Inflation Nowcasting: Frequently Asked Questions

These questions and answers accompany the technical working paper "Nowcasting U.S. Headline and Core Inflation" by Edward S. Knotek II and Saeed Zaman. See the paper for more details and references.

What is "nowcasting"?

Nowcasting is a combination of two terms: "now" and "forecasting." Economic indicators are often released with a considerable delay, but there is usually intense interest in knowing what a given indicator is before its official release. Nowcasting refers to the prediction of the present, the very recent past, and the very near future—in essence, trying to fill in that missing data point with an accurate estimate. Often, nowcasts take advantage of higher-frequency data that are informative about some other economic indicator to predict what the next release of the indicator of interest will say.

Why is inflation nowcasting useful?

Inflation is a topic of perennial interest because it influences the behavior and plans of consumers, businesses, financial markets, pension funds, governments—in essence, everyone in an economy. The Federal Reserve pays close attention to inflation because of its impact on the economy and because the Federal Open Market Committee has been explicitly tasked by Congress through the Federal Reserve Act to promote the goals of maximum employment and price stability, the so-called "dual mandate." When making decisions, consumers and businesses may have to forecast the inflation rate far into the future. For example, if a consumer is thinking about taking out a loan, it helps to know how quickly wages and prices will be rising during the life of the loan—after all, it will be much easier to service the loan with stronger wage and price growth. Unfortunately, inflation tends to be difficult to predict accurately. But some recent research finds that forecasts of inflation in the future can be improved by having more accurate estimates of near-term inflation, or nowcasts. These nowcasts serve as an important jumping-off point for thinking about where inflation will be in the future.

Which inflation measures do you nowcast?

We generate nowcasts for four inflation measures: (1) inflation in the price index for personal consumption expenditures, more commonly referred to as PCE inflation; (2) inflation in the PCE price index excluding food and energy, also known as core PCE inflation; (3) inflation in the consumer price index (CPI); and (4) inflation in the CPI excluding food and energy, also known as core CPI inflation. The model reports seasonally adjusted, month-over-month inflation rates in these four measures (expressed as nonannualized percent changes) and quarterly inflation rates in these four measures (expressed at seasonally adjusted annualized rates, or SAAR). The model also reports year-over-year inflation rates in these four measures (based on nonseasonally adjusted data for CPI inflation and core CPI inflation and seasonally adjusted data for PCE inflation and core PCE inflation).

What is so special about your inflation nowcasts?

Using historical comparisons from 1999 to 2015, inflation nowcasts from the Cleveland Fed's model have been highly accurate. Our model's nowcasts easily outperform a variety of nowcasts from alternative statistical models, especially over the course of a month or quarter. But perhaps even more notable is the model's performance compared with the best available benchmarks: nowcasts from surveys of professional forecasters. The model's nowcasts for CPI inflation tend to be more accurate than the consensus (average) nowcasts from the Blue Chip Economic Indicators survey. Its nowcasts for CPI and PCE inflation also tend to be more accurate than the median nowcasts from the Federal Reserve Bank of Philadelphia's Survey of Professional Forecasters (SPF). These results are somewhat surprising, because professional forecasters can and do use a range of models and judgment to capture the special factors that affect near-term inflation trends, and there is additionally evidence of the wisdom of the masses (or, more formally, the law of large numbers) when making forecasts. Meanwhile, the model's nowcasts for core inflation tend to be just as accurate as those in the SPF, but the model is available at a daily frequency, while the SPF is released only once per quarter.

Finally, we also compare the model's inflation nowcasting accuracy with publicly available nowcasts from the Greenbook (the in-depth analysis of economic conditions produced by the Federal Reserve Board of Governors staff for Federal Open Market Committee meetings). While Greenbook nowcasts tend to be highly accurate, they are only released to the public with a five-year lag. The Cleveland Fed's model nowcasts have historically tended to be quite similar to those in the Greenbook, but they are available in real time, without delay.

Are the estimates updated daily?

Yes, every business day around 10:00 a.m. Eastern time.

What is the timing of the CPI and PCE inflation data releases?

In the United States, the Bureau of Labor Statistics (BLS) typically releases the CPI for a given month around the middle of the following month (for example, the January CPI is released around mid-February). The Bureau of Economic Analysis (BEA) typically releases the other major measure of consumer prices, the PCE price index, around the end of the following month (for example, the January PCE price index is released around the end of February, after the CPI for January has been released).

Which high-frequency data in particular are used to produce the inflation nowcasts?

Daily Brent crude spot oil prices and weekly retail gasoline prices.

How many data series are used to produce the inflation nowcasts?

The nowcasts are produced using ten data series. Eight data series are monthly. Five of these monthly series come from the BLS: the CPI, the core CPI, the CPI for food, the CPI for food at home, and the CPI for gasoline. (All of the CPI series we use are seasonally adjusted.) Three of the monthly series come from the BEA: the PCE price index, the core PCE price index, and the PCE price index for food and beverages purchased for off-premises consumption. One data series is weekly: retail gasoline prices, released each Monday by the Energy Information Administration. One data series is daily: Brent crude spot oil prices, from the *Financial Times* or the Energy Information Administration.

Could you very briefly describe how the nowcasts are made?

In a nutshell, there are four parts to the model. The first part nowcasts (or, perhaps more accurately, forecasts) core inflation. To do so, we find that estimates based entirely on the recent past do a fairly good job. The second part forecasts food price inflation. Again, we primarily use the recent past to forecast future food price inflation. The third part nowcasts gasoline price inflation. We use a combination of current gasoline prices and current oil prices, under the assumption that today's oil prices are informative about where gasoline prices are likely to head in the future, and then we seasonally adjust the data. Finally, the fourth part combines the nowcasts and forecasts of core inflation, food price inflation, and gasoline price inflation to come up with nowcasts of inflation in either the CPI or PCE price index.

Why do the core inflation nowcasts change so little and so infrequently? Why do the headline inflation nowcasts move around so much and so often?

Nowcasts change when new data arrive that are different than what was expected. This can be either because of the frequency with which new data arrive or the volatility of that new data.

The core inflation nowcasts use very few data sources, which arrive infrequently; therefore, they only change infrequently. Nowcasts of core CPI inflation are based only on past core CPI inflation; thus, the nowcasts *can only change* when new CPI data are released (or when past data are revised). However, if the data come in exactly as expected, then the core CPI inflation nowcasts will not actually change. Nowcasts of core PCE inflation are primarily based on past core PCE inflation, but they can also be affected by the availability of CPI data: if we have one more core CPI observation than core PCE for some month t, then we map month t's core CPI observation to core PCE in month t. Thus, core PCE nowcasts *can only change* when we get new data on either the CPI or the PCE price index—but again, if the data come in exactly as expected, they actually will not change.

Core inflation also tends to be relatively stable, so revisions to core inflation nowcasts can be small when they do occur. By contrast, gasoline prices play an important role in headline CPI and PCE inflation nowcasts, and gasoline price nowcasts depend in turn on oil prices. Oil prices move around almost every day, and they can be quite volatile and thus difficult to predict. The headline inflation nowcasts can change when new CPI or PCE price index data are released; but between releases, they often move around based on incoming oil and gasoline price data.

How accurate are your nowcasts?

Nowcasts are a type of forecast and thus are always subject to considerable uncertainty and imprecision. After all, we are using a very small number of series (ten) and simple statistical techniques to make informed guesses about economic indicators before they are released. By contrast, the official data releases are usually the product of tens of thousands of observations, powerful computational techniques, and a very large number of hours worked. **On average**, the model's nowcasts tend to be more precise than a range of competing nowcasts, but the outperformance is not absolute. The accuracy of each individual nowcast depends on a range of factors, including the measure under consideration and the point in the month or quarter that the nowcast is being made. **On average**, nowcasts made with more information later in a month or quarter tend to be more accurate. But shocks are ever-present that can push inflation away from the nowcasted value, especially if they occur to inflation components that we do not try to model.

How do I compare the model's historical nowcasts with the actual data?

The charting tool on the nowcasting website allows you to plot a sequence of nowcasts for previous months or quarters along with the actual inflation release (for simplicity, we display the first available data release corresponding to the quarter or month being nowcasted). This allows the user to see how close or far the model's nowcasts were compared with what was actually released by the statistical agency (the BLS or the BEA). Typically, nowcasts early in the month or quarter are less precise than nowcasts made later in the month or quarter, because the early nowcasts are based on less information than later nowcasts.

What's coming next?

We are constantly seeking ways to improve this website and its functionality. One item on the to-do list is to provide an option that includes historical nowcasting error bands in the charts that vary with the exact information set available at a given point in time. These bands would start out wide at the beginning of a month or quarter and then shrink over time as more information accumulates.

Whom should I contact if I have any further questions, comments, or suggestions?

We welcome your questions, comments, and suggestions for ways to make this website more useful and to enhance our list of FAQs. While we cannot respond to every query, we plan to make updates to the site and the FAQs, and your input would be much appreciated. Email us at mailto:inflation.central@clev.frb.org.

Is the same model used in the nowcasting working paper and the nowcasting website (technical)?

Yes, the basic model is the same, but there is one difference. The paper uses a single series on food inflation (specifically, the CPI for food) in constructing both CPI and PCE inflation estimates. On the website, the PCE inflation estimates are constructed using a different measure of food inflation—inflation in the PCE price index for food and beverages purchased for off-premises consumption. This change is methodologically more consistent with how the PCE price index is constructed, and it would have been great to include it in the paper. But unfortunately, a long real-time history of this particular series is not available.

After deciding to include the PCE price index for food and beverages purchased for off-premises consumption, we added one more series to the model—the CPI for food at home. Monthly inflation in the CPI for food at home has historically been similar to monthly inflation in the PCE price index for food and beverages purchased for off-premises consumption. As a result, the CPI for food at home can provide some additional information at times about the PCE price index for food and beverages purchased for off-premises consumption.

Could you walk through the model (technical)?

For more on the model, see the working paper. Here, we very quickly go through the key elements of the model used to compute the inflation nowcasts, focusing more on practical implementation than technical details.

Let T denote quarters and t denote months, such that within some quarter T there are months t=1, 2, 3. Quarterly inflation π_T is usually measured at seasonally adjusted annualized rates as $\pi_T = 100[(P_T / P_{T-1})^4 - 1]$, where $P_T = (1/3)(P_{T,t=1} + P_{T,t=2} + P_{T,t=3})$. Monthly inflation is expressed in nonannualized terms as $\pi_t = 100(P_t / P_{t-1} - 1)$. Thus, if we have nowcasts or forecasts of monthly inflation rates, we can fill in the missing price levels for the months within a given quarter to compute quarterly inflation.

We denote nowcasts or forecasts for a variable x by using the notation $f(x)$. In what follows, we set J=12, τ=24, and τ_L=60.

Core CPI inflation. Suppose we have monthly core CPI inflation data through month $t-1$, $\pi_{t-1}^{\text{Core CPI}}$. We recursively forecast monthly core CPI inflation using J-month moving averages, so the forecast for month t is $f(\pi_t^{\text{Core CPI}}) = (1/J)\sum_{j=1,\dots,J} \pi_{t-j}^{\text{Core CPI}}$, and forecasts for months t+1, t+2, … follow.

Core PCE inflation. Suppose we have monthly core PCE inflation data through month $t-1$, $\pi_{t-1}^{\text{Core PCE}}$. If we have core CPI inflation through month t (*data*, not a forecast) but are missing core PCE inflation in month t, we bridge from core CPI inflation to core PCE inflation. That is, using a rolling window of length τ, we estimate the regression model

$\pi_t^{\text{Core PCE}} = \beta_0 + \beta_1 \pi_t^{\text{Core CPI}} + e_t$, then use the estimated coefficients and the actual reading on $\pi_t^{\text{Core CPI}}$ to obtain $f(\pi_t^{\text{Core PCE}})$. If we do not have more core CPI inflation data than we have core PCE inflation data, or if we've already bridged all the available core CPI inflation data, then we recursively forecast monthly core PCE inflation $f(\pi_{t+k}^{\text{Core PCE}})$, $k \geq 0$, using J-month moving averages.

CPI food inflation. Suppose we have monthly CPI food inflation through month $t-1$, $\pi_{t-1}^{\text{Food CPI}}$. As with core inflation, we recursively forecast monthly CPI food inflation $f(\pi_{t+k}^{\text{Food CPI}})$, $k \geq 0$, using J-month moving averages.

PCE food inflation. Suppose we have monthly inflation in the PCE price index for food and beverages purchased for off-premises consumption through month $t-1$, $\pi_{t-1}^{\text{Food Off Premises PCE}}$. If we have monthly inflation in the CPI for food at home (not just food) in month t, $\pi_t^{\text{Food at Home CPI}}$, we bridge it over to the PCE-equivalent concept by setting $f(\pi_t^{\text{Food Off Premises PCE}}) = \pi_t^{\text{Food at Home CPI}}$. If we do not have more CPI data than PCE data, or if we've already bridged it, we recursively forecast $f(\pi_{t+k}^{\text{Food Off Premises PCE}})$, $k \geq 0$, using J-month moving averages.

Gasoline inflation. Letting $P_t^{\text{Gasoline (NSA)}}$ denote the average of available weekly gasoline prices within a month t before any seasonal adjustment, we compute monthly gasoline inflation $\pi_t^{\text{Gasoline (NSA)}} = 100(P_t^{\text{Gasoline (NSA)}} / P_{t-1}^{\text{Gasoline (NSA)}} - 1)$. We seasonally adjust the series to make it useful: if $\pi_{t-j}^{\text{Gasoline CPI}}$ is monthly inflation in the (seasonally adjusted) CPI for gasoline in month $t-j$, we define the seasonal factor in month t to be $sf_t = (1/3)\sum_{j=1\text{ year, 2 years, 3 years}} (\pi_{t-j}^{\text{Gasoline (NSA)}} - \pi_{t-j}^{\text{Gasoline CPI}})$ and then use it to derive our measure of seasonally adjusted monthly gasoline inflation: $\hat{\pi}_t^{\text{Gasoline}} = \pi_t^{\text{Gasoline (NSA)}} - sf_t$. To augment our gasoline price data, let P_t^{Oil} be the average of available daily oil prices within a month t. Mechanically, we extend the monthly oil price series by one observation by setting P_{t+1}^{Oil} to the last available daily observation. Using a rolling window of length τ_L, we estimate the first-stage regression $P_t^{\text{Gasoline (NSA)}} = \alpha + \beta P_t^{\text{Oil}} + e_t$; let $g(P_t^{\text{Gasoline (NSA)}})$ denote the predicted gasoline prices based on the regression coefficients and the actual monthly oil prices. Next, we estimate the second-stage regression $[P_t^{\text{Gasoline (NSA)}} - g(P_t^{\text{Gasoline (NSA)}})] = a[P_{t-1}^{\text{Gasoline (NSA)}} - g(P_{t-1}^{\text{Gasoline (NSA)}})] + e_t$, using the same rolling window. We match the length of $g(P_{t+1}^{\text{Gasoline (NSA)}})$ to the length of the oil price series using the first-stage regression coefficients, and then derive forecasts of $f(P_{t+1}^{\text{Gasoline (NSA)}})$ based on the second-stage regression coefficient. Finally, we compute monthly inflation in the non-seasonally adjusted gasoline prices and then seasonally adjust the data as described above to produce a set of $f(\pi_{t+k}^{\text{Gasoline}})$, $k \geq 0$. Due to release lags, we typically have one or two more months of gasoline inflation nowcasts or forecasts than we have CPI and PCE inflation data.

CPI inflation. Suppose we have monthly CPI inflation data through month $t-1$, π_{t-1}^{CPI}. We estimate the regression model $\pi_t^{\text{CPI}} = \beta_0 + \beta_1 \pi_t^{\text{Core CPI}} + \beta_2 \pi_t^{\text{Food CPI}} + \beta_3 \pi_t^{\text{Gasoline}} + e_t$ using a rolling window of length τ. In nowcasting months t, t+1, etc., if we have an estimate of $f(\pi_{t+k}^{\text{Gasoline}})$ for some $k \geq 0$, we use it along with the coefficients we just estimated and the estimates of $f(\pi_{t+k}^{\text{Food CPI}})$ and $f(\pi_{t+k}^{\text{Core CPI}})$ produced earlier and to derive the nowcast of $f(\pi_{t+k}^{\text{CPI}})$. If we do not have an estimate of $f(\pi_{t+k}^{\text{Gasoline}})$, we recursively forecast $f(\pi_{t+k}^{\text{CPI}})$, $k \geq 0$, using J-month moving averages.

PCE inflation. Suppose we have monthly PCE inflation data through month $t-1$, π_{t-1}^{PCE}. If we have monthly CPI inflation *data* (not a forecast) through month t, we bridge it to nowcast PCE inflation by estimating the regression model $\pi_t^{\text{PCE}} = \beta_0 + \beta_1 \pi_t^{\text{CPI}} + e_t$ using a rolling window of length τ, then use the estimated coefficients and the actual reading on π_t^{CPI} to obtain $f(\pi_t^{\text{PCE}})$. Going forward for some month $t+k$, if we have an estimate of $f(\pi_{t+k}^{\text{Gasoline}})$, we first estimate the regression model $\pi_t^{\text{PCE}} = \beta_0 + \beta_1 \pi_t^{\text{Core PCE}} + \beta_2 \pi_t^{\text{Food Off Premises PCE}} + \beta_3 \pi_t^{\text{Gasoline}} + e_t$ using a rolling window of length τ, and then use those estimated coefficients along with estimates of $f(\pi_{t+k}^{\text{Gasoline}})$, $f(\pi_{t+k}^{\text{Food Off Premises PCE}})$, and $f(\pi_{t+k}^{\text{Core PCE}})$ produced earlier to derive the nowcast of $f(\pi_{t+k}^{\text{PCE}})$. Finally, if we do not have an estimate of $f(\pi_{t+k}^{\text{Gasoline}})$, we recursively forecast $f(\pi_{t+k}^{\text{PCE}})$ using J-month moving averages.