Federal Reserve Bank of Cleveland Working Paper Series

Nowcasting Inflation

Edward S. Knotek II and Saeed Zaman

Working Paper No. 24-06

March 2024

Suggested citation: Knotek II, Edward S. and Saeed Zaman. 2024. "Nowcasting Inflation." Working Paper No. 24-06. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-202406.

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

Nowcasting Inflation[*]

Edward S. Knotek II
Federal Reserve Bank of Cleveland

Saeed Zaman
Federal Reserve Bank of Cleveland

March 2024

Abstract

This chapter summarizes the mixed-frequency methods commonly used for nowcasting inflation. It discusses the importance of key high-frequency data in producing timely and accurate inflation nowcasts. In the US, consensus surveys of professional forecasters have historically provided an accurate benchmark for inflation nowcasts because they incorporate professional judgment to capture idiosyncratic factors driving inflation. Using real-time data, we show that a relatively parsimonious mixed-frequency model produces superior point and density nowcasting accuracy for headline inflation and competitive nowcasting accuracy for core inflation compared with surveys of professional forecasters over a long sample spanning 1999–2022 and over a short sample focusing on the period since the start of the pandemic.

Keywords: inflation, nowcasting, mixed-frequency models, survey nowcasts, real-time data

JEL classifications: C53, E3, E37

[*] This draft was prepared as a handbook chapter on nowcasting inflation; the final version will appear in the *Research Handbook of Inflation*. We thank Matthew Gordon for excellent research assistance. The views expressed herein are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Cleveland or the Federal Reserve System. Contact information: edward.knotek@clev.frb.org and saeed.zaman@clev.frb.org. Research Department, Federal Reserve Bank of Cleveland, P.O. Box 6387, Cleveland, OH 44101.

1. Introduction

Real-time tracking of inflation developments is important because inflation influences the behavior of everyone in an economy. When making decisions, consumers and businesses may have to forecast the inflation rate far into the future. Unfortunately, inflation tends to be difficult to predict accurately. However, some recent research finds that forecasts of inflation over the next several years can be improved by incorporating more accurate estimates of where inflation is likely to be in the near term.[1] These inflation "nowcasts" thus serve as an important jumping-off point for modeling how inflation is likely to behave over a longer period.

The recognition of the importance of good inflation nowcasts for multi-step forecast accuracy has coincided with the development of multiple mixed-frequency (MF) methods for simultaneously modeling data at different frequencies. Statistical agencies typically report official monthly inflation data with a lag of a few weeks following the end of the target month or quarter. However, a considerable amount of information that could have an influence on monthly and quarterly inflation readings is available at daily and weekly frequencies. Mixed-frequency dynamic factor models (MF-DFMs), mixed data sampling (MIDAS) regressions, mixed-frequency vector autoregressions (MF-VARs), customized MF approaches that take advantage of relationships among particular series, and a variety of machine learning techniques, using both a small number of indicators or big data sets, have been applied to combine both high- and low-frequency indicators into a unified framework to nowcast inflation.

Of the approaches proposed in the literature to nowcast US inflation, the customized MF model developed by Knotek and Zaman (2017), which relies on deterministic model switching (DMS) and a limited number of indicators, has been shown to be highly effective in producing timely and accurate nowcasts of inflation. This model underlies the inflation nowcasts produced and disseminated to the public each business day by the Federal Reserve Bank of Cleveland. Knotek and Zaman (2017; 2023a; 2023b) have shown that, historically, this DMS model has done quite well; in many cases, the model's nowcasts have been more accurate than those of common benchmarks from alternative statistical models, including from competing MF models, and from consensus inflation nowcasts from surveys of professional forecasters. Research by other authors has either indirectly or directly validated the competitive inflation nowcasting properties of this DMS model (e.g., Marsilli, 2017; Clark et al., 2022).

In this chapter, we illustrate both the real-time point and density inflation nowcasting accuracy of the Knotek and Zaman DMS model relative to surveys of professional forecasters. We do horseraces over a long sample spanning 1999:Q2 through 2022:Q4 and a shorter sample spanning the period since the onset of the COVID-19 pandemic, a period associated with very high economic uncertainty and volatile movements in economic variables, including inflation, where the ability to use professional judgment to capture special factors outside the scope of a particular model may have proven to be especially important for nowcasting inflation.

We evaluate the nowcasting accuracy of the DMS model and find that it has performed relatively well during both sample periods. Compared with professional forecasters, the DMS model has been relatively more accurate than the inflation nowcasts coming from the Blue Chip consensus (BC) and the Survey of Professional Forecasters (SPF). While nowcasting errors have increased in absolute size since the onset of the pandemic, the DMS model has tended to outperform survey estimates even during this recent period. These latter results are notable because they represent a true out-of-sample test for the DMS model in a period where the ability

[1] See Faust and Wright (2013); Krüger et al. (2017); Knotek and Zaman (2019).

of surveys to incorporate judgment, novel data sources, and different models could have been particularly beneficial.

A noteworthy feature of the DMS model is that the nowcasts are produced using only ten data series, of which eight are monthly, one weekly, and one daily. In contrast, most other contributions in nowcasting US inflation have considered larger information sets, with some approaches utilizing what can be called "Big Data." For example, Modugno (2013) applies a MF-DFM on eight monthly, four weekly, and fifteen daily variables to nowcast headline CPI inflation, but this expanded data set does not deliver more accurate inflation nowcasts compared with the DMS model (see Knotek and Zaman, 2017; 2023a). Clark et al. (2022) implement a state-of-the-art random forest machine learning model and apply it to a rich information set consisting of 150 variables. They find that the inflation nowcasting accuracy of their approach is comparable to that of the DMS model. Garciga, Knotek, and Zaman (2024) extend the information set of Knotek and Zaman (2017) by including pricing information contained in the regional Federal Reserve Bank surveys and find that doing so yields improvements in nowcast accuracy, but those gains were limited to the recent period. Aparicio and Bertolotto (2020) show that there are potentially some nowcasting gains coming from Big Data via scraped online retail prices, although one cannot reject the null for the US of equal predictive nowcasting accuracy with models that use offline fuel prices instead. Overall, it is not clear from the evidence thus far that greater amounts of information and/or more sophisticated statistical or data science approaches universally deliver more accurate inflation nowcasts.

The focus of this survey chapter is on nowcasting inflation in the US. In reviewing the relevant literature on mixed-frequency methods, we include techniques and papers that have focused on nowcasting inflation in other countries as well. This chapter borrows heavily from the results in Knotek and Zaman (2017; 2023a; 2023b) and proceeds as follows. Section 2 describes the data flow and inflation nowcasting process. Section 3 discusses mixed-frequency methods for nowcasting inflation. Section 4 discusses details on surveys of professional forecasters and their inflation nowcasts. Section 5 presents empirical results from horseraces comparing the nowcasting accuracy of the DMS model with the surveys of professional forecasters. Section 6 concludes.

2. The Data Flow and Inflation Nowcasting

In the United States, monthly inflation is usually reported in non-annualized terms as $\pi_t = 100(P_t / P_{t-1} - 1)$, where P_t is the price level in month t. Using these monthly price levels, statistical agencies compute year-over-year inflation rates as $\pi_{t,t-12} = 100(P_t / P_{t-12} - 1)$ and quarterly inflation rates π_T measured at seasonally adjusted annualized rates as $\pi_T = 100[(P_T / P_{T-1})^4 - 1]$, where P_T denotes the price level in quarter T, which is the average of the three monthly price levels in that quarter: $P_T = (1/3)(P_{T,t=1} + P_{T,t=2} + P_{T,t=3})$.[2] Following Knotek and Zaman (2017; 2023a; 2023b), we maintain consistency with this method of computing inflation in our nowcasting exercises in Section 5: we keep track of available monthly price levels and then nowcast or forecast the missing monthly readings of a given quarter to construct quarterly inflation rates.

[2] This formula is consistent with the way that the US Bureau of Economic Analysis reports quarterly inflation rates.

The nowcasting literature that has adhered to the real-time data flow has shown that the timing of data releases is important for constructing and assessing the accuracy of nowcasts. The US Bureau of Labor Statistics (BLS) releases the Consumer Price Index (CPI) around the middle of the following month; e.g., the May CPI is released around mid-June. The Bureau of Economic Analysis (BEA) typically releases the other major measure of consumer prices, the personal consumption expenditures (PCE) price index, around the end of the following month; e.g., the May PCE price index is released around the end of June, after the CPI for May has already been released. While the contents and coverage of the two price indexes differ, the CPI has predictive content over the PCE price index during the interim period before the latter is released.

Before the release of the CPI, there is a variety of available higher-frequency data (e.g., daily oil and raw commodity prices, daily financial variables, weekly gasoline prices) or monthly indicators released in a more timely manner (e.g., survey data on manufacturing and non-manufacturing price indexes compiled by the Institute for Supply Management and by the regional Federal Reserve Banks, producer price index data series) with predictive content for CPI and PCE inflation—and for their main components: core, food, and energy inflation—that could be helpful in the production of timely and accurate inflation nowcasts. Furthermore, in recent years, increasing amounts of (unstructured) high-frequency information in the form of online prices, scanner data, credit card transaction data, and the like have become available, with research showing that these series provide some benefits for nowcasting inflation and especially for certain components of aggregate inflation, such as food inflation and goods inflation. One of the distinguishing features of many contributions in the inflation nowcasting literature is the specific high-frequency data that are shown either to help improve the accuracy of inflation nowcasts over traditional methods or to provide timelier nowcasts. We highlight those contributions in Section 3.

A common finding across papers on nowcasting inflation is that accounting for high-frequency gasoline prices is crucial to obtaining accurate nowcasts of headline inflation (e.g., Modugno, 2013; Knotek and Zaman, 2017; Marsilli, 2017; Clark et al., 2022). This is because gasoline prices, which are heavily influenced by oil prices, usually dominate near-term fluctuations in consumer energy prices in the US, and even though energy consumption is a small share of total spending and thus receives a small weight in the aggregate price indexes, the volatility of this series is a key factor behind month-to-month fluctuations in headline inflation. The fact that gasoline and oil prices are available at a higher frequency than monthly can be used to nowcast monthly gasoline price inflation, which, in turn, delivers improved accuracy of the inflation nowcasts.

Figure 1, which is taken from Knotek and Zaman (2017), provides a visual illustration of the timing of the data flow over the course of a quarter that an inflation nowcasting framework focused on nowcasting both the monthly and the quarterly inflation rates could utilize to produce timely and accurate inflation nowcasts. Shown is the daily (e.g., oil prices), weekly (e.g., gasoline prices), and monthly (e.g., CPI and PCE prices) data flow for a representative quarter (e.g., Q1). The numbers in the circles are representative dates (cases) over the course of a quarter featuring distinct information sets. The data release lags imply that at the very beginning of a quarter (Case 1), the last available quarterly inflation reading would have been from two quarters earlier. For example, at the beginning of January, neither the CPI nor the PCE price index for December would be available. Thus, we will not have the complete data for the fourth quarter. As a result, the last available quarterly inflation reading would be from the third quarter of the preceding year. The quarterly data for the fourth quarter of the preceding year would become available in mid-January for the CPI (Case 3) and late January for the PCE price index (Case 5). For the PCE price index, at the beginning of January, we typically have the third estimate of PCE price inflation for the third quarter of the preceding year. Because the information flow determines what data were available in real time at any particular point in time for nowcasting, it is important to match information sets between competing models or between a model and a survey of professional forecasters to ensure that both are on an equal and accurate footing when conducting nowcasting horseraces. In our exercises below, we follow the real-time data flow whenever possible.

Figure 1: Data Flow Timing

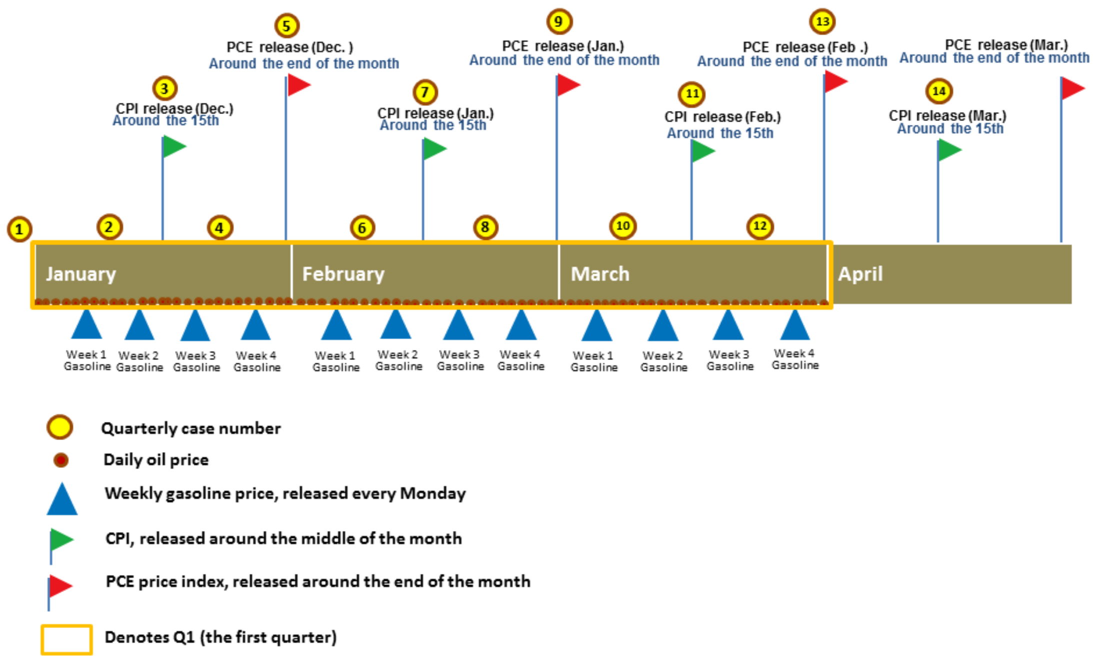

Notes: This figure is taken from Knotek and Zaman (2017).

3. Nowcasting Methods for Inflation

This section discusses the MF methods commonly used to nowcast inflation. These methods include the deterministic model switching (DMS) framework of Knotek and Zaman (2017), mixed data sampling (MIDAS), MF dynamic factor models (MF-DFMs), MF vector autoregressions (MF-VARs), and machine learning and Big Data approaches. To economize on space, we briefly discuss each of these methods, except for the DMS framework, for which we provide more details because of its use in the empirical section later in the chapter.

3.1. Deterministic Model Switching (DMS)

At its core, the DMS model developed by Knotek and Zaman (2017; 2023a; 2023b) follows a parsimonious approach to nowcasting inflation. First, it relies on a small number of carefully chosen data series to inform nowcast estimates. Second, it combines simple univariate and multivariate regression techniques. Third, it imposes time-varying weights on disaggregate and aggregate variables in nowcasting the aggregate. The switches between modeling techniques and weights occur deterministically based on the available information set at a point in time, thereby taking advantage of the information flow to improve nowcasting accuracy. For example, disaggregate information is utilized for nowcasting the aggregate only when this information is available and informative, resulting in time-varying weights.

The model takes the general form

$$\mathbf{A}_{s(t)}\mathbf{Z}_t = \mathbf{B}_{s(t)} + \mathbf{C}_{s(t)}\mathbf{X}_t + \sum_{j=1}^{J}\mathbf{D}_{j,s(t)}\mathbf{Z}_{t-j} + \boldsymbol{\varepsilon}_{s(t)}, \tag{1}$$

where $\mathbf{Z}_t$ is an $n \times 1$ vector of aggregates, $\mathbf{X}_t$ is an $m \times 1$ vector of disaggregates that are informative over $\mathbf{Z}_t$, and $\boldsymbol{\varepsilon}_{s(t)} \sim N(\mathbf{0}, \boldsymbol{\Sigma})$. The coefficient matrices **A**, **B**, **C**, and $\mathbf{D}_j$ are $n \times n$, $n \times 1$, $n \times m$, and $n \times n$, respectively, and vary depending on the available information set, denoted $s(t)$; in particular, **C** and $\mathbf{D}_j$ measure the weights put on the disaggregates and lagged aggregates, respectively. This model structure permits the use of information from diverse sources.

3.1.1. Nowcasting Core Inflation in DMS

With $\mathbf{Z}_t = [\pi_t^{\text{Core CPI}}, \pi_t^{\text{Core PCE}}]'$ as the aggregate of interest—where $\pi_t^{\text{Core CPI}}$ and $\pi_t^{\text{Core PCE}}$ are the monthly core CPI inflation rate and core PCE inflation rate in month t, respectively—then $\mathbf{X}_t = \mathbf{0}$ in equation (1). Because CPI releases precede PCE releases, we can at times take advantage of this timing mismatch. If we have $\pi_t^{\text{Core CPI}}$ but we only have $\pi_{t-1}^{\text{Core PCE}}$, we bridge core CPI to core PCE to nowcast the as-yet-unreleased month t core PCE inflation rate. Conditional on being in this state, the time-varying weights in equation (1) become

$$\mathbf{A}_{s(t)} = \begin{bmatrix} 0 & 0 \\ a_{21} & 1 \end{bmatrix}, \ \mathbf{B}_{s(t)} = \begin{bmatrix} 0 \\ b_2 \end{bmatrix}, \ \mathbf{D}_{j,s(t)} = \mathbf{0} \ \forall j. \tag{2}$$

The coefficients in equation (2) are estimated over a window of length τ to nowcast $\hat{\pi}_t^{\text{Core PCE}}$, where we use "^" to denote a nowcasted or forecasted value.

In all other cases, we rely on the spirit of Atkeson and Ohanian (2001), who find that inflation over the previous four quarters is a difficult forecasting benchmark to beat, and we forecast monthly core inflation $\hat{\mathbf{Z}}_t$ using recursive 12-month moving averages, by fixing

$$\mathbf{A}_{s(t)} = \mathbf{I}_2, \ \mathbf{B}_{s(t)} = \mathbf{0}, \ \mathbf{D}_{j,s(t)} = (1/12)\mathbf{I}_2, \ J = 12. \tag{3}$$

Thus, if we have data through time $t-1$ on $\mathbf{Z}_{t-1} = [\pi_{t-1}^{\text{Core CPI}}, \pi_{t-1}^{\text{Core PCE}}]'$, we use equation (3) to recursively generate forecasts for time t, $t+1$, …. If we have data through time t on $\pi_t^{\text{Core CPI}}$ but only through time $t-1$ on $\pi_{t-1}^{\text{Core PCE}}$, we first use equation (2) to nowcast $\hat{\pi}_t^{\text{Core PCE}}$ and then use equation (3) to recursively generate forecasts for time $t+1$, $t+2$, …, where $\hat{\pi}_t^{\text{Core PCE}}$ is included as an observation in taking the moving average. In this way, the arrival of a new core CPI reading affects its own forecast; but because PCE release dates lag behind CPI release dates, the arrival of core PCE inflation has no impact on core CPI. The arrival of a new core CPI reading also affects the nowcast for core PCE inflation for that month, and this nowcast in turn affects the core PCE inflation forecast for future months through the recursion. Once core PCE inflation data come out for that month, the forecast for core PCE inflation is potentially affected again, and the process resets.

3.1.2. Nowcasting Headline Inflation in DMS

In addition to core prices, food prices and energy prices are other key disaggregates for headline inflation. In theory, high-frequency futures and spot market prices for raw food items could have predictive content over monthly consumer food inflation, denoted π_t^{Food}, and serve as useful disaggregate indicators $\mathbf{X}_t$ for food inflation. In practice, raw food prices are a small determinant of most retail food prices. We follow the principle of parsimony and, assuming we have data through month $t-1$, forecast monthly food inflation as $\hat{\pi}_t^{\text{Food}} = (1/12)\sum_{j=1}^{12} \pi_{t-j}^{\text{Food}}$ and then recursively forecast $\hat{\pi}_{t+k}^{\text{Food}}$, $k=1,2,\ldots$[3]

Energy prices offer a contrast to food prices because gasoline prices dominate fluctuations in US consumer energy prices, and gasoline prices are heavily influenced by oil prices. Gasoline prices and oil prices are available at a higher frequency than monthly and can be used to nowcast monthly gasoline price inflation after seasonal adjustment, denoted $\hat{\pi}_t^{\text{Gasoline}}$, which can be used as one of the disaggregate variables in nowcasting headline inflation. The combination of data release lags, higher-than-monthly-frequency data on oil and gasoline prices, and the methodology we propose implies that we will typically have one or two more months of gasoline inflation nowcasts or forecasts, $\hat{\pi}_{t+k}^{\text{Gasoline}}$, $k\geq 0$, than we have data on monthly CPI and PCE inflation. The inclusion of these nowcasts $\hat{\pi}_t^{\text{Gasoline}}$ as disaggregates in $\mathbf{X}_t$ plays an important role in producing highly accurate nowcasts for headline inflation. In the interest of brevity, we refer readers to Knotek and Zaman (2017) for complete details on the construction of the gasoline inflation nowcasts.

With nowcasts of our four disaggregates, we can construct nowcasts and forecasts of headline inflation rates using the model in equation (1) and weights that vary deterministically with the available information set. Let $\mathbf{Z}_t = [\pi_t^{\text{CPI}}, \pi_t^{\text{PCE}}]'$ be the aggregate of interest, where π_t^{CPI} and π_t^{PCE} are the month t CPI inflation rate and PCE inflation rate, respectively. The vector of relevant disaggregates for headline inflation is

$$\mathbf{X}_t = \left[\pi_t^{\text{Core CPI}}, \pi_t^{\text{Core PCE}}, \pi_t^{\text{Food}}, \pi_t^{\text{Gasoline}}\right]'. \tag{4}$$

For states in which we have π_t^{CPI} but not π_t^{PCE}, we again bridge the headline CPI reading to headline PCE; the time-varying weights are:

[3] The CPI and the PCE price index treat food differently, suggesting there could potentially be two separate disaggregate series. In the CPI, the food index encompasses both food at home and food away from home, and the core CPI excludes all food. In the PCE price index, food and beverages purchased for off-premises consumption are classified as nondurable goods and are excluded from the core PCE price index. Food services and accommodations are classified as services and are included in the core PCE price index. This change took effect with the BEA's 2009 comprehensive revisions. Because of the limited availability of real-time data on food inflation measures, we consider a single food series—the CPI for food—that is used as a disaggregate measure for both CPI and PCE inflation.

$$\mathbf{A}_{s(t)} = \begin{bmatrix} 0 & 0 \\ a_{21} & 1 \end{bmatrix},\ \mathbf{B}_{s(t)} = \begin{bmatrix} 0 \\ b_2 \end{bmatrix},\ \mathbf{C}_{s(t)} = \mathbf{0},\ \mathbf{D}_{j,s(t)} = \mathbf{0} \text{ for all } j \tag{5}$$

and can be estimated over a window of data of length τ to produce a nowcast $\hat{\pi}_t^{\text{PCE}}$. For states in which we have $\hat{\pi}_t^{\text{Gasoline}}$, we pair that nowcast with the forecasts of $\hat{\pi}_t^{\text{Food}}$, $\hat{\pi}_t^{\text{Core CPI}}$, and $\hat{\pi}_t^{\text{Core PCE}}$ generated earlier to complete the vector $\hat{\mathbf{X}}_t$. The time-varying weights

$$\mathbf{A}_{s(t)} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix},\ \mathbf{B}_{s(t)} = \begin{bmatrix} b_1 \\ b_2 \end{bmatrix},\ \mathbf{C}_{s(t)} = \begin{bmatrix} c_{11} & 0 & c_{13} & c_{14} \\ 0 & c_{22} & c_{23} & c_{24} \end{bmatrix},\ \mathbf{D}_{j,s(t)} = \mathbf{0} \text{ for all } j \tag{6}$$

are estimated over a window of data of length τ and we can then forecast $\hat{\mathbf{Z}}_t = [\hat{\pi}_t^{\text{CPI}}, \hat{\pi}_t^{\text{PCE}}]'$ using $\hat{\mathbf{X}}_t$. Incoming high-frequency data on oil prices and gasoline prices that affect $\hat{\pi}_t^{\text{Gasoline}}$ will affect $\hat{\mathbf{X}}_t$ and headline inflation nowcasts through equation (6). And in states for which we lack $\hat{\pi}_t^{\text{Gasoline}}$ and thus do not have the complete disaggregate vector $\hat{\mathbf{X}}_t$, we use recursive 12-month moving averages by fixing

$$\mathbf{A}_{s(t)} = \mathbf{I}_2,\ \mathbf{B}_{s(t)} = \mathbf{0},\ \mathbf{C}_{s(t)} = \mathbf{0},\ \mathbf{D}_{j,s(t)} = (1/12)\mathbf{I}_2,\ J = 12\,. \tag{7}$$

As with the procedure for core inflation set out earlier, nowcasts or forecasts of headline inflation can enter the recursion in equation (7) if they are part of the 12-month window. Also as with core inflation, because PCE inflation is released after CPI inflation, we only bridge from CPI to PCE if there is an additional CPI reading. Hence, nowcasts or forecasts for headline CPI inflation are determined by equation (6) or (7), whereas nowcasts or forecasts of headline PCE inflation are determined by equation (5), (6), or (7) depending on the available information set.

3.2. Mixed Data Sampling (MIDAS)

Since their introduction by Ghysels et al. (2004), MIDAS models have become increasingly popular in nowcasting applications, including for inflation. A MIDAS model is a reduced-form regression that involves regressing a variable sampled at low frequency (e.g., monthly inflation) on its lags and high-frequency indicators and their lags (e.g., daily oil prices, weekly gasoline prices). To prevent parameter proliferation because of the large number of coefficients of the high-frequency regressors, MIDAS works with distributed lag polynomial operators that impose some strong restrictions on coefficients to reduce the estimation to a smaller number of parameters. The model is estimated using nonlinear least squares.

The MIDAS model with leads for inflation at time t+h, π_{t+h}, takes the form

$$\pi_{t+h} = \alpha_{(h)} + \sum_{j=0}^{P(M)-1} \chi_{j+1,(h)}\ \pi_{t-j} + \sum_{j=0}^{P(M)-1} \gamma_{j+1,(h)}\ Z_{t-j} +$$
$$\beta_h \sum_{j=0}^{P(HF)-1} \omega_{P(HF)-j}\left(\theta_{(h)}^{HF}\right) X_{P(HF)-j,t+1}^{HF} + e_{t+h} \tag{8}$$

where Z includes other monthly variables; $P(M)$ is the number of lags of the monthly regressors; and $P(HF)$ is the number of high-frequency observations, $X_{1,t+1}^{HF}$,….., $X_{P(HF),t+1}^{HF}$ in month t+1 (i.e., the target nowcast month). The coefficients are independently estimated for each forecast horizon (h). The assumption $\sum_{j=0}^{P(HF)-1} \omega_{P(HF)-j}\left(\theta_{(h)}^{HF}\right) = 1$ helps identify β_h.

Monteforte and Moretti (2013) apply a MIDAS model to nowcast euro area inflation. In their case, monthly inflation is regressed on high-frequency financial variables. Knotek and Zaman (2017) implement a MIDAS model to nowcast US inflation as an alternative to the DMS method. In their application, daily oil and weekly gasoline prices act as high-frequency leads, while monthly inflation is the lower-frequency variable. Breitung and Roling (2015) develop a nonparametric MIDAS model to nowcast German inflation rates using daily commodity price data. Marsilli (2017) uses a MIDAS model to nowcast US inflation using monthly unemployment rate data and high-frequency oil price data.

3.3. Dynamic Factor Model (DFM)

Building on Giannone et al. (2008), mixed-frequency dynamic factor models (MF-DFMs) are widely used for nowcasting macroeconomic variables. Modugno (2013) implements a MF-DFM to generate inflation point nowcasts from a data set comprising monthly, weekly, and daily data by extracting a common factor at a daily frequency via the estimation method of Bańbura and Modugno (2014). At the trading-day τ frequency (i.e., the point in time when our nowcasts are made), the DFM takes the form

$$y_\tau = Cf_\tau + \varepsilon_\tau,\ \varepsilon_\tau \sim N(0,\Sigma) \tag{9}$$

where y_τ is a vector of observations of mixed frequencies, C is a matrix of loadings, ε_τ is a vector of idiosyncratic components, and f_τ is a vector of unobserved common components that follows

$$B\, f_\tau = A(L) f_{\tau-1} + \upsilon_\tau, \upsilon_\tau \sim N(0,\mathrm{Q}) \tag{10}$$

where B and $A(L)$ are coefficient matrices that capture factor dynamics. The estimated latent daily factor(s) aggregate to weekly and monthly factors, which are used to construct nowcast estimates for monthly variables, including inflation.

With monthly, weekly, and daily data, $y_\tau = [y_\tau{}^M, y_\tau{}^W, y_\tau{}^D]'$, we have three corresponding factors, $f_\tau = [f_\tau{}^M, f_\tau{}^W, f_\tau{}^D]'$, each of dimension $r\times1$. The monthly factor(s) f_τ^M and the weekly factor(s) f_τ^W are a function of the daily factor(s) $f_\tau{}^D$. Thus equations (9) and (10) can be written as:

$$\begin{bmatrix} y_\tau^M \\ y_\tau^W \\ y_\tau^D \end{bmatrix} = \begin{bmatrix} C_M & 0 & 0 \\ 0 & C_W & 0 \\ 0 & 0 & C_D \end{bmatrix} \begin{bmatrix} f_\tau^M \\ f_\tau^W \\ f_\tau^D \end{bmatrix} + \begin{bmatrix} \varepsilon_\tau^M \\ \varepsilon_\tau^W \\ \varepsilon_\tau^D \end{bmatrix} \tag{11}$$

and

$$\begin{bmatrix} 1 & 0 & -1 \\ 0 & 1 & -1 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} f_\tau^M \\ f_\tau^W \\ f_\tau^D \end{bmatrix} = \begin{bmatrix} \Theta_\tau^M & 0 & 0 \\ 0 & \Theta_\tau^W & 0 \\ 0 & 0 & A_D \end{bmatrix} \begin{bmatrix} f_{\tau-1}^M \\ f_{\tau-1}^W \\ f_{\tau-1}^D \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \upsilon_\tau^D \end{bmatrix} \tag{12}$$

The matrices C_M, C_W, and C_D are the loadings for the monthly, weekly, and daily variables. Θ_τ^M and Θ_τ^W are time-varying coefficients: Θ_τ^M is equal to zero the day after the release of the monthly data and is equal to one elsewhere; similarly, Θ_τ^W is equal to zero the day after the release of the weekly data and is equal to one elsewhere.

Assuming that the monthly variables and weekly variables in our system at any time τ represent a stock (i.e., a snapshot), accordingly the monthly first difference (or growth rate) and weekly first difference (or growth rate) of those variables can be formed by summing up their respective daily first differences (or growth rates). To forecast arbitrarily far into the future, the daily factors are forecast via the transition equation (12) and are translated to daily nowcasts and aggregated to weekly and monthly nowcasts via equation (11).

Amstad and Fischer (2009) apply a MF-DFM on a monthly and daily data set consisting of 434 variables to nowcast Swiss core CPI inflation. Liang et al. (2020) use a MF-DFM to nowcast China's PPI inflation. Knotek and Zaman (2023a) consider a combination of MF-DFMs (and, separately, combinations of DMS and MIDAS models) to nowcast core and headline US inflation. Mariano and Ozmucur (2020) consider a combination of MF-DFMs (and other mixed-frequency models) to nowcast Philippine inflation. Carvalho (2020) also considers a combination of MF-DFMs applied to daily online retail prices and financial data to nowcast monthly consumer price inflation in Brazil. In his model, a factor at a daily frequency is extracted from many online retail price indices and financial variables and used to compute daily nowcasts of CPI inflation.

3.4. Mixed-Frequency VARs (MF-VARs)

VARs are the workhorse models in macroeconomics and macroeconomic forecasting. However, when it comes to nowcasting, they are less widely used than other methods mentioned thus far. This is because they are relatively computationally demanding and have less success in nowcasting horseraces than other parsimonious methods, especially for inflation. Schorfheide and Song (2015) is a prominent contribution to the use of MF-VARs for nowcasting and forecasting macroeconomic variables that are observed at mixed frequencies; i.e., some are observed at a quarterly frequency (e.g., real GDP, consumption) and some at a monthly frequency (e.g., inflation, the unemployment rate). Schorfheide and Song used their MF-VAR to produce quarterly nowcasts and forecasts of inflation, real GDP, and nine other macroeconomic and financial variables. Their MF-VAR implementation, which they estimate with Bayesian methods, entails specifying a state-space representation in which the state transition equations follow a VAR at monthly frequency and the measurement equations define the observed series as a function of potentially latent monthly variables, which are stacked in the state vector. McCracken et al. (2021) and Cimadomo et al. (2022) provide alternative proposals for handling mixed-frequency data in VARs that do not rely on state-space representation and are computationally less demanding. However, empirical applications of these proposals focus on nowcasting real GDP and not inflation. Recently, Aliaj et al. (2023) have proposed a Lasso-regularized MF-VAR (MF-VAR with machine learning method) that works reasonably well in nowcasting euro area HICP inflation. This is because the approach is based on a shrinkage method that can effectively handle large VARs by automatically excluding the information in the VAR that is deemed irrelevant.

3.5. Machine Learning Methods and Big Data

An increasing number of recent papers on inflation nowcasting use high-frequency information extracted from massive nontraditional data sets, including online pricing data, offline household and retail scanner data, and credit card transactions. These data fall broadly under the umbrella of Big Data, since they feature all three "Vs" of "volume, velocity, and variety," properties that define Big Data, according to Cimadomo et al. (2022). As Powell et al.

(2018) discuss, a key obstacle to working with these types of data is that they are “extremely messy” and unstructured and require tremendous effort to reconcile with standard official CPI statistics. In some cases, individuals or private companies have processed the online data into high-frequency online price indices that can be used by empirical researchers for forecasting and other purposes. For example, PriceStats took over the Billion Prices Project (BPP) and provides online price indices for various countries, including the US, at daily to weekly frequency.[4] Aparicio and Bertolotto (2020) provide evidence that aggregates of huge volumes of scraped online prices may hold some promise for nowcasting and short-term inflation forecasting beyond “offline” fuel prices for ten advanced economies, with large and statistically significant gains for nowcasting quarterly US inflation. That said, for nowcasting the current month’s CPI inflation rate in the US, one cannot reject the null of equal predictive nowcasting accuracy with models that use offline fuel prices instead of online prices. Similarly, Carvalho (2020) shows the usefulness of online retail price indices in nowcasting CPI inflation in Brazil in an MF-DFM model setting. Alvarez and Lein (2020) nowcast Swiss CPI inflation using online price indices (based on web-scraped pricing information) and real-time expenditure weights computed from daily debit card transaction data. Researchers have also used online prices to follow inflation developments in specific categories such as food inflation (e.g., for Poland: Jaworski, 2021; Macias et al. , 2022; for Turkey: Soybilgen et al., 2023).

Beck et al. (2023) use millions of weekly household scanner data entries to create 180 scanner-based price indices at a weekly frequency to nowcast monthly German inflation. The authors fit the mixed-frequency machine learning model (sparse-group LASSO) to these indices to nowcast six main disaggregates of aggregate inflation and use the respective nowcasts and official weights to compute the nowcast of aggregate inflation. Their results indicate that scanner-based indices have can significantly improve nowcasting accuracy for the components of inflation and for aggregate inflation.

Overall, research shows that online prices may help improve inflation nowcasting accuracy when included as additional predictors in standard models or when embedded in machine learning models.

Relatedly, advances in computing are allowing researchers to explore other unstructured data sets to help nowcast inflation. For example, Zheng et al. (2023) apply natural language processing tools to textual data to nowcast inflation in China.

There is also a growing body of research on applying machine learning methods to a rich data set of many (standard) predictors in nowcasting inflation. Schnorrenberger et al. (2024) compare linear machine learning methods (i.e., penalized regression methods: Elastic Net, the LASSO and Ridge, and sparse-group LASSO) to nonlinear machine learning methods (tree-based methods: random forest, Bayesian additive regression trees, generalized random forest, and local linear forest) in nowcasting CPI inflation in Brazil. They find that linear ML methods (e.g., penalized regression methods) generate more accurate nowcasts of monthly inflation compared to nonlinear tree-based algorithms. This finding is consistent with the results of Joseph et al. (2022) and Garcia et al. (2017), who find that linear ML methods generate more accurate short-term inflation forecasts than nonlinear ML methods.

These findings probably explain the inflation nowcast performance of the enhanced random forest model of Clark et al. (2022), which could be viewed as combining the benefits of linear and nonlinear ML methods. The novelty of their model is to allow for a linear relationship between the target variable and predictors at each node of the tree, as opposed to the standard

[4] For information on the BPP, see Cavallo and Rigobon (2016).

approach of just taking the average of the observations of the target variable in that node. They show that doing so leads to more accurate inflation nowcasts, whose accuracy outperforms simple benchmark models, and rivals the DMS method of Knotek and Zaman in a real-time comparison. This finding that a state-of-the-art ML method using 150 data series of different frequencies performs just as well as the parsimonious DMS model, which relies on a select few variables, is noteworthy. On the one hand, it suggests that the proposed model of Clark et al. (2022) is a very flexible approach, since it can effectively parse the rich information set to provide accurate inflation nowcasts that rival a hard-to-beat benchmark model. On the other hand, this finding instills confidence in the continuing use of the parsimonious approach of Knotek and Zaman.

4. Survey Nowcasts

In the US, various surveys of professional forecasters provide estimates of inflation nowcasts. These include the Philadelphia Fed's Survey of Professional Forecasters (SPF), Blue Chip Economic Indicators (BC), the National Association for Business Economics Outlook Survey, Bloomberg Economics, and Consensus Economics. With the exception of the SPF, the others are maintained by private entities and are accessible through a fee-based subscription. Of these surveys, the SPF and BC are the most widely used, in part due to their much longer history and in part due to an extensive body of research documenting their competitive forecasting properties. In particular, Faust and Wright (2009; 2013) have shown that professional forecasters' inflation nowcasts tend to outperform those from statistical models. Faust and Wright (2013) suggest that subjective nowcasts may hold a distinct advantage through their ability to "add expert judgment" to models (p. 20). Because of their documented nowcasting performance, they are a common benchmark against which to test other inflation nowcasting models. We briefly discuss the SPF and BC, including some details that will be useful in the next section, where we compare the accuracy of the surveys' inflation nowcasts to that of the DMS model.

4.1. Survey of Professional Forecasters (SPF)

The SPF is a publicly available survey that is published quarterly and is released around the middle of the second month of the quarter. The Federal Reserve Bank of Philadelphia publishes the survey dates; these dates are about one week prior to the release date, which means that SPF nowcasts of current quarter inflation are typically made before the first monthly CPI reading for the quarter is released. In the nowcasting horseraces in the next section, we match information sets that would have been available to the professional forecasters with the model's information set. The SPF reports quarterly nowcasts for all four inflation measures considered in this paper. It has a long history of reporting CPI forecasts, going back to 1981:Q3. The SPF started reporting core CPI, headline PCE, and core PCE inflation in 2007:Q1. In Section 5, we conduct nowcast comparison exercises beginning in 1999:Q2 for the CPI, which is when our real-time data set begins, and beginning in 2007:Q1 for the other three measures. In all cases, we end the comparisons in 2022:Q4. Following past research, we use the SPF median nowcasts to eliminate outliers.

4.2. Blue Chip Economic Indicators (BC)

Blue Chip Economic Indicators is a monthly survey of roughly 50 business economists conducted by Wolters Kluwer Legal and Regulatory Solutions. The survey asks respondents to provide quarterly nowcasts and forecasts of major US economic indicators, including three of the four inflation measures we consider: quarterly CPI, PCE, and core PCE inflation. However, the history of the latter two inflation measures is short, since BC started reporting them only recently, in June of 2020. For the nowcasts of quarterly CPI inflation, BC has a much longer history, with the first such nowcasts issued in March 1980 for 1980:Q1, but, as noted above, we will conduct horseraces beginning with 1999:Q2. For each variable, BC reports the consensus quarterly nowcast, which is an equal-weighted average of participants' forecasts. The BC survey is typically released around the 10th of each month, but the survey is conducted over an earlier two-day period that has historically often been indicated in the release. In the next section, we match this timing in our horserace.

5. Nowcasting US Inflation: Surveys vs. DMS

In this section, we compare the nowcasting accuracy of the headline and core inflation estimates coming from the DMS model with the accuracy of nowcasts coming from BC and the SPF. We evaluate both the point and the density nowcast accuracy over a long sample spanning 1999:Q2 (2000:Q4 for the density comparison) through 2022:Q4 and a shorter sample spanning the period since the onset of the COVID-19 pandemic, a period associated with very high economic uncertainty and volatile movements in economic variables, including inflation.

To preview, the evaluation results indicate that the DMS model has been relatively more accurate than the inflation nowcasts coming from the BC consensus and the SPF. While nowcasting errors have increased in absolute size since the onset of the pandemic, the DMS method has tended to outperform survey estimates even during the recent period. These results are noteworthy because when making forecasts, professional forecasters use a range of models and expert judgment to capture the special factors driving near-term inflation developments.

5.1. Point Nowcasting Accuracy Comparison

5.1.1. Point Nowcast Comparison with Blue Chip

Considering the timing of the BC survey and the publication of the CPI data, we compare BC nowcasting accuracy with that of the DMS model at four different points in time for each quarter. For example, nowcasts of the second quarter are collected in the April, May, June, and July surveys. The July BC survey data are released about one to two weeks before the BLS releases all the data needed to compute quarterly CPI inflation.

Table 1 reports the root mean squared error (RMSE) accuracy comparison for the sample period spanning 1999:Q2 through 2022:Q4 (top panel). We also report results for a short sample spanning 2020:Q1 through 2022:Q4 (bottom panel), to investigate inflation nowcasting performance since the start of the pandemic.[5] The results indicate the following.

[5] The "true" actual quarterly annualized inflation rates are computed using the third monthly estimate of PCE and CPI prices as the actual value, except for 2022:Q4, which uses the second estimate for November 2022 and the first estimate for December 2022.

First, as we move from month 1 (at the very beginning of the quarter) through month 4 (the survey from the month immediately following the quarter, which is released right before the quarterly CPI data are available), we see monotonic reductions in RMSEs for both our DMS model and the BC consensus. This improved nowcast accuracy is due in part to accumulating inflation data as we move from month 1 to month 4. Second, our model's nowcasts have been more accurate on average than BC nowcasts across all four months, as demonstrated by smaller RMSEs. Third, the magnitude of the errors experienced since the onset of the pandemic is notably larger, as evidenced by comparing RMSEs between the long and short evaluation samples, which is consistent with inflation becoming more volatile and more difficult to forecast. Nevertheless, our DMS model's inflation nowcasts have tended to be more accurate on average during this period, too.

Table 1: CPI Point Nowcasting Comparisons with the Blue Chip Consensus

Nowcast Evaluation Sample: 1999:Q2–2022:Q4				
	Blue Chip survey conducted in:			
	Month 1	Month 2	Month 3	Month 4
DMS Model RMSE	1.861	1.281	0.539	0.262
Blue Chip RMSE	1.964	1.523	0.854	0.425
Ratio, Blue Chip MSE to DMS MSE	1.114	1.412*	2.512***	2.631***
Nowcast Evaluation Sample: 2020:Q1–2022:Q4				
	Blue Chip survey conducted in:			
	Month 1	Month 2	Month 3	Month 4
DMS Model RMSE	2.968	2.511	0.955	0.321
Blue Chip RMSE	3.213	2.630	1.347	0.683
Ratio, Blue Chip MSE to DMS MSE	1.172	1.097	1.988	4.531

NOTES: Comparisons are matched based on Blue Chip survey dates; for example, when nowcasting the first quarter, month 1 would refer to the Blue Chip survey date in January, month 2 would be February's date, and month 3 would be March's date. The Blue Chip survey in month 4 (e.g., April) is conducted prior to the availability of CPI inflation data for the previous quarter and is the final nowcast. Quarterly inflation rates are seasonally adjusted annualized percent changes, so errors are expressed in annualized percentage points. *, **, and *** denote rejection of the null of equal predictive ability for the DMS model compared with each alternative model at the 10%, 5%, and 1% level, respectively, based on the Giacomini–White test. The exercise uses real-time data from 1999:Q2 through 2022:Q4. For the 2020:Q1–2022:Q4 sample, we do not compute statistical significance because the length of the sample size is so short.

To provide a visual illustration of recent quarterly performance, Figure 2 plots the profiles of the absolute nowcast errors for CPI inflation from our DMS model and the BC consensus for the short sample since the beginning of the COVID-19 pandemic in 2020:Q1. The four panels in the figure correspond to months 1 through 4 of each quarter. Looking at the figure, two observations immediately stand out. First, moving from panels (a) through (d), the magnitude of the absolute errors decreases (as can be seen by the changing scale of the y-axis). Second, it is generally the case that DMS model nowcasts were more accurate than the BC consensus, with the notable exception of very recent quarters.

Figure 2: CPI Point Nowcasting Comparison with the Blue Chip Consensus: Absolute Errors

5.1.2. Point Nowcast Comparison with the SPF

Table 2 reports the results when we compare the point nowcast accuracy of our DMS model to that of the SPF for all four inflation measures. When evaluated over the long sample, our model's nowcasts for both headline CPI and PCE inflation outperform the accuracy of the SPF nowcasts by 0.39 percentage point and 0.25 percentage point on average, respectively. When evaluated over the short sample that includes the pandemic and the recent surge in inflation, the magnitudes of the accuracy gains from our model compared with the SPF are significantly larger, at 0.65 percentage point for headline CPI and 0.41 percentage point for headline PCE.

Table 2: Point Nowcasting Comparisons with the Survey of Professional Forecasters

Nowcast Evaluation Sample: 1999:Q2–2022:Q4				
	CPI	Core CPI	PCE	Core PCE
DMS Model RMSE	1.118	0.708	0.874	0.563
SPF RMSE	1.510	1.062	1.128	0.780
Ratio, SPF MSE to DMS MSE	1.824***	2.251	1.667**	1.918
Nowcast Evaluation Sample: 2020:Q1–2022:Q4				
	CPI	Core CPI	PCE	Core PCE
DMS Model RMSE	1.988	1.229	1.343	0.801
SPF RMSE	2.636	2.166	1.756	1.505
Ratio, SPF MSE to DMS MSE	1.758	3.106	1.708	3.527

NOTES: Real-time comparisons are based on the SPF survey dates. SPF expectations for each quarter are the median value. Quarterly inflation rates are seasonally adjusted annualized percent changes, so errors are expressed in annualized percentage points. *,**, and *** denote rejection of the null of equal predictive ability for the DMS model compared with each alternative model at the 10%, 5%, and 1% level, respectively, based on the Giacomini–White test. The CPI exercise uses real-time data from 1999:Q2 through 2022:Q4. The core CPI, PCE, and core PCE exercises use real-time data from 2007:Q1 (the first available SPF estimates) through 2022:Q4. For the 2020:Q1–2022:Q4 sample, we do not compute statistical significance because the length of the sample size is so short.

In the case of core inflation, our model has also tended to outperform the SPF nowcasts on average. This is true over the long sample and since the start of the pandemic.[6]

Figure 3 provides a visual illustration of nowcast performance since the onset of the COVID-19 pandemic. The four panels in the figure plot the absolute nowcast errors for CPI, core CPI, PCE, and core PCE inflation measures from our model and the SPF. Early in the pandemic and as inflation began to move up in 2021 and early 2022, the inflation nowcasts from our model tended to be more accurate than those from the SPF; however, that relative performance reversed at the end of 2022 as inflation started showing signs of easing.

[6] The nowcast improvements for core inflation are not statistically significant according to the Giacomini-White test.

Figure 3: Nowcasting Comparison with the Survey of Professional Forecasters: Absolute Errors

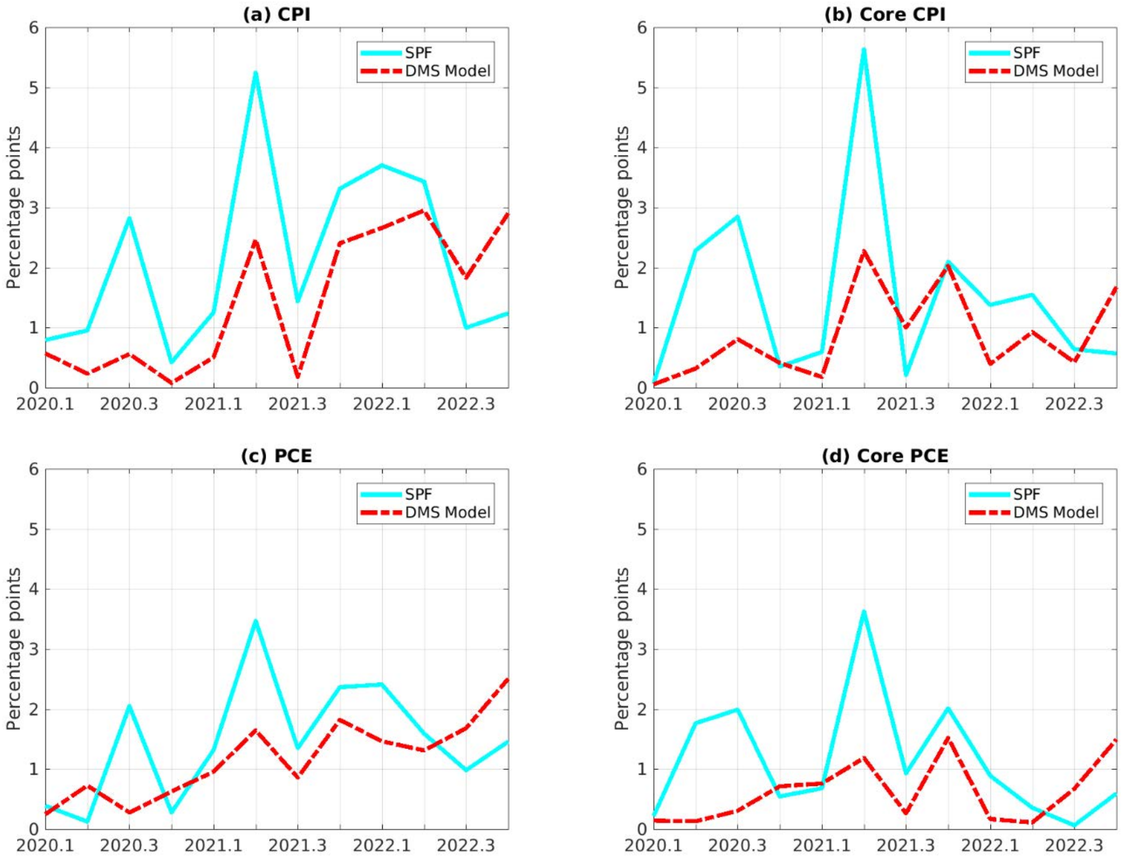

5.2. Density Nowcasting Accuracy Comparison

Most of the research on inflation nowcasting has focused on the production and evaluation of point nowcasts. There is growing recognition of the importance of the uncertainty around the point estimates, i.e., the range of potential inflation outcomes and their probability of occurring—the density nowcast. Knotek and Zaman (2023a) provide a comprehensive contribution to the production and evaluation of density nowcasts of inflation coming from various mixed-frequency models and their combinations, and using a variety of combination methods. Their results indicate that the density nowcasts coming from the DMS mixed-frequency model and from combinations of density nowcasts from a rich set of DMS model specifications are generally well-calibrated and tend to outperform other density nowcasting approaches. In the interest of space, we omit a full discussion of performance metrics and here focus on comparing density nowcasting performance for the DMS model and for a combination of DMS models with estimated density nowcasts from the SPF for our two sample periods over our four inflation measures.

For details on the computation of density nowcasts from the DMS model and DMS combinations, we refer the reader to Knotek and Zaman (2023a).[7] We construct the estimated

[7] The DMS combination we show in this paper is based on averaging the densities from 108 DMS model specifications. Knotek and Zaman (2023a) also consider a "grand combination" across mixed-frequency model

SPF density nowcasts using a normal distribution, whose mean is set equal to the median SPF point nowcast and whose variance is set to match the variance of the historical errors of the SPF point nowcasts over a short rolling window.[8] Prior work has documented that estimates of survey density nowcasts based on historical errors have historically been a good benchmark, especially for inflation.[9]

Table 3 reports the density accuracy results, based on log scores and relative log scores, from the out-of-sample density nowcasting horserace between our DMS model and the SPF and between the DMS combination and the SPF. The evaluation period runs from 2000:Q4 through 2022:Q4 for CPI inflation, and 2007:Q1 through 2022:Q4 for core CPI inflation, PCE inflation, and core PCE inflation. The results show that density nowcasts from the DMS model are substantially more accurate than the SPF-based benchmark density nowcasts, as indicated by significantly higher (less negative) log scores, and the gains are statistically significant. We obtain similar results for the DMS combination, whose density nowcasting performance is similar, based on log scores, to the performance of the single DMS model. For core CPI inflation and core PCE inflation, the density accuracy of the DMS model (and the DMS combination) is competitive with the SPF, suggesting that the normality assumption embedded in the estimated SPF density nowcasts is a reasonable strategy. Over the short evaluation period beginning with the pandemic when inflation volatility surged, the DMS model and the DMS combination produce large gains in density nowcast accuracy relative to the SPF-based benchmark.

classes that performs well for inflation density nowcasting. Because we focus on the DMS model in this paper, we do not use the grand combination. For more on the benefits of density combinations, see, e.g., Hall and Mitchell (2007).

[8] Results are robust to the use of mean SPF responses.

[9] Krüger, Clark, and Ravazzolo (2017) and Tallman and Zaman (2020) document competitive nowcasting performance, including a good calibration fit of the density nowcasts of inflation constructed through this simple procedure. The procedure's use of a short rolling window in computing the variance of the historical errors is a simple and convenient way to incorporate the changing variance of the density estimates. As discussed in Knotek and Zaman (2023a), the Federal Open Market Committee uses historical forecast errors to provide an estimate of the uncertainty surrounding the outlook in the Summary of Economic Projections (see Reifschneider and Tulip, 2019). While the SPF does provide some density forecasts by combining individual respondents' density forecasts, we favor the historical errors approach because Clements (2018) shows that the survey projection's second moments are inferior to simple statistical models. In addition, the SPF reports fixed-event density forecasts only for core PCE inflation and core CPI inflation, which limits their comparability to our results to the fourth quarter of each year.

Table 3: Density Nowcasting Comparisons with the Survey of Professional Forecasters

Nowcast Evaluation Sample: 2000:Q4–2022:Q4				
	CPI	Core CPI	PCE	Core PCE
SPF log score (SPF LS)	-2.070	-2.111	-1.594	-1.370
DMS Model log score	-1.459	-0.929	-1.170	-0.837
Relative, SPF LS – DMS Model LS	-0.610***	-1.182*	-0.424***	-0.534
DMS Combination log score	-1.452	-0.971	-1.145	-0.809
Relative, SPF LS – DMS Combination LS	-0.618***	-1.139*	-0.449***	-0.562
Nowcast Evaluation Sample: 2020:Q1–2022:Q4				
	CPI	Core CPI	PCE	Core PCE
SPF log score (SPF LS)	-4.150	-5.850	-2.546	-4.209
DMS Model log score	-2.131	-2.028	-1.722	-1.370
Relative, SPF LS – Model LS	-2.020	-3.822	-0.824	-2.840
DMS Combination log score	-2.080	-1.913	-1.841	-1.356
Relative, SPF LS – DMS Combination LS	-2.070	-3.937	-0.705	-2.853

Notes: The DMS model and DMS combination use real-time data available through the SPF survey date for each quarter. The SPF density nowcasts are based on historical forecast errors; see the text for details. The CPI exercise uses real-time data from 2000:Q4 through 2022:Q4. The core CPI, PCE, and core PCE exercises use real-time data from 2007:Q1 (the first available SPF estimate) through 2022:Q4. The DM type test reports the results of a test for equal predictive accuracy based on testing whether the constant term in the regression of the differences in the log score on the constant is statistically different from zero. For the 2020:Q1–2022:Q4 sample, we do not compute statistical significance because the length of the sample size is so short.

6. Conclusion

In recent years, the literature on inflation nowcasting has been rapidly developing. At its core, interest in inflation nowcasting comes from a realization that more accurate nowcasts, while valuable for their own sake, also help improve multi-step inflation forecasts. Progress on inflation nowcasting reflects a combination of advances in statistical and data science techniques to handle data sources at multiple frequencies and, in a growing number of cases, large and expanding amounts of non-traditional high-frequency data, such as online price indices, household and retailer scanner data, credit and debit card transaction data, and textual data. In this chapter, we briefly summarize the different methods being used to nowcast inflation.

We then illustrate the particular usefulness for inflation nowcasting of a specific parsimonious mixed-frequency model developed in Knotek and Zaman (2017) that features deterministic model switching, or DMS. Using real-time data, we evaluate and compare the DMS model's inflation nowcasting accuracy to two well-known surveys of professional forecasters. Across both a long sample spanning 1999–2022 and a short sample focusing on the period since the start of the pandemic, the DMS model has outperformed the surveys for nowcasting CPI inflation and PCE inflation, while the DMS model's nowcasting performance has been competitive with that of the surveys for core CPI and core PCE inflation.

Recent work using machine learning techniques opens the possibility of using many and different data series to nowcast inflation. For example, Clark et al. (2022) develop and apply one state-of-the-art machine learning method to 150 macroeconomic variables of different frequencies and document real-time nowcasting performance that is similar to that of the DMS model. Scraped online data also appear to hold promise for nowcasting and near-term forecasting of inflation. Future research should look to build on these frameworks and data sources to produce more accurate inflation nowcasts.

References

Aliaj, Tesi, Milos Ciganovic, and Massimiliano Tancioni. 2023. "Nowcasting inflation with Lasso-regularized vector autoregressions and mixed frequency data," *Journal of Forecasting* 42(3): 464-80. https://doi.org/10.1002/for.2944

Alvarez, Santiago E., and Sarah M. Lein. 2020. "Tracking Inflation on a daily basis," *Swiss Journal of Economics and Statistics* 156:18. https://doi.org/10.1186/s41937-020-00062-w

Amstad, Marlene, and Andreas M. Fischer. 2009. "Are Weekly Inflation Forecasts Informative?" *Oxford Bulletin of Economics and Statistics* 71(2): 237-52. https://doi.org/10.1111/j.1468-0084.2008.00523.x

Aparicio, Diego and Manuel I. Bertolotto. 2020. "Forecasting inflation with online prices," *International Journal of Forecasting* 36(2): 232-47. https://doi.org/10.1016/j.ijforecast.2019.04.018

Atkeson, Andrew, and Lee E. Ohanian. 2001. "Are Phillips Curves Useful for Forecasting Inflation?" Federal Reserve Bank of Minneapolis *Quarterly Review*, Winter, 2-11. https://doi.org/10.21034/qr.2511

Bańbura, Marta, and Michele Modugno. 2014. "Maximum Likelihood Estimation of Factor Models on Datasets with Arbitrary Pattern of Missing Data," *Journal of Applied Econometrics* 29(1): 133–60. https://doi.org/10.1002/jae.2306

Beck, Günter W., Kai Carstensen, Jan-Oliver Menz, Richard Schnorrenberger, and Elisabeth Wieland. 2023. "Nowcasting consumer price inflation using high-frequency scanner data: Evidence from Germany," Deutsche Bundesbank Discussion Paper No. 34/2023.

Breitung, Jörg, and Christoph Roling. 2015. "Forecasting Inflation Rates Using Daily Data: A Nonparametric MIDAS Approach," *Journal of Forecasting* 34(7): 588–603. https://doi.org/10.1002/for.2361

Carvalho, Igor da Silva. 2020. "Nowcasting CPI using online retail prices: Forecasting combination of dynamic factor models," FGV EMAp Dissertation. https://hdl.handle.net/10438/29681

Cavallo, Alberto, and Roberto Rigobon. 2016. "The Billion Prices Project: Using Online Prices for Measurement and Research," *Journal of Economic Perspectives* 30(2): 151-78. https://doi.org/10.1257/jep.30.2.151

Cimadomo, Jacopo, Domenico Giannone, Michele Lenza, Francesca Monti, and Andrej Sokol. 2022. "Nowcasting with large Bayesian vector autoregressions," *Journal of Econometrics* 231(2): 500-19. https://doi.org/10.1016/j.jeconom.2021.04.012

Clark, Todd E., Seton Leonard, Massimiliano Marcellino, and Philipp Wegmüller. 2022. "Weekly Nowcasting US Inflation with Enhanced Random Forests," Mimeo.

Clements, Michael P. 2018. "Are Macroeconomic Density Forecasts Informative?" *International Journal of Forecasting* 34(2): 181–98. https://doi.org/10.1016/j.ijforecast.2017.10.004

Faust, Jon, and Jonathan H. Wright. 2009. “Comparing Greenbook and Reduced Form Forecasts Using a Large Realtime Dataset,” *Journal of Business & Economic Statistics* 27(4): 468-79. https://doi.org/10.1198/jbes.2009.07214

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

Garcia, M. G., M.C. Medeiros, and G. F. R. Vasconcelos. 2017. “Real-time inflation forecasting with high-dimensional models: The case of Brazil,” *International Journal of Forecasting*, 33(3):679–93. https://doi.org/10.1016/j.ijforecast.2017.02.002

Garciga, Christian, Edward S. Knotek II, and Saeed Zaman. 2024. “Nowcasting U.S. Inflation: The Usefulness of Regional Fed Bank Surveys,” Manuscript (in preparation).

Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov. (2004) “The MIDAS Touch: Mixed Data Sampling Regression Models,” Unpublished manuscript.

Giannone, Domenico, Lucrezia Reichlin, and David Small. 2008. “Nowcasting: The Real-Time Informational Content of Macroeconomic Data,” *Journal of Monetary Economics* 55(4): 665–76. https://doi.org/10.1016/j.jmoneco.2008.05.010

Hall, Stephen G., and James Mitchell. 2007. “Combining Density Forecasts,” *International Journal of Forecasting* 23(1): 1–13. https://doi.org/10.1016/j.ijforecast.2006.08.001

Hendry, David F., and Kirstin Hubrich. 2011. “Combining Disaggregate Forecasts or Combining Disaggregate Information to Forecast an Aggregate,” *Journal of Business & Economic Statistics* 29(2): 216–27. https://doi.org/10.1198/jbes.2009.07112

Jaworski, Krystian. 2021. "Measuring food inflation during the COVID-19 pandemic in real time using online data: A case study of Poland," *British Food Journal* 123(13): 260-80. https://doi.org/10.1108/BFJ-06-2020-0532

Joseph, Andreas, Galina Potjagailo, Eleni Kalamara, Chiranjit Chakraborty, and George Kapetanios. 2022. “Forecasting UK inflation bottom up,” Bank of England Staff Working Paper No. 915.

Knotek, Edward S., II, and Saeed Zaman. 2017. “Nowcasting U.S. Headline and Core Inflation.” *Journal of Money, Credit and Banking* 49(5): 931–68. https://doi.org/10.1111/jmcb.12401

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

Knotek, Edward S., II, and Saeed Zaman. 2023a. “Real-Time Density Nowcasts of US Inflation: A Model Combination Approach.” *International Journal of Forecasting* 39(4): 1736-60. https://doi.org/10.1016/j.ijforecast.2022.04.007

Knotek, Edward S., II, and Saeed Zaman. 2023b. “A Real-Time Assessment of Inflation Nowcasting at the Cleveland Fed,” Federal Reserve Bank of Cleveland *Economic Commentary* No. 2023-06. https://doi.org/10.26509/frbc-ec-202306

Krüger, Fabian, Todd E. Clark, and Francesco Ravazzolo. 2017. "Using Entropic Tilting to Combine BVAR Forecasts With External Nowcasts," *Journal of Business & Economic Statistics* 35(3): 470–85. https://doi.org/10.1080/07350015.2015.1087856

Liang, Run, Fa Wang, and Jiawen Xu. 2021. "Nowcasting China's PPI inflation using low-frequency and mixed-frequency dynamic factor models." https://doi.org/10.13140/RG.2.2.25151.64165

Macias, Paweł, Damian Stelmasiak, and Karol Szafranek. 2023. "Nowcasting food inflation with a massive amount of online prices." *International Journal of Forecasting* 39(2): 809-826. https://doi.org/10.1016/j.ijforecast.2022.02.007

Mariano, Roberto S., and Suleyman Ozmucur. 2021. "Predictive Performance of Mixed-Frequency Nowcasting and Forecasting Models (with Application to Philippine Inflation and GDP Growth)," *Journal of Quantitative Economics* 19: 383–400. https://doi.org/10.1007/s40953-021-00276-6

Marsilli, Clément. 2017. "Nowcasting US Inflation Using a MIDAS Augmented Phillips Curve," *International Journal of Computational Economics and Econometrics* 7(1/2): 64. https://doi.org/10.1504/IJCEE.2017.080649

McCracken, Michael W., Michael T. Owyang, and Tatevik Sekhposyan. 2021. "Real-Time Forecasting and Scenario Analysis Using a Large Mixed-Frequency Bayesian VAR," *International Journal of Central Banking* 18(5): 327-67.

Modugno, Michele. 2013. "Now-Casting Inflation Using High Frequency Data." *International Journal of Forecasting* 29(4): 664–75. https://doi.org/10.1016/j.ijforecast.2012.12.003

Monteforte, Libero, and Gianluca Moretti. 2013. "Real-Time Forecasts of Inflation: The Role of Financial Variables," *Journal of Forecasting* 32(1): 51–61. https://doi.org/10.1002/for.1250

Powell, Ben, Guy Nason, Duncan Elliott, Matthew Mayhew, Jennifer Davies, and Joe Winton. 2018. "Tracking and Modelling Prices Using Web-Scraped Price Microdata: Towards Automated Daily Consumer Price Index Forecasting," *Journal of the Royal Statistical Society Series A: Statistics in Society* 181(3): 737–56 https://doi.org/10.1111/rssa.12314

Reifschneider, David, and Peter Tulip. 2019. "Gauging the Uncertainty of the Economic Outlook Using Historical Forecasting Errors: The Federal Reserve's Approach," *International Journal of Forecasting* 35(4): 1564–82. https://doi.org/10.1016/j.ijforecast.2018.07.016

Schorfheide, Frank, and Dongho Song. 2015. "Real-Time Forecasting With a Mixed-Frequency VAR," *Journal of Business & Economic Statistics* 33(3): 366-80. https://doi.org/10.1080/07350015.2014.954707

Schnorrenberger, Richard, Aishameriane Schmidt, Guilherme Valle Moura. 2024. "Harnessing Machine Learning for Real-Time Inflation Nowcasting," Manuscript.

Soybilgen, Barış, M. Ege Yazgan, and Hüseyin Kaya. 2023. "Nowcasting Turkish Food Inflation Using Daily Online Prices," *Journal of Business Cycle Research* 19: 171-90. https://doi.org/10.1007/s41549-023-00084-2

Tallman, Ellis W., and Saeed Zaman. 2020. "Combining Survey Long-Run Forecasts and Nowcasts with BVAR Forecasts Using Relative Entropy," *International Journal of Forecasting* 36(2): 373–98. https://doi.org/10.1016/j.ijforecast.2019.04.024

Zheng, Tingguo, Xinyue Fan, Wei Jin, and Kuangnan Fang. 2023. "Words or Numbers? Macroeconomic Nowcasting with Textual and Macroeconomic Data," *International Journal of Forecasting*. https://doi.org/10.1016/j.ijforecast.2023.05.006