Sectoral Information and the Real-Time Measurement of Trend Inflation: Evidence from Korea

Hyun Hak Kim*

August 31, 2026

Abstract

Korea's pandemic inflation cycle—a surge that carried headline inflation to a three-decade high in 2022, a disinflation, and a renewed firming through 2026—is the country's first complete cycle since the 1990s, and it settles questions that a sample ending mid-surge cannot. We estimate trend inflation with the unobserved-components model with stochastic volatility and outliers of Stock and Watson (2016b), on headline and core inflation from 1985 and on twelve CPI sectors from 1990, extended so that the frequency of outliers is itself a time-varying state estimated by Pólya–Gamma augmentation. The surge moved the trend rather than producing a string of outliers: the twelve-sector trend rises from 1.0 percent before the pandemic to 4.8 percent in 2022, retraces about half of that, and ends at 2.6 percent with a 68-percent band above the 2-percent target; led by processed food, eating out and petroleum, it leaves an above-target trend concentrated in market services. The estimated outlier frequency has fallen by half to two-thirds since the 1990s and does not rebound, though we cannot say why: the observables usually invoked for that decline fail to explain it once they are estimated inside the model. Sectoral information also buys detection time, identifying the movement two quarters before a univariate model does under vintage-consistent seasonal adjustment, cutting real-time estimation error by more than forty percent, and delivering the lowest forecast RMSE at every horizon among the ten models compared together with the only well-calibrated predictive densities. Two properties of the estimator account for that gain: it comes from heterogeneity between the trend and the transitory loading vectors and vanishes when the two are collinear, and the seasonal-adjustment procedure enters only through the covariance of what it revises at the endpoint.

Keywords: trend inflation; disaggregated prices; unobserved components; stochastic volatility; real-time estimation; density forecasting
JEL Classification: C22, C53, E31

---

* Professor, Department of Economics, Kookmin University, Seoul, Korea.
hyunhak.kim@kookmin.ac.kr. This research is supported by the Bank of Korea and by National Research Foundation of Korea grant No. 2018S1A5A8029423.

1 Introduction

For most of the decade before the pandemic, the question about Korean trend inflation was why it was so low. Headline inflation drifted below the Bank of Korea's target after 2012, and unobserved-components estimates placed the persistent component of inflation near one percent with little sign of movement. The pandemic ended that discussion and supplied something econometrically rare: a complete inflation cycle, from a surge that carried headline inflation to a three-decade high in 2022, through a disinflation, to a renewed firming at the end of our sample in 2026. This paper uses that cycle to ask three questions. Did the surge move the trend, or was it transitory? Could the movement have been detected in real time? And what is disaggregated, sector-level information worth—for detection, and for point and density forecasting of inflation?

We answer these questions with the unobserved-components model with stochastic volatility and outliers (UCSVO) of Stock and Watson (2016b), estimated on Korean data in univariate form for headline and core inflation (1985–2026) and in multivariate form on twelve sectors of the CPI (1990–2026). The model is built for exactly the identification problem the surge poses: it decomposes inflation into a drifting trend, a transitory component with time-varying volatility, and occasional outliers, so that "persistent shift," "noisy quarter," and "one-off spike" are competing explanations the data can choose among. The multivariate version adds a common factor in both components with drifting loadings, and aggregate trend inflation is the weighted sum of sector trends. We extend the model in one direction: the outlier *frequency*, a fixed parameter in Stock and Watson (2016b), is allowed to drift as a logit random walk estimated by Pólya–Gamma augmentation, so that whether an economy's propensity to generate inflation outliers has changed becomes an estimable question rather than an assumption. The data-augmentation machinery we use for this is not it-

self new—Han and Irie (2026) apply the same logit-random-walk and Pólya–Gamma construction to the probability of *zero* price changes in disaggregated CPI data—but its application to the outlier frequency of an unobserved-components model, and the economic question it answers, are.

Four findings emerge. First, the surge was a trend movement. The twelve-sector aggregate trend rises from 1.0 percent on the eve of the pandemic to 4.8 percent in 2022, retraces only about half of the rise, and stands at 2.6 percent at the end of the sample with a 68-percent band that excludes the 2-percent target. Notably, the estimated outlier frequency—which our extension turns into a time-varying object—has fallen by half to two-thirds since the 1990s and does *not* rebound during the pandemic: the surge did not look like outliers to the model, and the economy's propensity to generate inflation outliers appears to have declined secularly. We are unable to say why. Neither the shrinking weight of volatile items in the consumption basket nor the volatility of imported prices survives as an explanation once it is estimated inside the model rather than correlated with it after the fact (Section 4.4). Second, the sector composition of the trend changed. The surge was led by processed food, eating out, and petroleum, and the above-target trend remaining in 2026 is concentrated in market services; the rental-housing sector, whose trend tracked the aggregate closely before 2018, decoupled after 2022. Third, sector information buys detection time. In a pseudo-real-time exercise with quarterly re-estimation, the multivariate model's confident detection of an above-target trend comes in 2021Q4, two quarters before the univariate model's, and its real-time estimation error is more than 40 percent smaller. Fourth, the multivariate trend has the lowest forecast RMSE at every horizon among the ten models compared—including official core inflation, which is its closest rival in calm periods but fails in the surge—and its predictive *densities* are the only well-calibrated ones we evaluate, beating an AR(4) benchmark's

CRPS by 16 to 18 percent, with the practically useful asymmetry that future outlier states should be excluded from the univariate predictive distribution but retained in the multivariate one, where they carry the tails that the sector cross-section strips from the center.

These results speak to three strands of literature. The first is the statistical trend-inflation literature descending from Stock and Watson (2007) and Stock and Watson (2016b), including multi-country applications (Garnier et al., 2015), the recent precision-sampler implementation for the United States and the euro area of Dovì (2026), and, for Korea, Kim and Eo (2022), who applied the multivariate UCSVO to Korean expenditure-category data and diagnosed the early surge through 2022Q1. Relative to that work our contribution is the completed cycle and what it enables: the resolution of the transitory-or-trend question, the real-time detection analysis, and the forecast evaluation, none of which are possible on a sample that ends mid-surge. A complementary literature decomposes Korean inflation along a different dimension: Kim and Kwark (2024) use a multilevel factor model of G7 and Korean inflation to separate global from country-specific components and find that global factors have dominated since the mid-2000s, conspicuously so during the pandemic surge. Their decomposition is cross-sectional and international, ours is temporal and domestic—permanent versus transitory within one national basket—and the two fit together: Korea's exposure to volatile world prices has not diminished, yet the frequency with which that exposure produces outliers in headline inflation has fallen by half—a divergence that Section 4.4 documents and finds considerably harder to account for than it first appears.[1] The second strand asks what sectoral disaggregation contributes to understanding aggregate trend inflation: Eo et al. (2023) distinguish

[1]Earlier Korean work on trend inflation includes Yu and Kim (2003) using structural VARs and Choi (2016); a larger literature forecasts Korean inflation without extracting a trend, e.g., Kim (2018).

goods and services trends in the United States, Ahn and Luciani (2026) separate common from idiosyncratic inflation in disaggregated data, and the Federal Reserve Bank of New York's Multivariate Core Trend (Almuzara and Sbordone, 2024) has made a seventeen-sector trend decomposition an operational policy statistic. Qin and Qu (2026) make the underlying statistical point explicitly: cross-sectional aggregation attenuates the signal about latent states, so disaggregated data sharpen state inference—which is what our real-time results quantify for an actual policy decision. Our sector results give this program a small-open-economy data point in which the services-led trend re-emerges after a food-and-energy-led surge. The third strand is the post-pandemic reassessment of inflation forecasting. Watson (2026) shows for the United States that univariate time series models missed the surge in real time while models incorporating sectoral information did substantially better; our detection and density results are, to our knowledge, the first evidence on that message outside the United States, and extend it from point to density forecasting.

Beyond these findings, the paper offers a methodological caution of independent interest: the seasonal adjustment of disaggregated data is part of the estimator rather than a preliminary. Adjusting once on the full sample, or with a method whose endpoint factors are unstable, materially overstates how much a sectoral model would have known in real time (Sections 3 and 5.3). The variance decomposition between common-trend and transitory volatility, by contrast, turns out to be stable across sample lengths once the seasonal treatment is held fixed (Section 5.3).

Two of these points are given formal content. In the linear-Gaussian version of the model we show that the real-time gain from disaggregation comes from heterogeneity between the trend and the transitory loading vectors and disappears when the two are collinear, and that the effect of a seasonal-adjustment procedure on a real-time exercise runs entirely through the covariance of what that procedure revises at the

endpoint (Section 2.3). Evaluated at our estimates, the first accounts for the bulk of the measured real-time gain and the second for why the choice between adjustment methods matters as much as it does.

The remainder of the paper proceeds as follows. Section 2 presents the univariate and multivariate models, two results on what disaggregation and seasonal adjustment do to a real-time estimate, and the construction of predictive distributions, and Section 3 describes the data. Section 4 reports the estimates, first on the full sample and then in real time; Section 5 evaluates point and density forecasts and collects the robustness checks. Section 6 concludes.

2 Model

2.1 Univariate UCSVO model

We estimate trend inflation with the unobserved-components model with stochastic volatility and outliers (UCSVO) of Stock and Watson (2016b). The model decomposes inflation into a persistent component—the trend—and a serially uncorrelated transitory component, allows the volatility of both to drift, and permits occasional outliers in the transitory component:

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

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

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

$$\Delta \ln \sigma^2_{\varepsilon,t} = \gamma_\varepsilon\,\nu_{\varepsilon,t}, \tag{4}$$

$$\Delta \ln \sigma^2_{\Delta\tau,t} = \gamma_{\Delta\tau}\,\nu_{\Delta\tau,t}, \tag{5}$$

where $(\eta_{\varepsilon,t}, \eta_{\tau,t}, \nu_{\varepsilon,t}, \nu_{\Delta\tau,t})$ are *iid* $N(0, I_4)$. The trend follows a martingale whose innovation standard deviation $\sigma_{\Delta\tau,t}$ evolves as a log random walk, as does the transitory volatility $\sigma_{\varepsilon,t}$. The outlier indicator s_t equals one with probability $1 - q_t$ and, with probability q_t, is drawn from a uniform grid on $[2, 10]$, scaling up the transitory shock. This device lets the model distinguish a genuinely persistent movement from a one-off spike—a distinction at the heart of the "transitory or trend" question the pandemic surge poses.

Our one departure from Stock and Watson (2016b) is to let the outlier *frequency* itself drift. In the original model the outlier probability is a constant p; we allow

$$q_t = \frac{1}{1 + e^{-\psi_t}}, \qquad \psi_t = \psi_{t-1} + \sigma_\psi \, \zeta_t, \qquad \zeta_t \sim \textit{iid } N(0, 1), \tag{6}$$

so that the propensity of the inflation process to generate outliers is an estimated, slowly moving feature of the data rather than a fixed parameter. The motivation is substantive: whether the pandemic surge revived the outlier-generating regime of the 1990s, or whether outliers have become structurally rarer, is precisely the kind of question a constant p assumes away. Estimation adds one Gibbs block: given the outlier indicators $o_t = \mathbf{1}\{s_t > 1\}$, the log-odds path $\psi_{1:T}$ is drawn by forward-filtering backward-sampling after a Pólya–Gamma augmentation (Polson et al., 2013), and σ_ψ^2 is updated from an inverse-Gamma prior; Han and Irie (2026) use the same construction for the probability of zero price changes in disaggregated CPI data.[2]

Priors follow Stock and Watson (2016b), adapted to quarterly Korean data. The volatility-drift parameters γ_ε and $\gamma_{\Delta\tau}$ are uniform on a five-point grid over $[0,\ 0.2/\sqrt{np}]$,

[2]The Pólya–Gamma draws use the truncated infinite-sum representation with 200 terms. The prior for ψ_0 is centered on the logit of the constant-p prior mean (one outlier per two years) with unit variance, and the inverse-Gamma prior for σ_ψ^2 has mean 0.15^2 per quarter. The constant-p model is the limit $\sigma_\psi \to 0$; estimated trends are numerically indistinguishable across the two (correlations of 0.999 in the univariate model and 0.9987 in the multivariate one), so the extension refines the outlier block without disturbing the decomposition.

with np the number of periods per year. Initial states $(\tau_0, \ln\sigma_{\varepsilon,0}, \ln\sigma_{\Delta\tau,0})$ are $N(0, \kappa I_3)$ with $\kappa = 10^6$. For the outlier block, the constant-p benchmark places a Beta prior on p centred on one outlier every two years with a prior sample of ten years; in the time-varying specification that same centring is used for ψ_0. The posterior is simulated by Gibbs sampling using the mixture-of-normals approximation of Kim et al. (1998); computational details follow Stock and Watson (2016a).

Aggregate price indices, by averaging over many items, rarely display extreme outliers; sector indices do. The outlier device therefore matters most in the multivariate framework, to which we now turn.

2.2 Multivariate UCSVO model

The multivariate model adds common latent factors to both components and lets their loadings drift, extending the dynamic factor model of del Negro and Otrok (2008) with outliers and with the permanent–transitory distinction. For sector $i = 1, \ldots, n$,

$$\pi_{i,t} = \alpha_{i,\tau,t}\,\tau_{c,t} + \alpha_{i,\varepsilon,t}\,\varepsilon_{c,t} + \tau_{i,t} + \varepsilon_{i,t}, \tag{7}$$

$$\tau_{c,t} = \tau_{c,t-1} + \sigma_{\Delta\tau,c,t}\,\eta_{\tau,c,t}, \qquad \tau_{i,t} = \tau_{i,t-1} + \sigma_{\Delta\tau,i,t}\,\eta_{\tau,i,t}, \tag{8}$$

$$\varepsilon_{c,t} = \sigma_{\varepsilon,c,t}\,s_{c,t}\,\eta_{\varepsilon,c,t}, \qquad \varepsilon_{i,t} = \sigma_{\varepsilon,i,t}\,s_{i,t}\,\eta_{\varepsilon,i,t}, \tag{9}$$

$$\alpha_{i,\tau,t} = \alpha_{i,\tau,t-1} + \lambda_{i,\tau}\,\xi_{i,\tau,t}, \qquad \alpha_{i,\varepsilon,t} = \alpha_{i,\varepsilon,t-1} + \lambda_{i,\varepsilon}\,\xi_{i,\varepsilon,t}, \tag{10}$$

with each of the four volatilities evolving as an independent log random walk as in (4)–(5), and all disturbances *iid* standard normal. Subscript c denotes the common factors: $\tau_{c,t}$ is the trend movement shared across sectors and $\varepsilon_{c,t}$ the common transitory shock, while $\tau_{i,t}$ and $\varepsilon_{i,t}$ are sector-specific. Outliers are permitted in the common and each sector-specific transitory shock; the common outlier probability $q_{c,t}$ follows the logit random walk (6), while the sector-specific probabilities p_i remain constant,

keeping the time variation in outlier frequency where it is economically interpretable—in the aggregate-shock process—and the parameter count manageable.

The trend of sector i is $\tau_t^{(i)} = \alpha_{i,\tau,t}\tau_{c,t} + \tau_{i,t}$, and aggregate trend inflation weights the sector trends by (exogenous) expenditure shares $w_{i,t}$:

$$\tau_t = \sum_{i=1}^{n} w_{i,t} \left(\alpha_{i,\tau,t}\, \tau_{c,t} + \tau_{i,t}\right). \tag{11}$$

Priors for the γ's and outlier probabilities are as in the univariate model. Factor loadings have prior $N(0, \kappa_1^2\, \iota\iota' + \kappa_2^2 I_n)$ with $(\kappa_1, \kappa_2) = (20, 0.4)$ following Stock and Watson (2016b): uninformative about the average loading, restrictive about its cross-sectional spread. The loading drift parameters λ have Gamma priors following del Negro and Otrok (2008). Common-factor scales are identified by normalizing the initial variances of the common states.

2.3 What the cross-section buys, and what seasonal adjustment costs

Two of the claims this paper makes about real-time estimation—that disaggregated data identify a trend movement sooner and more precisely, and that the seasonal adjustment has to be redone at every origin—are usually supported by simulation or by example. In the linear-Gaussian version of the model they can be stated exactly. Doing so is worth the space, because the statements say *when* the claims should be expected to hold, and the conditions turn out to be checkable in our estimates. We record the two results here and prove them in the online supplement.

Consider the benchmark obtained from (7)–(10) by holding the loadings fixed at $\alpha_{i,\tau,t} = \alpha_i$ and $\alpha_{i,\varepsilon,t} = \beta_i$, holding the four volatilities constant, fixing the weights at $w_{i,t} = w_i$, and switching the outliers off ($s_t \equiv 1$). Collect the trends in $x_t =$

$(\tau_{c,t}, \tau_{1,t}, \ldots, \tau_{n,t})'$ and the sector inflation rates in π_t. The benchmark is then the linear-Gaussian state-space model

$$x_t = x_{t-1} + \xi_t, \qquad \pi_t = A x_t + \nu_t, \qquad A = [\,\alpha \;\; I_n\,], \tag{12}$$

with

$$\begin{aligned} \mathrm{Var}(\xi_t) &= Q = \mathrm{diag}\big(\sigma^2_{\Delta\tau,c},\, \sigma^2_{\Delta\tau,1}, \ldots, \sigma^2_{\Delta\tau,n}\big), \\ \mathrm{Var}(\nu_t) &= R = \sigma^2_{\varepsilon,c}\,\beta\beta' + \mathrm{diag}\big(\sigma^2_{\varepsilon,1}, \ldots, \sigma^2_{\varepsilon,n}\big), \end{aligned}$$

where $\alpha = (\alpha_1, \ldots, \alpha_n)'$ and $\beta = (\beta_1, \ldots, \beta_n)'$. Aggregate trend inflation (11) is the linear combination $\tau_t = a' x_t$ with $a = A'w = (\sum_i w_i \alpha_i,\, w_1, \ldots, w_n)'$.

It is worth being clear about what the comparison between the two models is and is not. Aggregating the observation equation gives $y_t \equiv w'\pi_t = \tau_t + w'\nu_t$, so the aggregate satisfies a univariate local-level model exactly, with trend innovation variance $q = a'Qa$ and transitory variance $h = w'Rw$. The univariate specification of Section 2.1 is therefore not a misspecified rival of the multivariate one: it is the correctly specified model for the series it observes. What separates the two is the information set, and nothing else.

Proposition 1 (Disaggregation and real-time trend uncertainty) *Let P_n^{dis} and P^{agg} denote the steady-state variances of τ_t conditional on $\pi_{1:t}$ and on $y_{1:t}$ respectively.*

(i) $P^{agg} = \frac{1}{2}\big(-q + \sqrt{q^2 + 4qh}\big)$, which is strictly increasing in both q and h, and $P_n^{dis} \le P^{agg}$.

(ii) If the n-sector panel is obtained from a finer m-sector panel by a known aggregation matrix, then $P_m^{dis} \le P_n^{dis}$.

(iii) Fix a date t and let n grow with granular weights, $\max_i w_i \to 0$ and $\sum_i w_i^2 \to 0$, and with the loadings and the idiosyncratic variances bounded. Write the cross-sectional second-moment matrix of the loading pairs as

$$M_n = n^{-1} \sum\nolimits_i (\alpha_i, \beta_i)'(\alpha_i, \beta_i).$$

If M_n converges to a nonsingular limit, then $\mathrm{Var}(\tau_t \mid \pi_{1:t}) \to 0$. If instead the loadings are collinear, $\beta = c\,\alpha$ for a scalar $c \neq 0$, then $\mathrm{Var}(\tau_t \mid \pi_{1:t})$ converges to the filtered variance of $\bar{\alpha}_w \tau_{c,t}$ in the local-level model with signal variance $q_c = \sigma^2_{\Delta\tau,c}$ and noise variance $c^2\sigma^2_{\varepsilon,c}$, which is strictly positive; here $\bar{\alpha}_w = \sum_i w_i \alpha_i$.

Part (i) is the ordering one expects, but it is worth having in this form because it makes the size of the gain a function of objects the model estimates. Part (iii) is the substantive one. It says that the value of disaggregation comes from *heterogeneity between the two loading vectors*, not from the number of series. Many sectors help because the sectors that carry the common trend are not the sectors that carry the common transitory shock, so the cross-section can be projected in two directions and the two common components separated. If the two loading vectors were proportional, no amount of disaggregation would separate them: the cross-section would reveal only $\tau_{c,t} + c\,\varepsilon_{c,t}$, and the residual uncertainty about the trend would be bounded below by an expression in the common transitory variance $\sigma^2_{\varepsilon,c}$. Disaggregation is worth doing to the extent that the two loading profiles differ.

That condition is satisfied comfortably in our estimates. The cross-sectional correlation between the posterior mean loadings $\hat{\alpha}_{i,\tau}$ and $\hat{\alpha}_{i,\varepsilon}$ is 0.49 on average over the sample and 0.52 over 2015–2026, and the 2×2 second-moment matrix M_n has a condition number of about 8—far from singular. Evaluating the benchmark at the posterior means of the loadings and volatilities, the steady-state real-time standard

error of the aggregate trend under disaggregation is 0.60, 0.54 and 0.62 times its aggregated counterpart at 2019Q4, 2022Q2 and 2026Q2: a reduction of 38 to 46 percent. The reduction we actually measure in Section 4.5, in the root mean squared error of the real-time estimate against the final one, is 47 percent. These are different objects—a steady-state filtered standard error and a realized RMSE against a smoothed benchmark—but they are the same size, which suggests the mechanism in Proposition 1 accounts for most of what the real-time exercise finds.

The second result concerns the data the filter is fed. Sector price indices are published without seasonal adjustment; write the observed series as $x_{i,t} = \pi_{i,t} + s_{i,t}$, where $s_{i,t}$ is the seasonal component and $\pi_{i,t}$ is the object the model in (7) describes. The seasonal component is not observed but estimated, and the estimate depends on the vintage: let $\hat{s}_{t|v} = E[s_t \mid \mathcal{X}_v]$ denote the estimate formed from data through v, where $\mathcal{X}_v$ is the information in the vintage-v sample and $\mathcal{X}_t \subseteq \mathcal{X}_T$. Define the real-time error $u_t = s_t - \hat{s}_{t|t}$, the full-sample error $e_t = s_t - \hat{s}_{t|T}$, and the endpoint revision $r_t = \hat{s}_{t|T} - \hat{s}_{t|t}$, with $\Omega = \mathrm{Var}(r_t)$.

Proposition 2 (Seasonal adjustment in a real-time design) *Under the projection assumption stated above,*

(i) $\mathrm{Cov}(e_t, r_t) = 0$ *and therefore* $\mathrm{Var}(u_t) = \mathrm{Var}(e_t) + \Omega$.

(ii) A design that re-adjusts at every origin faces the observation covariance $R + \mathrm{Var}(u_t)$*, while a design that adjusts once on the full sample and then filters through* t *faces* $R + \mathrm{Var}(e_t)$*. The two differ by exactly* Ω.

(iii) The filtered variance of τ_t *is nondecreasing in the observation covariance in the positive semi-definite order. Hence the full-sample design reports a filtered variance no larger than any genuine real-time procedure can attain; in the benchmark evaluated below, with* $\Omega \succ 0$*, the inequality is strict.*

The content of part (ii) is that the entire effect of the seasonal-adjustment method on the real-time problem runs through one object, the covariance of what the method revises at the endpoint. Two filters that revise their endpoint factors by the same amount are interchangeable here, however differently they are constructed; a filter that revises more is worse in real time by exactly that much, and adjusting once on the full sample hides exactly that much. This is why Section 3 measures the size of these endpoint revisions rather than comparing the two adjustment procedures on goodness of fit: Ω is the statistic the problem depends on.

It is also large enough to matter. Estimating Ω from the vintage panel described in Section 3—the covariance across sectors of the difference between what each vintage's adjustment assigned to its last four quarters and what the full-sample adjustment assigns to the same quarters—and evaluating the benchmark as before, a design that adjusts once on the full sample would understate the real-time standard error of the aggregate trend by 5 to 7 percent under X-13 and by 33 to 37 percent under a robust STL decomposition. The first number is why our vintage-consistent results in Section 4.5 are close to what a full-sample adjustment would have suggested; the second is why, as Section 5.3 reports, the same exercise run on STL-adjusted vintages loses most of the advantage that a full-sample STL adjustment appears to deliver.

Both propositions describe the linear-Gaussian benchmark, not the model we estimate: they hold the loadings and volatilities fixed and switch the outliers off. They are meant to identify mechanisms and to supply magnitudes against which the estimated results can be read, and on both counts the benchmark turns out to be informative. The estimates themselves come from the full specification of Sections 2.1 and 2.2.

2.4 Predictive distributions

The forecasting exercise in Section 5 requires predictive distributions of average inflation over the next h quarters, $\pi_{t+h}^{(h)} = (400/h)\ln(P_{t+h}/P_t)$, using only information through t. We construct them by model-consistent simulation. For each retained posterior draw at origin t, we iterate the state equations forward $j = 1, \ldots, H$: the log variances follow their random walks with the drawn γ's, the trends follow (2) (or (8)) with the simulated volatilities, transitory shocks are drawn from the outlier mixture with the drawn outlier probabilities, and—in the multivariate model—factor loadings are held at their date-t values and sector inflation is aggregated with the date-t weights.[3] Averaging the simulated $\pi_{t+1}, \ldots, \pi_{t+j}$ yields one draw of $\pi^{(j)}$; the collection across posterior draws is the predictive distribution.

We also evaluate a variant that sets $s_{t+j} = 1$ in the simulation, excluding future outliers from the predictive distribution while retaining them in estimation. The comparison speaks to whether the outlier states should be treated as forecastable dispersion or as non-forecastable noise, and Section 5.2 shows the answer differs sharply between the univariate and multivariate models.

2.5 Estimation settings

All models are estimated by MCMC with 5,000 burn-in draws and 50,000 post-burn-in draws thinned by 10, so posterior summaries use 5,000 retained draws.[4] The real-

[3]Holding the loadings fixed is a simplification, since (10) lets them drift. We checked it directly by re-running twelve origins (2018Q1–2020Q4) with the loadings simulated forward using their estimated random-walk innovations alongside the fixed-loading version. The two predictive distributions are practically identical: the ratio of predictive standard deviations is 1.000, 1.002 and 1.004 at $h = 1, 4$ and 8, the largest difference in any predictive standard deviation is 0.008 percentage points, and mean CRPS differs by at most 0.001. The loadings simply drift too slowly over eight quarters to matter for the predictive distribution.

[4]Chains with materially shorter burn-in can distort the variance decomposition between common-trend and transitory volatility on this dataset even when the trend path itself is little affected; results are stable across seeds at the settings above, and the online supplement reports inefficiency factors.

time analysis of Section 4.5 re-estimates the models on expanding samples; there we use 5,000 burn-in draws and 1,000 retained draws (thinned by 5) per origin for the multivariate model and 2,000 retained draws for the univariate models. Throughout Sections 4 and 5, trend estimates are reported from the constant-p specification and outlier-frequency estimates from the time-varying-q one, which is the only specification in which the latter exist. As noted in Section 2.1, the two deliver trends that are numerically indistinguishable, so nothing that follows turns on the choice.

3 Data

Our price data are drawn from the Bank of Korea's Economic Statistics System (ECOS). For the univariate models we use three quarterly series: the headline Consumer Price Index (CPI), the official core index excluding agricultural products and petroleum (CPI-xao), and the official core index excluding food and energy (CPI-xfe). The headline and CPI-xao series are available from 1985Q1 and CPI-xfe from 1990Q1; all series run through 2026Q2 and are expressed on the 2020=100 basis. Quarterly indices are averages of the monthly indices, and inflation is measured as the quarter-on-quarter log growth rate at an annual rate, $\pi_t = 400 \ln(P_t/P_{t-1})$. Figure 1 plots the three inflation series. As in the U.S. data studied by Stock and Watson (2016b), the two core measures track the headline series closely at low frequencies while removing much of its high-frequency variation; unlike the U.S. case, however, Korean inflation displays two complete disinflation episodes—after the 1998 crisis and after 2012—and, at the end of the sample, the surge and partial normalization that are the focus of this paper.

For the multivariate model we use the twelve-sector decomposition of the CPI by item characteristics (the Korean statistical agency's "special classification"), available from 1990Q1. The sectors, listed in Table 1, partition the headline basket exhaus-

tively: an agricultural group including livestock and fisheries products, six industrial-product groups, electricity–gas–water, and four service groups (rents, public services, eating out, and other personal services). This classification, rather than the COICOP expenditure-purpose classification, matches the grouping used by the Bank of Korea in its own analysis of underlying inflation and keeps administered-price-heavy categories (public services, electricity–gas–water) separate from market services, which proves important for the sectoral decomposition of the trend in Section 4.3. Table 1 also reports the CPI weights of the twelve sectors at three vintages. The weight shifts are themselves informative about the changing composition of Korean consumption: the farm-sector weight falls from 189 per 1,000 in 1990 to 76 currently, while the weight of personal services excluding eating out rises from 84 to 195.

Table 1: Disaggregate sectors and CPI weights

	CPI weight (per 1,000)		
Sector	1990	2017	2022–
Agricultural, livestock and fisheries	189.4	77.9	75.6
Processed food	83.4	70.1	82.7
Durable goods	90.8	64.8	73.0
Textile products	83.3	54.6	43.8
Publications	16.9	6.2	4.2
Petroleum products	19.1	46.6	46.6
Other industrial products	109.6	83.5	88.0
Electricity, gas and water	19.9	44.4	33.7
Rental for housing	118.7	93.2	99.1
Public services	125.8	145.1	120.0
Eating out	58.7	125.8	138.0
Personal services excl. eating out	84.4	187.8	195.3
Sum	1,000.0	1,000.0	1,000.0

Note: Special-classification (item-characteristics) sectors from ECOS table 901Y010. "Other industrial products" aggregates the current pharmaceutical, cosmetics, and other-industrial subgroups with fixed weights so that its coverage matches the pre-2020-basis definition. The 1990 and 2017 columns report the vintages used for those periods in estimation; the last column reports current ECOS weights, applied from 2022Q1.

Three construction details deserve mention; a replication script that downloads the raw ECOS series and reproduces every step accompanies the paper. First, the sector price indices are seasonally adjusted with X-13ARIMA-SEATS, applied to each monthly sector index before quarterly averaging; headline and core indices, whose remaining seasonality is mild, enter unadjusted. The choice of method is not incidental here. X-13 extends each series with regARIMA forecasts before estimating the seasonal factors, which stabilises those factors at the end of the sample—precisely the observations on which a real-time trend estimate depends. Measuring this directly, we re-adjusted every series using only the data available at a sequence of past dates and compared the result with the full-sample adjustment over the last four quarters of each vintage: the standard deviation of the revisions is 1.8 percentage points under X-13 against 5.8 under a robust STL decomposition. Because the real-time exercise of Section 4.5 re-adjusts the data at every origin, that difference matters, and we use X-13 throughout. Second, two sectors require splicing. The electricity–gas–water index is published as a separate special-classification category only from 2010; we extend it back to 1990 using the growth rates of the corresponding series in the discontinued pre-2020-basis classification. The other-industrial-products group is constructed as a fixed-weight composite of the current pharmaceutical (15.9), cosmetics (11.1), and other-industrial (61.0) subgroups so that its coverage matches the broader pre-2020-basis definition; before 2000, when the two subgroups are unavailable, it is likewise extended with the growth rates of the discontinued series. Third, for the farm sector we use the full agricultural–livestock–fisheries aggregate, whose weight the twelve-sector partition requires for the sectors to sum to the headline basket.[5]

Expenditure weights enter the multivariate model through the aggregation of sec-

[5]Pairing the price index of the narrower agricultural-products subgroup with the weight of the full farm aggregate would be a natural shortcut, but the distinction matters empirically: the livestock and fisheries components contribute sizable idiosyncratic episodes—the 2010–11 vegetable-price and foot-and-mouth episodes among them—that the narrower series omits.

tor trends. We use the official vintages through 2018, carry the 2017 vintage through 2021, and apply the current ECOS weights from 2022Q1 onward. Because the model aggregates with normalized shares, remaining differences between this step function and the exact annual weight revisions are second order for the aggregate trend.

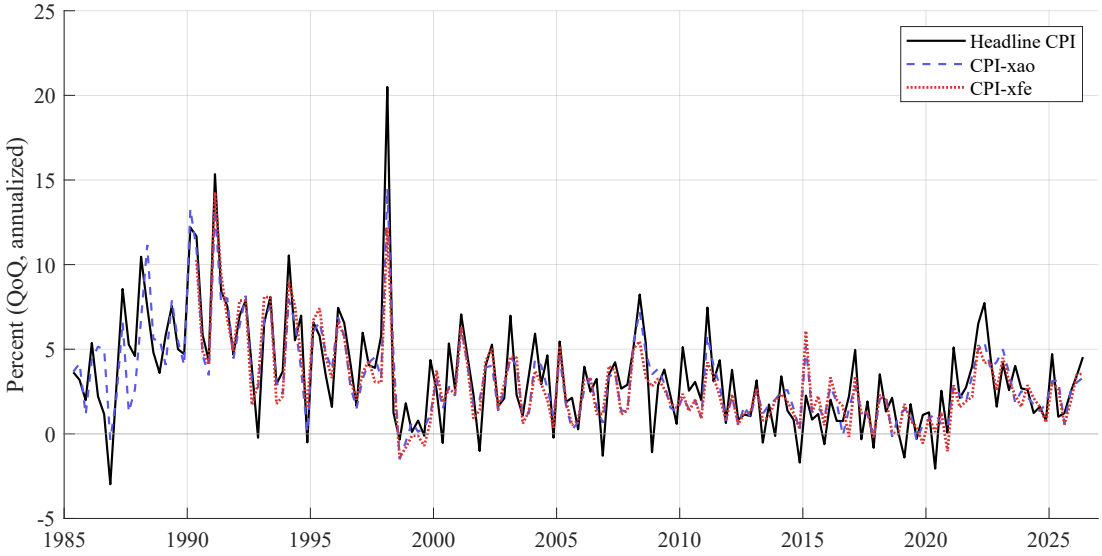


Figure 1: Headline and core inflation in Korea, 1985–2026

Note: Quarter-on-quarter log growth at annual rates. CPI-xao excludes agricultural products and petroleum; CPI-xfe excludes food and energy. Quarterly averages of monthly indices, 2020=100.

4 Trend Inflation in Korea: Full-Sample and Real-Time Estimates

4.1 Univariate estimates

We begin with the univariate UCSVO estimates for headline inflation, estimated over 1986Q1–2026Q2 (CPI-xfe, which starts later, over 1991Q1–2026Q2). The smoothed headline trend traces the familiar Korean disinflation—from about 8 percent in the late 1980s to roughly 3 percent in the 2000s and to just above 1 percent by the late 2010s—and then the episode that motivates this paper. On the eve of the pan-

demic the trend stood at 1.17 percent (2019Q4), with a 68-percent credible band of $[0.48, 1.86]$. It rose through 2021, peaked at 3.85 percent $[2.87, 4.83]$ in 2022Q2, declined to 2.36 percent by late 2024, and—rather than settling—climbed again to 3.04 percent $[2.19, 3.88]$ at the end of the sample. The core-based trends tell the same story with a lag and smaller amplitude: the CPI-xfe trend peaks at 3.28 percent and ends at 2.77 percent $[2.12, 3.43]$. By any of the three measures, the posterior places most of its end-of-sample mass above the Bank of Korea's 2-percent target.

Two features of the decomposition deserve emphasis. First, the model attributes the surge to the trend and to drifting volatility, not to outliers: the posterior mean of the outlier scale factor s_t peaks at 1.4 between 2020 and 2026 and has a median of 1.0 over that period, against 7.1 in 1998Q1. The model flags the Asian crisis unambiguously and the pandemic not at all. The pandemic surge, in other words, did not look like a spike to the model; it looked like a persistent shift. Second, the trend did not return to its pre-pandemic level. The post-2022 disinflation retraced roughly half of the rise, and the renewed increase over 2025–26 leaves the headline trend nearly two percentage points above its 2019 level.

4.2 Multivariate estimates: the aggregate trend

The twelve-sector model sharpens this picture considerably. Figure 2 plots the aggregate trend from equation (11) together with the univariate trend and realized headline inflation. The multivariate trend starts the pandemic period lower—0.97 percent $[0.70, 1.24]$ in 2019Q4—rises faster, and peaks higher, at 4.8 percent on a plateau running through 2022. It then falls to 2.02 percent $[1.74, 2.29]$ by 2024Q2 and edges back up to 2.55 percent $[2.19, 2.88]$ at the end of the sample. The 68-percent band excludes the 2-percent target at the sample end. Pooling the cross-section pays a large precision dividend: the multivariate band is roughly ± 0.30 percentage points

from 2015 onward against ±0.67 for the univariate trend, echoing the finding of Stock and Watson (2016b) for the United States that sectoral information tightens trend inference.

Figure 3 plots the volatility of the common trend and common transitory components. The common-trend volatility $\sigma_{\Delta\tau,c,t}$ declines through the 1990s and 2000s—the familiar secular stabilization—but the pandemic interrupts it: from 0.61 in 2018 it rises to 0.99 in 2022, peaking at 1.29, before easing to 0.70 by 2026, still well above its pre-pandemic level. The path over 1991–2018 is not an artefact of the longer window: re-estimating on a sample that ends before the pandemic reproduces it almost exactly (Section 5.3).

4.3 Sector trends

Where did the trend movement come from? Figure 4 plots the sector trends $\tau_t^{(i)}$ for the sectors that dominate the answer. The 2021–22 surge was led by processed food, eating out, and petroleum: at the surge peak their sector trends reach 9.3, 8.5, and 6.7 percent respectively, and their co-movement loads onto the common trend factor. This composition—food-related categories plus energy—confirms, on a completed cycle, the early diagnosis of Kim and Eo (2022), who identified food and transport-related categories as the drivers of the initial surge through 2022Q1.

The most striking sample dependence concerns rents. Over 1990–2018 the trend of the rental-housing sector tracked the aggregate trend closely, and on that sample alone it would look like the key sector for understanding Korean trend inflation. In the extended sample the rental trend collapses: it rises only briefly to about 1.6 percent in 2022 and then settles near 1.3 percent, well below the aggregate trend, through the end of the sample. The reversal coincides with the downturn of the *jeonse*-based rental market after 2022 and cautions against building structural interpretations on

any single sector's comovement.[6] As of 2026Q2 the composition has shifted again. Weighting sector trends by their expenditure shares, the four service sectors account for 1.53 of the 2.55 percentage points of aggregate trend inflation, with personal services excluding eating out alone contributing 0.83 points on a sector trend of 4.2 percent. Petroleum still carries a high sector trend of 5.7 percent but, with a weight of under 5 percent, contributes only 0.27 points. The above-target trend that remains at the end of the sample is therefore primarily a services phenomenon, whereas the surge that created it was led by food and energy.

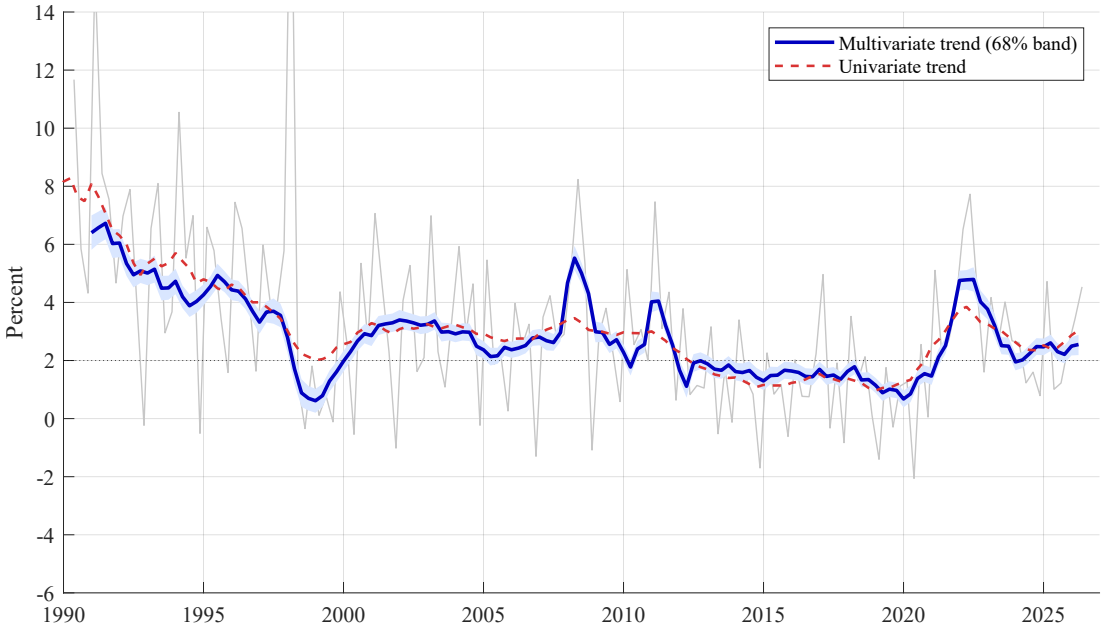


Figure 2: Trend inflation in Korea, 1990–2026Q2

Note: Posterior means; shaded area is the 68-percent band of the multivariate aggregate trend. Gray line is realized headline inflation (QoQ, annualized). Dotted line marks the 2-percent target.

[6]Rental prices in the Korean CPI mix *jeonse* (lump-sum deposit) and monthly-rent contracts, whose relative weights and pricing have shifted substantially since the 2020 revisions to the housing lease acts. A structural treatment of this margin is beyond our scope.

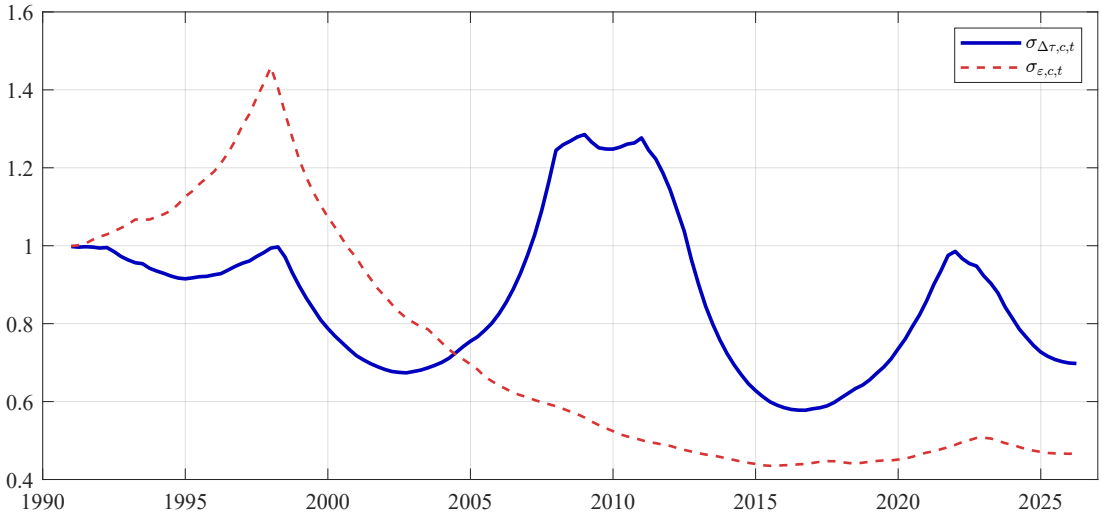


Figure 3: Common trend and transitory volatility

Note: Posterior means of the common-factor volatilities in the twelve-sector model, in normalized units.

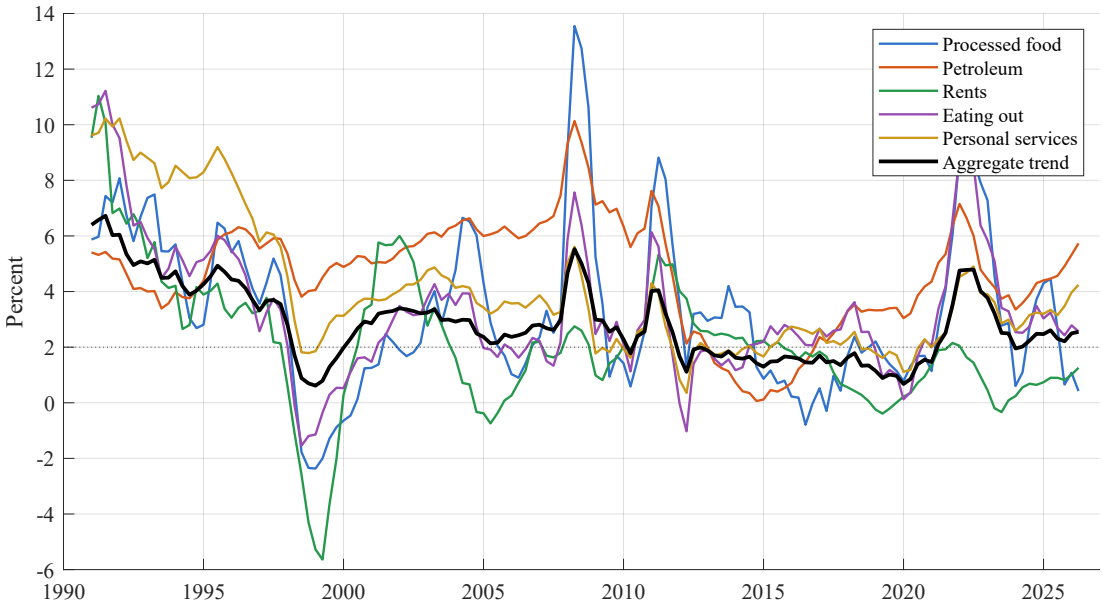


Figure 4: Sector trend inflation, selected sectors

Note: Posterior means of $\tau_t^{(i)} = \alpha_{i,\tau,t}\tau_{c,t} + \tau_{i,t}$ for selected sectors; black line is the aggregate trend.

4.4 The declining frequency of outliers

The time-varying outlier probability of equation (6) turns the informal observation that "the model flags no outliers after 2019" into an estimated structural feature. Figure 5 plots the posterior of q_t for headline inflation and of the common-component probability $q_{c,t}$ in the twelve-sector model. Both tell the same story. For headline inflation, the outlier probability falls from about 5 percent per quarter in the early 1990s to 1.5 percent from the mid-2010s onward; for the common component it falls by about two thirds, from roughly 10 percent around the 1997–98 crisis to 3.5–4 percent after 2015. Most striking is what does *not* happen: neither path rebounds during the pandemic. The 2021–22 surge, for all its violence, did not revive the outlier-generating regime of the 1990s—its shocks were large but persistent, exactly the configuration the trend and volatility states absorb. This is the model-internal counterpart of the decomposition results above, and it sharpens their interpretation: Korean inflation has not merely had fewer outliers lately; the estimated propensity to generate them has declined secularly, and a major global shock did not reverse the decline.[7]

What lies behind the decline? Table 2 compares five-year averages of q_t with three candidates. The first is composition: the share of the CPI basket in volatile items—farm products, petroleum, and utility prices—fell from 0.205 in the late 1990s to 0.156 at the end of the sample, and it tracks q_t closely, with a correlation of 0.84 across the five-year averages. The second and third are external shocks: the volatility of imported-goods price inflation and of the won/dollar rate. These correlate less well (0.45 and 0.54), and for a telling reason—imported-price volatility has *not* declined, averaging 19.1 in the late 1990s and 19.2 in 2025–26. Korea's exposure to volatile

[7]Estimated trends are numerically unchanged relative to the constant-p specification (correlations 0.999 in both models), and the paths are reproduced across independent chains. The posterior for σ_ψ remains close to its prior mean, as expected with roughly 150 binary observations; Section 5.3 shows that the shape of the q_t path is nonetheless robust to that prior.

world prices is much as it was, which is consistent with the evidence of Kim and Kwark (2024) that global factors have become more, not less, important for Korean inflation since the mid-2000s.

The composition reading is the tempting one, and we do not adopt it, because it does not survive being taken seriously. A correlation across seven five-year averages of series that both trend is weak evidence, and the natural way to do better is to put the candidates inside the model rather than alongside it. That is cheap here: because the Pólya–Gamma augmentation of Section 2.1 already makes the outlier-frequency equation conditionally Gaussian, replacing (6) by

$$\operatorname{logit} q_t = \psi_t + x_t'\beta, \qquad \psi_t = \psi_{t-1} + \sigma_\psi \zeta_t, \tag{13}$$

with x_t the standardised covariates, costs one additional Gibbs block, a weighted Gaussian regression for β.[8] Table 3 reports β for four specifications: the univariate model with ψ_t held constant, with the random walk retained, and with the won/dollar rate added on the shorter sample it is available for, and the common component of the multivariate model.

The coefficient on the basket share is indistinguishable from zero in all four: $+0.15$, -0.09, $+0.18$ and -0.08, with between 40 and 55 percent of the posterior mass below zero. This is not a matter of an unhelpful prior; accounting for the full decline in q_t would take $\beta \approx 0.56$, which a $N(0,1)$ prior accommodates comfortably. It is a matter of information. Over 162 quarters at an average q_t near 0.025 the headline series delivers a handful of outlier events, and a simulation with a known β confirms that at that event count the sampler recovers the truth but with 68-percent bands of roughly ± 0.4.

[8] Rolling headline inflation volatility, the fourth column of Table 2, is deliberately not a regressor: outliers are part of realised headline volatility, so its coefficient would not be interpretable. The two regressors used are close to orthogonal, with a correlation of 0.08.

The cross-section can supply the information the aggregate cannot. Pooling the twelve sectors and estimating $\text{logit}\, q_{i,t} = a_i + x_t'\gamma$, with a sector intercept in place of the constant p_i of the baseline, raises the number of outlier events from a handful to roughly 170, and the bands narrow accordingly. Here the basket share does get a positive coefficient, $+0.34$ with 93 percent of the mass above zero, and it would account for the whole decline: the implied cross-sector mean outlier probability falls from 0.150 to 0.063 over the sample, and holding the share fixed at its 1991 value removes that decline entirely.

Even this will not carry a causal reading, for a reason that can be shown directly. Over 1991–2026 the volatile-item share correlates -0.96 with a linear time trend, so a coefficient on it is close to a coefficient on a trend. Adding a standardised linear trend to the pooled design settles the matter: the trend takes a coefficient of -0.99 with 99 percent of its mass below zero, and the basket share's coefficient collapses from $+0.34$ to -0.33, no longer distinguishable from zero on the positive side. The share was standing in for the passage of time. Imported-price volatility, whose correlation with a trend is -0.06, behaves in the opposite way: its coefficient is unaffected by the trend and in fact sharpens, from $+0.18$ to $+0.21$ with 96 percent of the mass positive. The pooled design can therefore detect a covariate effect when one is present, which is what makes the null on the basket share informative; what no design here can do is separate a trending regressor from the trend itself.

A third check, which needs no estimation at all, runs against the composition story outright. If the aggregate frequency fell because weight moved away from outlier-prone sectors, then reweighting the estimated sector probabilities should reproduce the fall. It does the opposite: $\sum_i w_{i,t}\hat{p}_i$ *rises* by 5 percent over the sample. The sector that lost by far the most weight is farm products, down 114 per thousand, and it has the lowest estimated idiosyncratic outlier probability of the twelve, 0.048;

petroleum and utilities, which the volatile-item share also counts, gained weight. Whatever produced the decline in q_t operated on all sectors at once, not through the reallocation of weight among them.

We therefore report the decline as what the evidence supports: a robust feature of the data—reproduced across chains, priors, model variants and both seasonal-adjustment methods—whose source we are unable to identify with the observables available. Korea's inflation process has become markedly less prone to outliers over three decades, and a major global shock did not reverse that; the reason is a question this paper poses rather than answers.

Table 2: The decline in the outlier frequency and its candidate correlates

Period	q_t	Volatile-item basket share	Import-price volatility	Won/dollar volatility	Headline volatility
1995–99	0.036	0.205	19.1	20.0	3.65
2000–04	0.026	0.207	27.5	27.6	3.71
2005–09	0.020	0.185	15.5	15.7	2.20
2010–14	0.016	0.182	18.0	22.1	2.18
2015–19	0.015	0.169	13.9	13.1	1.65
2020–24	0.015	0.161	17.5	12.5	2.15
2025–26	0.016	0.156	19.2	13.7	1.96
Correlation with q_t	—	+0.84	+0.45	+0.54	+0.90

Note: Five-year averages. q_t is the posterior mean outlier probability for headline inflation. The volatile-item share is the CPI weight of farm products, petroleum, and electricity–gas–water per unit. Volatilities are rolling twenty-quarter standard deviations of import price inflation, of won/dollar depreciation, and of headline inflation, all at annual rates. Correlations are computed across the seven five-year averages and are descriptive.

Table 3: The outlier frequency and observable covariates

Specification	Covariate	Mean	68% band		$P(\cdot < 0)$
Univariate, covariates only	Volatile-item basket share	+0.15	-0.44	+0.74	0.40
	Import-price volatility	+0.32	-0.13	+0.76	0.24
Univariate, covariates + random walk	Volatile-item basket share	-0.09	-0.77	+0.59	0.55
	Import-price volatility	+0.46	-0.05	+0.98	0.19
Univariate, three covariates	Volatile-item basket share	+0.18	-0.59	+0.96	0.41
	Import-price volatility	+0.01	-0.67	+0.67	0.50
	Won/dollar volatility	+0.36	-0.34	+1.10	0.32
Multivariate, common component	Volatile-item basket share	-0.08	-0.73	+0.57	0.54
	Import-price volatility	+0.60	+0.12	+1.09	0.12
Multivariate, pooled over sectors	Volatile-item basket share	+0.34	+0.09	+0.61	0.07
	Import-price volatility	+0.18	+0.05	+0.31	0.10
adding a linear trend	Volatile-item basket share	-0.33	-0.74	+0.08	0.78
	Import-price volatility	+0.21	+0.10	+0.32	0.04
	Linear time trend	-0.99	-1.37	-0.61	0.99

Note: Posterior summaries of β in (13) for the univariate specifications and of γ in logit $q_{i,t} = a_i + x_t'\gamma$ for the pooled sector design. Covariates are standardised, so coefficients are per standard deviation, and the equations contain no constant: the level is ψ_t or a_i. The prior is $N(0, I)$. The 68-percent band is the interval between the 16.7th and 83.3rd posterior percentiles. The pooled design uses the twelve sector outlier indicators, roughly 170 events, against a handful in the aggregate specifications.

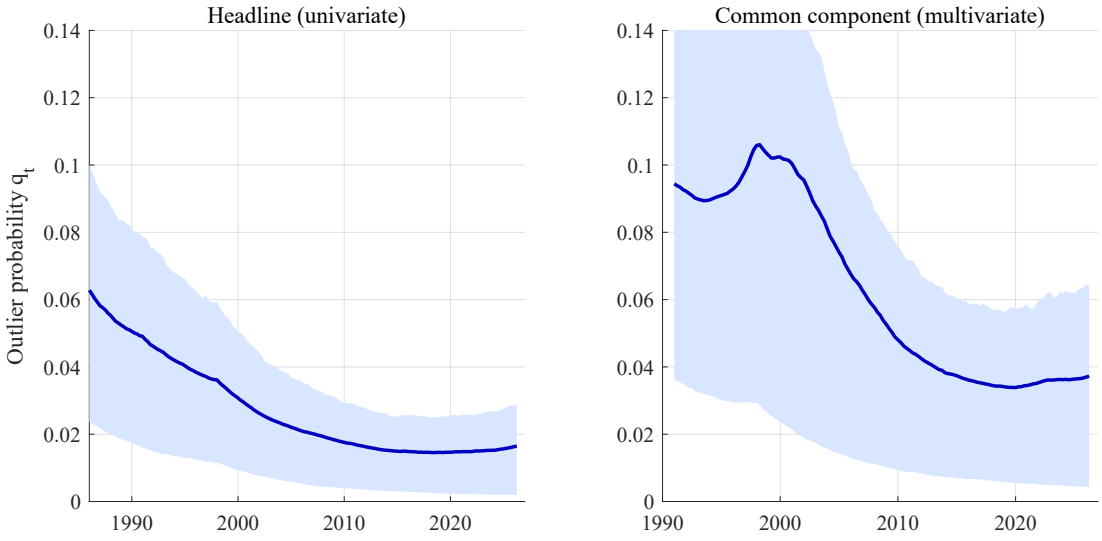


Figure 5: Time-varying outlier probability

Note: Posterior mean and 68-percent bands of q_t (left: univariate model, headline inflation) and $q_{c,t}$ (right: common transitory component of the twelve-sector model), from the logit-random-walk extension of equation (6).

4.5 Real-time estimates

The estimates reported so far use the full sample and thus answer a retrospective question. The question that mattered in 2021—is trend inflation rising?—had to be answered with the data available at the time. To replicate that problem we re-estimate both models on expanding samples, with quarterly origins from 2000Q1 through 2026Q2, and record at each origin the filtered estimate of the trend at the sample endpoint. This is a pseudo-real-time design: it uses the current data vintage rather than historical vintages, but CPI data are essentially unrevised apart from base-year relinking, so the informational constraint that matters—not seeing the future—is the binding one. One part of the data preparation is genuinely forward-looking, however, and we treat it accordingly: the seasonal factors. Rather than adjusting the sector series once on the full sample, we re-run X-13ARIMA-SEATS at every origin using only the months available then, so that each origin sees the seasonally adjusted data a forecaster would actually have had. Weight vintages are likewise applied as of their availability.

Figure 6 plots the real-time trend estimates against the final (smoothed) multivariate estimate, and Table 4 summarizes the detection timeline. Three results stand out. First, *sector information buys about two quarters of detection time.* The real-time multivariate mean crosses 2 percent at origin 2021Q2 and its 16.7th percentile—a natural threshold for a confident call—crosses at 2021Q4. The univariate model reaches the same milestones at 2021Q4 and 2022Q2. By the confident-call criterion, a policymaker relying on headline inflation alone learns in mid-2022 what the sector cross-section revealed in late 2021. It is notable that the Bank of Korea's first post-pandemic rate increase, in August 2021, preceded even the multivariate model's confident-detection date; judged against these model-based signals, the tightening cycle does not look late.

Second, the multivariate real-time estimate is much closer to the final estimate throughout: its RMSE against the smoothed trend over 2000–2024 is 0.398 percentage points against 0.757 for the univariate model—a reduction of nearly one half—and its real-time bands are less than half as wide. Both filtered estimates overshoot at the peak, as the filtered-versus-smoothed asymmetry implies: the multivariate estimate reaches 5.18 in real time at 2022Q3 against a final estimate of 4.79, and the univariate one reaches 4.94 against a final 3.85. The difference is that the multivariate estimate tracks the final path closely on both legs of the cycle, whereas the univariate one also lags the run-up.

Third, the disinflation is confirmed earlier as well. The multivariate model's 83.3rd percentile falls back below 3 percent at origin 2024Q1, three quarters before the univariate model reaches the same point. The cross-section helps at both turning points, not only on the way up—which matters for a central bank deciding when to stop tightening as much as when to start.

Finally, the end of the sample is itself a real-time reading. As of 2026Q2 the multivariate real-time trend stands at 2.51 percent with a 68-percent band of $[2.14, 2.85]$: above target, though history suggests such end-of-sample readings will be revised as the future arrives.

Table 4: Real-time detection of the trend rise

	Multivariate	Univariate
Real-time mean first exceeds 2%	2021Q2	2021Q4
16.7th percentile first exceeds 2%	2021Q4	2022Q2
Real-time peak (final estimate at peak)	5.18 (4.79)	4.94 (3.85)
83.3rd percentile back below 3%	2024Q1	2024Q4
RMSE, real-time vs. final, 2000–2024	0.398	0.757
Real-time trend, 2026Q2 [16.7, 83.3]	2.51 [2.14, 2.85]	3.01 [2.19, 3.87]

Note: Expanding-window estimation with quarterly origins 2000Q1–2026Q2, with the sector series re-adjusted at every origin using only the data available then; "real-time" refers to the filtered estimate of the endpoint trend at each origin. RMSE compares real-time and full-sample (smoothed) posterior means over origins through 2024Q4.

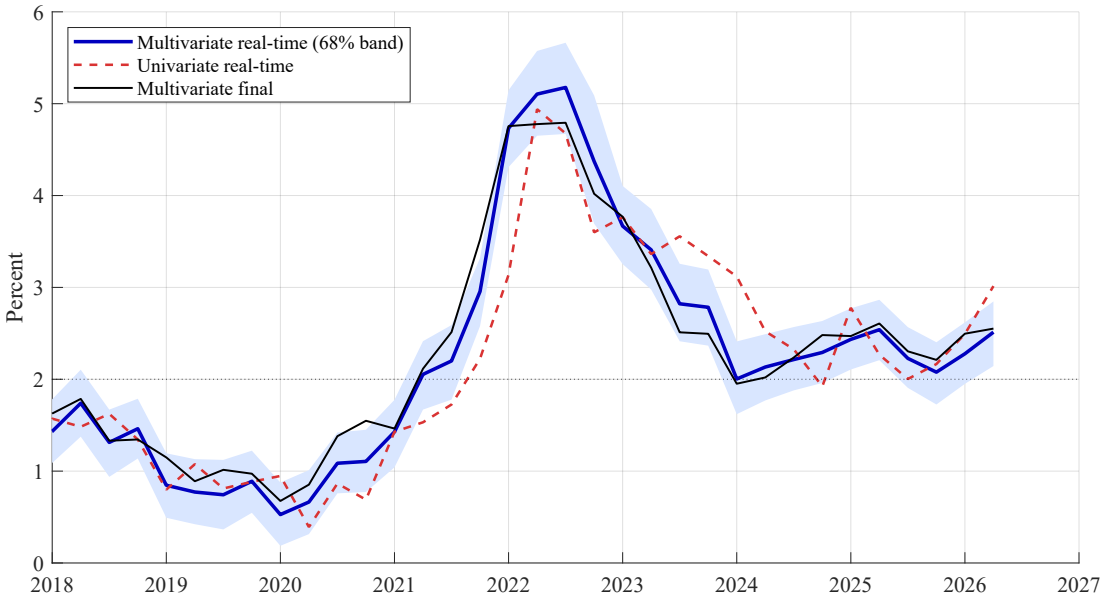


Figure 6: Real-time versus final trend estimates, 2018–2026

Note: Real-time series plot, at each origin, the filtered endpoint estimate using data through that origin; the black line is the full-sample smoothed multivariate trend. Shaded area is the 68-percent real-time band of the multivariate model.

5 Forecasting Performance and Robustness

If the trend estimates of Section 4.5 are useful in real time, they should forecast. The natural reading of the model is that the current trend estimate *is* the forecast

of average inflation at any horizon, since the trend is a martingale and the transitory component has mean zero. We evaluate this reading against standard benchmarks using the same expanding-window design as before: at each origin t from 2000Q1, we forecast average headline inflation over the next h quarters, $\pi_{t+h}^{(h)}$, for $h = 1, 4, 8$.

5.1 Point forecasts

Table 5 reports root mean squared errors for ten forecasts. Five are standard time series benchmarks: the real-time multivariate and univariate trends, the random-walk forecast of Atkeson and Ohanian (2001) (the average of the last four quarters of inflation), a direct-projection AR(4) re-estimated at each origin, and the expanding sample mean. Two are the official core measures, entered the way policy institutions use them—the current four-quarter average of core inflation as a forecast of future headline inflation—which is the comparison at the centre of Stock and Watson (2016b). The last three use the same twelve-sector information as the multivariate model but exploit it in conventional ways: a factor-augmented regression on the first three principal components of the sector inflation panel, an elastic net, and a random forest, each estimated recursively on current and lagged sector inflations. If the value of the sector cross-section were simply predictive content, a flexible learner given the same predictors should capture it.

Three messages emerge. First, over the full evaluation sample the multivariate trend has the lowest RMSE at every horizon. The margin over the naive benchmarks is statistically significant (Diebold–Mariano $p \leq 0.01$ against both the univariate trend and the Atkeson–Ohanian benchmark at every horizon), as is the margin over the two machine-learning forecasts at $h = 4$ (1.20 against 1.60 for the elastic net and 1.65 for the random forest, $p = 0.02$ for both). The state-space structure extracts something from the cross-section that flexible regressions on the same data do not:

the sector trends, not the sector inflation rates, are what forecasts. The margin is widest in the disinflation subsample, where the multivariate trend's RMSE at $h = 4$ is 0.26 against 0.38 for the nearest competitor.

Second, core inflation is a much tougher benchmark than the naive ones, and the comparison with it is instructive rather than decisive. Excluding food and energy is itself an effective way of estimating the persistent component, and in the calm two decades it works very well: at $h = 4$ and $h = 8$ the CPI-xfe benchmark attains RMSEs of 1.00 and 0.83 against 1.05 and 0.94 for the multivariate trend. What separates the two is the surge, where core deteriorates sharply (2.27 at $h = 4$ against 1.88) because the food and energy prices it discards were precisely where the persistent movement began. Over the full sample the multivariate trend is therefore ahead (1.20 against 1.28 at $h = 4$; 1.22 against 1.24 at $h = 8$), though the differences are not statistically significant ($p = 0.24$ and 0.60). The honest summary is that a model-based trend and a well-chosen exclusion index perform comparably in normal times, and that the model's advantage appears when the composition of inflation changes—which is when a central bank most needs it to.

Third, a caution. In the surge subsample at the longer horizons, the mean-reverting benchmarks—the AR(4) and even the unconditional sample mean—beat both trend forecasts, in the sample-mean case by a wide margin. The reason is instructive rather than embarrassing: the trend genuinely rose, but the *level* of inflation reverted toward it from above faster than the trend itself moved, so a forecast anchored on a lower long-run level looked good over 2022–23 even though it embodied the wrong reading of the trend. Detection and point forecasting are different loss functions; a model can be right about the trend and still lose an RMSE race at the peak of a surge.

5.2 Density forecasts

Point RMSE understates what the model produces: a full predictive distribution. Table 6 evaluates the simulated densities of Section 2.4 with the continuous ranked probability score (CRPS) and PIT-based coverage. The benchmarks are densities built around the AR(4) and Atkeson–Ohanian point forecasts. Because a Gaussian benchmark could be accused of being a straw man, we report the AR(4) both with Gaussian errors and with a recursive residual bootstrap that inherits whatever skewness and excess kurtosis the in-sample residuals display.

The multivariate densities win, and they win in the dimension that matters for a policy institution: calibration.[9] At $h = 4$ and $h = 8$ their central 68-percent interval covers the realization 71 and 70 percent of the time—essentially exact—whereas the univariate densities are markedly over-dispersed (coverage 0.82–0.86), a conservatism that survives even when future outliers are excluded from the simulation; block-bootstrap tests of PIT uniformity do not reject for either model, reflecting their limited power at these sample sizes, so we weight the coverage evidence. At $h = 1$ the multivariate density is somewhat too narrow (coverage 0.61), the one horizon at which the sector cross-section over-sharpens the center. On CRPS, the multivariate densities improve on the AR(4) benchmark by 18 percent at $h = 4$ and 16 percent at $h = 8$ over the full sample, and dominate every subsample except—consistently with the point results—the surge at longer horizons. Bootstrapping the benchmark's errors rather than assuming normality does not overturn this: the bootstrap AR(4) improves on its Gaussian counterpart at $h = 1$ but stays behind the multivariate density (CRPS 1.15 against 1.10), and it is no better at $h = 4$ and worse at $h = 8$ (0.82 and 0.87 against 0.67 and 0.69). Formal tests locate the advantage precisely:

[9] Density results are computed under the time-varying-outlier-frequency specification of Section 2.1, with q_t (and $q_{c,t}$) simulated forward along their random walks. Results under the constant-p specification are nearly identical (CRPS differences in the second decimal), as are the trend point forecasts, whose posterior means have correlation 0.999 across the two specifications.

on CRPS differentials the multivariate model is significantly better than the AR(4) at $h = 1$ over the full sample ($p = 0.010$, and $p = 0.030$ on the Amisano–Giacomini log-score comparison) and in the calm decades at every horizon ($p = 0.036$, 0.003 and 0.021), while the AR(4) is significantly better in the post-2020 window at the longer horizons ($p = 0.076$ at $h = 4$ and 0.004 at $h = 8$). The two regimes offset, so the pooled full-sample difference reaches significance only at the shortest horizon.

The outlier device plays opposite roles in the two models' predictive distributions. In the univariate model, excluding future outliers from the simulation improves both CRPS and calibration: with only one series, the outlier states mainly widen an already-too-wide density. In the multivariate model the prescription reverses: excluding outliers leaves the center of the density—sharpened by the cross-section—too confident, so that coverage falls short (central-68 coverage of 0.52–0.66) and the log score deteriorates badly in the surge, when realizations land in the tails. Sector information narrows the middle of the predictive distribution; someone must still hold the tails, and in this model that job belongs to the outlier states. We read this as a practical guideline for users of UCSVO-class models: report the multivariate density with its outlier component intact.

Table 5: Out-of-sample point forecast accuracy (RMSE)

	Multi	Uni	AO	AR4	Mean	xFE	xAO	FAAR	ENet	RF	n
Horizon $h = 1$											
Full, 2001–2026	**1.98**	2.21	2.09	2.04	2.58	2.08	2.10	2.06	2.12	2.08	101
Calm, 2001–2019	**2.01**	2.22	2.08	2.04	2.73	2.04	2.11	2.11	2.18	2.09	76
Surge, 2020–2023	**2.06**	2.42	2.34	2.25	2.27	2.48	2.33	2.18	2.15	2.25	16
Disinflation, 2024–2026	1.47	1.66	1.67	1.55	1.72	1.54	1.51	**1.38**	1.42	1.72	9
DM p (vs. Multi, full)	—	0.00	0.01	0.27	0.00	0.11	0.01	0.18	0.08	0.15	
Horizon $h = 4$											
Full, 2001–2026	**1.20**	1.39	1.32	1.40	1.97	1.28	1.34	1.38	1.60	1.65	98
Calm, 2001–2019	1.05	1.23	1.16	1.43	2.14	**1.00**	1.17	1.38	1.63	1.66	76
Surge, 2020–2023	1.88	2.13	2.03	1.47	**1.26**	2.27	2.12	1.62	1.69	1.87	16
Disinflation, 2024–2026	**0.26**	0.66	0.66	0.58	1.08	0.48	0.38	0.42	0.55	0.46	6
DM p (vs. Multi, full)	—	0.01	0.01	0.18	0.00	0.24	0.01	0.25	0.02	0.02	
Horizon $h = 8$											
Full, 2001–2026	**1.22**	1.40	1.42	1.38	1.90	1.24	1.31	1.35	1.46	1.43	94
Calm, 2001–2019	0.94	1.15	1.14	1.37	2.07	**0.83**	1.01	1.33	1.40	1.42	76
Surge, 2020–2023	2.13	2.25	2.35	1.50	**0.92**	2.40	2.30	1.54	1.72	1.56	16
DM p (vs. Multi, full)	—	0.01	0.00	0.24	0.01	0.60	0.07	0.33	0.13	0.27	

Note: Forecast target is average headline inflation over the next h quarters; forecasts use real-time (expanding-window) information with origins from 2000Q1. xFE and xAO are the current four-quarter averages of the two official core inflation measures. FAAR is a direct projection on the first three principal components of the twelve-sector inflation panel; ENet and RF are an elastic net and a random forest on current and lagged sector inflations, all estimated recursively. Bold marks the lowest RMSE in each row. DM rows report two-sided Diebold–Mariano p-values against the multivariate trend forecast on the full sample, with Newey–West standard errors.

Table 6: Density forecast accuracy and calibration

	Uni	Uni (no-out)	Multi	Multi (no-out)	AR4	AR4 boot	AO
CRPS, $h = 4$							
Full, 2001–2026	0.82	0.81	**0.67**	**0.67**	0.82	0.82	0.81
Calm, 2001–2019	0.75	0.73	0.59	**0.58**	0.84	0.85	0.75
Surge, 2020–2023	1.30	1.31	1.18	1.24	**0.84**	**0.85**	1.21
Disinflation, 2024–2026	0.49	0.48	0.33	**0.27**	0.45	0.40	0.54
PIT central-68 coverage	0.86	0.86	0.71	0.61	—	—	—
PIT uniformity, bootstrap p	0.92	—	0.69	—	—	—	—
CRPS diff. vs. AR4, HAC p	0.68	—	0.07	—	—	—	—
AG log-score vs. AR4, p	0.17	—	0.13	—	—	—	—
CRPS, $h = 8$							
Full, 2001–2026	0.85	0.83	**0.69**	**0.69**	0.82	0.87	0.84
Calm, 2001–2019	0.73	0.71	0.54	**0.53**	0.81	0.88	0.72
Surge, 2020–2023	1.45	1.46	1.43	1.50	0.90	**0.89**	1.43
PIT central-68 coverage	0.85	0.82	0.70	0.66	—	—	—
PIT uniformity, bootstrap p	0.67	—	0.93	—	—	—	—
CRPS diff. vs. AR4, HAC p	0.81	—	0.24	—	—	—	—
AG log-score vs. AR4, p	0.41	—	0.56	—	—	—	—

Note: CRPS computed from 2,000 predictive draws (model densities and the bootstrap benchmark) or analytically (Gaussian benchmarks); bold marks the best model in each row. "AR4 boot" replaces the Gaussian error density with a recursive residual bootstrap. "No-out" excludes future outlier states from the predictive simulation. PIT central-68 coverage is the fraction of realizations inside the central 68-percent predictive interval over the full sample (target 0.68).

5.3 Robustness

The results above rest on a number of choices—how finely the basket is split, how long the sample is, how the sector series are seasonally adjusted, how the sampler is configured, and how the new outlier-frequency state is given a prior. We take each in turn.

Coarser aggregation. Re-estimating the multivariate model on a seven-sector aggregation (farm products; all industrial products; electricity–gas–water; rents; public services; eating out; other personal services) yields an aggregate trend with correlation 0.987 to the twelve-sector baseline and the same cycle: 0.96 percent in 2019Q4,

a peak of 4.8 percent in 2022, and 2.0 percent in 2024Q2. Two differences are instructive. The coarser model reaches essentially the same peak (4.76 against 4.79) but comes off it faster, so its 2022 average is lower (4.36 against 4.59): averaging within the broad industrial group smooths some of the food-and-energy co-movement that the twelve-sector model assigns to the common trend. And it ends the sample higher (2.84 against 2.55 percent), since it cannot separate the sectors whose trends have normalised from those that have not. The common-trend volatility paths have correlation 0.988.

Sample length. Because the surge is the dominant event in the sample, one might worry that it reshapes the decomposition of the earlier years. It does not. Re-estimating the twelve-sector model on 1991Q1–2018Q2, a window that ends before the pandemic, yields an aggregate trend with correlation 0.999 to the full-sample estimate over the common window and a common-trend volatility path with correlation 0.985; the 2010 value of that volatility, the point at which the two windows might be expected to differ most, is 1.16 on the short sample against 1.25 on the full one. The short-sample estimate is lower at its own endpoint (0.46 in 2018Q2 against 0.62), which is the usual filtered-versus-smoothed difference rather than a change in regime. The decomposition is a property of the data, not of where the sample ends.

Seasonal adjustment and the real-time exercise. Section 3 reported that vintage and full-sample seasonal factors differ far less under X-13 than under a robust STL decomposition (revision standard deviations of 1.8 against 5.8 percentage points over the last four quarters of each vintage). That gap is not a curiosity: it determines how much of a sectoral trend model's apparent real-time performance is real. Re-running the real-time exercise on STL-adjusted vintages shrinks the detection advantage over the univariate model from two quarters to one and erases the

forecasting advantage entirely, whereas with X-13—which extends each series with regARIMA forecasts before estimating the seasonal factors—the results of Sections 4.5 and 5 are those reported. Users of sectoral trend indicators should therefore treat the endpoint behaviour of the adjustment procedure as part of the estimator, not as a data-preparation detail.

Sampler settings. At the settings of Section 2.5, results are stable across seeds. With materially shorter burn-in the variance decomposition between common-trend and transitory volatility can be distorted even though the trend path changes little, which is why we run 5,000 burn-in draws throughout. The online supplement reports inefficiency factors for the specifications we estimate.

Does the data prefer the time-varying specification? Because the extension adds a state, it should be asked whether the data want it. Following the predictive-Bayes-factor practice of Carriero et al. (2024), Table 7 accumulates the difference in one-step-ahead log predictive likelihoods between the time-varying and constant-p specifications across the real-time origins. The evidence favors the time-varying model in both cases, more clearly in the multivariate one: the cumulative log Bayes factor is +1.13 for the univariate model and +2.50 for the multivariate, where twice the log factor of 5.0 sits in the range conventionally read as positive but not strong evidence. The mean increment is positive in both, though not distinguishable from zero at conventional levels (HAC t of 1.80 and 1.63). What the split by period adds is that the multivariate gain is concentrated where the economics would put it: +1.56 of the +2.50 accumulates from 2020 onward, against +0.63 over the two calm decades. The specification earns its extra state when outlier behavior is actually in question.

We do not want to lean on this more than it will bear. A log factor of 2.5 over 105 origins is support, not proof, and the extension is not needed to forecast better:

the point forecasts of Section 5 are essentially the same either way. Its value is that it makes the outlier frequency an object the data can speak about at all—and, as Section 4.4 shows, they say something economically interesting. That the substantive results of Sections 4 and 5 do not hinge on the specification choice is itself reassuring.

Table 7: Predictive model comparison: time-varying versus constant outlier frequency

	Cumulative log predictive Bayes factor			Mean increment	
	Full	Calm	2020–	Value	HAC t
Univariate	+1.13	+0.72	+0.57	+0.0107	+1.80
Multivariate (12 sectors)	+2.50	+0.63	+1.56	+0.0238	+1.63

Note: Cumulative differences in one-step-ahead log predictive likelihoods (time-varying minus constant p) over real-time origins 2000Q1–2026Q1; positive values favor the time-varying specification. "Calm" covers origins through 2019Q4. The HAC t-statistic tests a zero mean increment.

The prior on the outlier-frequency drift. Our extension introduces one new parameter, the innovation standard deviation σ_ψ of the log-odds random walk (6). With roughly 150 quarterly observations of a rare binary event, σ_ψ is only weakly identified: its posterior tracks whatever prior we impose, giving posterior means of 0.050, 0.148, 0.293, and 0.193 under priors centered at 0.05, 0.15, 0.30, and 0.20 (Table 8). The question that matters for Section 4.4 is therefore not whether σ_ψ is pinned down—it is not—but whether the *shape* of the estimated q_t path survives changes in the prior. It does. Figure 7 shows that the four paths are the same object viewed at different degrees of smoothing: correlations with the baseline path range from 0.993 to 0.999, and the estimated trend is untouched (correlations of 0.9999 or better). The prior governs amplitude, not shape. Every specification puts q_t at its highest in the early 1990s and much lower from the mid-2010s onward, with the decline running from 0.035 to 0.029 under the tightest prior and from 0.047 to 0.008 under the loosest. And in none of them does the pandemic restore the earlier regime:

comparing 2015–19 with post-2020 averages, q_t is flat under the tight prior, 4 percent higher under the baseline and diffuse priors, and 18 percent higher under the loose prior—but that largest relative change is a rise from 0.7 to 0.8 percent per quarter, an order of magnitude below the early-1990s level.[10]

Table 8: Sensitivity of the outlier frequency to the prior on σ_ψ

	σ_ψ		Posterior mean q_t			Correlation with baseline	
Prior for σ_ψ^2	prior	posterior	1990s	2015–19	2020–	q_t	τ_t
IG(10, 0.0225)	0.05	0.050	0.035	0.029	0.029	0.993	0.9999
IG(10, 0.20) *(baseline)*	0.15	0.148	0.041	0.015	0.015	1.000	1.0000
IG(10, 0.81)	0.30	0.293	0.047	0.007	0.008	0.996	0.9999
IG(3, 0.08)	0.20	0.193	0.043	0.012	0.013	0.999	0.9999

Note: Univariate model, headline inflation, 1986Q1–2026Q2, 50,000 retained draws per specification. "Prior" reports the implied prior mean of σ_ψ; q_t columns report posterior means averaged over the indicated periods. Correlations are computed over the full estimation sample against the baseline specification.

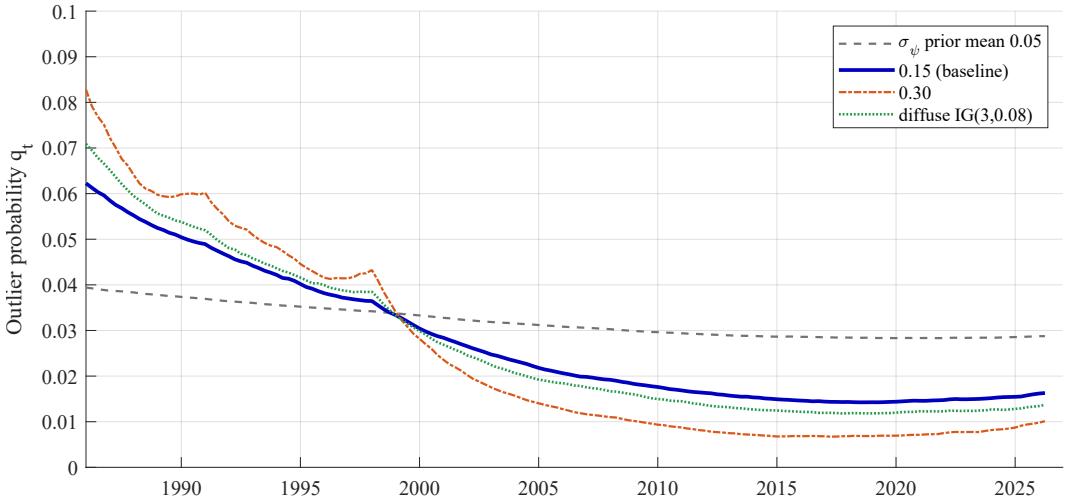


Figure 7: Estimated outlier probability under alternative priors

Note: Posterior mean of q_t for headline inflation under the four priors of Table 8.

[10]Two implementation details. The log-odds path is bounded to $q_t \in [6 \times 10^{-6}, 0.88]$; the bound is never approached in any specification reported here (baseline log-odds lie between -5 and -3) and serves only to prevent the drift to $-\infty$ that an unbounded logit random walk can display in long stretches without outliers. The Pólya–Gamma draws use a 200-term truncation of the infinite-sum representation, and results are unchanged when the truncation is doubled.

6 Conclusion

Korea's pandemic inflation cycle rewrote what a trend-inflation model has to say about the country. Estimated through 2026Q2, the multivariate UCSVO places the trend at 2.6 percent—a full band-width above target and roughly three times its pre-pandemic level—after a surge that the model reads unambiguously as a trend movement rather than a string of outliers, led by food-related sectors and petroleum and sustained, at the end of the sample, by market services. Letting the outlier frequency itself drift sharpens the reading: the estimated propensity to generate inflation outliers has fallen by half to two-thirds since the 1990s and did not revive during the pandemic. That decline is robust across chains, priors, model variants and both seasonal-adjustment methods, but we cannot say what produced it: neither basket composition nor external price volatility survives as an explanation when it is estimated inside the model, and external prices are as volatile as ever. The episode also delivered a verdict on methodology: sector-level information detected the trend movement two quarters before headline inflation alone did, cut the real-time estimation error by more than 40 percent, and produced the only correctly calibrated predictive densities in our comparison—outperforming not just naive benchmarks but factor-augmented and machine-learning forecasts built from the very same sector panel. For a central bank, the practical implication is concrete: a sectoral trend model of the Multivariate-Core-Trend type is worth maintaining as a real-time statistic, its density—with the outlier component retained—is worth reporting alongside its point estimate, and its filtered readings at the peak of a surge should be read with the overshoot that Section 4.5 documents in mind.

The analysis has limits that chart the next steps. Our expenditure weights apply official vintages as a step function rather than the exact annual revisions; the real-time exercise uses the current data vintage rather than historical ones; and the evidence

for the time-varying outlier frequency, while economically interpretable, is supported by predictive fit only moderately and mainly in the pandemic period. More fundamentally, twelve sectors are a coarse window on a basket of several hundred items: letting the data select which items carry the common trend, with shrinkage on the loadings at the item level, is the natural extension of the program begun here.

Data Availability Statement

The data that support the findings of this study are openly available from the Bank of Korea's Economic Statistics System (ECOS) at `https://ecos.bok.or.kr`. Every series used is public and can be retrieved through the ECOS open API: consumer price indices by special classification (statistic code 901Y010) and by COICOP division (901Y009), import price indices (401Y015), and the won–dollar exchange rate (731Y006). Seasonal adjustment uses X-13ARIMA-SEATS, distributed without charge by the U.S. Census Bureau at `https://www.census.gov/data/software/x13as.html`.

The replication package—code that downloads the raw series, constructs the estimation samples, estimates every model reported in the paper, and regenerates all tables and figures—will be deposited in a public repository upon acceptance and is available from the author in the meantime. Running the download scripts requires the user's own ECOS API key, which the scripts read from the environment variable `ECOS_API_KEY`; keys are issued free of charge on registration.

References

Ahn, H. J. and Luciani, M. (2026). Common and Idiosyncratic Inflation. *Journal of Applied Econometrics*, 41(2):156–168.

Almuzara, M. and Sbordone, A. M. (2024). Measurement and Theory of Core Inflation. Staff Report 1115, Federal Reserve Bank of New York.

Atkeson, A. and Ohanian, L. E. (2001). Are Phillips curves useful for forecasting inflation? *Quarterly Review*, 25(1):2–11.

Carriero, A., Clark, T. E., Marcellino, M., and Mertens, E. (2024). Addressing COVID-19 Outliers in BVARs with Stochastic Volatility. *The Review of Economics and Statistics*, 106(5):1403–1417.

Choi, J. (2016). Forecasting Medium-and Long-term Inflation in U.S. by Using Inflation Expectation. *Kukje Kyungje Yongu*, 22(2):55–82.

del Negro, M. and Otrok, C. (2008). Dynamic factor models with time-varying parameters: measuring changes in international business cycles. Staff Reports 326, Federal Reserve Bank of New York.

Dovì, M.-S. (2026). Multivariate Trend Inflation in the United States and the Euro Area. *Journal of Applied Econometrics*, pages 1–6. Replication section.

Eo, Y., Uzeda, L., and Wong, B. (2023). Understanding Trend Inflation Through the Lens of the Goods and Services Sectors. *Journal of Applied Econometrics*, 38(5):751–766.

Garnier, C., Mertens, E., and Nelson, E. (2015). Trend Inflation in Advanced Economies. *International Journal of Central Banking*, 11(4):65–136.

Han, G. and Irie, K. (2026). Zero-Inflated Stochastic Volatility Model for Disaggregated Inflation Data with Exact Zeros. *Journal of the Royal Statistical Society Series C: Applied Statistics*. Advance access; DOI 10.1093/jrsssc/qlag026.

Kim, D. and Eo, Y. (2022). Inflation in Korea after COVID-19: Evidence from Disaggregate CPI Inflation Rates by Expenditure Category. *The Korean Economic Forum*, 15(2):1–20. In Korean.

Kim, H. H. (2018). Looking into the black box of the korean economy: The sparse factor model approach. *Journal of Asia-Pacific Economy*, 23(1):1–16.

Kim, S., Shephard, N., and Chib, S. (1998). Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models. *Review of Economic Studies*, 65(3):361–393.

Kim, Y. J. and Kwark, N.-S. (2024). Changes in Inflation Dynamics in Korea: Global Factor, Country Factor and Their Propagation. *Oxford Bulletin of Economics and Statistics*, 86(6):1373–1408.

Polson, N. G., Scott, J. G., and Windle, J. (2013). Bayesian Inference for Logistic Models Using Pólya-Gamma Latent Variables. *Journal of the American Statistical Association*, 108(504):1339–1349.

Qin, A. and Qu, Z. (2026). Regime-Switching Models for Disaggregated Data. Technical report, arXiv:2606.08398.

Stock, J. H. and Watson, M. W. (2007). Why Has U.S. Inflation Become Harder to Forecast? *Journal of Money, Credit and Banking*, 39(s1):3–33.

Stock, J. H. and Watson, M. W. (2016a). Appendix to Core Inflation and Trend Inflation. Appendix, Princeton University.

Stock, J. H. and Watson, M. W. (2016b). Core Inflation and Trend Inflation. *Review of Economics and Statistics*, 98(4):770–784.

Watson, M. W. (2026). Forecasting the Covid Surge in Inflation. Working Paper 35435, National Bureau of Economic Research.

Yu, B.-C. and Kim, B.-H. (2003). The properties of estimated core inflation. *Journal of Industrial Economics and Business*, 16(2):195–208.