

NBER WORKING PAPER SERIES

LEARNING ABOUT THE LONG RUN

Leland Farmer
Emi Nakamura
Jón Steinsson

Working Paper 29495
<http://www.nber.org/papers/w29495>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2021

We thank Abhi Gupta and Ethan McClure for excellent research assistance. We thank Assaf Ben-Shoham, Daniel Benjamin, John Campbell, Gary Chamberlain, Anna Cieslak, Nicolas Chopin, Gopi Gaswami, Michael Johannes, David Laibson, Omiros Papaspiliopoulos and seminar participants at various institutions for valuable comments and discussions. We thank the Alfred P. Sloan Foundation, the Smith Richardson Foundation, and the Bankard Fund for Political Economy at the University of Virginia for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Leland Farmer, Emi Nakamura, and Jón Steinsson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Learning About the Long Run
Leland Farmer, Emi Nakamura, and Jón Steinsson
NBER Working Paper No. 29495
November 2021
JEL No. E37,E47,G12

ABSTRACT

Forecasts of professional forecasters are anomalous: they are biased, forecast errors are autocorrelated, and forecast revisions predict forecast errors. Sticky or noisy information models seem like unlikely explanations for these anomalies: professional forecasters pay attention constantly and have precise knowledge of the data in question. We propose that these anomalies arise because professional forecasters don't know the model that generates the data. We show that Bayesian agents learning about hard-to-learn features of the data generating process (low frequency behavior) can generate all the prominent aggregate anomalies emphasized in the literature. We show this for two applications: professional forecasts of nominal interest rates for the sample period 1980-2019 and CBO forecasts of GDP growth for the sample period 1976-2019. Our learning model for interest rates also provides an explanation for deviations from the expectations hypothesis of the term structure that does not rely on time-variation in risk premia.

Leland Farmer
Department of Economics
College and Graduate School of Arts & Science
University of Virginia
PO Box 400182
Charlottesville, VA 22904-418
farmer.leland@gmail.com

Jón Steinsson
Department of Economics
University of California, Berkeley
671 Evans Hall
Berkeley, CA 94720
and NBER
jsteinsson@berkeley.edu

Emi Nakamura
Department of Economics
University of California, Berkeley
685 Evans Hall
Berkeley, CA 94720
and NBER
enakamura@berkeley.edu

1 Introduction

For almost half a century, the assumption that people form rational expectations has dominated economic modelling in macroeconomics and finance. During this time, a substantial empirical literature has formulated and evaluated tests of rational expectations. One finding from this literature has been that even professional forecasters consistently fail such tests. Professional forecasts seem to suffer from a long list of “anomalies.” For example, they are biased, forecast errors are autocorrelated, and forecast revisions predict future forecast errors.

A related literature has tested the expectations hypothesis of the term structure. If the expectations hypothesis holds, yields on long-term bonds are the bond market’s forecast of future short rates (modulo a constant risk premium). Empirical tests of the expectations hypothesis fail spectacularly (e.g., [Campbell and Shiller, 1991](#)). One reaction to this finding is that risk premia in the bond market are time varying ([Wachter, 2006](#); [Bansal and Shaliastovich, 2013](#)). An alternative view is that this finding reflects forecasting anomalies among bond traders ([Froot, 1989](#)).¹

The traditional reaction to forecasting anomalies in macroeconomics is that they imply that professional forecasters are irrational, i.e., that forecasters are not making efficient use of the information available to them ([Mincer and Zarnowitz, 1969](#); [Friedman, 1980](#); [Nordhaus, 1987](#); [Maddala, 1991](#); [Croushore, 1998](#); [Schuh, 2001](#)). Recent behavioral work develops this perspective (e.g. [Bordalo et al., 2020](#)). An alternative reaction is that these anomalies result from information frictions ([Mankiw et al., 2003](#); [Coibion and Gorodnichenko, 2012, 2015](#)). The most prominent models of information frictions in macroeconomics are sticky information models ([Mankiw and Reis, 2002](#)) and noisy information models ([Sims, 2003](#); [Woodford, 2003](#)). These models seem eminently plausible for households and firms. Arguably, they are less well suited to explain the behavior of professional forecasters (and bond traders). Professional forecasters read the news every day and have no trouble observing the relevant data precisely (i.e., without noise).

In this paper, we consider another explanation. Standard tests of rational expectations impose the very strong assumption that agents know the model that generates the variables that are being forecast (parameter values and all). In reality, nobody knows the correct model of the world. Since professional forecasters don’t know the correct model of the world, they use incoming data to learn about how the world works. But such learning can fundamentally change the dynamics of even perfectly rational Bayesian forecasts. This idea has been recognized by researchers at least

¹See also [Bekaert, Hodrick, and Marshall \(2001\)](#), [Piazzesi, Salomao, and Schneider \(2015\)](#), [Cieslak \(2018\)](#), [Xu \(2019\)](#), and [Nagel and Xu \(2021\)](#).

since [Friedman \(1979\)](#).²

Realistic learning models are difficult to solve. As a consequence, early work on learning used relatively simple models. But in such models, Bayesian learning occurs quickly, suggesting that rational learning can't explain forecasting anomalies that persist over multiple decades. Structural breaks have sometimes been invoked as a reason why learning might persist over long periods of time. But such arguments have only been made informally.

Bayesian learning can, however, be extremely slow in richer, more realistic models ([Johannes, Lochstoer, and Mou, 2016](#)). Consider, for example, models with multiple unobserved components some stationary and other containing a unit root. A key property of such models is that the long-run trajectory of a variable may move quite independently from the short-run dynamics of that variable (if the short-run dynamics are dominated by the stationary components). This means that the quarter-to-quarter dynamics of the variable may be quite uninformative about its longer-run properties. Since information about low-frequency properties accumulates slowly, learning about the long run can be extremely slow. In such models, several different parameter combinations may yield similar fit for the high-frequency behavior of the series but may have very different implications about the low-frequency behavior of the series. We show that in such cases it can take many decades to learn the true parameters.³

We develop two applications of these ideas, one for forecasting nominal interest rates and another for forecasting real GDP growth. In each case, we endow Bayesian forecasters with an unobserved components model and initial beliefs about the parameters of this model. Each period, these agents use real-time U.S. data to update their beliefs about the parameters and states of the model. They then forecast the variable in question and we assess whether the resulting forecasts are "anomalous."

Our main result is that we are able to match all the main forecasting anomalies emphasized in the prior literature for both interest rates and real GDP when forecasters are endowed with "reasonable" initial beliefs. In addition, we construct long-term yield data from our model-generated forecasts of nominal interest rates assuming that the expectations hypothesis holds. We then run a battery of standard tests of the expectations hypothesis on these data. The model-generated yield data fail the tests of the expectations hypothesis in exactly the same way as do real-world bond

²Other important papers that emphasize this idea include [Lewis \(1989b,a\)](#), [Barsky and De Long \(1993\)](#), [Timmermann \(1993\)](#), [Lewellen and Shanken \(2002\)](#), [Brav and Heaton \(2002\)](#), [Cogley and Sargent \(2005\)](#), [Collin-Dufresne, Johannes, and Lochstoer \(2016\)](#), [Johannes, Lochstoer, and Mou \(2016\)](#), and [Guo and Wachter \(2019\)](#).

³[Collin-Dufresne, Johannes, and Lochstoer \(2016\)](#) and [Kozlowski, Veldkamp, and Venkateswaran \(2020\)](#) develop models with rare events in which Bayesian learning is slow.

yields. Notably, our sample period is roughly 40 years and we endow our Bayesian agents with data back to the early post-WWII period. Even though they learn for quite a few decades, agents' forecasts continue to display anomalies.⁴

Despite being quite simple, some of the parameters of our unobserved components models are very difficult to learn, even over multiple decades. As we discuss above, it is hard to tell apart stochastic processes with similar implications for short-run dynamics even if they have very different implications for long-run dynamics. As a consequence, Bayesian learning can be very slow. Yet, the long-run dynamics are essential for predictions several years into the future. It is for this same reason that unit root tests have low power in "short sample" (short often being many decades).

Since learning is slow, agents' initial beliefs matter for a long time. If our ability to match forecast anomalies with "rational learning" relied on unreasonable initial beliefs (e.g., very tight initial beliefs at implausible parameter values), it would not be very interesting. We show, however, that we can match the anomalous features of the forecast data for "reasonable" initial beliefs. In particular, the initial beliefs we endow agents with are quite dispersed. They also accord well with historical experience prior to our sample period. For example, our agents place small weight in 1951 on the possibility that the nominal interest rate has a large random walk component. This is consistent with the fact that (outside of war) the U.S. had been on a gold (or silver) standard almost continuously from its founding until that point in time and interest rates had therefore been quite stable at low frequencies. The large and persistent rise and fall in nominal interest rates that occurred subsequently was far outside of what had been experienced up to that point in history.

A potential concern with our results is that perhaps we are able to match the forecast anomalies we emphasize because we endow agents with a misspecified model. To address this concern and understand better what features of agents' initial beliefs drive our results, we conduct a Monte Carlo simulation of our model for nominal interest rates. In this case, we know the true data generating process and thus know that the agents in our model are not learning using a misspecified model. We show that when initial beliefs are centered on parameters that imply too little persistence in interest rates relative to the truth, our model generates the kinds of anomalies we find in the data.⁵ In contrast, if initial beliefs are centered on parameters that imply too much

⁴For interest rates, we focus on the mean forecast from the Survey of Professional Forecasters. An important related literature characterizes and seeks to explain the behavior of individual forecasts relative to the mean, as well as the behavior of forecast dispersion (see, e.g., [Andrade et al., 2016](#); [Angeletos, Huo, and Sastry, 2020](#); [Bordalo et al., 2020](#); [Cao et al., 2021](#); [Crump et al., 2021](#); [Singleton, 2021](#); and references therein)

⁵This result is similar in spirit to results in [Gourinchas and Tornell \(2004\)](#) about exchange rates.

persistence, our model generates anomalies in the opposite direction (negatively autocorrelated forecast errors and over-reaction rather than under-reaction in [Coibion and Gorodnichenko \(2015\)](#) regressions. etc.). If initial beliefs happen to be exactly centered on the true values in our Monte Carlo, no anomalies arise.

In the Monte Carlo simulations, we know what the truth is. When it comes to the real world, there is no way of knowing what the truth is without learning, and learning about the long run is extremely slow. In our Monte Carlo simulation, a decade is a “blink in the eye” in terms of learning about key parameters of in our model. Even after agents have been learning for 70 years, they are still very far from the truth and are inching towards the truth extremely slowly. These results illustrate how, rational expectations tests can be very misleading even when run over “long” periods of time.

Whether anomalies arise from Bayesian learning about parameters depends crucially on the nature of the data. If the fluctuations in the variables of interest are homoscedastic and not very persistent, agents will learn model parameters relatively quickly and none of the issues described above will arise. In contrast, models in which learning has been shown to be important include long-run risk models and models with disasters ([Croce, Lettau, and Ludvigson, 2015](#); [Collin-Dufresne, Johannes, and Lochstoer, 2016](#); [Kozlowski, Veldkamp, and Venkateswaran, 2020](#); [Bidder and Dew-Becker, 2016](#)). [Giacoletti, Laursen, and Singleton \(2018\)](#) also study a learning model, but consider a case where some model parameters are assumed to be known, leading to a substantially simplified learning rule. We focus on the challenges arising from the agents’ lack of knowledge of the model parameters which makes the learning problem much harder and learning much slower. Our work also relates to a rich literature on boundedly rational learning in macroeconomics (e.g., [Evans and Honkapohja, 2001](#); [Sargent, 2001](#); [Eusepi and Preston, 2011, 2018](#)). [Ben-David, Graham, and Harvey \(2013\)](#) provide evidence for Bayesian learning among firm CFOs.

The paper proceeds as follows. Section 2 describes our data. Section 3 reviews forecasting anomalies for interest rates and real GDP data. Section 4 presents our model and results for nominal interest rates. Section 5 presents our model and results for real GDP growth. Section 6 presents Monte Carlo simulation exercises aimed to shed light on why our results turn out the way they do. Section 7 concludes.




Figure 1: SPF Forecasts of the 3-Month T-Bill Rate

Note: The black solid line is the 3-month T-bill rate. Each short gray line with five circles represents the SPF forecasts made in a particular quarter about the then present quarter (first circle) and following four quarters (subsequent four circles).

2 Data

The paper discusses two applications, one to interest rate forecasting and the other to real GDP forecasting. We describe the data we use for these two applications in turn.

2.1 Interest Rate Data and Forecasts

The forecast data we use for the 3-month Treasury Bill (T-Bill) rate come from the Survey of Professional Forecasters (SPF) conducted by the Federal Reserve Bank of Philadelphia. Our sample period for these forecasts is 1981Q3 to 2019Q4. The SPF is a quarterly survey sent out to a rotating panel of forecasters. We use the mean forecast across forecasters. Figure 1 plots these forecasts.

The survey is sent out near the end of the first month of each quarter. The forecast therefore roughly coincides with the BEA's advance report of the national income and product accounts. Survey response deadlines are in the second to third week of the second month of the same quarter. Survey respondents are asked to provide nowcasts and one to four quarter ahead forecasts of the quarterly average 3-month T-Bill secondary market rate. The timing of these forecasts is as follows: the nowcast pertains to the quarterly average rate at the end of the quarter when the survey is

received, and the subsequent forecasts pertain to quarterly averages for each of the following four quarters.




Figure 2: The 3-Month T-Bill Rate

The data we use on the 3-month T-Bill secondary market rate is from the Board of Governors of the Federal Reserve System . Our sample period for this series is 1951Q2 to 2019Q4. Figure 2 plots the series. To be consistent with the forecast data, we use quarterly averages of the daily interest rate. We also use daily estimates of the zero-coupon yield curve from [Liu and Wu \(2020\)](#). Liu and Wu estimate the zero-coupon yield curve for bonds of maturity 1 month to 30 years (360 months) dating back to June 1961. We convert these data to quarterly data by computing the average yield in a quarter. Our sample period for these zero coupon bond yields is 1961Q3 to 2019Q4.

2.2 Real GDP Growth Data and Forecasts

The real GDP growth forecasts we analyze are from the Congressional Budget Office (CBO). Our sample period for these forecasts is 1976 to 2019. The CBO releases its annual economic outlook at the beginning of each year, where it provides projections for current and future real economic growth. Since 1996, the CBO has made projections out to a horizon of 11 years. Before that, they made projections out to a horizon of 6 years. The CBO forecasts the annual average level of real output over each calendar year. Growth rates are then computed as percentage changes in these




Figure 3: CBO Forecasts of Real GDP Growth

Note: The black solid line is the 2021Q1 vintage of real GDP growth from 1976 to 2019. The broken black line is the initial release of GDP growth at each point in time. Each short gray line with seven circles represents the initial release of real GDP for the previous year (first circle) and the CBO forecasts made in a particular year about GDP growth in the following six years (subsequent six circles).

average levels across years. Up to and including their 1992 report, the CBO forecast real Gross National Product (GNP). Since then, they have forecast real Gross Domestic Product (GDP). For expositional simplicity, we refer to these as real GDP forecasts throughout the paper.

The data we use on actual real GDP growth is from the Philadelphia Federal Reserve Bank’s Real-Time Data Set. This source publishes monthly vintages of real-time real output back to November of 1965. Most vintages contain data back to 1947Q1. However, a few vintages are missing data before 1959Q3, which limits our sample period as we discuss in greater detail in section 5.

3 Forecasting Anomalies

As we discuss in the introduction, the forecasts of professional forecasters exhibit a number of “anomalies”—i.e., patterns that previous researchers have argued suggest deviations from forecast rationality. Here we document a number of such anomalies for professional forecasts of the 3-month nominal T-bill rate and real GDP growth. We also document deviations from the expect-

tations hypothesis of the term structure—which may be interpreted as forecast anomalies on the part of the bond market. The facts we document in this section will be key empirical targets we seek to match with our models later in the paper.

The null hypotheses we consider below constitute tests of forecast rationality given two assumptions: 1) that forecasters aim to minimize the mean squared error of their forecasts, implying that optimal forecasts are equivalent to conditional expectations ($F_t y_{t+h} = E_t y_{t+h}$), and 2) that forecasters know the true model of the world. For the 3-month T-bill, we focus on forecast horizons of one to four *quarters*. For real GDP growth, however, we focus on forecast horizons of one to four *years*. These different forecasting horizons reflect differences in the horizons at which the forecast anomalies are most striking for the 3-month T-bill vs. real GDP growth.

Bias

A straightforward prediction of full-information rational expectations models is that forecasts should be unbiased at all horizons. Let y_t be the variable to be forecast, and let $F_t y_{t+h}$ denote the h -period ahead forecast of y_t given time t information. Define the forecast error as $e_{t+h|t} \equiv y_{t+h} - F_t y_{t+h}$. The forecasting bias can then be estimated using the following regression:

$$e_{t+h|t} = \alpha + u_{t+h}, \tag{1}$$

with $\alpha = 0$ indicating that forecasts are unbiased at a given horizon h .

Panel A of Table 1 displays our estimates of α for the 3-month T-bill rate and real output growth. Our estimates indicate that professional forecasts of the T-bill rate display negative bias—the truth being lower than the forecast on average—at all horizons and the magnitude of this bias increases with the horizon. At the 4-quarter forecast horizon, SPF forecasters overestimate the true T-bill rate by an average of 0.7 percentage points. These biases are statistically significant at the 1% level at all horizons. In contrast, there is little evidence of statistically significant bias in CBO forecasts of GDP growth at the horizons we study.

Autocorrelated Forecast Errors

Another prediction of full-information rational expectations models is that forecast errors should be serially uncorrelated. To assess this prediction, we consider the following regression of h -period ahead forecast errors on their own past value h periods earlier (i.e., we consider the correlation of

Table 1: Forecast Anomalies

	Forecast Horizon				
	1	2	3	4	5
<i>Panel A: Bias</i>					
T-Bill	-0.18*** (0.06)	-0.34*** (0.11)	-0.52*** (0.16)	-0.70*** (0.20)	–
GDP Growth	0.27 (0.21)	-0.27 (0.34)	-0.54 (0.44)	-0.62 (0.47)	-0.52 (0.44)
<i>Panel B: Autocorrelation</i>					
T-Bill	0.30* (0.16)	0.27** (0.11)	0.24** (0.11)	0.13 (0.12)	–
GDP Growth	0.22 (0.16)	0.16 (0.17)	0.11 (0.14)	0.08 (0.18)	0.08 (0.11)
<i>Panel C: Mincer-Zarnowitz</i>					
T-Bill	0.97 (0.02)	0.94* (0.04)	0.90** (0.05)	0.86** (0.06)	–
GDP Growth	0.94 (0.12)	0.60 (0.31)	0.03*** (0.31)	-0.42*** (0.22)	-0.43*** (0.36)
<i>Panel D: Coibion-Gorodnichenko</i>					
T-Bill	0.23* (0.13)	0.34** (0.15)	0.62*** (0.18)	–	–
GDP Growth	0.08 (0.11)	0.00 (0.23)	0.50 (0.52)	-1.63*** (0.41)	-1.46*** (0.45)

Note: The forecast horizons for the T-Bill are quarters, while the forecast horizons for the GDP growth are years. Stars represent significance relative to the following hypotheses: $\alpha = 0$ for bias, $\beta = 0$ for autocorrelation, $\beta = 1$ for Mincer-Zarnowitz, $\beta = 0$ for Coibion-Gorodnichenko. P-values are computed using Newey-West standard errors with lag length selected as $L = \lceil 0.75 \times T^{1/3} \rceil$. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

contiguous, non-overlapping h -period forecasts):

$$e_{t+h|t} = \alpha + \beta e_{t|t-h} + u_{t+h}. \quad (2)$$

In a full-information setting, forecast rationality implies that $\alpha = 0$ and $\beta = 0$, i.e., there should be no bias and forecast errors should not be predictable by known information (the time t forecast error).

Panel B of Table 1 reports our estimates of β from equation (2). SPF forecasts of the T-bill display substantial positive autocorrelation. The 1-quarter forecast has an autocorrelation of 0.30. This falls to 0.24 at three quarters. These estimates are statistically significantly different from zero, especially at horizons two and three. CBO forecasts of GDP growth also display positive

autocorrelation. But in this case the autocorrelation is smaller and not statistically significantly different from zero.

Mincer-Zarnowitz Regressions

A classic test of forecast rationality proposed by [Mincer and Zarnowitz \(1969\)](#) investigates the intuitive prediction that the truth should on average move one-for-one with a rational forecast: when the forecast rises by 1%, on average, the realized value should also rise by 1%. This prediction can be analyzed using the regression

$$y_{t+h} = \alpha + \beta F_t y_{t+h} + u_{t+h}. \quad (3)$$

In a full-information setting, forecast rationality implies that $\alpha = 0$ and $\beta = 1$, i.e., there should be no bias and realized values should move one-for-one with forecasts.

Panel C of [Table 1](#) reports our estimates of β from (3). In this case, it is the GDP growth forecasts that display substantial deviations from the null of forecast rationality. While the estimate of β for the 1-year ahead forecast is close to one, it falls sharply at longer horizons. For the 3-year ahead forecast, we estimate a β close to zero. In other words, actual GDP growth is no more likely to be high when it was forecast to be high three years earlier than when it was forecast to be low three years earlier. For the 4-year and 5-year ahead forecast, we estimate negative values (high forecasted growth predicts low growth on average). These three estimates are strongly statistically significantly different from one. In contrast, our estimate of β for the T-bill forecasts are close to one. They are somewhat below one and the difference is statistically significant. But the difference is much less stark than in the case of GDP forecasts.

Coibion-Gorodnichenko Test

Another property of rational forecasts under full information is that they should not underreact or overreact to new information. [Coibion and Gorodnichenko \(2012\)](#) propose the following regression to assess this:

$$e_{t+h|t} = \alpha + \beta (F_t y_{t+h} - F_{t-1} y_{t+h}) + u_{t+h}.$$

Forecast rationality in a full-information setting implies that $\alpha = 0$ and $\beta = 0$. $F_t y_{t+h} - F_{t-1} y_{t+h}$ is known at time t and forecast errors should not be predictable by known information. If $\beta > 0$, the

forecasts are said to suffer from “underreaction.” In this case, an increase in the forecast predicts a situation where the new forecast is still too low on average, i.e., didn’t increase enough. If $\beta < 0$, the forecasts are said to suffer from “overreaction.”

Panel D of Table 1 reports our estimates of β from (3). In this case, we see opposite anomalies for the two applications we consider. For the T-bill forecasts, we see evidence of underreaction: we estimate positive values for β rising from 0.22 at the 1-quarter horizon to 0.64 at the 3-quarter horizon. For GDP growth forecasts, however, we estimate neither over- nor underreaction at short horizons. At the 4-year and 5-year horizons, however, we estimate negative values of β indicating overreaction.

Failures of the Expectations Hypothesis

The expectations hypothesis of the term structure implies that the yield on an n -period bond should equal the average expected values of yields on 1-period bonds over the lifetime of the n -period bond, up to a constant risk premium. This should hold regardless of the process followed by the short rate. Following [Campbell and Shiller \(1991\)](#) and others, we can test this implication with the following regression:

$$\frac{1}{n} \sum_{i=0}^{n-1} y_{t+i}^{(1)} - y_t^{(1)} = \alpha + \beta(y_t^{(n)} - y_t^{(1)}) + u_t, \quad (4)$$

where $y_t^{(n)}$ denotes the yield of a n -period bond at time t . The expectations hypothesis implies that when the yield spread between short-term and long-term bonds ($y_t^{(n)} - y_t^{(1)}$) is high, short-term bond yields will rise in the future (the dependent variable will be large). Specifically, the expectations hypothesis implies that $\beta = 1$. Early papers estimating equation (4) include [Fama \(1984\)](#) and [Fama and Bliss \(1987\)](#).

The first row in Table 2 presents our estimates of β in equation (4) for bonds of maturity 2 to 40 quarters. Consistent with a large earlier literature, we find that the null hypothesis of $\beta = 1$ is resoundingly rejected at short horizons with β being estimated to be smaller than one. As the horizon grows, our estimate of β rises closer to one, but remains below one for all horizons we consider.

Another implication of the expectations hypothesis of the term structure is that at times when the yield spread is unusually high the yield on long bonds will rise. One intuition for this is that returns must be equalized (modulo a constant) for short-term and long-term bonds. If the yield

Table 2: Failures of the Expectations Hypothesis

	Long Horizon n						
	2	3	4	8	12	20	40
Future Short Rates	-0.01*** (0.22)	0.11*** (0.22)	0.18*** (0.21)	0.39** (0.25)	0.57 (0.27)	0.74 (0.23)	0.71 (0.22)
Change in Long Rate	-1.02*** (0.43)	-0.91*** (0.54)	-1.03*** (0.57)	-1.29*** (0.61)	-1.61*** (0.65)	-2.04*** (0.67)	-2.75*** (0.86)

Note: The sample period is from 1961Q3 to 2019Q4. The top row reports estimates of β from regression (4). The bottom row reports estimates of β from regression (5). In both cases, the horizon n is listed at the top of the table. Stars represent significance relative to the hypothesis that $\beta = 1$. P-values are computed using Newey-West standard errors with lag length selected as $L = \lceil 0.75 \times T^{1/3} \rceil$. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

spread is high, then the long-bond yield needs to rise to reduce the return on the long bond so that it can be equal to that of the short bond. Another intuition is that the high yield spread implies that the short yield will rise over the life of the long bond. As time passes, the relatively low current short rate will then drop out of the sum of future short rates that determines the long yield (according to the expectations hypothesis). As this happens, the sum increases and so the long yield should increase.

We can test this implication of the expectations hypothesis with the following regression:

$$y_{t+1}^{(n-1)} - y_t^{(n)} = \alpha + \beta \left(\frac{1}{n-1} \right) (y_t^{(n)} - y_t^{(1)}) + u_t. \quad (5)$$

It is straightforward to show that the expectations hypothesis implies $\beta = 1$. Early papers estimating equation (5) include Shiller (1979), Shiller, Campbell, and Schoenholtz (1983), and Campbell and Shiller (1991).

The second row of Table 2 presents our estimates of β in equation (5). Consistent with earlier research, we find large deviations from the null implied by the expectations hypothesis. Rather than being equal to one, we estimate values for β around negative one at short horizons and even larger negative values at longer horizons. This means that when the yield spread is large the long rate has tended to fall rather than rise as the expectation hypothesis implies that it should. The conventional interpretation of this result is that it implies large predictable excess returns on the long bond when the yield spread is high. Our results later in the paper question whether these excess returns were truly predictable.

The previous literature has identified a number of potential econometric issues associated with

these tests of the expectations hypothesis. One issue is that, in regression (5), the long-term yield appears in the dependent variable with a negative sign and in the regressor with a positive sign. As a consequence, measurement error in the long yield will bias the estimated coefficient downward and may even result in a negative estimate. [Campbell and Shiller \(1991\)](#) use instrumental variables techniques to assess whether measurement error is the cause of the negative estimates but find that the negative coefficients are quite robust. A second issue is small sample bias. This issue was emphasized for regressions (4) and (5) by [Bekaert, Hodrick, and Marshall \(1997\)](#). We conduct Monte Carlo analysis in section 6 based on our model from section 4. This analysis does find evidence of some small sample biases. But the quantitative magnitude of these biases is small.

4 Learning about Nominal Interest Rates

Traditional tests of forecast rationality evaluate the joint hypothesis that agents form conditional expectations rationally, and that they know the true model that generates the data. Our goal is to assess whether we can explain the forecast anomalies documented in section 3 by relaxing the assumption that forecasters know the true model, while maintaining the assumption that agents form expectations rationally given their beliefs about the model that generates the data. To this end, we consider agents who update their beliefs about how the world works using Bayesian learning and then form real-time Bayesian forecasts.

Our first application is to learning about the 3-month T-bill rate (short rate). We begin by presenting the model we assume the agents use to learn about and forecast the short rate. We then describe the details of how they learn and forecast. Finally, we compare the resulting forecasts with the SPF forecasts and longer-term yields.

4.1 An Unobserved Components Model for the Nominal Short Rate

Following [Kozicki and Tinsley \(2001\)](#), we propose a “shifting end-point” model for the short rate.⁶ Specifically, the model we assume agents use to learn about and forecast the short rate is:

$$y_t = \mu_t + x_t \tag{6}$$

$$\mu_t = \mu_{t-1} + \sqrt{\gamma}\sigma\eta_t, \quad \eta_t \sim N(0, 1), \tag{7}$$

$$x_t = \rho x_{t-1} + \sqrt{1-\gamma}\sigma\omega_t, \quad \omega_t \sim N(0, 1), \tag{8}$$

Here, the short rate y_t is modelled as the sum of two unobserved components: a permanent random walk component μ_t and a transitory AR(1) component x_t . The transitory component x_t is assumed have mean 0 and persistence ρ . Shocks to μ_t and x_t are independent, normally distributed. The total variance of these two innovations to y_t conditional on time $t - 1$ information is σ^2 . The share of the variance of these innovations that is attributable to shocks to the permanent component μ_t is assumed to be γ , with the complementary share $1 - \gamma$ attributable to the transitory component x_t . We refer to this model as an unobserved components model (UC).

To gain some intuition about the implications of the model, consider the h -period forecast of the short rate assuming the unobserved components at time t and parameters of the model are known:

$$E_t y_{t+h} = \mu_t + \rho^h x_t \tag{9}$$

This shows that μ_t corresponds to the long run forecast of the short rate (as $h \rightarrow \infty$), while x_t captures short run deviations of the short rate from this long run forecast. The expectations hypothesis implies that the yield on an n -period zero coupon bond is

$$y_t^{(n)} = c^{(n)} + \frac{1}{n} \sum_{h=0}^{n-1} E_t y_{t+h} = c^{(n)} + \mu_t + \frac{1}{n} \sum_{h=0}^{n-1} \rho^h x_t \tag{10}$$

where $c^{(n)}$ denotes the constant risk premium on n -period bonds. This expression shows that μ_t represents a level factor for the term structure of bond yields, while the slope and curvature of the term structure is governed by x_t .

⁶See also [van Dijk et al. \(2014\)](#), [Cieslak and Povala \(2015\)](#), [Bauer and Rudebusch \(2020\)](#), [Bianchi, Lettau, and Ludvigson \(2020\)](#), and [Crump et al. \(2021\)](#).

4.2 Bayesian Learning and Forecasting about the Nominal Short Rate

We assume that agents do not know the value of the unobserved components (states) μ_t and x_t . We, furthermore, assume that they do not know the value of the parameters ρ , γ , and σ . We endow them with initial beliefs about these unknown states and parameters and data on the short rate. We assume that they use Bayes Law to update their beliefs about the states and parameters over time and then in each period construct a forecast of future short rates based on their then current beliefs. More specifically, we start the agents off with initial beliefs in 1951Q2. The agents then use data on the short rate from 1951Q2 onward to update their beliefs. Starting in 1961Q3 they perform “online” forecasting of the short rate. In other words, each quarter they forecast the short rate based their beliefs at that point in time.

The world did not begin in 1951Q2. So, why do we not use data going further back in time? The reason for this is that the monetary policy regime in the U.S. was fundamentally different before 1951Q2. In March 1951, the U.S. Treasury and the Federal Reserve reached an agreement – commonly referred to as the Treasury-Fed Accord – to separate government debt management and monetary policy (Romero, 2013). From 1942 until the Accord, the Federal Reserve had abdicated its monetary independence by committing to fix the short rate at a low value to aid the financing of WWII and manage the massive government debt left after WWII. Before 1942, the U.S. had for the most part been on a gold (or silver) standard. Rather than model these fundamentally different monetary regimes explicitly, we start our analysis at the time of the Treasury-Fed Accord and simply endow agents with initial beliefs at that date (which presumably reflect information gleaned from the prior history).

We use a Gibbs Sampling algorithm (augmented with random walk Metropolis-Hastings steps when needed) to sample from the posterior distribution of the model parameters and the latent states at each time period t . We describe this algorithm in more detail in Appendix A. Armed with an estimate of agent’s belief distribution for the unknown parameters and states in a time period t , we use our unobserved components model to forecast the future evolution of the short rate. We describe the algorithm we use to do this in Appendix B. We do this for each quarter starting in 1961Q3, which is the first quarter for which we have zero-coupon yield curve data.

We assume that agents make their forecasts on the final day of each quarter. This implies that they have access to the average level of the interest rate in that quarter and their “nowcast” is the true realized interest rate for the quarter. This is an approximation: in reality, the SPF forecasters

only have information up to the second to third week of the second month of the quarter as we discuss above.

The short rate was constrained by the zero lower bound (ZLB) towards the end of our sample period. We define the period when the target federal funds rate was at or below 25 basis points as the ZLB period. This corresponds to 2009Q1 to 2015Q4 in our sample period. We view this as a period when the desired short rate is censored (but for simplicity follows the same process as before). Our approximation to Bayesian learning for this period is to assume that agents do not update their beliefs about the parameters (ρ , γ , and σ) but that they continue to filter the hidden states (μ_t and x_t) using the parameter estimates from 2008Q4. Full learning then resumes in 2016Q1. This short-cut allows us to avoid substantial additional complications which we believe are unlikely to materially affect our results.

4.3 Initial Beliefs about the Nominal Short Rate

A basic premise of our analysis is that Bayesian learning is slow in certain settings. This implies that initial beliefs will matter. If learning is fast, beliefs converge quickly to the truth and initial beliefs quickly cease to matter. If learning is slow, beliefs will not converge quickly to the truth and initial beliefs will continue to influence later beliefs non-trivially for a long time – as long as it takes for beliefs to converge to the truth.

In our setting, learning about the parameters ρ , γ , and σ is slow, while learning about the states μ_t and x_t is reasonably fast. Our choice of initial beliefs about μ_t and x_t , therefore, does not matter for our results as long as they are reasonable. We assume that initial beliefs about μ_t in 1951Q2 are $N(y_{1951Q12}, 1)$ and initial beliefs about x_t in 1951Q2 are $N(0, 1)$. These initial conditions are assumed to have a correlation of -1 due to the form of the observation equation (6).

For ρ , γ , and σ we specify initial beliefs in 1951Q2 of the following form:

$$\rho \sim N(\mu_\rho, \sigma_\rho^2), \quad \gamma \sim \mathcal{B}(\alpha_\gamma, \beta_\gamma), \quad \sigma^2 \sim \mathcal{IG}(\alpha_{\sigma^2}, \beta_{\sigma^2}),$$

where \mathcal{B} denotes a beta distribution and \mathcal{IG} denotes an inverse-gamma distribution. As we discuss above, these initial beliefs encode professional forecaster's understanding of how the world works as of 1951Q2, based on prior history. What precisely forecasters believed in 1951 is not known to us. Our approach is to search over the space of initial beliefs specified above to see if we can find an initial belief which can rationalize the forecast anomalies we document in section 3. If

we can find a belief (or perhaps more than one different beliefs) that can rationalize the forecast anomalies reasonably well, then we ask whether any of these beliefs can be viewed as a reasonable initial beliefs for professional forecasters to have in 1951Q2. If so, we conclude that the forecast anomalies we have documented can be explained by Bayesian learning and Bayesian forecasting and are therefore not necessarily evidence of forecaster irrationality.

To keep our analysis manageable, we fix the initial beliefs for σ by setting $\alpha_{\sigma^2} = 1.25$ and $\beta_{\sigma^2} = 0.5625$. This belief distribution is plotted in the bottom panel of Figure 4.⁷ This leaves four parameters: $\mu_{\rho}, \sigma_{\rho}^2, \alpha_{\gamma}, \beta_{\gamma}$. We search over the space of these parameters to find beliefs that match the forecast anomalies as well as possible. Specifically, for each point in this space, we construct forecasts as described above and estimate the forecasting regressions discussed in section 3. We then minimize an unweighted average of the difference between the regression coefficients from the regressions based on model-generated forecasts and the regression coefficients we estimated in section 3 based on real-world data. Appendix C provides more detail.

The top two panels of Figure 4 plot the initial belief distributions for ρ and γ that minimize the objective function discussed above. The belief distribution for ρ is concentrated on moderately large values. It is centered at 0.60 and has a standard deviation of 0.12. With a $\rho = 0.60$ the half-life of innovations to x_t is roughly four months. The belief distribution for γ is concentrated on relatively small values. It has a mode of 0.07 and a standard deviation of 0.06. This implies that forecasters believed in 1951Q2 that most of the variation in the short rate was due to transitory fluctuations of moderate persistence (i.e., an x_t with a ρ around 0.60) rather than permanent fluctuations (μ_t).

Are the initial belief distributions plotted in Figure 4 reasonable? We argue they are for two reasons. First, these belief distributions are quite dispersed. If we could only match the forecast anomaly results with extremely concentrated initial belief distributions, we would conclude that our model could not rationalize the forecast anomalies. This is not the case. The initial belief distributions we estimate put substantial mass on a wide range of parameter combinations, a sufficiently wide range that we think they constitute plausible beliefs forecasters might have had in 1951Q2.

Second, the belief that γ was relatively small is arguably consistent with the history of interest rates prior to 1951Q2. Outside of war, the United States had been on a gold standard (or silver

⁷This belief distribution has a mode of 0.25. The standard deviation of the distribution is undefined for values of $\alpha_{\sigma^2} \leq 2$. Our choice of $\alpha_{\sigma^2} = 1.25$ is thus a very dispersed initial belief.




Figure 4: Marginal Initial Beliefs Distributions: T-bill Rate Model

Note: Each panel plots the initial beliefs held in 1951Q2 by agents in our T-bill rate model for each of the three model parameters: ρ , γ , and σ^2 respectively.

standard) almost continuously from its founding (and England had been on a gold or bimetallic standard for hundreds of years before that). Over this long time span, interest rates had been quite stable at low frequencies with most variation being rather transient (due to seasonal cycles and financial crises). Given this history, it seems reasonable that forecasters' beliefs were heavily skewed towards believing that most fluctuations in interest rates would be relatively transient. The long upward march of interest rates in the 1960s, 70s, and early 80s and subsequent downward march of interest rates since then was completely without parallel in history. It seems unlikely that forecasters in 1951Q2 would put much weight on such an unprecedented sequence of events occurring.

4.4 Model's Fit to the Data

Figure 5 offers a visual depiction of the fit of the model's forecasts to the data. The top panel plots SPF forecasts of the short rate (the same data as is plotted in Figure 1). The bottom panel plots the forecasts generated by our model with the initial beliefs discussed above. Our model captures the fact that SPF forecasters tend to predict that the short rate will "mean revert" slowly towards a "normal" value, i.e., something close to its average value over the past business cycle. More specifically, when rates are falling, SPF forecasters tend not to expect further decreases but




Figure 5: Forecasted T-bill Rate: Data vs. Model

Note: The black solid line is the 3-month T-bill rate. Each short gray line with five circles represents forecasts made in a particular quarter about the then present quarter (first circle) and following four quarters (subsequent four circles). In the top panel, these forecasts are SPF forecasts. In the bottom panel, these forecasts are mean forecasts generated from the UC model estimated in real-time.

instead expect gradual increases; when rates are increasing, they tend to underestimate the speed of further increases.

The visual fit of the model-implied forecasts to the data is particularly strong in the first 25 years of our sample period. Consider, for example the easing cycles of 1984-1986, 1988-1994, and 2001-2004. In each of these episodes, our model matches the data extremely well. The same is true during the tightening cycles of 1986-1988, 1994-1995, and 1998-2000. Since 2004, there are periods when our model's fit to the data is somewhat worse. Interestingly, these tend to be episodes when the Federal Reserve has successfully used forward guidance to inform market participants about its policy intentions. One example of this is the period 2004-2007. Over this period, the SPF forecasters correctly anticipate the speed of Fed tightening, while our model delivers smaller expected increases. This is the period when the Fed used the term "measured pace" to signal to the market that they would increase interest rates by 0.25% at each FOMC meeting for quite some time. A second example is the period from 2012-2015, when the Fed explicitly stated that they would keep the short rate at 0.25% for several years. Our model does not incorporate this forward guidance and therefore fails to capture its effect on SPF forecasts.

Table 3 presents results for the forecast anomaly regressions we analyze in section 3 for our

Table 3: T-Bill Rate Forecast Anomalies: Model vs. Data

	Forecast Horizon			
	1	2	3	4
<i>Panel A: Bias</i>				
SPF	-0.18*** (0.06)	-0.34*** (0.11)	-0.52*** (0.16)	-0.70*** (0.20)
UC Model	-0.15** (0.06)	-0.27** (0.12)	-0.39** (0.17)	-0.50** (0.21)
<i>Panel B: Autocorrelation</i>				
SPF	0.30* (0.15)	0.27** (0.11)	0.24** (0.11)	0.13 (0.12)
UC Model	0.36** (0.16)	0.39*** (0.12)	0.35*** (0.11)	0.23** (0.11)
<i>Panel C: Mincer-Zarnowitz</i>				
SPF	0.97 (0.02)	0.94* (0.04)	0.90** (0.05)	0.86** (0.06)
UC Model	0.96* (0.02)	0.93** (0.03)	0.88** (0.05)	0.83*** (0.06)
<i>Panel D: Coibion-Gorodnichenko</i>				
SPF	0.23* (0.13)	0.34** (0.15)	0.62*** (0.18)	–
UC Model	0.38** (0.17)	0.55* (0.32)	0.88*** (0.34)	–

Note: The forecast horizons are quarters. Stars represent significance relative to the following hypotheses: $\alpha = 0$ for bias, $\beta = 0$ for autocorrelation, $\beta = 1$ for Mincer-Zarnowitz, $\beta = 0$ for Coibion-Gorodnichenko. P-values are computed using Newey-West standard errors with lag length selected as $L = \lceil 0.75 \times T^{1/3} \rceil$. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

model-generated data (rows labelled “UC model”) and compares these with analogous results for the real-world data (rows labelled “SPF”). Our model’s ability to match the forecast anomalies seen in the real-world data is remarkable. For all four types of regressions and at all horizons, our model matches the magnitude and statistical significance of the real-world estimates quite closely. Specifically, our model generates a negative bias that increases in size with the horizon, as in the data; our model generate autocorrelation in forecast errors of about 0.35 at horizons one through three and much less at horizon four, as in the data; our model yields Mincer-Zarnowitz coefficients slightly below one and decreasing with horizon, as in the data; and our model generates underreaction that grows with horizon, as in the data.

Table 4 presents results for the expectations hypothesis regressions we discuss in section 3

Table 4: Failures of the Expectations Hypothesis: Model vs. Data

	Long Horizon n						
	2	3	4	8	12	20	40
<i>Panel A: Future Short Rates</i>							
Data	-0.01*** (0.22)	0.11*** (0.22)	0.18*** (0.21)	0.39** (0.25)	0.57 (0.27)	0.74 (0.23)	0.71 (0.22)
UC Model	-0.19*** (0.43)	0.01** (0.45)	0.14* (0.46)	0.58 (0.51)	0.83 (0.47)	0.94 (0.36)	0.95 (0.34)
<i>Panel B: Change in Long Rate</i>							
Data	-1.02*** (0.43)	-0.91*** (0.54)	-1.03*** (0.57)	-1.29*** (0.61)	-1.61*** (0.65)	-2.04*** (0.67)	-2.75*** (0.86)
UC Model	-1.08*** (0.72)	-1.11*** (0.73)	-1.14*** (0.74)	-1.28*** (0.80)	-1.42*** (0.87)	-1.72*** (1.00)	-2.36** (1.56)

Note: The sample period is from 1961Q3 to 2019Q4. The top panel reports estimates of β from regression (4). The bottom panel reports estimates of β from regression (5). In both cases, the horizon n is listed at the top of the table. Stars represent significance relative to the hypothesis that $\beta = 1$. P-values are computed using Newey-West standard errors with lag length selected as $L = \lceil 0.75 \times T^{1/3} \rceil$. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

based on model-generated data and compares these results with those based on real-world data. Again, our model's ability to match the results based on real-world data is impressive. For the future-short-rate regressions in Panel A, we estimate β coefficients close to zero at short horizons, as in the data. The estimates then rise for longer-term bonds as they do for the data. For the change-in-long-rate regressions in Panel B, we estimate β coefficients that are negative at all horizons and increasingly so as the horizon increases. Quantitatively, our estimates are close to -1 at short horizons and rise to -2 at long horizons. This is all quite consistent with the the pattern of the real-world estimates.

Table 4 shows that our model provides an explanation for why the long rate has tended to fall when the yield spread is large rather than rise as full-information rational expectations models imply it should. In our model, this occurs because learning about the interest rate regime has tended to yield larger movements in the long rate than the forces captured by full-information rational expectations models. In particular, when the yield spread is large, agents in our model tend to revise downward their estimate of the long-run level of the short rate (μ_t) by enough to offset the forces pushing the long rate up.




Figure 6: Parameter Estimates: T-bill Rate Model

Note: Each panel plots the evolution of beliefs about one of the three UC model parameters: ρ , γ , and σ . The black solid line is the mean and the dotted black lines are the 5th and 95th percentiles of the posterior distribution for the parameter in question. Recall that we only update beliefs about these parameters every fourth quarter.

4.5 Parameter and State Estimates

Figure 6 plots the evolution of the mean of the posterior distributions of ρ , γ , and σ along with 90% credible intervals between 1961Q3 and 2019Q4. The mean estimate of ρ is around 0.75 early in the sample. It then gradually rises over time and is around 0.9 towards the end of the sample. The mean estimate of γ is initially low at 0.1. It rises to around 0.2 in the early 1980s and gradually declines until very late in the sample when it rises sharply. The mean estimate of σ is around 0.4 early in the sample. It rises sharply during the Volcker disinflation and gradually decreases after the early 1980s.

Figure 7 plots the mean estimates of the states μ_t and x_t over the course of the sample. The solid black line is the mean of the “real-time” filtering distribution, i.e., the belief distribution about μ_t and x_t conditional on data up to time t , while the solid gray line is the mean of the “ex-post” smoothing distributions, i.e., the belief distribution about μ_t and x_t conditional on data up to 2019Q4. The broken black lines plot 90% credible intervals for the real-time filtering distribution.

It is interesting to compare the real-time filtering distribution and the ex-post smoothing dis-




Figure 7: State Estimates: T-bill Rate Model

Note: Each panel corresponds to one of the two UC hidden state variables: μ_t and x_t respectively. The black solid line is the posterior mean of the real-time filtering distributions, the dotted black lines are the 5th and 95th percentiles of the posterior real-time filtering distributions, and the solid gray line is the posterior mean of the ex-post smoothing distributions for the corresponding parameter.

tribution. For μ_t , the real-time filtering distribution is consistently below the ex-post smoothing distribution from the beginning of our sample until the early 1980s and then consistently above from the early 1980s until very late in our sample. This reflects the fact that in real time the agents in our model underestimate the persistence of the run-up of interest rates in the 1960s and 70s, and again underestimate the persistence of the fall in interest rates after the early 1980s. Ex post, agents revise their view of history and conclude that both the run-up and fall in interest rates was more persistent than they believed at the time. This helps explain the persistent downward drift of long rates in the 1980s at a time when the yield spread was high. A full information rational expectations model would predict that long rates should rise during this period. But persistent learning about how to interpret the rise and fall of interest rates more than counteracted this.

5 Learning about the Real GDP Growth

Our second application is to learning about real GDP growth. As in section 4, we begin by presenting the model we assume agents use to learn about and forecast GDP growth. We then describe the details of how they learn and forecast. Finally, we compare the resulting forecasts with the

CBO forecasts we discussed in section 3.

5.1 An Unobserved Components Model for GDP

The model we assume agents use to learn about and forecast real GDP is:

$$y_t = z_t + x_t \tag{11}$$

$$\Delta z_t = \mu + \sqrt{\gamma}\sigma u_t, \quad u_t \sim N(0, 1), \tag{12}$$

$$x_t = \rho_1 x_{t-1} + \rho_2 x_{t-2} + \sqrt{1 - \gamma}\sigma v_t, \quad v_t \sim N(0, 1), \tag{13}$$

where y_t denotes quarterly log real GDP. This model posits that real GDP is the sum of two components: a difference stationary component z_t and a trend-stationary component x_t . The difference stationary component z_t is assumed to follow a random walk with drift μ . The trend-stationary component x_t is assumed to follow a mean zero AR(2) process with autoregressive coefficients ρ_1 and ρ_2 . The conditional standard deviation of y_t is denoted σ . The share of innovations to y_t that hit the difference-stationary component z_t is γ , with the complementary share $1 - \gamma$ hitting the trend-stationary component x_t . The parameter γ therefore governs “how big” the random walk component of GDP is (Cochrane, 1988). We refer to this model as an unobserved components (UC) model.

5.2 Bayesian Learning and Forecasting about GDP

As in the interest rate application discussed in section 4, we assume that agents in the model do not know the value of the unobserved components (states) z_t and x_t or parameters μ , ρ_1 , ρ_2 , σ , and γ . We start the agents off with an initial belief distribution about these unknown states and parameters in 1959Q3. The agents then observe (real-time) data on GDP and update their beliefs about the states and parameters using Bayes Law. Starting in 1976Q1 they perform annual “online” forecasting of GDP. This corresponds to the first period for which CBO forecasts are available.

The reason we start the agents off with an initial belief distribution in 1959Q3 is that we do not have a complete set of real-time GDP vintages that go further back than this date. The initial beliefs in 1959Q3 thus embody all knowledge agents in the model have been able to glean from earlier data about the statistical process that GDP follows. We discuss the choice of these initial beliefs in more detail in section 5.3 below.

We assume that agents have access to the first release of Q4 GDP for the prior year (the BLS’s “advance release” for that quarter) when they forecast. This is meant to approximate the information set the CBO has access to when it forecasts GDP each year. The CBO’s forecasts (contained in its Economic Outlook report) are typically released in January or February of each year. While this is usually before the advance release of Q4 GDP for the previous year, much of the underlying data that is used to construct the advance release has been made public at this point. This implies that the Q4 advance release can be predicted fairly accurately at this point. We therefore think that endowing our model agents with the Q4 advance release is the best way to approximate the information set of the CBO at the time it constructs its annual GDP forecast.

The parameters of the model and latent state estimates are updated every 4 quarters to line up with the timing of when the CBO constructs forecasts. We describe the algorithm we use to update agent’s beliefs in Appendix D. Armed with estimates of agent’s beliefs, we use our unobserved components model to construct forecasts of GDP growth. We describe the algorithm we use to do this in Appendix E.

5.3 Initial Beliefs about GDP

As in the interest rate application in section 4, learning about the parameters in our model for GDP is slow. This implies that agent’s initial beliefs about these parameters will matter. In contrast, learning about the states z_t and x_t is reasonably fast implying that initial beliefs about these states is less consequential. We assume that agents’ initial beliefs about z_t and x_t in 1959Q3 are $z_t \sim N(y_{1959Q3}, 0.01^2)$ and $x_t \sim N(0, 0.01^2)$.

We specify initial beliefs for the parameters in 1959Q3 of the following form

$$\begin{aligned} \rho_1 + \rho_2 &\sim N(\mu_\rho, \sigma_\rho^2) \mathcal{I}(\rho_1, \rho_2), & \rho_2 &\sim N(\mu_{\rho_2}, \sigma_{\rho_2}^2) \mathcal{I}(\rho_1, \rho_2), \\ \gamma &\sim \mathcal{B}(\alpha_\gamma, \beta_\gamma), & \mu &\sim N(\mu_\mu, \sigma_\mu^2), & \sigma &\sim \mathcal{IG}(\alpha_\sigma, \beta_\sigma). \end{aligned}$$

where $\mathcal{I}(\rho_1, \rho_2)$ is an indicator function which is 1 if the x_t process is stationary and 0 otherwise. For more detail, see Appendix D.

We fix $\mu_\mu = 0.01$ and $\sigma_\mu = 0.01$, corresponding to a prior for average annual long run growth of 4%. We fix $\alpha_\sigma = 7.0625$ and $\beta_\sigma = 0.0014$ corresponding to a prior mean for σ^2 of 0.015^2 and standard deviation of 0.01. That leaves 6 parameters to estimate to fit the forecast anomalies presented in section 2, which we denote $\theta = (\mu_\rho, \sigma_\rho, \mu_{\rho_2}, \sigma_{\rho_2}, \alpha_\gamma, \beta_\gamma)'$. We do this in a similar




Figure 8: Marginal Initial Beliefs Distributions: Real GDP Growth Model

Note: Each panel plots the initial beliefs held in 1959Q3 by agents in our model for the following five parameter combinations: $\rho_1 + \rho_2$, ρ_2 , γ , μ , and σ^2 .

fashion to what we do in the interest rate application in section 4. Appendix F provides details.⁸

The resulting initial beliefs are plotted in Figure 8. We view these as reasonable initial beliefs in that they are quite dispersed. For example, the initial belief distribution on $\rho_1 + \rho_2$ puts substantial weight on values between 0.7 and 1. This range spans cases where the transitory component x_t has a modest half-life of less than a year and cases where it is very persistent. Likewise, the initial beliefs for γ is centered close to 0.5 and has high variance. The initial belief for ρ_2 embeds a belief that the transitory component of GDP is hump-shaped. But again, this distribution has substantial variance.

5.4 Model's Fit to the Data

Figure 9 offers a visual depiction of the fit of the forecasts that our model generates to the data. The top panel plots CBO forecasts of real GDP growth (the same data as is plotted in Figure 3).

⁸We place some bounds on the values of parameters that can be chosen in this estimation. Namely, we restrict the standard deviation of the initial beliefs on $\rho_1 + \rho_2$, ρ_2 , and γ to be greater than or equal to 0.05. For the initial belief distribution for γ , we additionally put an upper bound on the standard deviation of 0.15 and restrict the mode of the distribution to be less than 0.6. The latter restriction imposes that agents believe at least 40% of the variation in output comes from trend-stationary fluctuations. These restrictions are useful to avoid estimating what we consider unreasonably dogmatic initial beliefs and capture the belief that a significant fraction of output fluctuations are trend-stationary.

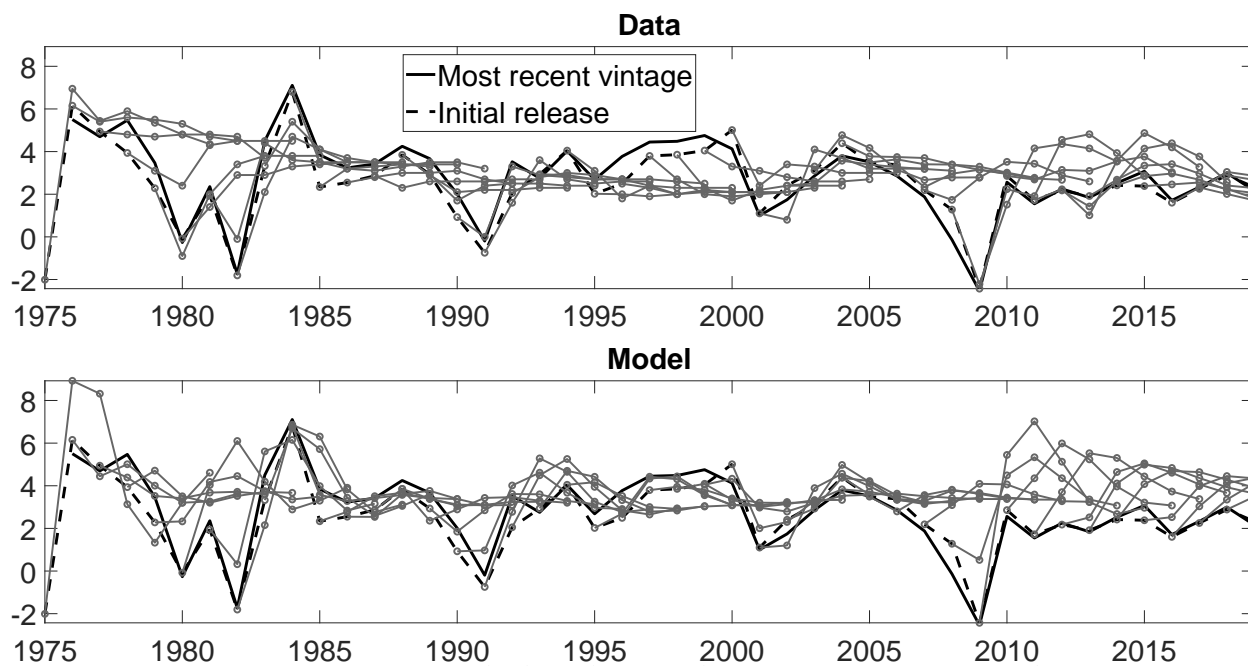


Figure 9: Forecast Whisker Plots: Real Economic Output Growth

Note: The black solid line is the most recent vintage of GDP growth estimates. The dashed black line is the initial release of GDP growth for each period. Each short gray line with seven circles represents forecasts made in a particular year about that year (first circle) and following six years (subsequent six circles). In the top panel, these forecasts are CBO forecasts. In the bottom panel, these forecasts are mean forecasts generated from the UC model estimated in real-time.

The bottom panel plots the forecasts generated by our model with the initial beliefs discussed above. The model is able to match the broad characteristics of CBO forecast errors. For example, the model matches the large forecast errors the CBO made in the early 2010s when it forecast that the economy would grow unusually fast after the Great Recession but growth turned out to be more modest. Also, the model generates persistent forecast errors in the late 1990s when growth was high for several years but the CBO persistently forecast lower growth.

Table 5 presents results for the forecast anomaly regressions we analyze in section 3 for our model-generated data (rows labelled “UC model”) and compares these with analogous results for the real-world data (rows labelled “CBO”). Our model is able to match the anomalies in the CBO forecasts quite well. The most spectacular anomaly in the case of the CBO forecasts is for the Mincer-Zarnowitz regressions in Panel C. These start off close to one at the one-year horizon but fall to zero at the three-year horizon and to roughly -0.4 at the four and five-year horizons. Our model is able to match this pattern quite well. The model also yields positively autocorrelated forecast errors, overreaction at long horizons in the Coibion-Gorodnichenko regression, and negative bias. For almost all of the anomaly statistics, the model estimate is not statistically significantly different from the data estimate. However, the point estimates from the model are not in

Table 5: Real GDP Forecast Anomalies: Model vs. Data

	Forecast Horizon				
	1	2	3	4	5
<i>Panel A: Bias</i>					
CBO	0.27 (0.21)	-0.27 (0.34)	-0.54 (0.44)	-0.62 (0.47)	-0.52 (0.44)
UC Model	-0.65** (0.27)	-1.65*** (0.40)	-1.36*** (0.38)	-0.85** (0.38)	-0.66* (0.39)
<i>Panel B: Autocorrelation</i>					
CBO	0.22 (0.16)	0.16 (0.17)	0.11 (0.14)	0.08 (0.18)	0.08 (0.11)
UC Model	0.39** (0.17)	0.31* (0.16)	0.23** (0.11)	0.06 (0.12)	-0.05 (0.07)
<i>Panel C: Mincer-Zarnowitz</i>					
CBO	0.94 (0.12)	0.60 (0.31)	0.03*** (0.31)	-0.42*** (0.22)	-0.43*** (0.36)
UC Model	0.84 (0.14)	0.35*** (0.20)	0.34** (0.33)	-0.38*** (0.26)	-0.98*** (0.63)
<i>Panel D: Coibion-Gorodnichenko</i>					
CBO	0.08 (0.11)	-0.00 (0.23)	0.50 (0.52)	-1.63*** (0.41)	-1.46*** (0.45)
UC Model	0.06 (0.10)	-0.76 (0.56)	-0.11 (0.29)	-0.78 (0.56)	-1.22** (0.50)

Note: The forecast horizons are years. Stars represent significance relative to the following hypotheses: $\alpha = 0$ for bias, $\beta = 0$ for autocorrelation, $\beta = 1$ for Mincer-Zarnowitz, $\beta = 0$ for Coibion-Gorodnichenko. P-values are computed using Newey-West standard errors with lag length selected as $L = \lceil 0.75 \times T^{1/3} \rceil$. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

all cases as closely aligned with the data as in the interest rate application discussed in section 4. For example, the model yields a larger negative bias than the data, especially at horizons two and three.

5.5 Parameter Estimates

Figure 10 plots the evolution of the mean of the posterior distributions of the five parameters of our model for GDP along with 90% credible intervals over the period 1976 and 2019. Perhaps the most striking feature of Figure 10 is how little beliefs about the parameters change over time. We do see that σ trends downward by a modest amount, likely reflecting the Great Moderation. Also, ρ_2 trends modestly upward. But ρ_1 , γ , and μ change very little. This lack of change presumably

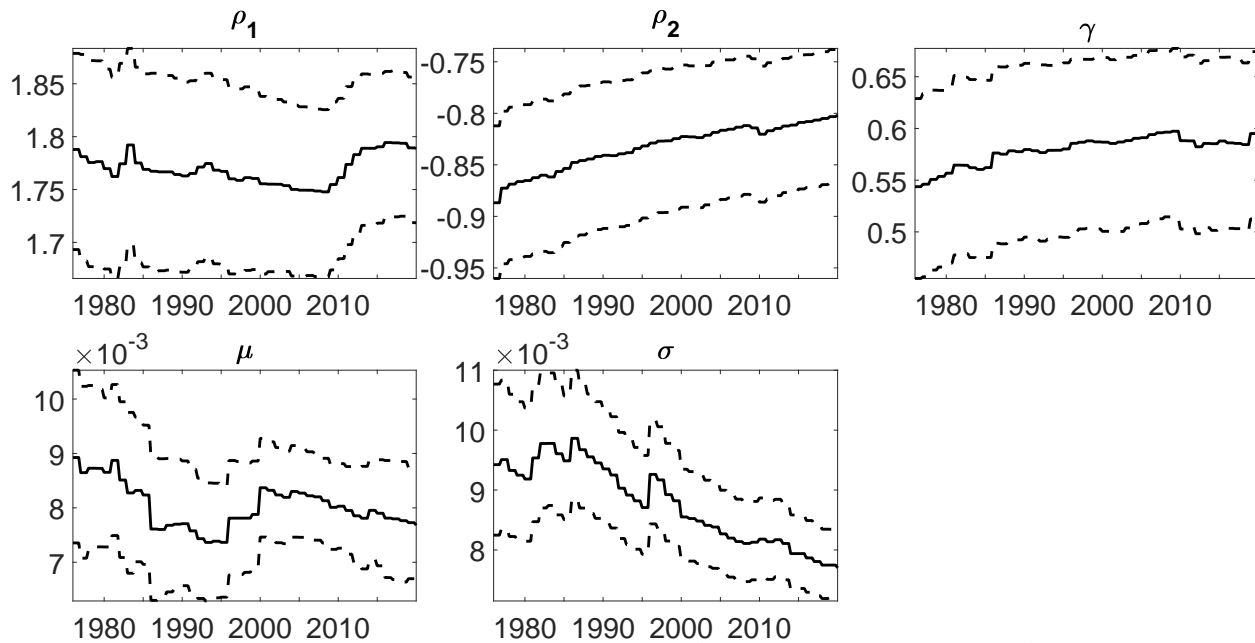


Figure 10: Parameter Estimates: Real Economic Output Growth

Note: Each panel plots the evolution of beliefs about one of the five UC model parameters: ρ_1 , ρ_2 , γ , μ , and σ . The black solid line is the mean and the dotted black lines are the 5th and 95th percentiles of the posterior distribution for the parameter in question. Recall that we only update beliefs about these parameters every fourth quarter.

reflects a combination of two things. First, it may be that some of the parameters are close to their true values. Second, for those parameters that are further away from their true values, little information can be gleaned from the data about their true values resulting in posterior beliefs being little changed even over a 40 year period. This is perhaps not surprising given how difficult it is to distinguish between difference-stationary time series and persistent but trend-stationary time series.

6 Why Does it Work?

To understand better why it is that our Bayesian learning model can match the forecast anomalies that we document in section 3, we now simulate data from the model we use in section 4 and assess how learning occurs in this model. Relative to the analysis earlier in the paper, in this section, we know the true data generating process. We can therefore assess how long it takes agents to learn the truth and how initial beliefs that differ in various ways from the truth affect results from the forecasting regressions we consider in section 3.

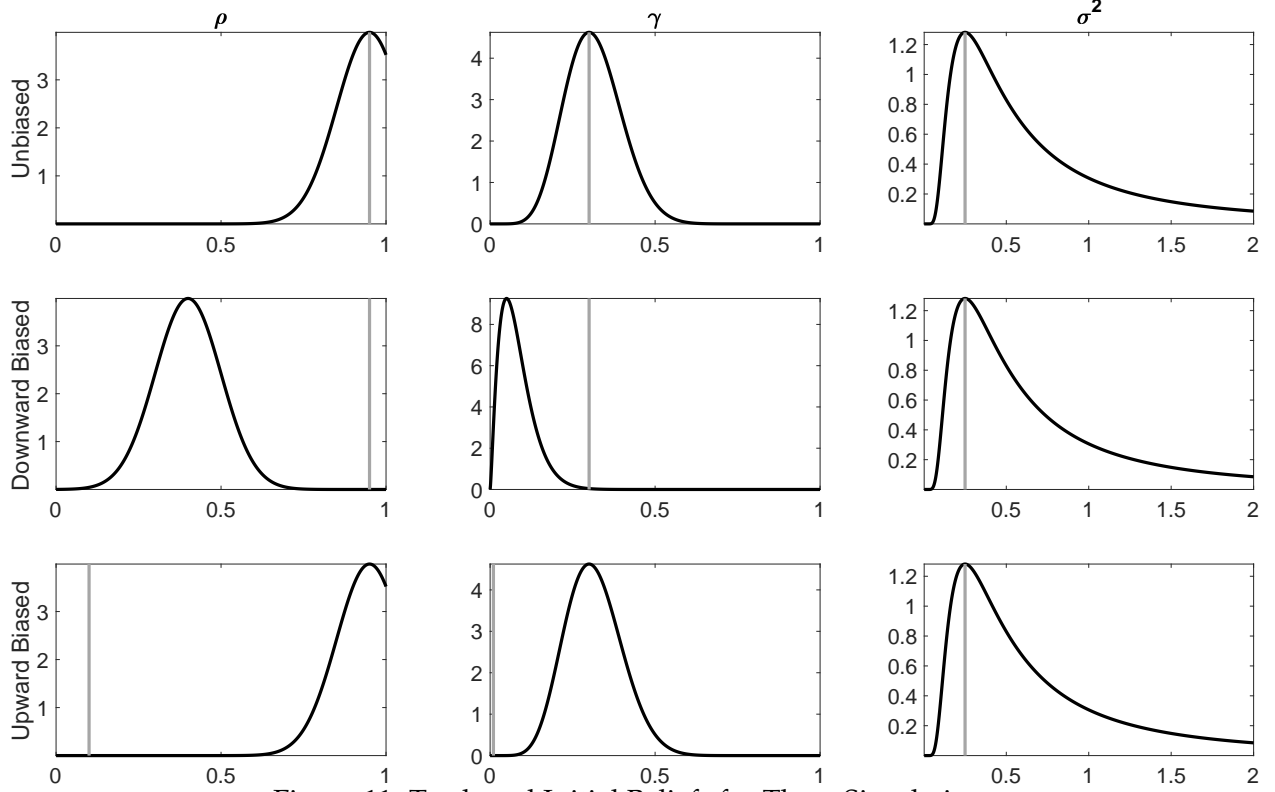


Figure 11: Truth and Initial Beliefs for Three Simulations

Note: The figure plots the truth (gray vertical line) and initial belief distribution (black line) for ρ (left column), γ (middle column), and σ (right column) for the three cases we consider. The first row of figures is the Unbiased Initial Beliefs case, the middle row is the Downward-Biased Initial Beliefs case, and the bottom row is the Upward-Biased Initial Beliefs case.

Recall that the model we use for the short rate in section 4 is:

$$y_t = \mu_t + x_t \tag{14}$$

$$\mu_t = \mu_{t-1} + \sqrt{\gamma}\eta_t, \quad \eta_t \sim N(0, 1), \tag{15}$$

$$x_t = \rho x_{t-1} + \sqrt{1 - \gamma}\sigma\omega_t, \quad \omega_t \sim N(0, 1). \tag{16}$$

We present results for three cases which we refer to as a case of unbiased initial beliefs, downward-biased initial beliefs, and upward-biased initial beliefs. Figure 11 plots the true parameter values (gray vertical lines) and initial belief distributions (black lines) for these three cases. A more detailed description follows:

- **Unbiased Initial Beliefs:** In this case, we set the true parameters to values $\rho = 0.95$, $\gamma = 0.3$, and $\sigma = 0.5$. These values are close to the mean of the belief distribution we estimate from the real-world data in the second half of our sample. We assume that agents in the model

have an initial belief distribution with the property that the mode of the belief distribution for each parameter is equal to the truth:

$$\rho \sim N(0.95, 0.01), \quad \gamma \sim \mathcal{B}(9.052, 19.788), \quad \sigma^2 \sim \mathcal{IG}(1.25, 0.5625).$$

- **Downward-Biased Initial Beliefs:** In this case, we again set the true parameters to values $\rho = 0.95$, $\gamma = 0.3$, and $\sigma = 0.5$. We however assume that agents in the model have an initial belief distribution with the property that the modes of the belief distributions for ρ and γ are smaller than the truth:

$$\rho \sim N(0.4, 0.01), \quad \gamma \sim \mathcal{B}(2.34, 26.5), \quad \sigma^2 \sim \mathcal{IG}(1.25, 0.5625).$$

- **Upward-Biased Initial Beliefs:** In this case, we set the true parameters to values $\rho = 0.1$, $\gamma = 0.01$, and $\sigma = 0.5$. We then assume that agents in the model have an initial belief distribution with the property that the modes of the belief distributions for ρ and γ are larger than the truth:

$$\rho \sim N(0.95, 0.01), \quad \gamma \sim \mathcal{B}(9.052, 19.788), \quad \sigma^2 \sim \mathcal{IG}(1.25, 0.5625).$$

The reason why we choose different true values for this case is that the true value of ρ used in the other two cases is sufficiently large that it is difficult to illustrate the effects of beliefs that are upward biased relative to this truth.

For each of these three sets of assumptions, we simulate 500 samples of the same length as the short rate data we use in section 4, i.e., 275 periods corresponding to the sample period from 1951Q2 to 2019Q4. For each of these simulated data series, we then perform the same exercise as we did in section 4. Given their initial beliefs, the agents in the model learn about the parameters of the model using the short rate series and Bayes Law. They then construct Bayesian forecasts. The length of the sample period for the Bayesian forecasts is the same as for the real-world data. We then run the same forecast rationality and expectations hypothesis tests on the resulting data as we did on the real-world data in section 4.

Tables 6 and 7 present the results from this analysis. Table 6 presents results on autocorrelation of forecast errors, the Mincer-Zarnowitz test, and Coibion-Gorodnichenko tests of over- and

Table 6: Forecast Anomalies in Simulated Data

	Forecast Horizon			
	1	2	3	4
<i>Panel A: Autocorrelation</i>				
Unbiased Initial Beliefs	0.01 (0.08) 1.00	0.00 (0.09) 1.00	-0.00 (0.11) 0.99	-0.01 (0.13) 0.84
Downward-Biased Initial Beliefs	0.16 (0.09) 0.93	0.19 (0.10) 0.78	0.19 (0.12) 0.61	0.18 (0.14) 0.33
Upward-Biased Initial Beliefs	-0.34 (0.06) 1.00	-0.32 (0.07) 1.00	-0.28 (0.08) 1.00	-0.26 (0.08) 1.00
<i>Panel B: Mincer-Zarnowitz</i>				
Unbiased Initial Beliefs	0.96 (0.03) 0.58	0.92 (0.05) 0.62	0.88 (0.08) 0.57	0.83 (0.11) 0.53
Downward-Biased Initial Beliefs	0.98 (0.03) 0.27	0.95 (0.05) 0.40	0.90 (0.08) 0.45	0.85 (0.11) 0.47
Upward-Biased Initial Beliefs	0.37 (0.17) 1.00	0.33 (0.22) 1.00	0.34 (0.25) 1.00	0.35 (0.27) 0.99
<i>Panel C: Coibion-Gorodnichenko</i>				
Unbiased Initial Beliefs	0.01 (0.09) 0.99	0.01 (0.12) 0.99	0.01 (0.15) 1.00	-
Downward-Biased Initial Beliefs	0.18 (0.11) 0.66	0.32 (0.19) 0.55	0.41 (0.25) 0.79	-
Upward-Biased Initial Beliefs	-0.52 (0.10) 1.00	-0.55 (0.13) 1.00	-0.53 (0.17) 1.00	-

Note: The top number for each case is the mean estimate across simulations. The middle number for each case (in parentheses) is the standard deviation across simulations. The bottom number for each case is the fraction of simulations that give a smaller estimate than the real-world data.

underreaction, while Table 7 presents results on the two tests of the expectations hypothesis we consider in section 3. In each case, we report three statistics. The first is the mean estimated coefficient across the 500 simulations; the second statistic is the standard deviation of the estimated effects across simulations (in parentheses); and the third statistic is the fraction of simulations that give a smaller estimate than the estimate based on real-world data.

The main finding from this analysis is that the downward-biased initial beliefs simulation

Table 7: Failures of the Expectations Hypothesis in Simulated Data

	Long Horizon n						
	2	3	4	8	12	20	40
<i>Panel A: Future Short Rates</i>							
Unbiased Initial Beliefs	0.95 (0.64) 0.07	1.01 (0.63) 0.07	1.05 (0.66) 0.08	1.19 (0.72) 0.12	1.31 (0.76) 0.16	1.51 (0.82) 0.16	2.06 (1.03) 0.08
Downward-Biased Initial Beliefs	0.17 (0.19) 0.17	0.20 (0.21) 0.30	0.23 (0.22) 0.38	0.33 (0.29) 0.57	0.42 (0.33) 0.65	0.57 (0.40) 0.66	0.97 (0.56) 0.29
Upward-Biased Initial Beliefs	2.46 (0.21) 0.00	2.14 (0.16) 0.00	1.97 (0.14) 0.00	1.71 (0.09) 0.00	1.64 (0.07) 0.00	1.59 (0.06) 0.00	1.50 (0.05) 0.00
<i>Panel B: Change in Long Rate</i>							
Unbiased Initial Beliefs	0.90 (1.27) 0.07	0.93 (1.32) 0.08	0.95 (1.36) 0.07	1.01 (1.50) 0.06	1.08 (1.63) 0.05	1.20 (1.92) 0.03	2.08 (3.00) 0.03
Downward-Biased Initial Beliefs	-0.66 (0.38) 0.17	-0.69 (0.40) 0.28	-0.74 (0.42) 0.24	-1.03 (0.52) 0.32	-1.39 (0.64) 0.34	-2.04 (0.91) 0.51	-3.62 (1.84) 0.66
Upward-Biased Initial Beliefs	3.91 (0.42) 0.00	4.13 (0.42) 0.00	4.38 (0.42) 0.00	5.59 (0.51) 0.00	6.90 (0.62) 0.00	9.56 (0.86) 0.00	13.77 (1.62) 0.00

Note: The top number for each case is the mean estimate across simulations. The middle number for each case (in parentheses) is the standard deviation across simulations. The bottom number for each case is the fraction of simulations that give a smaller estimate than the real-world data.

roughly matches all of the anomalies we document in real-world data. This simulation yields positively autocorrelated forecast errors, underreaction in the Coibion-Gorodnichenko regression, values below one in the future-short-rate regression, and negative values for the change-in-long-rate regressions. In virtually all cases, the downward-biased initial beliefs simulation is quantitatively consistent with our real-world estimates of the anomalies in the sense that the real-world estimate is well within the 95% central probability mass of the distribution of estimates from the simulation.

In sharp contrast, the upward-biased initial belief simulation yields anomalies with the opposite sign from what we see in the real-world data. This simulation yields negatively autocorrelated forecast errors, overreaction in the Coibion-Gorodnichenko regression, and values above one in both the future-short-rate regression and the change-in-long-rate regression. In addition to this, the upward-biased initial beliefs simulation also yields a very different pattern from the




Figure 12: Parameter Learning with Downward-Biased Initial Beliefs

Note: The figure plots the evolution of beliefs about ρ (left panel) and γ (right panel) over time when agents start off with downward-biased initial beliefs. The solid gray line is the truth. The solid black line is the evolution over time of the the mean point estimate across simulation. Recall that the point estimate in a particular simulation is the mean of the belief distribution of the parameter in question in that simulation. The broken black lines plot the evolution of the 90% and 10% quantiles of the distribution of point estimates across simulations.

real-world data for the Mincer-Zarnowitz regression, while the downward-biased initial beliefs simulation matches the real-world data for this regression as well.

Finally, the unbiased initial beliefs simulation yields results that are in most cases consistent with full-information rational expectations on average. It yields virtually no autocorrelation of forecast errors and a coefficient very close to zero in the Coibion-Gorodnichenko regressions (i.e., neither underreaction nor overreaction). For the expectations hypothesis regressions, it yields coefficients that are on average slightly larger than one at longer horizons. But the value one is not far from the middle of the distribution of coefficients across simulations.

From these results, we conclude that beliefs in society about interest rates in 1951 that underestimated the extent to which fluctuations in interest rates would be persistent relative to what turned out to be the case provide an explanation for the forecast anomalies and failures of the expectations hypothesis that we discuss in section 3. As we discuss earlier in the paper, such beliefs seem reasonable given the prior history of interest rate movements. Outside of war, the U.S. had been on a gold or silver standard and a run-up and run-down of interest rates such as was experienced from the 1960s to the 2000s had never before happened.

It is instructive to consider the speed of learning about the key parameters ρ and γ in the sim-




Figure 13: Learning about Persistence with Different Priors on γ

Note: The figure plots the evolution of beliefs about ρ for the downward-biased case (solid black line) and for a case that is the same as the downward-biased case except that γ is set (very close) to zero agents have very tight initial beliefs around the true value of γ (dashed black line). The solid gray line is the truth.

ulations with downward-biased initial beliefs. Figure 12 plots the evolution of beliefs about these parameters over time in the simulations. The gray line denotes the true value of the parameters. The solid black line plots the evolution of the mean point estimate of the parameters across simulations from 1961Q3 to 2019Q4. In 1961, the point estimates of both ρ and γ are substantially below the truth. Over time, both estimates rise, but this happens very slowly and both continue to be substantially below the truth at the end of the sample – when agents have been learning about these parameters for almost 70 years.

Figure 12 shows that it takes substantially longer than 70 years for the agents in our model to learn the true values of the parameters ρ and γ . One reason for this is that increases in ρ and γ both increase the persistence of fluctuations in the short rate. When agents revise upward their beliefs about the persistence of the short rate, they face the problem of whether the higher persistence is due to a more persistent x_t process (i.e., a higher ρ) or to a more volatile μ_t process (i.e., higher γ). Figure 13 compares the speed of learning about ρ in the downward-biased case with a case that is the same as the downward-biased case except that γ is set (very close) to zero and agents have a very tight initial belief distribution around the true value of γ – i.e., we turn off variation in μ_t and learning about γ . In this case, learning about ρ is much quicker.

Figure 13 shows that having two unobserved persistent components slows down learning.⁹ But even when variation in μ_t and learning about γ has been shut down, learning about ρ still takes quite a few decades. This illustrates that learning about the persistence of highly persistent time series processes is quite slow. Unit root tests have low power for similar reasons.

7 Conclusion

In this paper, we provide a new interpretation of well-known forecast anomalies of professional forecasters. We stress that tests of forecast rationality are joint tests of rationality and the notion that forecasters know the true model of the world. We relax the assumption that forecasters know the true model of the world and show that the anomalies can be explained via Bayesian learning of unobserved components models. Since the anomalies in question persist for decades, it is important that learning is slow in our setting. We show that learning is indeed extremely slow in the type of unobserved components model we consider. This implies that forecasters with reasonable initial beliefs that turn out not be centered on the truth result in forecast anomalies of the kind we observe in the data that persist for decades. We also perform a simulation exercise in which we know the true value of the parameters. We show in this exercise that reasonable initial belief errors can result in extremely persistent forecast anomalies as in the real world setting. Importantly, in this simulation exercise, we know that agents are using a correctly-specified model to learn and yet learning is extremely slow. At a big picture level, the argument we are making is that forecast anomalies arise for essentially the same reason that it is hard for econometricians to distinguish certain classes of models / parameters even with decades of data, e.g., it is hard to reject a unit root in many settings.

⁹Johannes, Lochstoer, and Mou (2016) develop similar results in a different setting.

A Bayesian Updating about Parameters and States for Interest Rates

In this appendix, we describe in more detail the Gibbs Sampling algorithm we use to sample from the joint posterior of the model parameters and latent states in the UC model for the short rate. Define the vector of the parameters and latent states of the model through date t as $\boldsymbol{\theta} \equiv (\rho, \gamma, \sigma, \boldsymbol{\mu}_{1:t}, \mathbf{x}_{1:t})'$. Let $p(\boldsymbol{\theta})$ denote the joint prior over the parameter vector $\boldsymbol{\theta}$. Let $L(\mathbf{y}_{1:t}|\boldsymbol{\theta})$ denote the likelihood function of the data through time t , given a set of parameters $\boldsymbol{\theta}$. Our goal is to sample from the posterior of the parameters given the data, $p(\boldsymbol{\theta}|\mathbf{y}_{1:t})$, where we know

$$p(\boldsymbol{\theta}|\mathbf{y}_{1:t}) \propto L(\mathbf{y}_{1:t}|\boldsymbol{\theta})p(\boldsymbol{\theta})$$

We assume functional forms for the initial beliefs as follows

$$\begin{aligned}\rho &\sim N(\mu_\rho, \sigma_\rho^2) \\ \gamma &\sim \mathcal{B}(\alpha_\gamma, \beta_\gamma) \\ \sigma^2 &\sim \mathcal{IG}(\alpha_{\sigma^2}, \beta_{\sigma^2})\end{aligned}$$

The initial beliefs for the states are given by

$$\begin{aligned}\mu_{1951Q2} &\sim N(y_{1951Q2}, 1) \\ x_{1951Q2} &\sim N(0, 1)\end{aligned}$$

where y_{1952Q2} denotes the 3-month Treasury bill rate in 1952Q2.

We start with an initial guess of the parameters $\boldsymbol{\theta}^{(0)} = (\rho^{(0)}, \gamma^{(0)}, \sigma^{(0)}, \boldsymbol{\mu}_{1:t}^{(0)}, \mathbf{x}_{1:t}^{(0)})'$. Given a draw of the parameters $\boldsymbol{\theta}^{(b)}$, we draw $\boldsymbol{\theta}^{(b+1)}$ as follows:

1. Draw $\rho^{(b+1)}|\gamma^{(b)}, \sigma^{(b)}, \boldsymbol{\mu}_{1:t}^{(b)}, \mathbf{x}_{1:t}^{(b)}, \mathbf{y}_{1:t}$. Given the other parameters, beliefs about ρ can be updated from the autoregression

$$x_t^{(b)} = \rho x_{t-1}^{(b)} + \sqrt{1 - \gamma^{(b)}} \sigma^{(b)} \omega_t$$

Define

$$\tilde{\sigma}_\rho^2 \equiv \left[\sigma_\rho^{-2} + \frac{\sum_{s=2}^t \left(x_{s-1}^{(b)} \right)^2}{(1 - \gamma^{(b)}) (\sigma^{(b)})^2} \right]^{-1}$$

$$\tilde{\mu}_\rho \equiv \tilde{\sigma}_\rho^2 \left[\frac{\mu_\rho}{\sigma_\rho^2} + \frac{\sum_{s=2}^t x_{s-1}^{(b)} x_s^{(b)}}{(1 - \gamma^{(b)}) (\sigma^{(b)})^2} \right]$$

The posterior of ρ is $N(\tilde{\mu}_\rho, \tilde{\sigma}_\rho^2)$ and thus we draw $\rho^{(b+1)} \sim N(\tilde{\mu}_\rho, \tilde{\sigma}_\rho^2)$.

2. Draw $\gamma^{(b+1)} | \rho^{(b+1)}, \sigma^{(b)}, \boldsymbol{\mu}_{1:t}^{(b)}, \mathbf{x}_{1:t}^{(b)}, \mathbf{y}_{1:t}$. There is no closed form expression for the posterior of γ . We therefore draw it using a random walk Metropolis-Hastings step. Specifically, we draw a proposal $\tilde{\gamma}^{(b+1)} \sim N(\gamma^{(b)}, \sigma_{\gamma,prop}^2)$ where $\sigma_{\gamma,prop}^2$ is a proposal variance chosen such that this step has between a 25 and 40% acceptance rate over the burn-in period. We then set $\gamma^{(b+1)} = \tilde{\gamma}^{(b+1)}$ with probability α_{b+1} , where

$$\alpha_{b+1} \equiv \frac{L\left(\mathbf{y}_{1:t} | \rho^{(b+1)}, \tilde{\gamma}^{(b+1)}, \sigma^{(b)}, \boldsymbol{\mu}_{1:t}^{(b)}, \mathbf{x}_{1:t}^{(b)}\right) p_\gamma\left(\tilde{\gamma}^{(b+1)}\right)}{L\left(\mathbf{y}_{1:t} | \rho^{(b+1)}, \gamma^{(b)}, \sigma^{(b)}, \boldsymbol{\mu}_{1:t}^{(b)}, \mathbf{x}_{1:t}^{(b)}\right) p_\gamma\left(\gamma^{(b)}\right)}$$

Otherwise we set $\gamma^{(b+1)} = \gamma^{(b)}$.

3. Draw $\sigma^{(b+1)} | \rho^{(b+1)}, \gamma^{(b+1)}, \boldsymbol{\mu}_{1:t}^{(b)}, \mathbf{x}_{1:t}^{(b)}, \mathbf{y}_{1:t}$. Given the other parameters, beliefs about σ can be updated from the two equations

$$\mu_t^{(b)} = \mu_{t-1}^{(b)} + \sqrt{\gamma^{(b+1)}} \sigma \eta_t$$

$$x_t^{(b)} = \rho^{(b+1)} x_{t-1}^{(b)} + \sqrt{1 - \gamma^{(b+1)}} \sigma \omega_t$$

Since η_t and ω_t are independent, these regression equations can be treated as two independent sources of information. Define

$$\tilde{\alpha}_{\sigma^2} \equiv \alpha_{\sigma^2} + (t - 1)$$

$$\tilde{\beta}_{\sigma^2} \equiv \beta_{\sigma^2} + \frac{\sum_{s=2}^t \left(\mu_s^{(b)} - \mu_{s-1}^{(b)} \right)^2}{2\gamma^{(b+1)}} + \frac{\sum_{s=2}^t \left(x_s^{(b)} - \rho^{(b+1)} x_{s-1}^{(b)} \right)^2}{2(1 - \gamma^{(b+1)})}$$

The posterior of σ^2 is $\mathcal{IG}(\tilde{\alpha}_{\sigma^2}, \tilde{\beta}_{\sigma^2})$ and thus we draw $(\sigma^{(b)})^2 \sim \mathcal{IG}(\tilde{\alpha}_{\sigma^2}, \tilde{\beta}_{\sigma^2})$.

4. Draw $\boldsymbol{\mu}_{1:t}^{(b+1)}, \boldsymbol{x}_{1:t}^{(b+1)} | \rho^{(b+1)}, \gamma^{(b+1)}, \sigma^{(b+1)}, \boldsymbol{y}_{1:t}$. This can be done using the standard Kalman filter and simulation smoother.

This algorithm is repeated to produce B draws from the posterior distribution of the parameters and states at each time t .

B Bayesian Forecasting of Interest Rates

The algorithm described in Appendix A yields B samples of the posterior of the states and parameters of our UC model at each point in time t . We index these samples by b as follows $\{\rho^{(b)}, \gamma^{(b)}, \sigma^{(b)}, \boldsymbol{\mu}_{1:t}^{(b)}, \boldsymbol{x}_{1:t}^{(b)}\}_{b=1}^B$. We then use the following algorithm to produce a real-time forecast distribution for the yield curve at time t :

1. For each $b = 1, \dots, B$

- (a) Simulate a path of shocks $\{\eta_{t+h}^{(b)}, \omega_{t+h}^{(b)}\}_{h=1}^H$ from the standard Normal distribution.
- (b) Starting from $h = 1$, construct a simulated path of the states over H subsequent periods using equations (7)-(8):

$$\begin{aligned}\mu_{t+h}^{(b)} &= \mu_{t+h-1|t}^{(b)} + \sqrt{\gamma^{(b)}} \sigma^{(b)} \eta_{t+h}^{(b)} \\ x_{t+h}^{(b)} &= \rho^{(b)} x_{t+h-1|t}^{(b)} + \sqrt{1 - \gamma^{(b)}} \sigma^{(b)} \omega_{t+h}^{(b)}\end{aligned}$$

- (c) Use the simulated states to construct the forecast distribution of the short rate $\{y_{t+h|t}^{(b)}\}_{h=1}^H$ where

$$y_{t+h|t}^{(b)} = \mu_{t+h|t}^{(b)} + x_{t+h|t}^{(b)}$$

2. The forecast of y_{t+h} given time t information is computed as

$$F_t y_{t+h} = \frac{1}{B} \sum_{b=1}^B y_{t+h|t}^{(b)}$$

The implied yield of a bonds of maturity n is given by

$$y_t^{(n)} = c^{(n)} + \frac{1}{n} \sum_{h=0}^{n-1} F_t y_{t+h}$$

We estimate the constant $c^{(n)}$ as the average level of the corresponding n -period bond yield in the data since it is not identified from the expectations hypothesis alone. Note that this estimate of the constant does not affect the results of the expectations hypothesis regression tests we run since it only affects the level of the n -period yield.

At the end of the estimation we are left with a sequence of model-implied 1 to H-quarter ahead forecasts $\{F_t y_{t+h}\}_{h=1}^H$ and model-implied yields $\{y_t^{(h)}\}_{h=1}^H$ for every quarter t from 1961Q3 to 2019Q4.

C Search over Initial Beliefs for Nominal Short Rate

Let $\theta = (\alpha_\rho, \beta_\rho, \alpha_\gamma, \beta_\gamma)'$. Let $\alpha = \{\alpha_h\}_{h=1}^H$ and $\beta = \{\beta_h\}_{h=1}^H$ denote vectors of estimated coefficients from the forecasting anomaly regressions for different horizons up through a maximum horizon of H using the SPF and yield curve data. Let $\hat{\alpha} = \{\hat{\alpha}_h\}_{h=1}^H$ and $\hat{\beta} = \{\hat{\beta}_h\}_{h=1}^H$ denote those same quantities estimated on the model implied forecasts and yields for a particular value of θ . Define the moment function as

$$\hat{m}(\theta) = \begin{bmatrix} \alpha_{bias} - \hat{\alpha}_{bias} \\ \alpha_{ar} - \hat{\alpha}_{ar} \\ \beta_{ar} - \hat{\beta}_{ar} \\ \alpha_{mz} - \hat{\alpha}_{mz} \\ \beta_{mz} - \hat{\beta}_{mz} \\ \alpha_{cg} - \hat{\alpha}_{cg} \\ \beta_{cg} - \hat{\beta}_{cg} \\ \beta_{sr} - \hat{\beta}_{sr} \\ \beta_{lr} - \hat{\beta}_{lr} \end{bmatrix} \quad (17)$$

The parameters are then estimated via the simulated method of moments (SMM) with an identity weighting matrix

$$\hat{\theta} = \operatorname{argmax}_{\theta} \|\hat{m}(\theta)\|_2 = \operatorname{argmax}_{\theta} \hat{m}(\theta)' \hat{m}(\theta)$$

Every evaluation of the moment function $\hat{m}(\theta)$ requires us to sample from the posterior of the UC model sequentially. Since this step is very computationally costly, we only re-estimate the model every 4 quarters rather than every quarter, and use a burn-in sample of 20,000 draws and keep the subsequent 20,000 draws rather than 50,000 for each of those quantities in our empiri-

cal specification. The global minimum is found using MATLAB's "particleswarm" optimization routine.

D Bayesian Updating about Parameters and States for GDP

Here we describe the initial beliefs and sampling algorithm for our GDP application. We assume that the initial belief of the CBO about the mean of the difference stationary component μ is Normal,

$$\mu \sim N(\mu_\mu, \sigma_\mu^2).$$

We assume that the CBO has independent Normal initial beliefs about the sum of the autoregressive parameters $\rho_1 + \rho_2$ and for the second autoregressive parameter ρ_2 . We truncate these initial belief distributions in such a way as to put zero weight on parameter combinations that result in the x_t component being non-stationary. We can write these initial belief distributions as

$$\begin{aligned}\rho_1 + \rho_2 &\sim N(\mu_\rho, \sigma_\rho^2) \mathcal{I}(\rho_1, \rho_2), \\ \rho_2 &\sim N(\mu_{\rho_2}, \sigma_{\rho_2}^2) \mathcal{I}(\rho_1, \rho_2).\end{aligned}$$

where $\mathcal{I}(\rho_1, \rho_2)$ is an indicator variable which is 1 for (ρ_1, ρ_2) combinations that result x_t being stationary and 0 otherwise.

This implies a joint initial belief distribution for ρ_1, ρ_2 the moments of which are

$$\begin{aligned}\mu_{\rho_1} &= \mu_\rho - \mu_{\rho_2}, \\ \sigma_\rho^2 &= \sigma_{\rho_1}^2 + \sigma_{\rho_2}^2 + 2\sigma_{\rho_1, \rho_2}, \\ \sigma_{\rho, \rho_2} &= \sigma_{\rho_1, \rho_2} + \sigma_{\rho_2}^2 = 0, \\ \sigma_{\rho_1, \rho_2} &= -\sigma_{\rho_2}^2, \\ \sigma_{\rho_1}^2 &= \sigma_\rho^2 + \sigma_{\rho_2}^2.\end{aligned}$$

In other words,

$$\begin{bmatrix} \rho_1 \\ \rho_2 \end{bmatrix} \sim N \left(\begin{bmatrix} \mu_\rho - \mu_{\rho_2} \\ \mu_{\rho_2} \end{bmatrix}, \begin{bmatrix} \sigma_\rho^2 + \sigma_{\rho_2}^2 & -\sigma_{\rho_2}^2 \\ -\sigma_{\rho_2}^2 & \sigma_{\rho_2}^2 \end{bmatrix} \right) \mathcal{I}(\rho_1, \rho_2).$$

We assume that the CBOs initial belief distribution about the the variance share γ of shocks to the trend component is a Beta distribution,

$$\gamma \sim \mathcal{B}(\alpha_\gamma, \beta_\gamma).$$

We assume that the CBOs initial belief distribution about the conditional variance σ^2 is an Inverse Gamma distribution,

$$\sigma^2 \sim \mathcal{IG}(\alpha_{\sigma^2}, \beta_{\sigma^2}).$$

Lastly, we assume that agents' initial beliefs about z_t and x_t in 1959Q3 are $z_t \sim N(y_{1959Q3}, 0.01^2)$ and $x_t \sim N(0, 0.01^2)$.

We start with an initial guess of the unknown parameters

$$\boldsymbol{\theta}^{(0)} = \left(\mu^{(0)}, \rho_1^{(0)}, \rho_2^{(0)}, \gamma^{(0)}, \sigma^{(0)}, \mathbf{z}_{1:t}^{(0)}, \mathbf{x}_{1:t}^{(0)} \right)'.$$

Given a draw of the parameters $\boldsymbol{\theta}^{(b)}$, we draw $\boldsymbol{\theta}^{(b+1)}$ as follows:

1. Draw $\mu^{(b+1)} | \rho_1^{(b)}, \rho_2^{(b)}, \gamma^{(b)}, \sigma^{(b)}, \mathbf{z}_{1:t}^{(b)}, \mathbf{x}_{1:t}^{(b)}, \mathbf{y}_{1:t}$. Given the other parameters, beliefs about μ can be updated from the equation for z_t :

$$\Delta z_t^{(b)} = \mu + \sqrt{\gamma^{(b)} \sigma^{(b)}} u_t.$$

Define

$$\begin{aligned} \tilde{\sigma}_\mu^2 &\equiv \left[\sigma_\mu^{-2} + \frac{t-1}{\gamma^{(b)} (\sigma^{(b)})^2} \right]^{-1}, \\ \tilde{\mu}_\mu &\equiv \tilde{\sigma}_\mu^2 \left[\frac{\mu_\mu}{\sigma_\mu^2} + \frac{\sum_{s=2}^t \Delta z_{s-1}^{(b)}}{\gamma^{(b)} (\sigma^{(b)})^2} \right]. \end{aligned}$$

The posterior of μ is $N(\tilde{\mu}_\mu, \tilde{\sigma}_\mu^2)$ and thus we draw $\mu^{(b+1)} \sim N(\tilde{\mu}_\mu, \tilde{\sigma}_\mu^2)$.

2. Draw $\rho_1^{(b+1)}, \rho_2^{(b+1)} | \mu^{(b+1)}, \gamma^{(b)}, \sigma^{(b)}, \mathbf{z}_{1:t}^{(b)}, \mathbf{x}_{1:t}^{(b)}, \mathbf{y}_{1:t}$. Given the other parameters, beliefs about ρ_1, ρ_2 can be updated from the equation for x_t :

$$x_t^{(b)} = \rho_1 x_{t-1}^{(b)} + \rho_2 x_{t-2}^{(b)} + \sqrt{(1 - \gamma^{(b)}) \sigma^{(b)}} v_t.$$

Define

$$\begin{aligned}\tilde{\Sigma}_\rho &\equiv \left[\Sigma_\rho^{-1} + \frac{\sum_{s=3}^t [x_{s-1}^{(b)}, x_{s-2}^{(b)}]' [x_{s-1}^{(b)}, x_{s-2}^{(b)}]}{(1 - \gamma^{(b)}) (\sigma^{(b)})^2} \right]^{-1}, \\ \tilde{\mu}_\rho &\equiv \tilde{\Sigma}_\rho \left[\Sigma_\rho^{-1} \mu_\rho + \frac{\sum_{s=3}^t [x_{s-1}^{(b)}, x_{s-2}^{(b)}]' x_s^{(b)}}{(1 - \gamma^{(b)}) (\sigma^{(b)})^2} \right].\end{aligned}$$

The posterior of $(\rho_1, \rho_2)'$ is $N(\tilde{\mu}_\rho, \tilde{\Sigma}_\rho)$ and thus we draw $(\rho_1^{(b+1)}, \rho_2^{(b+1)})' \sim N(\tilde{\mu}_\rho, \tilde{\Sigma}_\rho)$.

3. Draw $\gamma^{(b+1)} | \mu^{(b+1)}, \rho_1^{(b+1)}, \rho_2^{(b+1)}, \sigma^{(b)}, \mathbf{z}_{1:t}^{(b)}, \mathbf{x}_{1:t}^{(b)}, \mathbf{y}_{1:t}$. There is no closed form expression for the posterior of γ . We therefore draw it using a random walk Metropolis-Hastings step. Specifically, we draw a proposal $\tilde{\gamma}^{(b+1)} \sim N(\gamma^{(b)}, \sigma_{\gamma, prop}^2)$ where $\sigma_{\gamma, prop}^2$ is a proposal variance chosen such that this step has between a 25 and 40% acceptance rate over the burn-in period. We then set $\gamma^{(b+1)} = \tilde{\gamma}^{(b+1)}$ with probability α_{b+1} , where

$$\alpha_{b+1} \equiv \frac{L(\mathbf{y}_{1:t} | \mu^{(b+1)}, \rho_1^{(b+1)}, \rho_2^{(b+1)}, \tilde{\gamma}^{(b+1)}, \sigma^{(b)}, \mathbf{z}_{1:t}^{(b)}, \mathbf{x}_{1:t}^{(b)}) p_\gamma(\tilde{\gamma}^{(b+1)})}{L(\mathbf{y}_{1:t} | \mu^{(b+1)}, \rho_1^{(b+1)}, \rho_2^{(b+1)}, \gamma^{(b)}, \sigma^{(b)}, \mathbf{z}_{1:t}^{(b)}, \mathbf{x}_{1:t}^{(b)}) p_\gamma(\gamma^{(b)})}.$$

Otherwise we set $\gamma^{(b+1)} = \gamma^{(b)}$.

4. Draw $\sigma^{(b+1)} | \mu^{(b+1)}, \rho_1^{(b+1)}, \rho_2^{(b+1)}, \gamma^{(b+1)}, \mathbf{z}_{1:t}^{(b)}, \mathbf{x}_{1:t}^{(b)}, \mathbf{y}_{1:t}$. Given the other parameters, beliefs about σ can be updated from the two equations

$$\begin{aligned}\Delta z_t^{(b)} &= \mu^{(b+1)} + \sqrt{\gamma^{(b+1)}} \sigma u_t, \\ x_t^{(b)} &= \rho_1^{(b+1)} x_{t-1}^{(b)} + \rho_2^{(b+1)} x_{t-2}^{(b)} + \sqrt{1 - \gamma^{(b+1)}} \sigma v_t.\end{aligned}$$

Since u_t and v_t are independent, these regression equations can be treated as two independent sources of information. Define

$$\begin{aligned}\tilde{\alpha}_{\sigma^2} &\equiv \alpha_{\sigma^2} + (2t - 3)/2, \\ \tilde{\beta}_{\sigma^2} &\equiv \beta_{\sigma^2} + \frac{\sum_{s=2}^t \left(\Delta z_s^{(b)} - \mu^{(b+1)} \right)^2}{2\gamma^{(b+1)}} + \frac{\sum_{s=3}^t \left(x_s^{(b)} - \rho_1^{(b+1)} x_{s-1}^{(b)} - \rho_2^{(b+1)} x_{s-2}^{(b)} \right)^2}{2(1 - \gamma^{(b+1)})}.\end{aligned}$$

The posterior of σ^2 is $\mathcal{IG}(\tilde{\alpha}_{\sigma^2}, \tilde{\beta}_{\sigma^2})$ and thus we draw $(\sigma^{(b)})^2 \sim \mathcal{IG}(\tilde{\alpha}_{\sigma^2}, \tilde{\beta}_{\sigma^2})$.

5. Draw $\mathbf{z}_{1:t}^{(b+1)}, \mathbf{x}_{1:t}^{(b+1)} | \mu^{(b+1)}, \rho_1^{(b+1)}, \rho_2^{(b+1)}, \gamma^{(b+1)}, \sigma^{(b+1)}, \mathbf{y}_{1:t}$. This can be done using the stan-

standard Kalman filter and simulation smoother.

E Bayesian Forecasting of GDP

The algorithm described in Appendix D yields B samples of the posterior of the states and parameters of our UC model for GDP at each point in time t . We index these samples by b as follows $\left\{ \rho_1^{(b)}, \rho_2^{(b)}, \gamma^{(b)}, \mu^{(b)}, \sigma^{(b)}, z_{t|t}^{(b)}, x_{t|t}^{(b)}, x_{t-1|t}^{(b)} \right\}_{b=1}^B$. We then use the following algorithm to produce a real-time forecast distribution for the GDP at time t :

1. For each $b = 1, \dots, B$

- (a) Simulate a path of shocks $\left\{ u_{t+h}^{(b)}, v_{t+h}^{(b)} \right\}_{h=1}^H$ from the standard Normal distribution.
- (b) Starting from $h = 1$, construct a simulated path of the states over H subsequent periods using equations

$$\begin{aligned} z_{t+h|t}^{(b)} &= \mu^{(b)} + \sqrt{\gamma^{(b)} \sigma^{(b)}} u_{t+h}^{(b)}, \\ x_{t+h|t}^{(b)} &= \rho_1^{(b)} x_{t+h-1|t}^{(b)} + \rho_2^{(b)} x_{t+h-2|t}^{(b)} + \sqrt{1 - \gamma^{(b)} \sigma^{(b)}} v_{t+h}^{(b)}. \end{aligned}$$

- (c) Use the simulated states to construct $\left\{ y_{t+h|t}^{(b)} \right\}_{h=1}^H$ where

$$y_{t+h|t}^{(b)} = z_{t+h|t}^{(b)} + x_{t+h|t}^{(b)}.$$

2. The forecast of y_{t+h} given time t information is computed as

$$F_t y_{t+h} = \frac{1}{B} \sum_{b=1}^B y_{t+h|t}^{(b)}.$$

At the end of the estimation we are left with a sequence of model-implied 1 to H-quarter ahead forecasts $\{F_t y_{t+h}\}_{h=1}^H$ for every year t from 1976Q4 to 2019Q4.

We must perform a few additional steps to transform our forecasts to ones that are comparable to those produced by the CBO. The CBO publishes forecasts of growth in the average annual level of real output. We define the average annual level of real output over the year preceding quarter t as

$$\bar{Y}_t \equiv \frac{1}{4} \sum_{s=t-3}^t \exp(y_s).$$

As an example, in the CBO's economic outlook published in 1990, its 1-year ahead forecast of GDP growth is

$$100 \times \left(\frac{\bar{Y}_{1990Q4}}{\bar{Y}_{1989Q4}} - 1 \right).$$

Thus to convert the model's forecasts of quarterly log real GDP to average annual h -year ahead level forecasts, we apply the following transformation to the simulated forecast distribution

$$F_t \bar{Y}_{t+h} \equiv \frac{1}{B} \sum_{b=1}^B \left[\frac{1}{4} \sum_{s=t+4h-3}^{t+4h} \exp \left(F_t y_s^{(b)} \right) \right].$$

The associated forecasts of growth in average annual levels between year $t + h - 1$ and $t + h$ for $h = 1, \dots, H$ are

$$100 \times \left(\frac{F_t \bar{Y}_{t+h}}{F_t \bar{Y}_{t+h-1}} - 1 \right).$$

F Search over Initial Beliefs for GDP Growth

We denote $\theta = (\mu_\rho, \sigma_\rho, \mu_{\rho_2}, \sigma_{\rho_2}, \alpha_\gamma, \beta_\gamma)'$. Let $\alpha = \{\alpha_h\}_{h=1}^H$ and $\beta = \{\beta_h\}_{h=1}^H$ denote vectors of estimated coefficients from the forecasting anomaly regressions for different horizons up through a maximum horizon of H using the CBO data. Let $\hat{\alpha} = \{\hat{\alpha}_h\}_{h=1}^H$ and $\hat{\beta} = \{\hat{\beta}_h\}_{h=1}^H$ denote those same quantities estimated on the model implied forecasts and yields. Additionally, denote the t -statistics associated with these coefficients as $\{t_\alpha, t_\beta\} = \{t_{\alpha,h}, t_{\beta,h}\}_{h=1}^H$ for the data and $\{t_{\hat{\alpha}}, t_{\hat{\beta}}\} = \{t_{\hat{\alpha},h}, t_{\hat{\beta},h}\}_{h=1}^H$ for the model. Define the moment function as

$$\hat{m}(\theta) = \begin{bmatrix} \alpha_{bias} - \hat{\alpha}_{bias} \\ t_{\alpha,bias} - t_{\hat{\alpha},bias} \\ \beta_{ar} - \hat{\beta}_{ar} \\ t_{\beta,ar} - t_{\hat{\beta},ar} \\ \beta_{mz} - \hat{\beta}_{mz} \\ t_{\beta,mz} - t_{\hat{\beta},mz} \\ \beta_{cg} - \hat{\beta}_{cg} \\ t_{\beta,cg} - t_{\hat{\beta},cg} \end{bmatrix} \quad (18)$$

The parameters are then estimated via SMM with the following objective function

$$\hat{\theta} = \operatorname{argmax}_{\theta} \hat{m}(\theta)' \mathbf{W} \hat{m}(\theta)$$

where the elements of the objective function associated with the Mincer-Zarnowitz and Coibion-Gorodnichenko coefficients are given 3 times the weight of all other elements in \mathbf{W} . We also place bounds on the estimated parameters as described in footnote 8 in the main text. The estimated initial belief distributions are plotted in Figure 8.

Every evaluation of the moment function $\hat{m}(\theta)$ requires us to sample from the posterior of the UC model sequentially. Since this step is very computationally costly, we only re-estimate the model every 4 quarters rather than every quarter, and use a burn-in sample of 15,000 draws and keep the subsequent 15,000 draws rather than 50,000 for each of those quantities in our empirical specification. The global minimum is found using MATLAB's "particleswarm" optimization routine.

References

- ANDRADE, P., R. K. CRUMP, S. EUSEPI, AND E. MOENCH (2016): "Fundamental Disagreement," *Journal of Monetary Economics*, 83, 106–128.
- BANSAL, R. AND I. SHALIASTOVICH (2013): "A Long-run Risks Explanation of Predictability Puzzles in Bond and Currency Markets," *The Review of Financial Studies*, 26, 1–33.
- BARSKY, R. B. AND J. B. DE LONG (1993): "Why Does the Stock Market Fluctuate?" *Quarterly Journal of Economics*, 108, 291–311.
- BAUER, M. D. AND G. D. RUDEBUSCH (2020): "Interest Rates under Falling Stars," *American Economic Review*, 110, 1316–54.
- BEKAERT, G., R. J. HODRICK, AND D. MARSHALL (1997): "On Biases in Tests of the Expectations Hypothesis of the Term Structure of Interest Rates," *Journal of Financial Economics*, 42, 309–348.
- (2001): "Peso Problem Explanations for Term Structure Anomalies," *Journal of Monetary Economics*, 48, 241–270.
- BEN-DAVID, I., J. R. GRAHAM, AND C. R. HARVEY (2013): "Managerial Miscalibration," *Quarterly Journal of Economics*, 128, 1547–1584.
- BIANCHI, F., M. LETTAU, AND S. C. LUDVIGSON (2020): "Monetary policy and Asset Valuation," *Journal of Finance*.
- BIDDER, R. AND I. DEW-BECKER (2016): "Long-Run Risk Is the Worst-Case Scenario," *American Economic Review*, 106, 2494–2527.
- BORDALO, P., N. GENNAIOLI, Y. MA, AND A. SHLEIFER (2020): "Overreaction in Macroeconomic Expectations," *American Economic Review*, 110, 2748–2782.
- BRAV, A. AND J. B. HEATON (2002): "Competing Theories of Financial Anomalies," *Review of Financial Studies*, 15, 575–606.
- CAMPBELL, J. Y. AND R. SHILLER (1991): "Yield Spreads and Interest Rate Movements: A Bird's Eye View," *Review of Economic Studies*, 58, 495–514.
- CAO, S., R. K. CRUMP, S. EUSEPI, AND E. MOENCH (2021): "Fundamental Disagreement about Monetary Policy and the Term Structure of Interest Rates," Working Paper, University of Texas, Austin.
- CIESLAK, A. (2018): "Short-rate expectations and unexpected returns in treasury bonds," *The Review of Financial Studies*, 31, 3265–3306.
- CIESLAK, A. AND P. POVALA (2015): "Expected Returns in Treasury Bonds," *Review of Financial Studies*, 28, 2859–2901.
- COCHRANE, J. (1988): "How Big Is the Random Walk in GNP?" *Journal of Political Economics*, 96, 893–920.
- COGLEY, T. AND T. J. SARGENT (2005): "The Conquest of US Inflation: Learning and Robustness to Model Uncertainty," *Review of Economic Dynamics*, 8, 528–563.
- COIBION, O. AND Y. GORODNICHENKO (2012): "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy*, 120, 116–159.
- (2015): "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review*, 105, 2644–2678.
- COLLIN-DUFRESNE, P., M. JOHANNES, AND L. A. LOCHSTOER (2016): "Parameter Learning in General Equilibrium: The Asset Pricing Implications," *American Economic Review*, 106, 664–98.
- CROCE, M. M., M. LETTAU, AND S. C. LUDVIGSON (2015): "Investor Information, Long-run Risk,

- and the Term Structure of Equity," *Review of Financial Studies*, 28, 706–742.
- CROUSHORE, D. (1998): "Evaluating Inflation Forecasts," Working Paper, Federal Reserve Bank of Philadelphia.
- CRUMP, R. K., S. EUSEPI, E. MOENCH, AND B. PRESTON (2021): "The Term Structure of Expectations," in *Handbook of Economic Expectations*, ed. by R. Bachmann, G. Topa, and W. van der Klaauw, Amsterdam, Holland: Elsevier.
- EUSEPI, S. AND B. PRESTON (2011): "Expectation, Learning, and Business Cycle Fluctuations," *American Economic Review*, 101, 2844–2872.
- (2018): "The Science of Monetary Policy: An Imperfect Knowledge Perspective," *Journal of Economic Literature*, 56, 3–59.
- EVANS, G. W. AND S. HONKAPOHJA (2001): *Learning and Expectations in Macroeconomics*, Princeton, NJ: Princeton University Press.
- FAMA, E. (1984): "The Information in the Term Structure," *Journal of Financial Economics*, 13, 509–528.
- FAMA, E. F. AND R. R. BLISS (1987): "The Information in Long-Maturity Forward Rates," *American Economic Review*, 77, 680–692.
- FRIEDMAN, B. M. (1979): "Optimal Expectations and the Extreme Information Assumptions of 'Rational Expectations' Macromodels," *Journal of Monetary Economics*, 5, 23–41.
- (1980): "Survey Evidence on the 'Rationality' of Interest Rate Expectations," *Journal of Monetary Economics*, 6, 453–465.
- FROOT, K. A. (1989): "New Hope for the Expectations Hypothesis of the Term Structure of Interest Rates," *Journal of Finance*, 44, 283–305.
- GIACOLETTI, M., K. T. LAURSEN, AND K. J. SINGLETON (2018): "Learning and Risk Premiums in an Arbitrage-free Term Structure Model," Working Paper, Stanford University.
- GOURINCHAS, P.-O. AND A. TORNELL (2004): "Exchange Rate Puzzles and Distorted Beliefs," *Journal of International Economics*, 64, 303–333.
- GUO, H. AND J. WACHTER (2019): "'Superstitious' Investors," NBER Working Paper No. 25603.
- JOHANNES, M., L. A. LOCHSTOER, AND Y. MOU (2016): "Learning about Consumption Dynamics," *Journal of Finance*, 71, 551–600.
- KOZICKI, S. AND P. TINSLEY (2001): "Shifting endpoints in the term structure of interest rates," *Journal of Monetary Economics*, 47, 613–652.
- KOZLOWSKI, J., L. VELDKAMP, AND V. VENKATESWARAN (2020): "The Tail that Wags the Economy: Beliefs and Persistent Stagnation," *Journal of Political Economy*, 128, 2839–2879.
- LEWELLEN, J. AND J. SHANKEN (2002): "Learning, Asset-pricing Tests, and Market Efficiency," *Journal of Finance*, 57, 1113–1145.
- LEWIS, K. K. (1989a): "Can Learning Affect Exchange-Rate Behavior?: The Case of the Dollar in the Early 1980's," *Journal of Monetary Economics*, 23, 79–100.
- (1989b): "Changing Beliefs and Systematic Rational Forecast Errors with Evidence from Foreign Exchange," *American Economic Review*, 621–636.
- LIU, Y. AND J. C. WU (2020): "Reconstructing the Yield Curve," *Journal of Financial Economics*, forthcoming.
- MADDALA, G. S. (1991): "Survey data on Expectations: What Have We Learnt?" in *Issues in Contemporary Economics*, ed. by M. Nerlove, London, UK: Palgrave MacMillan, 319–344.

- MANKIW, N. G. AND R. REIS (2002): "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics*, 117, 1295–1328.
- MANKIW, N. G., R. REIS, AND J. WOLFERS (2003): "Disagreement about Inflation Expectations," *NBER Macroeconomics Annual*, 18, 209–248.
- MINCER, J. A. AND V. ZARNOWITZ (1969): "The Evaluation of Economic Forecasts," in *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance*, ed. by J. A. Mincer, New York, NY: NBER, 3–46.
- NAGEL, S. AND Z. XU (2021): "Dynamics of Subjective Risk Premia," Working Paper, University of Chicago.
- NORDHAUS, W. (1987): "Forecasting Efficiency: Concepts and Applications," *Review of Economics and Statistics*, 69, 667–674.
- PIAZZESI, M., J. SALOMAO, AND M. SCHNEIDER (2015): "Trend and Cycle in Bond Premia," Working Paper, Stanford University.
- ROMERO, J. (2013): "Treasury-Fed Accord," <https://www.federalreservehistory.org/essays/treasury-fed-accord>.
- SARGENT, T. J. (2001): *The Conquest of American Inflation*, Princeton, NJ: Princeton University Press.
- SCHUH, S. (2001): "An Evaluation of Recent Macroeconomic Forecast Errors," *New England Economic Review*, 2001, 35–36.
- SHILLER, R. J. (1979): "The Volatility of Long-Term Interest Rates and Expectations Theories of the Term Structure," *Journal of Political Economics*, 87, 1190–1219.
- SHILLER, R. J., J. Y. CAMPBELL, AND K. L. SCHOENHOLTZ (1983): "Forward Rates and Future Policy: Interpreting the Term Structure of Interest Rates," *Brookings Papers on Economic Activity*, 1983, 173–223.
- SIMS, C. A. (2003): "Implications of Rational Inattention," *Journal of Monetary Economics*, 50, 665–690.
- SINGLETON, K. (2021): "How Much "Rationality" Is There in Bond-Market Risk Premiums?" *Journal of Finance*, 76, 1611–1654.
- TIMMERMANN, A. G. (1993): "How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices," *Quarterly Journal of Economics*, 108, 1135–1145.
- VAN DIJK, D., S. J. KOOPMAN, M. VAN DER WEL, AND J. H. WRIGHT (2014): "Forecasting Interest Rates with Shifting Endpoints," *Journal of Applied Econometrics*, 29, 693–712.
- WACHTER, J. (2006): "A Consumption-Based Model of the Term Structure of Interest Rates," *Journal of Financial Economics*, 79, 365–399.
- WOODFORD, M. (2003): "Imperfect Common Knowledge and the Effects of Monetary Policy," *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, 25.
- XU, Z. (2019): "Expectation Formation in the Treasury Bond Market," Working Paper, City University of Hong Kong.