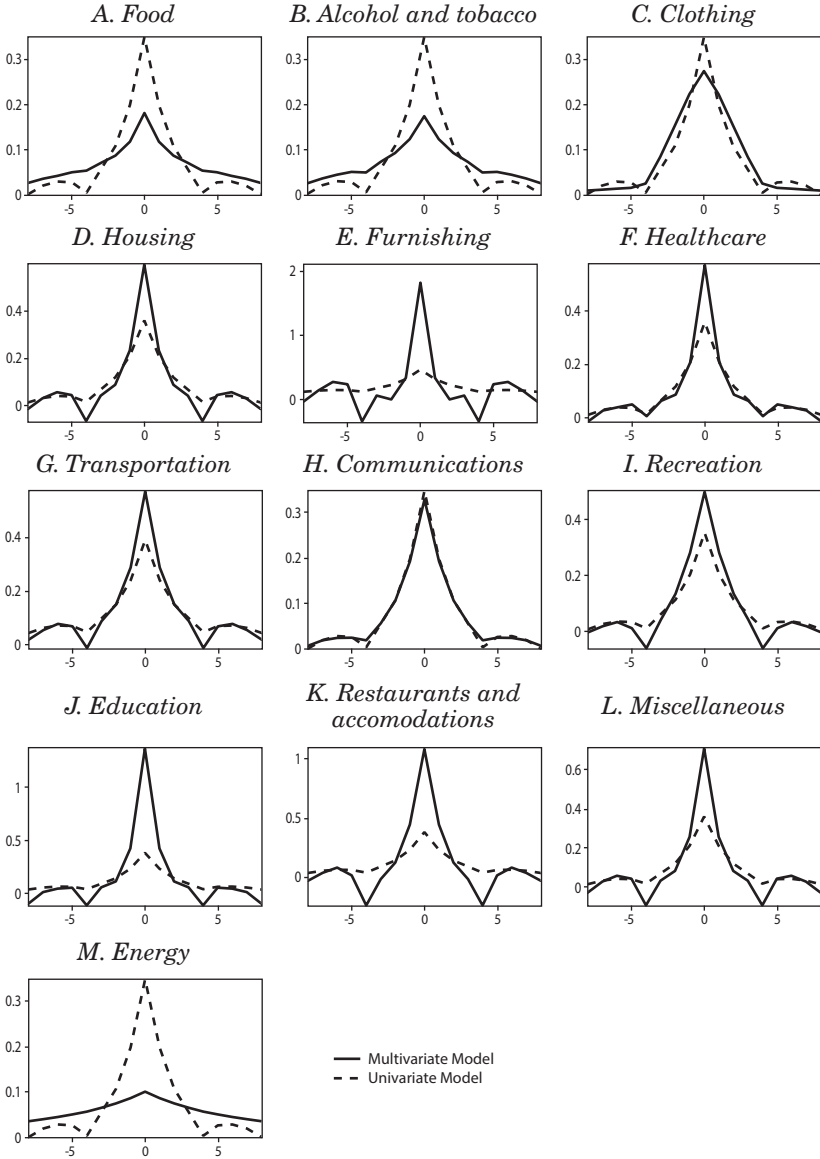
Figure 6. Trend Estimates from the 13-sector UCSV Model



Source: Authors' calculations.

Notes: These are the 12 HICP tier-two sectors, with energy excluded from the housing and transportation sectors, and shown separately as the 13th sector.

Figure 7. Weight Placed on $\pi_{i,t+j}$ for Estimating Aggregate τ_t

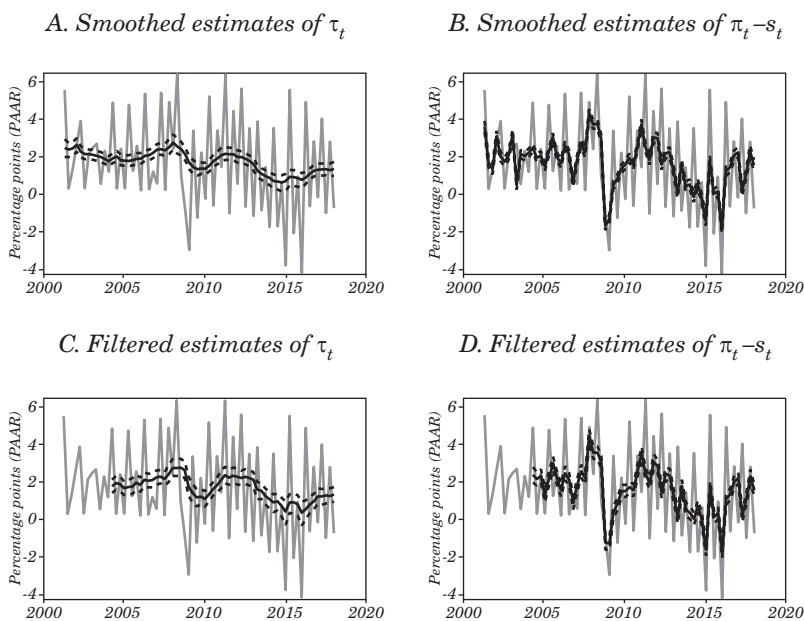


Source: Authors' calculations.

Notes: Values shown are the sum of the Kalman smoother weights on $\pi_{i,t+j}$ for estimating τ_t . The results in the first row are from the univariate model for aggregate inflation. Weights are normalized by expenditure shares, so the weights for all sectors sum to unity over all leads and lags.

This improved precision from the multivariate model can be seen in table 4 and figure 8, which show aggregate estimates constructed as share-weighted averages of the sectorial components. Comparing the error bands in table 4 with the corresponding error bands for the univariate model in table 2 shows a tightening of the bands for the multivariate model. For example, the multivariate errors bands for $\tau_{2018:Q1}$ are roughly 80 percent as wide as the univariate bands, and the multivariate error bands for $s_{2018:Q1}$ are roughly 60 percent as wide as the univariate bands.

Figure 8. Smoothed and Filtered Estimates from 13-sector Multivariate UCSV Model for Aggregate HICP Inflation



Source: Authors' calculations.

Notes: The values shown are the posterior median and 68-percent equal-tail posterior credible intervals for the dates shown. Aggregate values are computed as share-weighted averages of the sectorial values.

Table 4. Selected Results for Aggregate Inflation from the 13-sector UCSV Model*(a) Estimated trends from the multivariate model*

	2001:Q2	2009:Q4	2018:Q1
τ_t	2.43 (1.95 2.92)	1.35 (1.05 1.66)	1.32 (0.91 1.73)

(b) Estimated seasonal factors

	Q1	Q2	Q3	Q4
2002	-0.44 (-0.75 -0.14)	1.95 (1.69 2.22)	-1.48 (-1.71 -1.24)	0.07 (-0.18 0.32)
2009	-1.65 (-1.89 -1.40)	2.94 (2.69 3.17)	-2.13 (-2.37 -1.89)	0.89 (0.66 1.13)
2017	-2.42 (-2.71 -2.13)	3.87 (3.53 4.18)	-2.09 (-2.38 -1.81)	0.52 (0.19 0.88)

Source: Authors' calculations.


Notes: The values shown are the posterior median and 68-percent equal-tail posterior credible intervals for the dates shown. Aggregate values are computed as share-weighted averages of the sectorial values.

2.3 Different Levels of Disaggregation

The results presented thus far show that the 13-sector multivariate trend and seasonal estimates are more accurate than estimates that only use aggregate inflation. A natural question to ask is how much of these gains could be achieved by using a coarser disaggregation scheme, for example by using a three-sector decomposition of food, energy, and the aggregate of all of the other sectors. Using data for the U.S., Stock and Watson (2016) found that much of the gain from using a 17-sector decomposition of U.S. PCE inflation could be achieved by using this three-sector decomposition. Can similar gains be achieved from the euro-area HICP?

To answer this question, we estimated three additional multivariate UCSV models. The first is a two-sector model composed of energy and HICP excluding energy. The second is a three-sector decomposition composed of food, energy, and HICP excluding food and energy. The third is a four-sector decomposition that uses third-tier components to further decompose the non-food-and-energy HICP into goods and services. The two- and three-sector models are special cases of the 13-sector model; the four-sector model is not: as discussed above, the second-tier decomposition in the 13-sector model includes goods and services jointly in many of the sectors.

Figure 9. Estimates of Trend Inflation from the Various UCSV Models



Source: Authors' calculations.
 Notes: Values shown are full-sample posterior medians.

Table 5. Width of Credible Intervals, Final Quarter

Model	68% credible interval		90% credible interval	
	τ	$\pi - s$	τ	$\pi - s$
Univariate	1.09	1.00	1.89	1.77
2 sectors	0.96	0.81	1.63	1.39
3 sectors	0.87	0.68	1.53	1.19
4 sectors	0.82	0.67	1.48	1.19
13 sectors	0.82	0.62	1.45	1.09

Source: Authors' calculations.
 Notes: The values are the widths of 68-percent and 90-percent credible intervals for τ and $\pi - s$ for the final quarter in the sample (2018:Q1).

Figure 9 plots the estimates of trend inflation computed for each model. The estimated trends are generally similar, although there are noteworthy differences between the one- and multi-sector trends during 2009 and 2015.⁵ Table 5 summarizes the accuracy of these alternative models by showing the final quarter (2018:Q1) width of the 68-percent and 90-percent error bands for trend and seasonally adjusted inflation. Each decomposition yields marginal improvements,

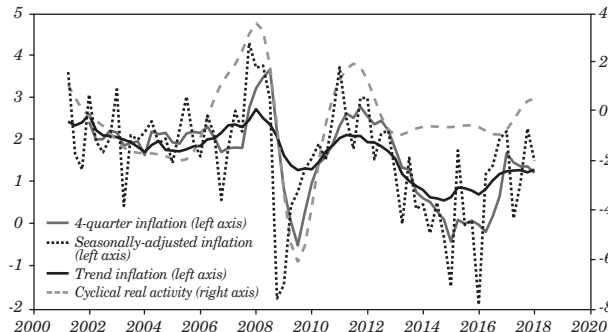
5. This paper has taken a multivariate approach to trend (and seasonal adjustment) of aggregate inflation by using sectorial inflation rates. Other series beyond sectorial inflation rates may also help identify trend inflation. Mertens (2016) provides an interesting application by using inflation expectations and nominal interest rates as additional indicators.

but much of the gain can be achieved by using the three-sector decomposition; this is consistent with the results for the U.S. reported in Stock and Watson (2016).

3. INFLATION AND REAL ACTIVITY

The multivariate estimates of trend inflation suggest a large variation in the trend level of inflation over the 2001–2018 sample period. Figure 10 shows how this variation in inflation was related to variation in real economic activity, where real activity is measured as an average of three coincident indicators for the euro area: the unemployment gap (inverted), capacity utilization, and the logarithm of industrial production, each band-pass filtered to isolate business-cycle variation (6–32 quarters) and standardized to have zero mean and unit variance. Over 2001–2018, changes in trend inflation closely mirrored changes in real activity: trend inflation increased to nearly 3 percent in early 2008 as activity was near its cyclical peak, fell by 1.5 percent during the 2009 recession, returned to 2 percent during the recovery, but fell again to under 1 percent as real activity weakened during 2013–2016.

Figure 10. Inflation and Real Activity



Source: Authors' calculations.

Notes: The trend and seasonally adjusted inflation values are the full-sample posterior medians from the 13-sector UCSV model. The cyclical activity index is the average of standardized band-pass filtered values of the unemployment gap (inverted), the capacity utilization rate, and the logarithm of industrial production, for a pass band of 6–32 quarters.

Table 6 presents correlations between the cyclical activity index and various measures of HICP inflation. The lowest correlation is with seasonally unadjusted quarterly inflation, and the highest (0.55) is with four-quarter inflation. As can be seen in figure 10, 4-quarter inflation falls sharply with economic activity in the 2009 recession, whereas trend inflation falls less, hence it has a somewhat lower correlation with the cyclical activity index. These correlations are all substantial and are consistent with a Phillips relation being present in euro-area inflation.

Table 6. Width of Credible Intervals, Final Quarter

<i>Inflation measure</i>	<i>Correlation</i>
Quarterly inflation	0.20
4-quarter inflation ($100\Delta\ln(Pt/Pt-4)$)	0.55
Seasonally adjusted HICP	0.42
Univariate trend	0.43
3-sector trend estimate	0.47
13-sector trend estimate	0.44

Source: Authors' calculations.

Notes: Seasonally adjusted HICP is the smoothed estimate of $\pi_t - s_t$ computed by using the univariate UCSV model. The three trend estimates are computed by using the UCSV model (univariate or multivariate, depending on the estimate). The cyclical activity index is the average of standardized band-pass filtered values of the unemployment gap (inverted), the capacity utilization rate, and the logarithm of industrial production, for a pass band of 6-32 quarters).

REFERENCES

- Ball, L.M. and S.G. Cecchetti. 1990. "Inflation and Uncertainty at Short and Long Horizons." *Brookings Papers on Economic Activity* 1: 215–54.
- Bryan, M.F. and S.G. Cecchetti. 1994. "Measuring Core Inflation." In *Monetary Policy*, edited by N. G. Mankiw. Chicago, IL: University of Chicago Press.
- Carter, C.K. and R. Kohn 1994. "On Gibbs Sampling for State Space Models." *Biometrika* 81: 541–53.
- Cecchetti, S., P. Hooper, B. Kasman, K. Schoenholtz, and M. Watson. 2007. "Understanding the Evolving Inflation Process." *U.S. Monetary Policy Report 1*.
- Cogley, T. and T.J. Sargent. 2015. "Measuring Price-Level Uncertainty and Instability in the United States, 1850-2012." *Review of Economics and Statistics* 97(4): 827–38.
- Cogley, T., T.J. Sargent, and P. Surico. 2015. "Price-Level Uncertainty and Instability in the United Kingdom." *Journal of Economic Dynamics and Control* 52: 1–16.
- Eckstein, O. 1981. *Core Inflation*. New York: Prentice Hall.
- Ehrmann, M., G. Ferrucci, M. Lenza, and D. O'Brien. 2018. "Measures of Underlying Inflation in the Euro Area." *ECB Economic Bulletin* No. 4.
- Gordon, R.J. (1975), "Alternative Responses of Policy to External Supply Shocks." *Brookings Papers on Economic Activity* 1: 183–206.
- Harvey, A.C. 1989. *Forecasting, Structural Time Series Models, and the Kalman Filter*. Cambridge MA: Cambridge University Press.
- Hausman, J.A. and M.W. Watson. 1985. "Errors in Variables and Seasonal Adjustment Procedures." *Journal of the American Statistical Association* 80(391): 531–40.
- Hillmer, S.C. and G.C. Tiao. 1982. "An ARIMA-Model-Based Approach to Seasonal Adjustment." *Journal of the American Statistical Association* 77: 63–70.
- Kang, K. H., C.J. Kim, and J. Morley. 2009. "Changes in U.S. Inflation Persistence." *Studies in Nonlinear Dynamics and Econometrics* 13(4): 1–21.
- Kim, C.J. and C.R. Nelson. 1999. *State-Space Models with Regime Switching: Classical and Gibbs Sampling Approaches with Applications*. Cambridge, MA: MIT Press.

- Kim, S., N. Shephard, and S. Chib. 1998. "Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models." *Review of Economic Studies* 65: 361–93.
- Maravall, A. 1995. "Unobserved Components in Economic Time Series." In *The Handbook of Applied Econometrics*, edited by M.H. Pesaran and M. Wickens. Oxford: Basil Blackwell.
- Mertens, E. 2016. "Measuring the Level and Uncertainty of Trend Inflation." *Review of Economic Statistics* 98(5): 950–67.
- Nerlove, M., D.M. Grether, and J.L. Carvalho. 1979. *Analysis of Economic Time Series*. New York, NY: Academic Press.
- Omori, Y., S. Chib, N. Shephard, and J. Nakajima. 2007. "Stochastic Volatility with Leverage: Fast and Efficient Likelihood Inference." *Journal of Econometrics* 140: 425–49.
- Stock, J.H. and M.W. Watson. 2007. "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Banking, and Credit* 39(1): 3–33. (Also see Erratum in vol. 39, No. 7, 2007)
- . 2016. "Trend and Core Inflation." *Review of Economics and Statistics* 98(4): 770–84.