Simple Ways to Forecast Inflation: What Works Best?

Brent H. Meyer and Mehmet Pasaogullari

There are many ways to forecast the future rate of inflation, ranging from sophisticated statistical models involving hundreds of variables to hunches based on past experience. We generate a number of forecasts using a simple statistical model and an even simpler estimating rule, adding in various measures thought to be helpful in predicting the course of inflation. Then we compare their forecast accuracy. We find that no single specification outperforms all others over all time periods. For example, the median and 16 percent trimmed-mean measures outperform all other specifications during the 1990s, and survey-based inflation expectations seem to do better during volatile periods.

Just about everybody pays attention to inflation and wonders when prices are going up, and by how much. Households and businesses need estimates of future prices to make well-informed decisions. Policymakers, whose job is to aid in those decisions by promoting stable prices, need accurate forecasts in order to monitor inflation and make course corrections when necessary.

To get a glimpse into the probable future, one can use a statistical model. In this *Commentary*, we investigate a few simple versions of these to forecast Consumer Price Index (CPI) inflation, along with some even-simpler rules of thumb. We start with univariate forecasting techniques. Then, in an effort to improve these forecasts, we investigate the forecasting properties of other variables that are thought to affect inflation—economic slack, underlying inflation, and survey measures of expected inflation. We compare the forecast accuracy of a number of different specifications with variants of all of these.

We find that there isn't just one dominant specification that outperforms all other forecast models in every time period. Also, over the past ten years, simple statistics—such as annual inflation rates in alternative price-change measures and inflation expectations obtained from surveys—turn out to be more informative than the statistical models we tested.

A Starting Point

Inflation tends to be a relatively persistent process, which means that current and past values should be helpful in forecasting future inflation. Applying that intuition, we construct two basic models that exploit information embedded in past values of CPI inflation. Each uses a different technique to forecast CPI inflation over the year ahead: One is based on regression analysis and the other is based on the naïve specification made popular by Atkeson and Ohanian (2001). Later, we add and switch out different variables and different ways of measuring these variables to get other specifications.

The first specification is a regression that forecasts one-year-ahead CPI inflation using lags of the CPI (specifically, past values of the quarterly annualized percent change in the CPI).[1] We estimate this regression in a recursive manner, starting with a sample that includes 40 quarters of data and adds an additional data point to the sample in each successive quarter.[2] This approach is equivalent to saying that the next year's inflation is a function of all past values of inflation up to 4 quarters before. The regression analysis figures out the parameters of that function.

Table 1. Accuracy of Forecasts Based on Past Inflation Only (in Root Mean Squared Errors)

	1960:Q2–1969:Q4	1970:Q1–1979:Q4	1980:Q1–1989:Q4	1990:Q1–1999:Q4	2000:Q1–2010:Q2	1960:Q2–1983:Q4	1984:Q1–2006:Q4	1984:Q1–2010:Q2	1995:Q1–2010:Q2	1960:Q2–2010:Q2
Regression with CPI inflation	0.98	2.18	2.31	0.96	2.39	2.00	1.27	1.79	2.00	1.89
Naïve forecast with CPI inflation	0.85	2.55	2.28	0.93	1.98	2.18	1.11	1.52	1.67	1.86

Note: The specification with the lowest root mean squared error (RMSE) in each time period is highlighted in green.

Table 2. Accuracy of Forecasts Based on Alternative Specifications (in Root Mean Squared Errors)

	1960:Q2–1969:Q4	1970:Q1–1979:Q4	1980:Q1–1989:Q4	1990:Q1–1999:Q4	2000:Q1–2010:Q2	1960:Q2–1983:Q4	1984:Q1–2006:Q4	1984:Q1–2010:Q2	1995:Q1–2010:Q2	1960:Q2–2010:Q2
Economic activity										
Regression with annualized growth rate of real GDP	0.93	2.09	2.15	0.88	2.39	1.85	1.21	1.78	2.00	1.82
Different measures of inflation										
Regression with core CPI inflation		3.12	2.51	1.00	1.46		1.12	1.29	1.28	
Naïve forecast with core CPI inflation	0.76	3.27	2.51	0.93	1.47	2.64	1.11	1.25	1.24	2.03
Regression with median CPI inflation			2.48	1.22	1.57		1.37	1.51	1.48	
Naïve forecast with median CPI inflation		2.91	2.02	0.78	1.50		0.98	1.19	1.29	
Regression with 16% trimmed-mean CPI inflation			2.34	1.11	1.66		1.23	1.47	1.47	
Naïve forecast with 16% trimmed-mean CPI inflation		2.84	2.07	0.77	1.57		0.97	1.21	1.33	
Inflation expectations										
Regression with UM inflation expectations				0.92	1.85				1.59	
Regression with SPF inflation expectations					1.37				1.21	
Naïve forecast with UM inflation expectations			1.51	0.88	1.74		0.92	1.32	1.49	
Naïve forecast with SPF inflation expectations				0.81	1.37		0.97	1.15	1.22	

Notes: For simplicity, we report only the root mean squared errors (RMSEs) from the highest-performing economic activity specification, given that the relative performance of this set of variables diminished after the 1980s. The specification with the lowest RMSE in each time period is highlighted in green.

The second specification forecasts one-year-ahead CPI inflation using a naïve specification, in which the forecast over the year ahead is simply the past four-quarter growth rate in the CPI. For example, the four-quarter growth rate in the CPI stands at 1.2 percent through the third quarter of 2010. Using the naïve technique, 1.2 percent becomes our forecast for inflation over the next four quarters (through the third quarter of 2011). This approach is equivalent to saying that inflation over the upcoming year is most likely to be what it was in the past year up to that point.

Because it is possible that the underlying inflation process has changed over time, we test the forecasting performance of these models over a variety of time periods. We first examine forecast accuracy by decade, starting in 1960. Next, since monetary policy changed in the 1980s, we break the data series into two time periods, one pre-1983 and one post-1983. The inflation process may have been altered following a period of disinflation in the early to mid-1980s—commonly referred to as the "Volcker-era" disinflation—after which both inflation and inflation expectations became less volatile. We also break out the 1984–2006 time period (excluding the last four years) to see what has happened to forecast accuracy over the most recent period, which includes the 2007–2009 recession. The last time period we examine is from 1995–2010, which allows us to examine measures of inflation expectations as predictors of inflation for time periods over which we have limited data.

To compare the accuracy of these specifications, we compute the root mean squared error (RMSE) statistic, a measure of forecast error, for each. A RMSE of 0 indicates a perfect forecasting performance, and positive values reflect deviations between the forecasted values and the realized values. The higher the RMSE, the higher the deviation between the forecasted values and the realized values on average. Table 1 reports the forecast accuracy for our backward-looking regression and naïve specification.

There are a couple of patterns to note from table 1. First, neither model consistently outperforms the other across different time periods, although the naïve method definitely has the upper hand. Second, the forecasting performance of these specifications, which depend only on past inflation, varies appreciably across different time periods and has deteriorated over the last four years.

The implication of this recent deterioration in forecasting performance is that inflation seems to be explained to a lesser extent by past inflation than it used to be. An explanation for that deterioration could be that the underlying inflation process has changed. Carlstrom, Fuerst, and Paustian (2007) suggest, for example, that inflation has become less persistent. The loss of explanatory power by lagged inflation could also be tied, in part, to the energy price shock in mid-2008. By the third quarter of 2008, the four-quarter growth rate in the CPI had jumped up to 5.3 percent (a 17-year high), only to fall below 0.0 percent a mere two quarters later. A dramatic swing like that had not been experienced in the recent past, and it probably contributed to a larger forecast error, since backward-looking measures could not have accounted for such extreme variation.

Can We Do Better?

We attempt to improve on the forecasts that depend only on past inflation by incorporating three other types of information into the basic models. First, we add different measures of economic activity into the regression.[3] The approach is common, and this kind of specification is sometimes referred to as a Phillips curve.

Second, we investigate measures of underlying or core inflation (such as the median CPI), statistics which attempt to lessen some of the volatility in the headline CPI, thereby extracting a more precise inflation trend.

Finally, we see if survey measures of inflation expectations have any useful predictive content.

The measures of economic activity that we add are those that are thought to improve inflation forecasts: real GDP, unemployment, industrial production, manufacturing production, and capacity utilization. Measures of economic activity are thought to be useful in forecasting inflation, with the underlying pace of expansion, or robustness of growth putting pressure on prices. For example, when output is rising at a fast pace or the unemployment rate is relatively low, prices in general often rise, leading to higher rates of inflation. Conversely, periods when growth is slow or the unemployment rate is high tend to be disinflationary.

One way to exploit this relationship when forecasting is to use "gaps"—or deviations from a trend—because they can indicate exceptional variation. Another is to look at growth rates, where faster growth rates tend to be associated with higher inflation rates and vice versa. We use deviations from trend and growth rates of the economic variables in the regressions.[4]

Measures of underlying inflation may more accurately uncover trend inflation than the headline measure, so they may be more useful in forecasting. In any given period there can be a substantial amount of noise in the overall CPI, which can arise from a myriad of issues including seasonal adjustment issues, measurement problems, and idiosyncratic price changes (such as excise tax increases). This noise can obscure the signal from past headline inflation and may lead to poor forecasts. Economists and forecasters frequently try to eliminate that noise in an attempt to uncover the underlying inflation trend, what is often called core inflation.

We test three measures of core CPI inflation in place of the past headline inflation. The first, and most common measure of underlying inflation, the core CPI, is the CPI

minus food and energy. As the name implies, this measure excludes two of the most historically volatile components—food and energy prices. However, Bryan and Cecchetti (1993) argue that volatility (noise) can arise from any component during a given period (usually a month), and by trimming (or excluding) the most volatile monthly price changes, a clearer signal of underlying inflation can be uncovered. Two such trimmed-mean statistics are the median CPI and the 16 percent trimmed-mean CPI, and we test these as well.

We incorporate the different measures of core inflation into both of our model types. In the regression specification, we use current and past values of the core inflation measures as explanatory variables.[5] We create alternative naïve specifications by incorporating the trailing four-quarter percent change in each core inflation measure.

Finally, we investigate two readily available survey measures of one-year-ahead inflation expectations, the median expectation from the University of Michigan's Survey of Consumers (hereafter UM) and the median expectation for CPI inflation from the Federal Reserve Bank of Philadelphia's Survey of Professional Forecasters (SPF).

These measures are intriguing as forecasting tools, as it is highly plausible that, given wage and price stickiness, individuals embed expectations of future inflation into their price-setting and decision-making behavior today. In fact, if you've read or listened to a Federal Reserve official lately, chances are you've probably heard something to the effect of "…inflation expectations matter." Indeed, central bankers' sensitivity to inflation expectations seems warranted, as it is theoretically possible that expectations can be self-fulfilling prophecies. However, we are mainly interested in the forecasting properties of these measures here.

Just as we do for the alternative measures of inflation, we incorporate the two measures of inflation expectations into our two basic models. The regression uses the current median expectation as an explanatory variable, and the naïve specifications simply assume the median expectation will be the future year-ahead rate of inflation.[6]

Which Gives the Best Forecast?
Table 2 details the RMSEs for the different specifications, reporting only the highest-performing measures from each type of additional variable (economic activity, core inflation, and inflation expectations) to save space.

Economic Activity. As was the case with the specifications that just depended on past inflation, the forecasting performance of activity measures varies markedly over time. Moreover, the inclusion of activity measures, particularly the annualized quarterly growth rate in real GDP, seemed to improve upon the specifications that depend just on past inflation only from the 1970s to the 1990s. For the 1970s,

activity measures are the best predictors among the remaining alternatives we examine. This result may be related to the relatively high volatility of that time period and may suggest that the loss of predictive power was due to the relative stability of the economy over the past 30 years or so (up until recently).

Core Inflation Measures. Generally, it appears that these measures of underlying inflation are useful when forecasting inflation. Moreover, it seems that the predictive ability of this set of specifications has improved recently relative to the techniques reported in table 1. During the 1980s the lowest RMSE is only 10 percent better than the specifications that just use past values of headline inflation, while over the past 10 years the lowest RMSE is roughly 25 percent better. Moreover, just paying attention to less noisy measures of underlying inflation, like the naïve forecasts from the four-quarter percent change in the median and 16 percent trimmed-mean CPI measures, tends to be more useful than a Phillips curve specification that includes lagged inflation and economic activity measures (except during the 1970s).[7] Also, it looks as if the naïve specifications of this type generally outperform the forecasts stemming from the estimated regressions.

Inflation Expectations. Similar to Ang, Bekaert, and Wei (2007), we find that survey measures of inflation expectations tend to outperform other, more standard, inflation-forecasting specifications. The most striking result is the relative performance of the naïve forecast constructed with the median expectation from the University of Michigan's Survey of Consumers. Over the 1980s, it had a RMSE of 1.51, roughly 25 percent better than the next closest RMSE of 2.02, which belongs to the naïve forecast from the median CPI. We find the timing of this superior performance particularly interesting, because it corresponds to the period during which the Fed started and pursued a strategy to reduce the rate of inflation.

Expectations measures might pick up on the 1980s disinflation better than other approaches because statistical models can't detect new directions that break with past relationships. Since statistical models exploit past relationships between variables, their forecasts are more persistent. At significant inflection points, these forecasts are persistently wrong, whereas individuals are free to use judgment to discern if those relationships have changed. In the early part of the 1980s, inflation was running in the double digits, and the Federal Open Market Committee, with Chairman Volcker at the helm, increased the federal funds rate to, at one point, nearly 20 percent in order to tame inflation. Individuals appear to have adjusted to the change in policy and started to expect lower, more stable inflation. Statistical methods could not detect this policy change as quickly, leading to the better forecasting performance of expectations measures over that time period.

Median year-ahead inflation expectations from the Survey of Professional Forecasters also tend to forecast relatively well, compared to other specifications we tested…and thank goodness! It might be a little embarrassing if professional forecasters, who not only employ more sophisticated inflation forecasting models than we've investigated here but also have years of experience to shape their judgment, ended up with relatively poor forecasts. That said, the naïve UM forecast did outperform the naive SPF forecast from 1984 to 2006, though that difference was negligible. However, inflation was relatively stable over that time period, and that could be driving those results. Perhaps more importantly, by including the most recent data that encompass the 2007-09 recession (a relatively volatile time period for headline CPI inflation), the RMSE for the naïve SPF forecast is about 15 percent more accurate than the forecast from the naïve University of Michigan specification.

Conclusion

While this exercise was rather simple, it did yield some interesting results. First, there doesn't seem to be a single specification that outperforms all others over all time periods. Second, "naïve" specifications (other than the naïve forecast using the headline CPI) seem to perform well compared to simple statistical models, and during some periods, forecast significantly better. For example, the naïve forecasts from the median and 16 percent trimmed-mean measures outperform all other specifications during the 1990s.

Finally, inflation expectations appear to forecast future inflation rather well, yielding the lowest RMSE in every time period for which we have data but one. Over some time periods, especially during the 1980s, a significant inflection point for inflation, expectations-based specifications forecast exceedingly well relative to the other specifications we tested.

References and Further Reading

"Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?" by Andrew Ang, Geert Bekaert, and Min Wei. 2007. *Journal of Monetary Economics*, 54(4): 1163–1212.

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

"Measuring Core Inflation," by Michael F. Bryan and Stephen G. Cecchetti. 1993. Federal Reserve Bank of Cleveland, working paper no. 93-04.

"Inflation Persistence, Inflation Targeting, and the Great Moderation," by Charles T. Carlstrom, Timothy S. Fuerst, and Matthias Paustian. 2007. Federal Reserve Bank of Cleveland, working paper no. 07-21.

"Forecasting Inflation," by James H. Stock and Mark W. Watson. 1999. *Journal of Monetary Economics*. 44(2): 293–335.

Footnotes

1. We estimate this regression using OLS with quarterly data on the CPI from 1947:Q1 to 2010:Q2. Following Stock and Watson (1999), we estimate the following regression: $\pi_{t+4}^4 - \pi_t = \alpha + \beta(L)(\pi_t - \pi_{t-1}) + \varepsilon_t$.

Here, π_{t+4}^4 is the four-quarter ahead annual inflation, π_t is the annualized quarterly inflation, and $\beta(L)$ is the lag polynomial operator. We include the current value and the first four lags in this lag polynomial.

2. Another possible approach is to estimate the model over rolling sample windows (typically 10 years or 40 quarters at a time). We tried this technique as well and obtained qualitatively similar results.

3. We alter that first regression slightly here by allowing the lags of inflation and lagged measures of economic activity to vary independently, choosing the "optimal" lag length using the Bayesian information criterion with a maximum lag length of four quarters. The regression equation is the following:
$\pi_{t+4}^4 - \pi_t = \alpha + \beta(L)(\pi_t - \pi_{t-1}) + \gamma(L)x_t + \varepsilon_t$.
Here, x_t is the activity measure.

4. The gaps for each measure of economic activity are obtained with the Hodrick-Prescott filter.

5. Here we are using a slightly different regression than our first regression, which included headline CPI, because of its performance:
$\pi_{t+4}^4 = \alpha + \lambda(L)\pi_t^* + \varepsilon_t$.
π_t^* is the different measure of inflation. We allow a maximum lag length of four quarters, with the "optimal" lag length chosen by the Bayesian information criterion.

6. The regression that uses expectations as the explanatory variable is the following:
$\pi_{t+4}^4 = \alpha + \beta_0\pi_{t+4|t}^e$.
$\pi_{t+4|t}^e$ is today's expectations for four quarters ahead.

7. This prompted us to test a set of Phillips curves that include these alternative underlying inflation measures and different measures of economic slack. Interestingly, we found that the inclusion of slack measures did little to improve upon the forecasting performance of the underlying inflation measures alone.
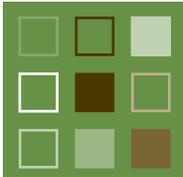
Brent H. Meyer is a senior economic analyst at the Federal Reserve Bank of Cleveland. Mehmet Pasaogullari is a research economist at the Bank. The views they express here are theirs and not necessarily those of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System or its staff.