JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 53, No. 6, Dec. 2018, pp. 2559–2586
COPYRIGHT 2018, MICHAEL G. FOSTER SCHOOL OF BUSINESS, UNIVERSITY OF WASHINGTON, SEATTLE, WA 98195
doi:10.1017/S0022109018000467

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

Monetary-Policy Rule as a Bridge: Predicting Inflation without Predictive Regressions

Jian Hua and Liuren Wu*

Abstract

A major issue with predicting inflation rates using predictive regressions is that estimation errors can overwhelm the information content. This article proposes a new approach that uses a monetary-policy rule as a bridge between inflation rates and short-term interest rates and relies on the forward-interest-rate curve to predict future interest-rate movements. The 2-step procedure estimates the predictive relation not through a predictive regression but far more accurately through the contemporaneous monetary-policy linkage. Historical analysis shows that the approach outperforms random walk out of sample by 30%–50% over horizons from 1 to 5 years.

I. Introduction

Forecasting inflation rates plays a central role in monetary-policy decision making (Svensson (2003), (2005)), but it is notoriously difficult to predict inflation rates with accuracy (Stock and Watson (2007), (2010)). Predictive regressions on inflation rates often fail to outperform the simple random-walk assumption out of sample (Faust and Wright (2013)). The poor out-of-sample performance of predictive regressions has also been observed for many financial series, such as interest rates (Duffee (2002), (2013)), exchange rates (Engel and West (2005)), and stock returns (Welch and Goyal (2008)). This article proposes a different approach to predicting inflation rates that drastically reduces the estimation error and improves the out-of-sample forecasting performance.

Several studies have discussed the inference issues and estimation biases with predictive regressions.[1] A major issue with predictive regressions in

*Hua, jian.hua@baruch.cuny.edu, Baruch College Zicklin School of Business; Wu (corresponding author), liuren.wu@baruch.cuny.edu, Baruch College Zicklin School of Business. The authors thank Tobias Adrian, David Backus, Mikhail Chernov, Todd Clark, Jennifer Conrad (the editor), Francis X. Diebold, Taeyoung Doh, Andrew Karolyi, Edith Liu, Eric Sims (the referee), Oleg Sokolinskiy, Byron Tsang, Jonathan Wright, and participants at Baruch College, Queens College, New York University, the 2013 Annual Society for Financial Econometrics (SoFiE) Conference, and the 2014 Triple Crown Conference for comments.

[1] See, for example, Elliott and Stock (1994) and Stambaugh (1999) on spurious relations and biased coefficient estimates in predictive regressions, Jansson and Moreira (2006) on invalid inferences, Campbell and Yogo (2006) on misspecification, and Phillips (2013) for a general overview of the pitfalls of predictive regressions.

out-of-sample tests is that, even with a truly predictive relation, estimation errors in the relation coefficients can overwhelm the information content of the predictive variables, leading to poor out-of-sample forecasting performance. This is particularly true for predicting the changes of highly persistent series, such as inflation rates, interest rates, and exchange rates, and for predicting financial security returns. The estimation errors on the coefficients of these weakly predictive relations tend to be large, more so for rolling-window estimation in out-of-sample tests. Even if a full-sample estimation generates statistically significant coefficients, the shorter sample size for the rolling-window estimation is likely to increase the estimation error and the sample variation, which deteriorates the out-of-sample performance.

To reduce the impact of estimation errors in out-of-sample tests, the literature has considered imposing strong priors and/or structural constraints. For example, Campbell and Thompson (2008) show that the forecasting performance for stock excess returns can be improved by imposing restrictions on the signs of coefficients and on the range of return forecasts based on our structural knowledge of the underlying relations. In predicting inflation rates, Faust and Wright (2013) find that although an unconstrained autoregressive regression cannot beat random walk out of sample, the specification can generate better performance with the autoregressive coefficient fixed ex ante, thus removing the issue of estimation errors.

This article proposes a new approach that predicts inflation rates without running predictive regressions. The approach constitutes a 2-step procedure. The first step is to estimate a contemporaneous relation between the realized year-over-year inflation rate and the yearly average of a short-term interest-rate series. The relation is motivated by monetary-policy rules (e.g., Taylor (1993)). As central banks adjust target interest rates to respond to inflation-rate variations, the two series exhibit strong comovements that can be identified robustly via a simple regression. The estimated relation serves as a bridge, transforming the task of inflation forecasting to short-term interest-rate forecasting. The second step is to predict the future path of the short-term interest rate using the forward-interest-rate curve, with a rolling moving-average bias correction that removes the average effect of the risk premium and convexity.

Combining the two steps leads to a simple relation that predicts inflation-rate changes with the forward-spot interest-rate differential of the corresponding maturity. One can in principle estimate this predictive relation with a standard predictive regression. By contrast, our 2-step procedure estimates the coefficient of the predictive relation far more accurately through the contemporaneous monetary-policy linkage.

Historical analysis of U.S. inflation-rate series shows that this new approach can outperform random walk out of sample by 30%–50% over horizons from 1 to 5 years. By comparison, estimating the same relation with standard predictive regressions leads to drastic out-of-sample deterioration and generates performance worse than the simple random-walk assumption. Applying the new approach to predict U.K. inflation rates generates similar superior performance. The superior performance is also robust to variations in specifications and performance metrics.

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

Many studies have used interest rates as predictors of future inflation rates (see, for example, Fama (1975), Mishkin (1990), Jorion and Mishkin (1991), Frankel and Lown (1994), Forni, Hallin, Lippi, and Reichlin (2003), Stock and Watson (2003), Diebold, Rudebusch, and Arouba (2006), and Ang, Bekaert, and Wei (2007)). Often, however, the estimation errors from the predictive regression completely overwhelm the information in the interest-rate curve, leading to poor out-of-sample forecasting performance. What our new approach achieves is a considerably more accurate identification of the predictive relation.

The remainder of the article is organized as follows: The next section lays out the structural basis for our new inflation-forecasting approach. Section III describes the data behavior and the empirical implementation details. Section IV discusses the empirical results. Section V extends the analysis to forecasting U.K. inflation rates. Section VI performs robustness analysis. Section VII concludes.

II. Estimate Predictive Relations without Predictive Regressions

Our construction of inflation forecasts relies on two key components, i) the strength of a monetary-policy rule that builds a bridge between inflation rates and short-term interest rates and ii) the information content in the forward-interest-rate curve to predict the future path of short-term interest-rate movements. We lay out the foundations for each component and illustrate how we combine the two components to predict inflation rates without running predictive regressions.

A. Monetary-Policy Rule as a Bridge between Inflation and Interest Rates

An economy's inflation is normally measured via some price index P_t, with its drift capturing the instantaneous expected inflation rate $\pi_t = (1/dt)\mathbb{E}_t[dP_t/P_t]$. Because the drift is not observable, one strives to make inferences based on the realized inflation rate over a certain historical period. A common choice is the year-over-year realized inflation rate, defined as the log price change over a year, $p_t \equiv \ln(P_t/P_{t-1})$. The price series are usually updated at either a monthly or quarterly frequency. The year-over-year realized inflation rate smoothes over the seasonality effects within a year and constitutes a more stable basis for long-run forecasting. Throughout the article, we focus on year-over-year inflation rates, and we use p_t to denote its value at time t while suppressing the notation for the 1-year time horizon to reduce notation cluttering.

One of the most important objectives of modern monetary policy is to control inflation (Woodford (2003)). Central banks alter their short-term interest-rate target in response to their forecasts of future inflation rates, explicitly or implicitly following some version of policy rules, such as the one suggested by Taylor (1993). One such policy rule can be written as follows:

$$r_t = \alpha + \beta\pi_t + x_t, \tag{1}$$

where β denotes the response of the interest-rate target r_t to per-unit shocks in the expected instantaneous inflation rate, and x_t denotes other policy considerations, such as output gap, unemployment rate, and past interest-rate level.

Our application reverts the policy rule and infers the expected inflation rate from the level of the interest-rate target by regressing the year-over-year realized inflation rate against the yearly average of a short-term interest-rate series:

$$p_t = a + br_t + e_t. \tag{2}$$

We take the year-over-year inflation rate as a noisy measurement of the average expected inflation rate over the time period $(t-1,t)$ and relate it to the average interest rate during this period. The reversal of the policy rule in equation (2) not only serves our purpose of inflation forecasting but also mitigates the impact of the errors-in-variables problem experienced when one regresses the policy rate against a noisy realized inflation rate as a proxy for expected inflation rate (Orphanides (2002)).

The particular formulation of equation (2) balances several considerations. An alternative translation of the policy rule in equation (1) would be to regress the next-period (say, monthly) realized inflation rate against the current short-term rate level, a specification examined by, for example, Fama (1975) as a test of the Fisher equation. Instead, equation (2) chooses to define the inflation rate over a longer, 1-year horizon to smooth out seasonality effects and chooses to use the yearly average of the daily interest-rate observations as the regressor to formulate a contemporaneous relation rather than a predictive relation. The literature has proposed many different versions of policy rules and has related the policy rate to both backward-looking and forward-looking inflation rates (Carlstrom and Fuerst (2000)), as well as lagged interest rates (Woodford (1999)). The annual averaging not only reduces the seasonality effect in the inflation-rate measurement but also captures a stronger average relation across different leads and lags.

The estimated contemporaneous linkage in equation (2) transforms the task of forecasting inflation rates to the task of forecasting short-term interest rates:

$$\widehat{p}_{t+h|t} - p_t = b(\widehat{r}_{t+h|t} - r_t), \tag{3}$$

where $\widehat{p}_{t+h|t}$ and $\widehat{r}_{t+h|t}$ denote the time t forecast of the time $(t+h)$ year-over-year inflation rate and interest rate, respectively, with h denoting the forecasting horizon.

There has been strong interest in the literature in estimating monetary-policy rules (e.g., Clarida, Gali, and Gertler (2000), Kim and Nelson (2006)), as well as debates on under what conditions the monetary-policy rules can be effectively identified (e.g., Sims (2008), Cochrane (2011)). The specification in equation (2) relies on strong assumptions on how the policy rule can be identified and on how omitted determinants of the monetary policy (e.g., employment and output gap) do not significantly alter the identified coefficient. The specification in equation (3) further assumes that the identified relation at time t holds at future times $t+h$. Whether these assumptions hold are subject to debate. Nevertheless, for the purpose of inflation prediction, we can regard equation (2) as a purely statistical relation motivated by empirical observations of strong comovements between interest rates and inflation rates, regardless of whether the comovements are induced by monetary-policy rules or not. The relation can serve as the bridge for inflation forecasting as long as it is strong, empirically identifiable, and stable.

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

B. Forecasting the Future Path of Interest Rates via the Forward-Rate Curve

We use the forward-interest-rate curve to forecast the future path of the short-rate movements. The forward curve includes contributions from three sources (Carr and Wu (2017)): i) the expectation of the future short rate, ii) the risk premium, and iii) the convexity effect. To use the forward-rate curve to predict future interest rates, one must find an effective way to separate the interest-rate expectation from the other two components.

The literature on interest-rate forecasting historically relies on the expectation hypothesis that the forward rate represents an expectation of the future short rate, thus assuming a risk premium of 0 and ignoring the convexity effect. Depending on how the expectation hypothesis is formulated, the literature proposes different forms of forecasting regressions. The one that is particularly relevant to our purpose is the forecasting regression based on forward interest rates (e.g., Fama (1984), Fama and Bliss (1987), and Mishkin (1988)):

$$r_{t+h} - r_t = c^h + d^h(f_t^h - r_t) + e_{t+h}. \tag{4}$$

The expectation hypothesis predicts that $c^h = 0$ and $d^h = 1$. Most studies find the slope estimate to deviate from 1 and interpret the deviation as evidence for a time-varying risk premium.

Although it is possible that the convexity component and the risk premium are time varying, it is difficult to identify their effects accurately. For the purpose of out-of-sample interest-rate forecasting, we propose a more conservative approach that captures the combined effects of the risk premium and the convexity via a simple moving average, thus avoiding a predictive regression altogether. We can write our forecasting specification as follows:

$$\widehat{r}_{t+h|t} - r_t = (f_t^h - r_t) - \overline{(f_t^h - r_t)}, \tag{5}$$

where we impose a slope of 1 and set the intercept to the negative of the rolling-window historical moving average of the forward-spot interest-rate differential, $\overline{(f_t^h - r_t)}$.

C. Predicting Inflation Rates without Running Predictive Regressions

Equations (2) and (5) form the basis of our inflation forecasting, with equation (2) building a linkage between inflation rates and short-term interest rates and equation (5) predicting future short rates using the current forward-rate curve with an average-bias correction. Given the contemporaneous coefficient estimate $\widehat{b}$ and the moving-average estimate of the forward-spot interest-rate differential, we can generate the inflation-rate forecasts ($\widehat{p}_{t+h|t}$) as

$$\widehat{p}_{t+h|t} - p_t = \widehat{b}\left(f_t^h - r_t - \overline{(f_t^h - r_t)}\right). \tag{6}$$

Equation (6) predicts future inflation-rate changes with the current forward-spot interest-rate differential. In principle, one could estimate this predictive relation via a standard predictive regression:

$$\widehat{p}_{t+h} - p_t = \eta + b(f_t^h - r_t) + e_{t+h}. \tag{7}$$

Our contribution lies not only in choosing the forward-spot interest-rate differential of the corresponding maturity as the predictive variable but also in proposing a new, more accurate way of determining the predictive-relation coefficients. Contemporaneous relations such as equation (2) are usually stronger than predictive relations such as equation (7). Coefficient identification is naturally stronger from a stronger relation.

III. Data and Implementation Details

We examine four major price indexes in the U.S. economy: the Consumer Price Index (CPI), the core CPI (CCPI), the personal consumption expenditure (PCE) deflator, and the core PCE deflator (CPCE). The CPI measures the average change in the prices of a basket of goods and services bought by a typical urban household. The PCE deflator measures the average change in the prices of a basket of goods and services purchased by a typical consumer. Their respective core measures exclude food and energy, the prices of which tend to be volatile. The price index time series are available from the Federal Reserve Bank of St. Louis, at a monthly frequency from June 1962 to Dec. 2016.

The spot- and forward-interest-rate curves are constructed from Treasury bond prices. The data sources and curve-construction details are described by Gurkaynak, Sack, and Wright (2007). The daily coefficient estimates for the extended Nelson and Siegel (1987) functional form are available from the Federal Reserve Board. Spot and forward interest rates can be directly computed from these coefficients based on the function form. In principle, one can construct spot rates at any maturity, but we find that the constructed spot rates with maturities of less than 3 months tend to exhibit stability issues. In our implementation, we use the 3-month spot rate as a proxy for the short rate r_t and compute h-period forwards of the 3-month rate from $h = 1$–5 years.

A. Summary Statistics

Table 1 reports the summary statistics of the year-over-year inflation rates for the four price indexes and the corresponding average spot and forward 3-month interest rates at selected forward horizons of 1–5 years. The inflation rates and the interest rates show similar average magnitude (mean), similar variation (standard deviation), and similar persistence (annual nonoverlapping autocorrelation).

Among the four inflation-rate measures, the inflation rate defined on PCE deflators shows smaller variation and higher persistence than that defined on the CPI index. The two classes of indexes have some subtle differences in their definitions. The CPI represents the price paid by urban customers, whereas the PCE deflator is a broader measure that covers both urban and rural customers. Furthermore, the PCE deflator is a chain-weighted index that captures shifting spending patterns, whereas the CPI is a fixed-weight index based on spending patterns several years ago. The broader base and the chain-weighting may contribute to the smaller variation and higher persistence of the PCE deflator. Within each index class, the core inflation rate shows smaller variation and higher persistence because of the exclusion of the more volatile energy and food component.

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

TABLE 1

Summary Statistics of Inflation Rates and Treasury Interest Rates

Table 1 reports summary statistics of inflation rates defined on the Consumer Price Index (CPI), core CPI (CCPI), personal consumption expenditure (PCE) deflator, and core PCE deflator (CPCE) in Panel A and the spot (0) 3-month and forward interest rates at forward horizons of 1, 2, 3, 4, and 5 years in Panel B. The inflation rates are computed from price indexes obtained from the Federal Reserve Bank of St. Louis, and the interest rates are computed from daily Nelson–Siegel (1987) model coefficients provided by the Federal Reserve on the U.S. Treasuries. The statistics are computed over 655 monthly observations from June 1962 to Dec. 2016. The inflation rates are computed as year-over-year log changes of the four price indexes, and the interest-rate series are averages of daily observations over the corresponding year. The statistics include the sample average (Mean), standard deviation (Std. dev.), minimum, maximum, yearly nonoverlapping autocorrelation (Auto), and cross-correlation (CC) with the 3-month spot interest rate.

	Panel A. Inflation Rates				*Panel B. Spot and Forward 3-Month Rates*					
Statistics	CPI	CCPI	PCE	CPCE	0	1	2	3	4	5
Mean	3.81	3.80	3.35	3.30	5.07	5.56	5.92	6.19	6.40	6.58
Std. dev.	2.73	2.46	2.36	2.09	3.23	3.17	3.00	2.83	2.69	2.58
Minimum	−1.98	0.60	−1.19	0.94	0.11	0.23	0.49	0.90	1.38	1.87
Maximum	13.62	12.76	10.95	9.74	15.53	14.11	13.80	13.55	13.41	13.35
Auto	0.75	0.82	0.80	0.88	0.89	0.93	0.94	0.95	0.95	0.95
CC	0.71	0.79	0.71	0.77	1.00	0.97	0.94	0.92	0.90	0.88

Comparing the spot 3-month interest rate with the forward rates at different forward horizons in Panel B of Table 1 shows an upward-sloping average term structure. Because the convexity effect drives the term structure downward-sloping, the upward-sloping average term structure is evidence for a positive bond risk premium. The longer-term forward rates also show less variation and more persistence.

The last row of Table 1 measures the cross-correlation of each series with the 3-month spot interest rate. The correlation estimates between the four inflation rates and the 3-month interest rate are high, ranging from 0.71 for the two raw measures to 0.77 for the core PCE inflation rate and 0.79 for the core CPI inflation rate. As expected, the correlation between forward and spot interest rates is even higher.

Figure 1 plots the time series of the four inflation rates and the corresponding average 3-month interest rate. The four inflation rates and the interest rate share similar trends. There has been extensive interest in understanding inflation trends, including inflation's dynamics (Clark and Doh (2011)) and its relations to monetary policy (Cogley and Sbordone (2008)) and the interest-rate term structure (De Graeve, Emiris, and Wouters (2008), Cieslak and Povala (2015)). Our approach amounts to inferring future inflation trends from the current interest-rate term structure.

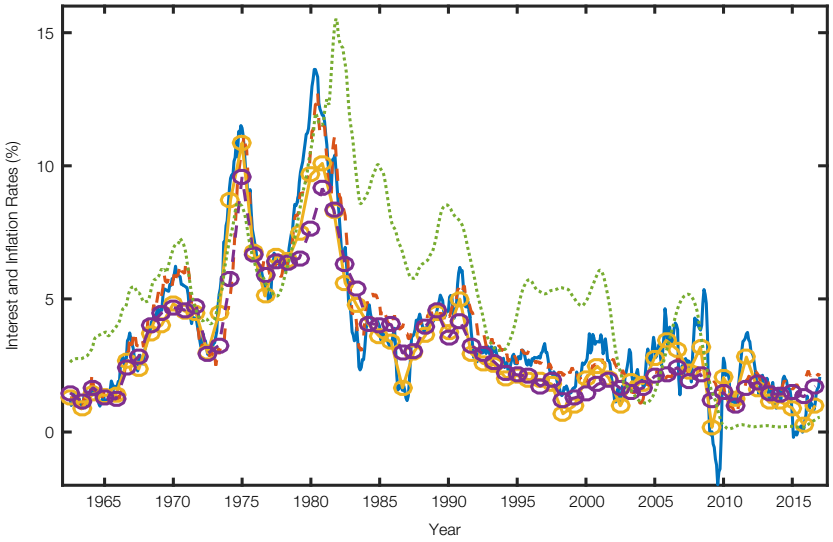
B. Rolling-Window Estimation and Out-of-Sample Prediction

The out-of-sample forecasting exercise relies on rolling-window estimation. At each date, we use a fixed 10-year rolling window to estimate i) the contemporaneous relation between the year-over-year inflation rate and the average 3-month interest rate and ii) the average difference between forward and spot interest rates. The choice of window size balances the need for a sample long enough to include several business cycles and enough inflation variation for effective identification with the need to allow for structural variations caused by potential regime changes.

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

FIGURE 1
Time Series of Inflation Rates and 3-Month Treasury Interest Rate

Figure 1 compares the time series of the year-over-year realized inflation rates computed from the four price indexes, the Consumer Price Index (CPI) (solid line), the core CPI (dashed line), the personal consumption expenditure (PCE) deflator (circle solid line), and the core PCE deflator (circle dashed line), with the yearly averages of the daily 3-month continuously compounded Treasury spot rate (dotted line).



From the estimated relations, we generate inflation-rate forecasts and compare the forecasting accuracy relative to the random-walk benchmark. The random-walk assumption defines the forecast of the future inflation rate as the most recent realized year-over-year inflation rate,

$$\widehat{p}_{t+h|t} = p_t. \tag{8}$$

By comparison, the time t out-of-sample inflation-rate forecast from our approach is given by

$$\widehat{p}_{t+h|t} = p_t + \widehat{b}_t \left(f_t^h - r_t - \overline{(f_t^h - r_t)} \right), \tag{9}$$

where $\widehat{b}_t$ denotes the time t rolling-window estimate of the slope coefficient, and $\overline{(f_t^h - r_t)}$ denotes the time t rolling-window estimate of the average forward-spot interest-rate differential.

C. Performance Measures and Test Statistics

With out-of-sample inflation-rate forecasts $\widehat{p}_{t+h|h}$, we define the ex post out-of-sample forecasting error e_{t+h} as

$$e_{t+h} = p_{t+h} - \widehat{p}_{t+h|t}. \tag{10}$$

We use the squared forecasting error to define the loss function and compare the loss of each forecast with the loss from the random-walk benchmark, the forecasting error of which is given by $e_{0,t+h} = p_{t+h} - p_t$.

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

We compute the mean squared forecasting error (MSFE) of each forecast:

$$\text{MSFE} = \frac{1}{T}\sum_{t=1}^{T} e_{t+h}^2, \tag{11}$$

with T denoting the number of out-of-sample observations. We define an R^2 measure against the random-walk benchmark:

$$R^2 = 1 - \text{MSFE}/\text{MSFE}_0, \tag{12}$$

where MSFE_0 denotes the mean squared forecasting error from the random-walk benchmark. The literature often uses the MSFE ratio to represent the relative performance. Although it contains the same information, the R^2 transformation has a more intuitive interpretation as the percentage improvement over the random-walk benchmark, positive when the forecast outperforms the random walk benchmark and negative when the forecast underperforms.

To test the statistical significance of the forecasting performance difference, we compute the Diebold–Mariano (DM) (1995) t-statistics on the squared forecasting error difference, $\delta_{t+h} = e_{0,t+h}^2 - e_{t+h}^2$,

$$\text{DM} = \frac{\overline{\delta}}{s_\delta}\left(\frac{T+1-2(h+1)+(h(h+1)/T)}{T}\right)^{0.5}, \tag{13}$$

where $\overline{\delta}$ denotes the sample mean of the difference, and s_δ denotes the Newey–West (1987) standard-error estimate, computed with a lag equal to the forecasting horizon. The statistics adjust for the small-sample-size bias according to Harvey, Leybourne, and Newbold (1997). Under the null hypothesis that each rolling-window-estimated model and the random-walk benchmark have equal finite-sample forecast accuracy,[2] Clark and McCracken (2012) find that the thus-computed DM test statistic can be compared with standard normal critical values.

IV. Empirical Findings

This section examines the estimated relation between inflation and short-term interest rates, the effectiveness of forward rates as predictors of future short rates, and the out-of-sample forecasting performance of future inflation rates from our approach.

A. Policy Linkage between Inflation Rates and Short-Term Interest Rates

We regress the year-over-year inflation rates on the 3-month interest rate to determine the linkage between the two:

$$p_t = a + br_t + e_t. \tag{14}$$

Panel A of Table 2 reports the full-sample regression estimates of the four inflation-rate series. The R^2 estimates of the regressions range from 50% for the

[2]See Inoue and Kilian (2004) for a discussion on the distinction between population and finite-sample predictive accuracy, and see Clark and McCracken (2012) for a thorough review of forecast performance evaluation.

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

TABLE 2

Linkage between Inflation Rates and Short-Term Interest Rates

Table 2 reports the coefficient (intercept $\hat{a}$ and slope $\hat{b}$) estimates, their Newey–West (1987) standard errors (in parentheses), and the R^2 values from regressing year-over-year realized inflation rates on 3-month Treasury interest rates. The regressions in Panel A are performed on the whole sample from June 1962 to Dec. 2016. Panels B and C report regression results over subsample periods. The Newey–West standard errors are computed with 36 lags. The CPI, CCPI, PCE, and CPCE indexes are defined in Table 1.

Price Index	$\hat{a}$		$\hat{b}$		R^2
Panel A. June 1962–Dec. 2016					
CPI	0.75	(0.45)	0.60	(0.13)	0.51
CCPI	0.74	(0.40)	0.60	(0.10)	0.63
PCE	0.72	(0.37)	0.52	(0.10)	0.50
CPCE	0.77	(0.30)	0.50	(0.06)	0.59
Panel B. June 1962–Sept. 1979					
CPI	−3.97	(0.71)	1.62	(0.14)	0.87
CCPI	−3.14	(0.63)	1.43	(0.11)	0.83
PCE	−3.03	(0.90)	1.37	(0.18)	0.79
CPCE	−2.03	(0.86)	1.15	(0.15)	0.74
Panel C. Oct. 1979–Dec. 2016					
CPI	0.89	(0.44)	0.49	(0.12)	0.54
CCPI	0.87	(0.39)	0.52	(0.10)	0.70
PCE	0.75	(0.38)	0.42	(0.09)	0.58
CPCE	0.77	(0.32)	0.42	(0.07)	0.72

PCE deflator series to 63% for the core CPI series. The slope estimates range from 0.5 for the core PCE series to 0.6 for the two CPI series. The estimates suggest that for each percentage change in the short-term interest rate, the expected inflation rate changes by 0.5–0.6 percentage points. The policy response is the reciprocal of the slope estimates ($\beta = 1/b$): Each percentage point of expected inflation change leads to 1.7–2 percentage points of action on the short-term interest rate, close to the suggestion by Taylor (1993).

Monetary policies in the United States experienced significant shifts during our sample period. Walsh (2003) describes the historical variation of Federal Reserve (Fed) policies based on differences in Fed operating procedures. Taylor (1999) and Clarida et al. (2000) divide the sample into two main subperiods based on the responsiveness of interest rates to expected inflation variation.[3] The first is the pre-Volcker period up until Sept. 1979. During this period, the Fed allowed nonborrowed reserves to adjust automatically to stabilize the federal funds rate within a narrow band around its target level. Such interest-rate targeting has been shown to be inflationary (Goodfriend (1987)). New theoretical developments (e.g., Woodford (1999), Clarida, Gali, and Gertler (1999)) advocate inflation-rate targeting instead, which requires that the interest rate responds more than the inflation-rate variation. Since Volcker, the Fed's operating procedures have experienced a series of changes, from targeting nonborrowed reserves in Oct. 1979 to targeting borrowed reserves after 1982 and then to directly targeting the Fed

[3] In other related literature, Goodfriend (2002) provides a detailed analysis of various phases of the U.S. monetary policy since 1987. Sims and Zha (2006) strive to identify regime shifts in U.S. monetary policies via statistical analysis. Boivin and Giannoni (2006) investigate the implications of changes in the structure of the U.S. economy for monetary-policy effectiveness over the pre- and post-1980 periods.

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

funds rate since Greenspan. The commonality underlying these different operating procedures is an enhanced interest-rate response to inflation variation and hence a more stable monetary policy.

Following Taylor (1999) and Clarida et al. (2000), we divide the sample into two subperiods and perform subperiod policy-rule estimation. The first subperiod is from the start of our sample to Sept. 1979, and the second is from Oct. 1979 to the present. Panel B of Table 2 reports the subsample estimates for the pre-Volcker period from June 1962 to Sept. 1979. The slope estimates are between 1.15 and 1.62, much higher than the full-sample estimates. From the policy-rule perspective, these estimates suggest that for each percentage-point increase in inflation, the short interest rate rises by 0.62–0.87 of a percentage point. This muted interest-rate response has been identified as a key reason for the observed high inflation and high inflation variation during this period. Panel C reports the subsample estimates from Oct. 1979 to Dec. 2016. The slope estimates are from 0.42 to 0.52, suggesting a much more responsive policy rule since Volcker. The strong response contributes to the declining and stabilizing inflation rate during the period.

Using the CPI index to define inflation, Clarida et al. (2000) estimate the interest-rate response to inflation at 0.68 for the pre-Volcker period and 2.14 after that. The reciprocals of these estimates are 1.47 and 0.47, close to our slope estimates of 1.62 for the 1962–1979 period and 0.49 for the period after. With the gross domestic product (GDP) deflator as the inflation measure, Taylor (1999) generates slightly higher responses at 0.813 for the 1960–1979 period and slightly lower responses at 1.533 for the 1987–1997 period. The patterns are nevertheless similar.

Figure 2 plots the 10-year rolling-window slope estimates for the relation. The estimates are above 1 in the 1970s but come below 1 starting in the early 1980s. Since Greenspan, monetary policy has been successful in containing the inflation rate to a low and stable level. As a result of this success, the slope coefficient estimates have stayed low during this time period.

B. Out-of-Sample Interest-Rate Prediction Based on the Forward Curve

We rely on the forward-interest-rate curve to generate the forecast on the future path of short-term interest-rate movements while removing the average bias induced by risk premiums and convexity effects via a simple moving-average correction,

$$r_{t+h|t} - r_t = (f_t^h - r_t) - \overline{(f_t^h - r_t)}. \tag{15}$$

Compared with the traditional expectation-hypothesis-based predictive regression,

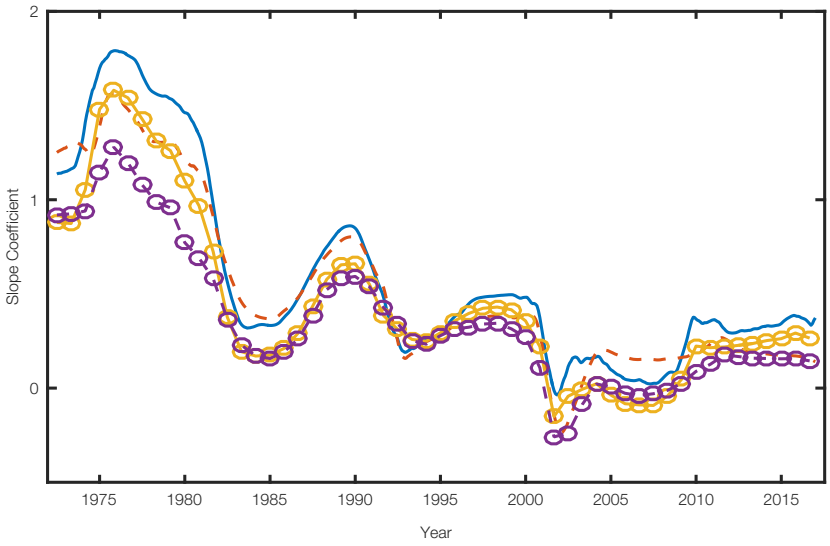
$$r_{t+h} - r_t = c^h + d^h(f_t^h - r_t) + e_{t+h}, \tag{16}$$

our approach amounts to imposing a slope of 1 and setting the intercept to the negative of the rolling-window moving average of the forward-spot interest-rate differential.

To see how our moving-average bias-removal approach in equation (15) compares with the predictive regression in equation (16) in out-of-sample

FIGURE 2
Rolling-Window Slope Estimates of the Inflation–Interest Rate Relation

Figure 2 compares the time series of the 10-year rolling-window estimates of the slope coefficient of the inflation–interest rate relation for the four inflation-rate series: Consumer Price Index (CPI) (solid line), core CPI (dashed line), personal consumption expenditure (PCE) deflator (circle solid line), and core PCE deflator (circle dashed line).



forecasting performance, we perform 10-year rolling-window estimations for both relations, from which we generate out-of-sample forecasts for future 3-month interest rates. Figure 3 plots the rolling-window expectation-hypothesis regression slope estimates ($\widehat{d}_t^h$) at forecasting horizons from 1 year (solid line) to 5 years (dashed line). The many dotted lines denote the slope estimates at intermediate horizons. The dash-dotted line denotes the null of the expectation hypothesis at 1. Although it is true that the slope estimates often deviate from the null value of 1, the deviations themselves show large sample variation, both above and below the null value, across different sample periods and forecasting horizons. The large variation highlights the inherent instability of the regression slope estimates.

Figure 4 plots the out-of-sample forecasting R^2 estimates from the two specifications at forecasting horizons from 1 to 5 years. The solid line shows the performance of the expectation hypothesis predictive regression in equation (16), which generates negative forecasting R^2 estimates for all forecasting horizons. Thus, the expectation-hypothesis-based predictive regressions cannot beat the simple random-walk hypothesis out of sample. By contrast, our moving-average bias-correction specification (dashed line) performs much better. The forecasting R^2 estimates are positive for forecasting horizons longer than 1.5 years and reach as high as 26% at the 4-year forecasting horizon. Therefore, regardless of whether the interest-rate risk premium is time varying or not, the expectation-hypothesis predictive regression itself is not stable enough to generate robust out-of-sample forecasting performance. Our simple moving-average bias-correction approach performs much better.

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

FIGURE 3

Rolling-Window Slope Estimates of the Expectation-Hypothesis Regressions

Figure 3 plots the 10-year rolling-window slope estimates on the expectation-hypothesis regression at different forecasting horizons. The solid line denotes the 1-year horizon. The dashed line denotes the 5-year horizon. The other dotted lines show estimates at intermediate horizons. The dash-dotted line represents the null value of the expectation hypothesis.

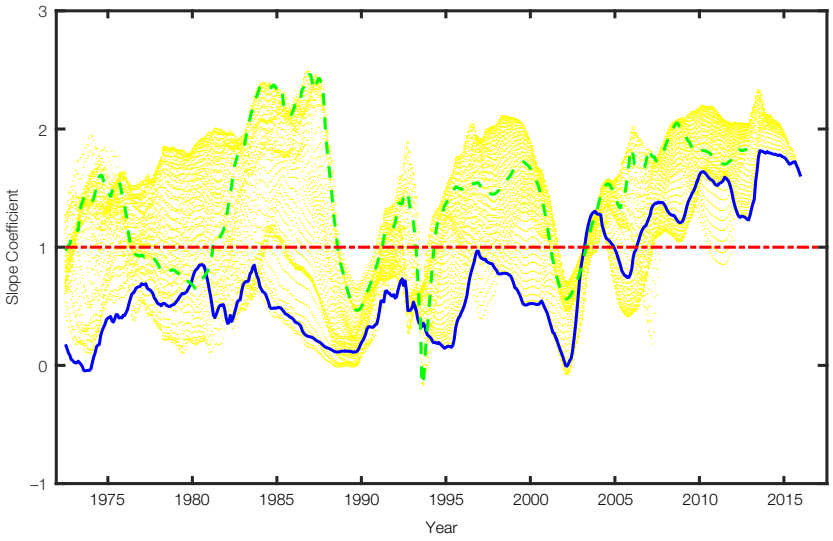
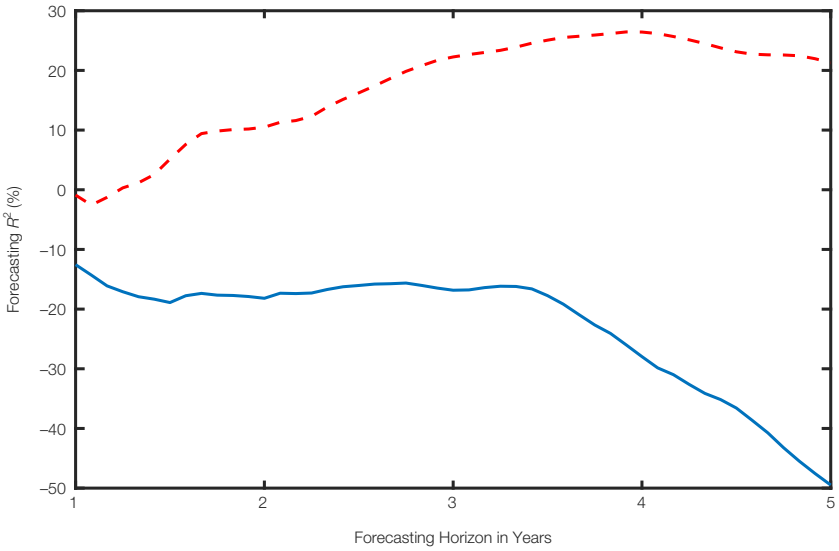


FIGURE 4

Out-of-Sample Forecasting R^2 for Future Interest Rates

Figure 4 compares the out-of-sample forecasting performance of the expectation-hypothesis (EH) regression (solid line) to the performance of our moving-average bias-correction approach (dashed line).



To examine the statistical significance of the forecasting performance, Table 3 reports the out-of-sample forecasting R^2 estimates and the DM t-statistics against the random-walk hypothesis at selected forecasting horizons of 1, 2, 3, 4, and 5 years over the out-of-sample period from June 1972 to Dec. 2016. The DM statistics for the expectation-hypotheses regressions are all negative and insignificant. By contrast, our average-bias-removal approach generates statistically significant outperformance over the random-walk benchmark at the 10% level for the 4-year forecasting horizon.

TABLE 3
Out-of-Sample Forecasting Performance for Future Interest Rates

Table 3 reports the out-of-sample forecasting performance on future interest rates at forecasting horizons of 1, 2, 3, 4, and 5 years, based on two approaches: the expectation-hypothesis (EH) regression (Panel A) and the moving-average (MV) bias removal (Panel B). The performance measures include both a forecasting R^2 measure and the Diebold–Mariano (DM) (1995) t-statistic defined on the squared forecasting error differences, which has standard normal critical values. The models are estimated with a 10-year-rolling window.

Horizon (years)	*Panel A. EH Regression*		*Panel B. MV Bias Correction*	
	R^2	DM	R^2	DM
1	−0.13	−1.14	−0.01	−0.08
2	−0.18	−0.79	0.10	0.69
3	−0.17	−0.35	0.22	1.59
4	−0.28	−0.42	0.26	1.85
5	−0.49	−0.71	0.21	1.33

C. Out-of-Sample Inflation Prediction without Running Predictive Regressions

We combine the rolling-window estimates on the contemporaneous relation between inflation rates and short-term interest rates in equation (14) with the short-rate prediction using average-bias-corrected forward rates in equation (15) to generate out-of-sample inflation-rate forecasts:

$$\widehat{p}_{t+h|t} - p_t \quad = \quad \widehat{b}_t \left(f_t^h - r_t - \overline{(f_t^h - r_t)} \right). \tag{17}$$

Figure 5 plots the out-of-sample forecasting R^2 estimates at different forecasting horizons for the four inflation-rate series. The forecasting performance is reasonably uniform across the four inflation-rate series and is strongly positive over long horizons. Over horizons from 1 to 5 years, the outperformance averages at 41%, 40%, 37%, and 28%, respectively, for the four inflation-rate series.

Table 4 reports both the R^2 estimates and the DM test statistics against the random-walk hypothesis at selected forecasting horizons of 1, 2, 3, 4, and 5 years. The DM statistics show strong statistical significance at 2- to 3-year horizons. At longer horizons, although the R^2 estimates remain high, the statistical significance declines as a result of the shortening of the effective sample size.

D. Performance Deterioration with Predictive Regressions

Equation (17) represents a very simple predictive relation, yet our out-of-sample forecasting exercise shows that it generates strong out-of-sample forecasting performance. A key contribution to the superior performance comes from the way we estimate the coefficients of the predictive relation, where the slope coefficient is determined by a strong contemporaneous relation, and the intercept is

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

FIGURE 5

Out-of-Sample Forecasting R^2 for Future Inflation Rates

Figure 5 plots the out-of-sample forecasting R^2 estimates across different forecasting horizons for the four inflation-rate series: Consumer Price Index (CPI) (solid line), core CPI (dashed line), personal consumption expenditure (PCE) deflator (circle solid line), and core PCE deflator (circle dashed line).

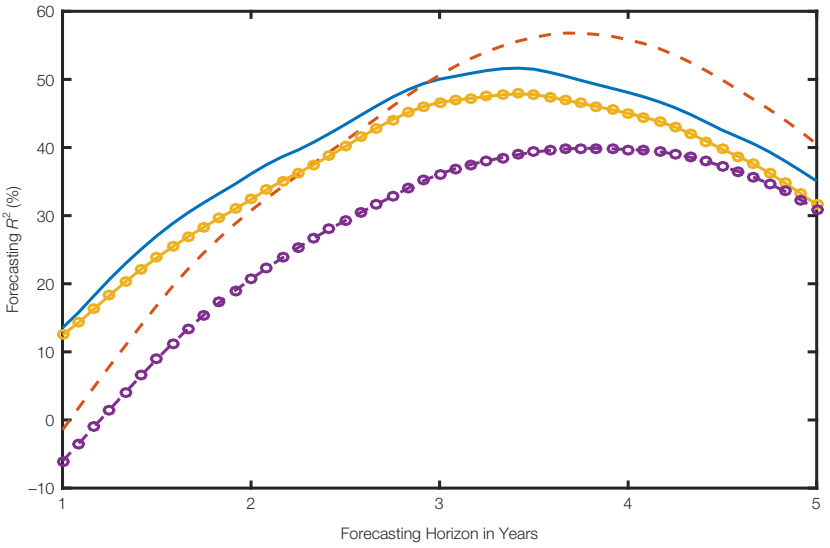


TABLE 4

Out-of-Sample Forecasting Performance for U.S. Inflation Rates without Predictive Regression

Table 4 reports the out-of-sample forecasting performance on four U.S. inflation-rate series at forecasting horizons of 1, 2, 3, 4, and 5 years under our approach without performing predictive regressions. The performance measures include both a forecasting R^2 measure and the Diebold–Mariano (DM) (1995) *t*-statistic defined on the squared forecasting error differences. The models are estimated with a 10-year-rolling window. The CPI, CCPI, PCE, and CPCE indexes are defined in Table 1.

Horizon (years)	CPI		CCPI		PCE		CPCE	
	R^2	DM	R^2	DM	R^2	DM	R^2	DM
1	0.14	1.39	−0.01	−0.11	0.12	1.13	−0.06	−0.40
2	0.36	2.08	0.31	1.65	0.32	1.94	0.21	1.18
3	0.50	1.82	0.51	1.75	0.47	1.78	0.36	1.54
4	0.48	1.53	0.56	1.52	0.45	1.47	0.40	1.35
5	0.35	1.05	0.41	1.12	0.32	1.03	0.31	1.03

determined by a simple moving average. This new approach drastically reduces the estimation error traditionally experienced in standard predictive regressions.

To highlight the strength of our approach, we estimate the same predictive relation using standard predictive regressions for comparison. We consider three ways of running the predictive regression, with varying degrees of structural constraints:

$$
\begin{aligned}
(18)\qquad \text{A:}\quad p_{t+h}-p_t &= c_t^h + b_t^h\left(f_t^h - r_t\right) + e_{t+h},\\
\text{B:}\quad p_{t+h}-p_t &= b_t^h\left(f_t^h - r_t - \overline{\left(f_t^h - r_t\right)}\right) + e_{t+h},\\
\text{C:}\quad p_{t+h}-p_t &= b_t\left(f_t^h - r_t - \overline{\left(f_t^h - r_t\right)}\right) + e_{t+h},
\end{aligned}
$$

Approach A imposes no structural constraints and directly estimates the intercept and slope of each predictive relation with a rolling-window predictive regression. Approach B removes the average forward-spot interest-rate bias through a moving average and estimates each predictive relation with a rolling-window regression without intercept. This specification acknowledges the existence of an average bias between the forward and spot interest rates, but it imposes a 0 mean on the inflation-rate changes. Approach C imposes the additional constraint that the relation between future inflation-rate changes and the corresponding average-bias-corrected forward-spot interest-rate differential is the same across all forecasting horizons and estimates the common slope coefficient by a stacked predictive regression without intercept.

Table 5 reports the out-of-sample forecasting performance at selected forecasting horizons from the three alternative estimation methods, with one panel for each type of regression. In each panel and for each inflation series, we report the out-of-sample forecasting R^2 defined against the random-walk benchmark, as well as two DM test statistics, one defined against the random-walk hypothesis (DMR) and the other defined against our proposed approach that estimates the predictive relation without running predictive regressions (DMN). A positive estimate for R^2 or, equivalently, for DMR suggests outperformance over the random-walk hypothesis, and a positive estimate for DMN indicates better predictive performance than our proposed method.

The results in Panel A of Table 5 are obtained following approach A of equation (18). The forecasting R^2 values are negative at both short and long horizons for all four series, and they only become positive at intermediate horizons.

TABLE 5

Forecasting Performance Deterioration Induced by Predictive Regression Estimation

Table 5 reports the out-of-sample forecasting performance for four U.S. inflation-rate series at forecasting horizons of 1, 2, 3, 4, and 5 years. The prediction is based on three alternative methods of estimation, one in each panel. The performance measures include both a forecasting R^2 measure and the Diebold–Mariano (DM) (1995) *t*-statistics, both against the random-walk hypothesis (DMR) and against our approach that does not include any predictive regression estimation (DMN). The models are estimated with a 10-year-rolling window. The CPI, CCPI, PCE, and CPCE indexes are defined in Table 1.

Horizon	CPI			CCPI			PCE			CPCE		
(years)	R^2	DMR	DMN	R^2	DMR	DMN	R^2	DMR	DMN	R^2	DMR	DMN
Panel A. Predictive Regression Estimation of Intercept and Slope												
1	−0.13	−1.18	−2.27	−0.30	−2.21	−1.29	−0.19	−1.66	−2.43	−0.37	−2.23	−1.13
2	0.04	0.13	−1.41	−0.19	−0.73	−2.10	−0.12	−0.48	−2.15	−0.35	−1.60	−2.03
3	0.13	0.22	−0.88	0.07	0.11	−0.91	0.04	0.08	−1.16	−0.05	−0.10	−1.24
4	0.05	0.15	−1.95	−0.03	−0.05	−1.37	−0.02	−0.06	−2.32	−0.06	−0.15	−1.57
5	−0.44	−1.77	−1.44	−0.84	−1.85	−1.56	−0.35	−1.47	−1.69	−0.43	−1.37	−1.50
Panel B. Predictive Regression Estimation of Slope without Intercept												
1	−0.07	−0.54	−1.70	−0.13	−1.35	−0.72	−0.12	−0.94	−1.98	−0.21	−1.68	−0.67
2	0.07	0.21	−1.20	−0.07	−0.24	−1.51	−0.09	−0.32	−1.74	−0.26	−1.02	−1.54
3	0.16	0.26	−0.75	0.11	0.19	−0.91	0.08	0.14	−0.89	−0.03	−0.06	−1.22
4	0.10	0.21	−1.35	−0.01	−0.01	−1.27	0.03	0.07	−1.33	−0.06	−0.12	−1.28
5	−0.18	−0.82	−2.32	−0.58	−1.18	−2.02	−0.09	−0.26	−2.07	−0.21	−0.50	−1.69
Panel C. Stacked Predictive Regression Estimation of Common Slope without Intercept												
1	−0.21	−1.02	−2.05	−0.27	−1.30	−1.55	−0.27	−1.21	−2.28	−0.39	−1.62	−1.90
2	0.09	0.26	−1.07	−0.04	−0.11	−1.46	−0.00	−0.01	−1.35	−0.16	−0.64	−1.69
3	0.27	0.58	−0.78	0.19	0.42	−1.13	0.20	0.41	−0.86	0.09	0.24	−1.06
4	0.30	0.64	−0.76	0.26	0.48	−1.03	0.25	0.55	−0.80	0.22	0.52	−0.82
5	0.19	0.39	−0.68	0.14	0.28	−0.92	0.20	0.41	−0.51	0.19	0.46	−0.57

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

The DM test statistics against the random-walk hypothesis (DMR) are either negative or insignificant. Not surprisingly, then, when we compute the DM statistics against our proposed approach (DMN), all the estimates are negative, and many significantly so.

By relying on the predictive regression only for the slope estimate but not for the intercept, Panel B of Table 5 shows some improvement as the forecasting R^2 estimates become either less negative or more positive. The performance improves further in Panel C when the predictive relations across all forecasting horizons are constrained to have the same slope coefficient. Nevertheless, even in this best-alternative case, none of the predictions significantly outperforms the random-walk hypothesis, and all of them perform worse than our approach, significantly so at short forecasting horizons.

What this exercise tells us is that even with the same predictive relation, the out-of-sample forecasting performance can still vary greatly depending on how the predictive relation is estimated. In particular, even if the predictive variable is truly informative, rolling-window predictive regressions tend to generate unstable coefficient estimates that lead to deteriorated out-of-sample forecasting performance. Therefore, to enhance out-of-sample forecasting power, it is important to not only search for informative predictive variables but also find robust ways of estimating the predictive relation. Although predictive regressions are often the most apparent way of estimating predictive relations, they are often the least likely to generate stable coefficient estimates and robust out-of-sample forecasting performance.

E. General Implications and Discussions

Although this article focuses on the particular application of inflation forecasting, the issue of poor out-of-sample forecasting performance for predictive regressions is general and is particularly severe for inherently weak predictive relations. In such cases, reformulating the problem to achieve relation identification in a stronger contemporaneous setting can be beneficial.

In many investment applications, return predictions are not generated based on predictive regressions but are based on constructing robust contemporaneous co-integrating relations. In equity markets, classic value investing strategies generate stock return predictions based on deviations between the market price and fundamental valuations. In fixed-income markets, although it is difficult to predict interest-rate changes based on autoregressive models or interest-rate factors, Bali, Heidari, and Wu (2009) show that the pricing residuals from a well-designed dynamic term-structure model can be used to predict short-term interest-rate movements and interest-rate swap returns. The estimation does not involve predictive regressions but relies on the fitting of a contemporaneous term-structure relation. In corporate credit markets, Carr and Wu (2011) construct a structural contemporaneous linkage between deep-out-of-the-money American put options on individual stocks and the company's credit default swap (CDS) spreads and show that deviations from the contemporaneous linkage predict future CDS and option movements. Bai and Wu (2016) show that deviations between CDS market quotes and fundamental-based CDS valuations can predict future market movements. Although Engle and Granger (1987) formulate a vector error-correcting forecasting

relation based on co-integration, practical applications often stop at the stage of identifying and estimating the contemporaneous co-integrating relation.

These applications are similar in spirit to our approach in identifying predictive relations by estimating a contemporaneous relation. The difference is that these applications rely on the estimated co-integrating relations to generate short-term error-correcting predictions, whereas our approach transforms the prediction of one variable to the prediction of another. The yield curve provides a natural foundation for generating long-term interest-rate predictions. Thus, any economic and financial series that show strong contemporaneous linkage with the short-term interest rates can be projected to the yield curve to generate long-term predictions. Inflation is one particular example. Other series that show strong contemporaneous linkages with short-term interest rates include, among others, mortgage rates, deposit rates, real growth, and unemployment rates. Indeed, policy agencies and financial institutions often use the yield curve as the basis for making macroeconomic and business projections.

Another way of turning predictive relations into contemporaneous relations is via professional surveys, which are available for many economic and financial series. Professional surveys have often been found to predict future movements better than most statistical regressions (e.g., Ang et al. (2007), Faust and Wright (2013)). The conventional thinking is to incorporate these survey numbers as yet another predictive variable (e.g., Aiolfi, Capistran, and Timmermann (2010)), but a potentially more promising application is to use these ex ante market forecasts to replace ex post realizations in identifying and/or recalibrating predictive relations. With this replacement, the predictive relations are transformed into a stronger contemporaneous relation. The identification accordingly becomes stronger, leading to potentially better out-of-sample predictions.

The benefit of using professional surveys goes beyond generating better out-of-sample forecasts. Asset-pricing theories generate implications for the contemporaneous relation between risks and expected excess returns. Empirical asset-pricing tests often replace ex ante return expectations with ex post realizations and attempt to estimate the contemporaneous pricing relation with a predictive regression. Ex post return realizations can deviate greatly from ex ante expectations. As a result, these predictive regressions often have little explanatory power and are a poor fit for testing asset-pricing theories. Using professional surveys instead, one can link the risk measures at a certain time to the actual market expectations at that time, thus generating stronger asset-pricing tests. In related literature, Brav, Lehavy, and Michaely (2005) and Wu (2018) propose the use of analysts' price targets on stocks to generate ex ante expected excess return estimators in testing classic risk–return relations. Kim and Orphanides (2012) incorporate interest-rate forecasts into dynamic term-structure model estimation to enhance the identification of the risk-premium behavior.

V. Predicting U.K. Inflation Rates

This section extends the forecasting analysis to the United Kingdom and examines whether our proposed methodology also generates superior forecasting performance for U.K. inflation rates.

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

We collect the U.K. retail price index (RPI) from the Office for National Statistics. The RPI measures the change in the cost of a representative sample of retail goods and services. The data sample is monthly from Jan. 1980 to Dec. 2014. We link the year-over-year inflation rates to the interest rates on U.K. government bonds. The Bank of England provides stripped daily spot interest rates from 1 to 20 years.[4] Based on the available data, we use the short end of the data, at 1-year maturity, to proxy the policy rate r_t and compute h-period forwards of the 1-year rate from 1 to 5 years.

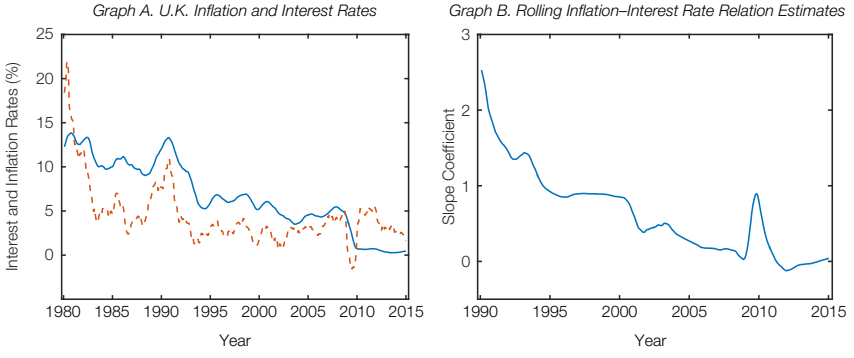
Figure 6 compares the time series of the U.K. sterling 1-year interest rate (solid line) with the time series of the year-over-year U.K. inflation rate (dashed line) in Graph A. The inflation rate reached over 20% in the early 1980s but has come down significantly and has been fluctuating within a narrow range since the early 1990s. The 1-year sterling interest-rate time series shows a similar downward trend. During this sample period, U.K. monetary policies experienced several structural transformations (Osborne (2013)). From 1976 to 1987, the U.K. monetary policy aimed to control various monetary aggregates. From 1987 to 1992, the policy was directed to target the exchange rate because U.K. interest rates were set to keep the value of sterling within a certain band relative to the German currency. In the face of large international capital flows and speculation, the monetary authorities could not maintain the exchange-rate target, resulting in some instability of both output and prices, as well as a sharp depreciation of sterling in 1992. Following the exit from the exchange-rate-targeting regime, the monetary policy started to target inflation. In 1997, the Monetary Policy Committee of the Bank of England was given operational independence, and inflation targeting has been the chief objective since then.

When we regress the U.K. realized inflation rate on the 1-year sterling interest rate over the entire sample period, we obtain a slope estimate of $b = 0.56$

FIGURE 6

Time Variation in the U.K. Inflation Rate, Interest Rate, and Their Relation

Graph A of Figure 6 plots the time series of the 1-year U.K. sterling interest rate (solid line) and the year-over-year U.K. inflation rate (dashed line) computed from the retail price index. Graph B plots the time series of the 10-year rolling-window slope estimates of the U.K. inflation–interest rate relation.



[4]Spot-rate data at shorter maturities are sparse, with large amounts of missing values.

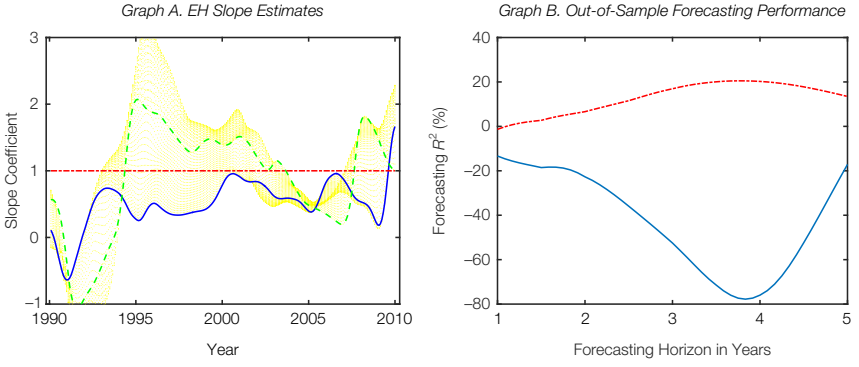
with a Newey and West (1987) standard deviation of 0.20 and an R^2 of 41%. The estimate corresponds to a strong monetary response of $\beta = 1/b = 1.78$. When we perform rolling-window estimation on the relation, the slope coefficient estimates evolve in line with the historical development of the U.K. monetary policy, as shown in Graph B of Figure 6. The rolling-window slope estimates are high and above 1 before 1992, suggesting that the U.K. monetary policy during this early period was inflationary. Since the introduction of inflation targeting, however, the coefficient estimates have come down to below 1.

To forecast U.K. sterling interest rates with the forward curve, we again compare the expectation-hypothesis-based predictive regression with our simple moving-average bias-removal approach. Graph A of Figure 7 plots the rolling-window expectation-hypothesis regression slope estimates over different forecasting horizons. The solid line denotes estimates on the 1-year-ahead forecasting relation, the dashed line denotes estimates on the 5-year-ahead forecasting relation, and the dotted lines denote estimates for intermediate horizons. Similar to the observation for the U.S. interest rates, the slope estimates often deviate strongly from the null value of 1, but the deviations themselves vary greatly over time and over different forecasting horizons, ranging from estimates as low as -1 to estimates as high as 3. Again, the large sample variation highlights the inherent instability of the estimates, making them unsuitable for out-of-sample forecasting application.

Graph B of Figure 7 plots the out-of-sample interest-rate forecasting R^2 estimates over different horizons for both the expectation-hypothesis regression approach (solid line) and our simple moving-average bias-removal approach (dashed line). The solid line stays far below 0, further highlighting the out-of-sample instability of the expectation-hypothesis regressions. By contrast, our moving-average bias-removal approach outperforms the random-walk benchmark over all horizons. The outperformance is over 20% at 3- to 4-year horizon.

FIGURE 7
U.K. Expectation-Hypothesis Regressions and Interest-Rate Forecasting

Graph A of Figure 7 plots the 10-year rolling-window slope estimates of the expectation-hypothesis forecasting regressions, with the solid line denoting the 1-year horizon, the dashed line denoting the 5-year horizon, and the other dotted lines representing intermediate horizons. The dash-dotted line represents the null value of the expectation hypothesis. Graph B compares the out-of-sample interest-rate forecasting R^2 estimates across different horizons based on the expectation-hypothesis (EH) regression (solid line) and our moving-average bias-correction approach (dashed line).



https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

Figure 8 plots the out-of-sample forecasting R^2 estimates for the U.K. inflation rate at different forecasting horizons. Our methodology works equally well for predicting U.K. inflation rates. The outperformance reaches as high as 40% around the 3- to 4-year forecasting horizon.

FIGURE 8
Out-of-Sample Forecasting R^2 for U.K. Inflation Rates

Figure 8 plots the out-of-sample forecasting R^2 estimates on U.K. inflation rates across different forecasting horizons.

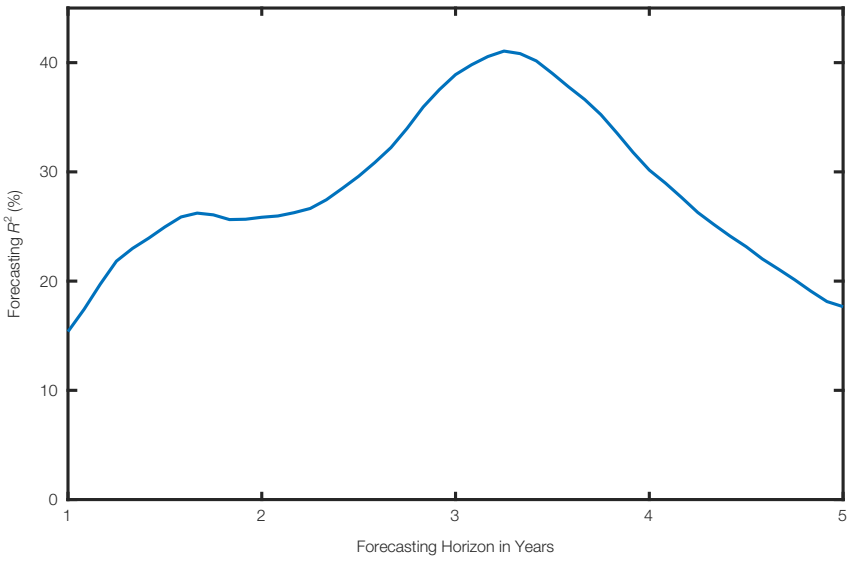


Table 6 reports the out-of-sample R^2 estimates and the DM t-statistics at selected forecasting horizons. The shorter sample makes the t-statistics weaker than those observed for the U.S. data. Nevertheless, the outperformance at the 1-year forecasting horizon over the random-walk benchmark is statistically significant at the 10% confidence level.

Several researchers (e.g., Welch and Goyal (2008), Kelly and Pruitt (2013)) stress the importance of testing the performance robustness of a methodology across different samples. The extension to U.K. inflation analysis shows that our proposed inflation-forecasting methodology works well not only over different forecasting horizons but also across different economies.

TABLE 6
Out-of-Sample Forecasting Performance for U.K. Inflation Rates

Table 6 reports the out-of-sample forecasting performance on the U.K. inflation-rate series defined on the U.K. retail price index at forecasting horizons of 1, 2, 3, 4, and 5 years. The performance measures include both a forecasting R^2 measure and the Diebold–Mariano (DM) (1995) t-statistic defined on the squared forecasting error differences, which has standard normal critical values.

	Horizon (years)				
	1	2	3	4	5
R^2	0.15	0.26	0.39	0.30	0.18
DM	1.66	1.06	0.86	0.65	0.34

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

VI. Robustness Analysis

This section examines the robustness of our proposed forecasting methodology in terms of specification variations and performance metrics.

A. Incorporating Monetary-Policy Inertia

Several researchers (e.g., Clarida et al. (2000), Levin, Wieland, and Williams (1999)) observe that the actual interest-rate response to economic conditions appears to show some level of gradualism or policy inertia.[5] In the presence of policy inertia, we can replace the actual interest rate r_t in the policy rule of equation (1) with an interest-rate target ($\bar{r}_t$),

$$\bar{r}_t = \alpha + \beta\pi_t + x_t, \tag{19}$$

and allow the actual interest rate r_t to move toward this target $\bar{r}_t$ gradually over time,

$$\bar{r}_t - r_t = \rho(r_t - r_{t-1}), \tag{20}$$

where $\rho > 0$ captures the degree of policy inertia.

In this section, we modify the original bridge relation in equation (2) to account for the effect of policy inertia,

$$p_t = a + b\bar{r}_t + e_t = a + br_t + c(r_t - r_{t-1}) + e, \tag{21}$$

where the second term adjusts the original bridge equation for the effect of policy inertia, with $c = b\rho$.

Table 7 reports the full-sample regression estimates of the extended specification. Compared with the results from the original specification in Table 2, the R^2 estimates become marginally larger, and the slope estimates $\widehat{b}$ become slightly smaller. The inertia coefficient $\widehat{c}$ estimates are large and significant for CPI and PCE but smaller and no longer statistically significant for the two core inflation measures.

TABLE 7

Linking Inflation Rates to Short-Term Interest Rates with Policy Inertia

Table 7 reports the coefficient estimates, their Newey–West (1987) standard errors (in parentheses), and the R^2 values from regressing year-over-year realized inflation rates on the 3-month Treasury interest-rate levels and the year-over-year changes. The Newey–West standard errors are computed with 36 lags. The CPI, CCPI, PCE, and CPCE indexes are defined in Table 1.

Price Index	$\widehat{a}$		$\widehat{b}$		$\widehat{c}$		R^2
CPI	1.15	(0.39)	0.53	(0.09)	0.51	(0.17)	0.58
CCPI	0.92	(0.37)	0.57	(0.08)	0.20	(0.16)	0.64
PCE	0.96	(0.33)	0.48	(0.07)	0.31	(0.14)	0.54
CPCE	0.84	(0.29)	0.49	(0.06)	0.06	(0.10)	0.59

[5]There are also active debates on the existence of monetary-policy inertia (Rudebusch (2002), (2006), Consolo and Favero (2009)) and under what conditions it is optimal to exhibit monetary-policy inertia (Woodford (1999), Sack (2000)).

Table 8 reports the out-of-sample forecasting performance for the 4 inflation series when the specification is adjusted for the monetary-policy inertia. The performance is similar to the original performance in Table 4, somewhat worse for CPI, CCPI, and PCE but slightly better for CPCE. The overall differences are small, suggesting that our year-over-year averages partially smooth out the short-term timing issue and that the long-term inflation predictions are not sensitive to the short-run adjustment behavior.

TABLE 8

Out-of-Sample Forecasting Performance for U.S. Inflation Rates under Policy Inertia

Table 8 reports the out-of-sample forecasting performance on four U.S. inflation-rate series at forecasting horizons of 1, 2, 3, 4, and 5 years when we adjust the specification to accommodate monetary-policy inertia. The performance measures include both a forecasting R^2 measure and the Diebold–Mariano (DM) (1995) t-statistic defined on the squared forecasting error differences. The models are estimated with a 10-year-rolling window. The CPI, CCPI, PCE, and CPCE indexes are defined in Table 1.

Horizon (years)	CPI		CCPI		PCE		CPCE	
	R^2	DM	R^2	DM	R^2	DM	R^2	DM
1	−0.11	−0.60	0.02	0.11	−0.09	−0.43	0.11	0.54
2	0.32	1.62	0.31	1.51	0.29	1.53	0.30	1.51
3	0.45	1.54	0.45	1.56	0.42	1.49	0.38	1.61
4	0.46	1.48	0.47	1.44	0.40	1.35	0.37	1.38
5	0.33	0.98	0.31	1.08	0.28	0.96	0.24	1.03

B. Choice of Rolling-Window Size

Hansen and Timmermann (2012) discuss how the out-of-sample predictive performance can vary greatly with the choice of rolling-window size. Rossi and Inoue (2012) propose to use a large range of window sizes to gain robustness. In our view, the choice of window size is not a purely statistical issue. Economic considerations are equally important. In our particular application, choosing a window size shorter than the average duration of a business cycle is likely to generate large and spurious variations in the estimated economic linkage between inflation and interest rates. Conversely, increasing the window size not only reduces the actual number of observations for the out-of-sample tests, thus reducing the statistical power of the tests, but also runs the risk of suffering from biases induced by structural regime changes.

To examine how the choice of window size affects the inflation-forecasting performance, Table 9 reports the out-of-sample forecasting exercise with different rolling-window sizes: 5-year window (Panel A), 15-year window (Panel B), and recursive (Panel C). With the 5-year window, we start the forecasts in June 1967 and hence obtain more observations for the statistical tests. Thus, although the R^2 estimates become slightly lower than those from the 10-year rolling-window estimation in Table 4, the DM statistics remain as strong in most cases. When we use a 15-year window, the number of out-of-sample observations becomes much smaller. Accordingly, the DM statistics become weaker. Nevertheless, the overall performance patterns are similar across the different window sizes. Using a recursive approach, we start the forecasts in June 1972, which gives the same 10-year initial window, and the R^2 estimates are similar to those from the 10-year rolling-window results, with somewhat lower DM statistics.

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

TABLE 9

Effects of Choice of Rolling-Window Size on Out-of-Sample Inflation Forecasting Performance

Table 9 reports the out-of-sample forecasting performance for four U.S. inflation-rate series at forecasting horizons of 1, 2, 3, 4, and 5 years when we perform the rolling-window estimation using different rolling-window sizes. The performance measures include both a forecasting R^2 measure and the Diebold–Mariano (DM) (1995) t-statistic defined on the squared forecasting error differences. The models are estimated with rolling windows of 5 years (Panel A) and 15 years (Panel B). Panel C uses an expanding window starting at 10 years. The CPI, CCPI, PCE, and CPCE indexes are defined in Table 1.

Horizon	CPI		CCPI		PCE		CPCE	
(years)	R^2	DM	R^2	DM	R^2	DM	R^2	DM
Panel A. 5-Year Rolling Window								
1	−0.02	−0.12	−0.07	−0.54	0.02	0.16	−0.07	−0.59
2	0.29	1.82	0.19	1.28	0.24	1.74	0.08	0.66
3	0.44	1.76	0.36	1.44	0.40	1.81	0.23	1.30
4	0.41	1.48	0.40	1.29	0.38	1.54	0.27	1.29
5	0.27	0.91	0.26	0.90	0.22	0.94	0.16	0.93
Panel B. 15-Year Rolling Window								
1	0.15	1.59	−0.01	−0.06	0.12	1.08	−0.13	−0.77
2	0.36	2.01	0.29	1.54	0.31	1.79	0.15	0.85
3	0.48	1.71	0.48	1.53	0.44	1.58	0.30	1.13
4	0.44	1.40	0.49	1.27	0.40	1.26	0.31	0.94
5	0.29	0.84	0.29	0.78	0.25	0.74	0.18	0.54
Panel C. Recursive Approach								
1	0.14	1.38	−0.05	−0.37	0.10	0.84	−0.21	−1.17
2	0.37	1.95	0.27	1.30	0.32	1.70	0.11	0.55
3	0.48	1.66	0.44	1.33	0.44	1.48	0.25	0.80
4	0.40	1.22	0.42	1.02	0.35	1.01	0.22	0.56
5	0.21	0.57	0.18	0.43	0.15	0.40	0.04	0.10

C. Policy-Rate Proxy

Our main analysis uses the 3-month spot interest rate to proxy the policy rate. The rates constructed at shorter maturities from the Nelson–Siegel (1987) coefficients show numeric instability. Nevertheless, we can investigate how the forecasting performance varies if we choose a longer maturity interest rate as the policy-rate proxy. Table 10 reports the out-of-sample forecasting performance for the 4 sets of inflation rates when we use 6-month (Panel A) and 12-month (Panel B) interest rates as the policy-rate proxy. Using either the 6-month or the 12-month interest rate as the workhorse of the inflation forecasting, we observe similar patterns in both R^2 and DM statistics.

D. Mean Absolute Error as an Alternative Loss Function

Although the squared forecasting error is the most commonly used metric for loss-function definition, there are many alternatives (Giacomini and White (2006)). To gauge the robustness of our results, this subsection considers the mean absolute error as an alternative loss measure. The relative mean absolute error (RMAE) measure is defined as the ratio of the mean absolute forecasting errors of our model against the random-walk benchmark. We compute the corresponding DM test statistics based on the absolute forecasting-error differences. Table 11 reports the RMAE ratio and the corresponding DM statistics. The overall pattern is quite similar to the R^2 measure. For CPI and PCE, our model statistically outperforms the random-walk benchmark at the 2- and 3-year horizons. For CCPI, the outperformance is statistically significant at the 3-year-ahead horizon.

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

TABLE 10

Effects of Benchmark-Rate Choice on Out-of-Sample Inflation Forecasting Performance

Table 10 reports the out-of-sample forecasting performance for four U.S. inflation-rate series at forecasting horizons of 1, 2, 3, 4, and 5 years. The performance measures include both a forecasting R^2 measure and the Diebold–Mariano (DM) (1995) *t*-statistic defined on squared forecasting error differences. Panel A reports the statistics using the 6-month interest rate as the benchmark short rate, and Panel B reports the statistics using the 1-year interest rate as the benchmark short rate. The forecasts are computed based on a 10-year rolling-window estimation. The CPI, CCPI, PCE, and CPCE indexes are defined in Table 1.

Horizon	CPI		CCPI		PCE		CPCE	
(years)	R^2	DM	R^2	DM	R^2	DM	R^2	DM
Panel A. 6-Month Interest Rate as Benchmark								
1	0.08	0.91	−0.09	−0.63	0.06	0.63	−0.15	−0.90
2	0.33	2.11	0.25	1.59	0.29	1.95	0.15	1.10
3	0.48	1.84	0.45	1.70	0.45	1.81	0.32	1.53
4	0.46	1.53	0.51	1.48	0.42	1.47	0.35	1.30
5	0.33	1.03	0.36	1.06	0.29	1.00	0.26	0.96
Panel B. 1-Year Interest Rate as Benchmark								
1	0.06	0.79	−0.09	−0.65	0.04	0.48	−0.14	−0.97
2	0.29	2.13	0.22	1.62	0.25	1.99	0.13	1.07
3	0.44	1.83	0.42	1.72	0.41	1.81	0.29	1.57
4	0.42	1.51	0.48	1.49	0.38	1.45	0.32	1.30
5	0.29	1.03	0.34	1.10	0.25	0.98	0.23	0.94

TABLE 11

Out-of-Sample Forecasting Performance for U.S. Inflation Rates Based on Mean Absolute Errors

Table 11 reports the out-of-sample forecasting performance on four U.S. inflation-rate series at forecasting horizons of 1, 2, 3, 4, and 5 years. The performance measures include both a forecasting relative mean absolute error (RMAE) measure, defined as the ratio of the mean absolute forecasting errors of each model against the random-walk benchmark, and the Diebold–Mariano (DM) (1995) *t*-statistic defined on the absolute error differences. The models are estimated with a 10-year-rolling window. The out-of-sample statistics are computed using data from June 1972 to Dec. 2016. The CPI, CCPI, PCE, and CPCE indexes are defined in Table 1.

Horizon	CPI		CCPI		PCE		CPCE	
(years)	RMAE	DM	RMAE	DM	RMAE	DM	RMAE	DM
1	0.95	1.05	1.04	−0.59	0.96	0.70	1.05	−0.72
2	0.84	2.27	0.90	1.23	0.87	1.82	0.98	0.29
3	0.76	2.13	0.77	1.90	0.79	1.79	0.88	1.31
4	0.78	1.56	0.76	1.53	0.82	1.26	0.84	1.06
5	0.92	0.70	0.84	1.16	0.92	0.75	0.94	0.48

VII. Concluding Remarks

Predictive regressions have been used extensively to determine the existence and statistical significance of many predictive relations; nevertheless, they have often been found to be ineffective in generating superior out-of-sample forecasting performance. A major issue with predictive regressions in out-of-sample tests is that even with a truly predictive relation, large estimation errors on the coefficients can overwhelm the information content. In this article, with regard to inflation-rate forecasting, we propose a new approach for estimating a predictive relation without resorting to standard predictive regressions. The approach drastically reduces the coefficient estimation error and significantly enhances the out-of-sample forecasting performance of a simple predictive relation.

Our forecasting specification contains two major elements: a contemporaneous relation between inflation rates and short-term interest rates motivated by monetary-policy rules and a forecasting specification of future interest rates using

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

the current forward-interest-rate curve. Combining the two elements generates a predictive relation that predicts future inflation-rate changes with the current forward-spot interest-rate differential. Although the use of the forward-interest-rate curve to predict future inflation rates is not new, the way we estimate the relation without resorting to predictive regressions is. Through our 2-step procedure, we effectively determine the slope of the predictive relation through the estimation of a contemporaneous relation. When we estimate the same predictive relation with standard predictive regressions, the out-of-sample predictive performance deteriorates significantly. Thus, to obtain superior forecasts, it is not only important to identify predictively informative variables but also equally important to find robust ways of estimating the predictive relation.

References

Aiolfi, M.; C. Capistran; and A. Timmermann. "Forecast Combinations." In *Oxford Handbook of Economic Forecasting*, M. P. Clements and D. F. Hendry, eds. Oxford, UK: Oxford University Press (2010).

Ang, A.; G. Bekaert; and M. Wei. "Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?" *Journal of Monetary Economics*, 54 (2007), 1163–1212.

Bai, J., and L. Wu. "Anchoring Corporate Credit Default Swap Spreads to Firm Fundamentals." *Journal of Financial and Quantitative Analysis*, 51 (2016), 1521–1543.

Bali, T.; M. Heidari; and L. Wu. "Predictability of Interest Rates and Interest-Rate Portfolios." *Journal of Business and Economic Statistics*, 27 (2009), 517–527.

Boivin, J., and M. P. Giannoni. "Has Monetary Policy Become More Effective?" *Review of Economics and Statistics*, 88 (2006), 445–462.

Brav, A.; R. Lehavy; and R. Michaely. "Using Expectations to Test Asset Pricing Models." *Financial Management*, 34 (2005), 31–64.

Campbell, J. Y., and S. B. Thompson. "Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?" *Review of Financial Studies*, 21 (2008), 1509–1531.

Campbell, J. Y., and M. Yogo. "Efficient Tests of Stock Return Predictability." *Journal of Financial Economics*, 81 (2006), 27–60.

Carlstrom, C. T., and T. S. Fuerst. "Forward-Looking versus Backward-Looking Taylor Rules." Working Paper 0009, Federal Reserve Bank of Cleveland (2000).

Carr, P., and L. Wu. "A Simple Robust Link between American Puts and Credit Protection." *Review of Financial Studies*, 24 (2011), 473–505.

Carr, P., and L. Wu. "Decomposing Long Bond Returns: A Decentralized Modeling Approach." Working Paper, New York University and Baruch College (2017).

Cieslak, A., and P. Povala. "Expected Returns in Treasury Bonds." *Review of Financial Studies*, 28 (2015), 2859–2901.

Clarida, R.; J. Gali; and M. Gertler. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature*, 37 (1999), 1661–1707.

Clarida, R.; J. Gali; and M. Gertler. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *Quarterly Journal of Economics*, 115 (2000), 147–180.

Clark, T. E., and T. Doh. "A Bayesian Evaluation of Alternative Models of Trend Inflation." Working Paper, Federal Reserve Bank of Cleveland and Federal Reserve Bank of Kansas City (2011).

Clark, T. E., and M. W. McCracken. "Advances in Forecast Evaluation." In *Handbook of Economic Forecasting*, Vol. 2, A. Timmermann and G. Elliott, eds. Amsterdam, The Netherlands: Elsevier (2012).

Cochrane, J. H. "Determinacy and Identification with Taylor Rules." *Journal of Political Economy*, 119 (2011), 565–615.

Cogley, T., and A. M. Sbordone. "Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve." *American Economic Review*, 98 (2008), 2102–2126.

Consolo, A., and C. A. Favero. "Monetary Policy Inertia: More a Fiction than a Fact?" *Journal of Monetary Economics*, 56 (2009), 900–906.

De Graeve, F.; M. Emiris; and R. Wouters. "A Structural Decomposition of the US Yield Curve." *Journal of Monetary Economics*, 56 (2008), 545–559.

Diebold, F. X., and R. S. Mariano. "Comparing Predictive Accuracy." *Journal of Business and Economic Statistics*, 13 (1995), 253–263.

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

Diebold, F. X.; G. D. Rudebusch; and S. B. Arouba. "The Macroeconomy and the Yield Curve: A Dynamic Latent Factor Approach." *Journal of Econometrics*, 131 (2006), 309–338.
Duffee, G. R. "Term Premia and Interest Rate Forecasts in Affine Models." *Journal of Finance*, 57 (2002), 405–443.
Duffee, G. R. "Forecasting Interest Rates." In *Handbook of Economic Forecasting*, Vol. 2, G. Elliott and A. Timmermann, eds. Amsterdam, The Netherlands: Elsevier (2013).
Elliott, G., and J. H. Stock. "Inference in Time Series Regression When the Order of Integration of a Regressor Is Unknown." *Econometric Theory*, 10 (1994), 672–700.
Engel, C., and K. West. "Exchange Rates and Fundamentals." *Journal of Political Economy*, 113 (2005), 485–517.
Engle, R. F., and C. W. Granger. "Cointegration and Error Correction: Representation, Estimation and Testing." *Econometrica*, 55 (1987), 251–276.
Fama, E. F. "Short-Term Interest Rates as Predictors of Inflation." *American Economic Review*, 65 (1975), 269–282.
Fama, E. F. "Forward and Spot Exchange Rates." *Journal of Monetary Economics*, 14 (1984), 319–338.
Fama, E. F., and R. R. Bliss. "The Information in Long-Maturity Forward Rates." *American Economic Review*, 77 (1987), 680–692.
Faust, J., and J. H. Wright. "Forecasting Inflation." In *Handbook of Economic Forecasting*, Vol. 2, G. Elliott and A. Timmermann, eds. Amsterdam, The Netherlands: Elsevier (2013).
Forni, M.; M. Hallin; M. Lippi; and L. Reichlin. "Do Financial Variables Help Forecasting Inflation and Real Activity in the Euro Area?" *Journal of Monetary Economics*, 50 (2003), 1243–1255.
Frankel, J. A., and C. S. Lown. "An Indicator of Future Inflation Extracted from the Steepness of the Interest Rate Yield Curve along Its Entire Length." *Quarterly Journal of Economics*, 109 (1994), 517–530.
Giacomini, R., and H. White. "Tests of Conditional Predictive Ability." *Econometrica*, 74 (2006), 1545–1578.
Goodfriend, M. "Interest-Rate Smoothing and Price Level Trend-Stationarity." *Journal of Monetary Economics*, 19 (1987), 335–348.
Goodfriend, M. "The Phases of U.S. Monetary Policy: 1987 to 2001." *Federal Reserve Bank of Richmond Economic Quarterly*, 88 (2002), 1–17.
Gurkaynak, R. S.; B. Sack; and J. H. Wright. "The U.S. Treasury Yield Curve: 1961 to the Present." *Journal of Monetary Economics*, 54 (2007), 2291–2304.
Hansen, P. R., and A. Timmermann. "Choice of Sample Split in Out-of-Sample Forecast Evaluation." Working Paper, UCSD and CREATES (2012).
Harvey, D. I.; S. J. Leybourne; and P. Newbold. "Testing the Equality of Prediction Mean Squared Errors." *International Journal of Forecasting*, 13 (1997), 281–291.
Inoue, A., and L. Kilian. "In-Sample or Out-of-Sample Tests of Predictability: Which One Should We Use?" *Econometric Reviews*, 23 (2004), 371–402.
Jansson, M., and M. J. Moreira. "Optimal Inference in Regression Models with Nearly Integrated Regressors." *Econometrica*, 74 (2006), 681–714.
Jorion, P., and F. Mishkin. "A Multicountry Comparison of Term-Structure Forecasts at Long Horizons." *Journal of Financial Economics*, 29 (1991), 59–80.
Kelly, B., and S. Pruitt. "Market Expectations in the Cross-Section of Present Values." *Journal of Finance*, 68 (2013), 1721–1756.
Kim, C.-J., and C. R. Nelson. "Estimation of a Forward-Looking Monetary Policy Rule: A Time-Varying Parameter Model Using Ex-Post Data." *Journal of Monetary Economics*, 53 (2006), 1949–1966.
Kim, D. H., and A. Orphanides. "Term Structure Estimation with Survey Data on Interest Rate Forecasts." *Journal of Financial and Quantitative Analysis*, 47 (2012), 241–272.
Levin, A.; V. Wieland; and J. C. Williams. "Robustness of Simple Monetary Policy Rules under Model Uncertainty." In *Monetary Policy Rules*, J. B. Taylor, ed. Chicago, IL: University of Chicago Press (1999).
Mishkin, F. "The Information in the Term Structure: Some Further Results." *Journal of Applied Econometrics*, 3 (1988), 307–314.
Mishkin, F. S. "What Does the Term Structure Tell Us about Future Inflation?" *Journal of Monetary Economics*, 25 (1990), 77–95.
Nelson, C. R., and A. F. Siegel. "Parsimonious Modeling of Yield Curves." *Journal of Business*, 60 (1987), 473–489.
Newey, W. K., and K. D. West. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55 (1987), 703–708.

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press

Orphanides, A. "Monetary-Policy Rules and the Great Inflation." *American Economic Review*, 92 (2002), 115–120.
Osborne, G. "Review of the Monetary Policy Framework." Report, HM Treasury (2013).
Phillips, P. "Pitfalls and Possibilities in Predictive Regressions." SoFie Conference Presentation, Yale University (2013).
Rossi, B., and A. Inoue. "Out-of-Sample Forecast Tests Robust to the Choice of Window Size." *Journal of Business and Economic Statistics*, 30 (2012), 432–453.
Rudebusch, G. D. "Term Structure Evidence on Interest Rate Smoothing and Monetary Policy Inertia." *Journal of Monetary Economics*, 49 (2002), 1161–1187.
Rudebusch, G. D. "Monetary Policy Inertia: Fact or Fiction?" *International Journal of Central Banking*, 2 (2006), 85–135.
Sack, B. "Does the Fed Act Gradually? A VAR Analysis." *Journal of Monetary Economics*, 46 (2000), 229–256.
Sims, C. A., and T. Zha. "Were There Regime Switches in the U.S. Monetary Policy?" *American Economic Review*, 96 (2006), 54–81.
Sims, E. R. "Identification and Estimation of Interest Rate Rules in the New Keynesian Model." Working Paper, University of Notre Dame (2008).
Stambaugh, R. F. "Predictive Regressions." *Journal of Financial Economics*, 54 (1999), 375–421.
Stock, J. H., and M. Watson. "Why Has Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking*, 39 (2007), 3–33.
Stock, J. H., and M. Watson. "Modeling Inflation after the Crisis." In *Proceedings—Economic Policy Symposium, Jackson Hole*, Kansas City, MO: Federal Reserve Bank of Kansas City (2010).
Stock, J. H., and M. W. Watson. "Forecasting Output and Inflation: The Role of Asset Prices." *Journal of Economic Literature*, 41 (2003), 788–829.
Svensson, L. E. O. "What Is Wrong with Taylor Rules? Using Judgement in Monetary Policy through Targeting Rules." *Journal of Economic Literature*, 41 (2003), 426–477.
Svensson, L. E. O. "Monetary Policy with Judgement: Forecast Targeting." *International Journal of Central Banking*, 1 (2005), 1–54.
Taylor, J. B. "Discretion vs. Policy Rules in Practice." *Carnegie-Rochester Conference Series on Public Policy*, 39 (1993), 195–214.
Taylor, J. B. "A Historical Analysis of Monetary Policy Rules." In *Monetary Policy Rules*, J. B. Taylor, ed. Chicago, IL: University of Chicago Press (1999).
Walsh, C. E. *Monetary Theory and Policy*, 2nd ed. Cambridge, MA: MIT Press (2003).
Welch, I., and A. Goyal. "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction." *Review of Financial Studies*, 21 (2008), 1455–1508.
Woodford, M. "Optimal Monetary Policy Inertia." *The Manchester School*, 67 (1999), 1–35.
Woodford, M. *Interest and Prices*. Princeton, NJ: Princeton University Press (2003).
Wu, L. "Estimating Risk-Return Relations with Analysts Price Targets." *Journal of Banking and Finance*, 93 (2018), 183–197.

https://doi.org/10.1017/S0022109018000467 Published online by Cambridge University Press