Learning and Risk Premiums in an Arbitrage-free Term Structure Model*

Marco Giacoletti Kristoffer T. Laursen Kenneth J. Singleton †

May 20, 2018

Abstract

We study the evolution of risk premiums on US Treasury bonds from the perspective of a *real-time Bayesian learner* $\mathcal{RA}$ who updates her beliefs using a dynamic term structure model. Learning about the historical dynamics of yields led to substantial variation in $\mathcal{RA}$'s subjective risk premiums. Moreover, she gained substantial forecasting power by conditioning her learning on measures of disagreement among professional forecasters about future yields. This gain was distinct from the (much weaker) forecasting power of macroeconomic information. $\mathcal{RA}$'s views about the pricing distribution of yields remained nearly constant over time. Her learning rule outperformed consensus forecasts of market professionals, particularly following U.S. recessions.

*We are grateful for feedback from Darrell Duffie, Lars Lochstoer, Jules Van Binsbergen, and seminar participants at Columbia University, the University of Michigan, University of Southern California, and Washington University-St. Louis.

†Marco Giacoletti is at the Marshall School of Business, University of Southern California, marco.giacoletti@marshall.usc.edu; Kristoffer T. Laursen is at AQR Capital Management, LLC; and Kenneth J. Singleton is at the Stanford University, Graduate School of Business, and NBER, kenneths@stanford.edu

1 Introduction

Market participants trading US treasury bonds need to form *prospective* views in *real-time* on bond expected returns, while facing structural changes in the economy (e. g. in policy, regulatory or political environments). However, most of the literature that studies risk compensation in bond markets uses *retrospective within-sample* measures of risk premiums,[1] and mostly relies on the assumption that the parameters governing the evolution of the risk factors in the economy are fixed and known by market participants.

In this paper we instead take the perspective of an agent $\mathcal{RA}$ who is forming beliefs in real-time about the distribution of future bond yields. $\mathcal{RA}$ is modeled as a Bayesian learner who forms beliefs about the parameters governing the dynamics of the key factors driving future excess returns, while taking account of possible structural changes in these parameter. We document that real-time learning materially changes estimates of bond risk premiums, relative to those based on retrospective full-sample estimates, and we explore the nature of the conditioning information that lends precision to $\mathcal{RA}$'s learning rule.

The yield curve is a high-dimensional object, and restrictions based on the cross section of bond yields have proven to be crucial in generating accurate estimates of risk premiums. Accordingly, Bayesian learning is explored within a dynamic term structure model (*DTSM*) in which the priced risks are the first three principal components (PCs or $\mathcal{P}$) of bond yields. Specifically, $\mathcal{RA}$ prices $\mathcal{P}$-sensitive payoffs using a stochastic discount factor (SDF) that reflects learning within a conditionally Gaussian framework, under the presumption that her conditioning information is informative about future yields.

Initially, $\mathcal{RA}$'s conditioning information is set equal to the history of the priced risks $\mathcal{P}$. This benchmark case is premised on investors recognizing that the yield portfolios $\mathcal{P}$ are effective predictors of future bond returns.[2] A striking empirical finding that emerges from this benchmark is that $\mathcal{RA}$ effectively knows the risk-neutral ($\mathbb{Q}$) drift of the PC's of bond yields and treats its parameters as virtually constant over time within her real-time Bayesian learning scheme. We show formally that, within our restricted arbitrage-free *DTSM*, this implies that $\mathcal{RA}$'s learning rule simplifies to a version of constant-gain learning.

As a reference point for $\mathcal{RA}$'s expected excess returns, we use the consensus beliefs among professional forecasters from the Blue-Chip Financial Forecasters (BCFF) panel, which is a unique survey run monthly among professional forecasters at financial and non-financial institutions. The forecasters included in the survey provide their forecasts for key macroeconomic variables and a wide cross-section of bond yields. Notably, $\mathcal{RA}$'s real-time learning rule conditioned only on $\mathcal{P}$ gives substantially smaller out-of-sample RMSEs across

[1]See, for examples, Duffee (2002), Ludvigson and Ng (2010), and Joslin, Priebsch, and Singleton (2014).

[2]See Cochrane and Piazzesi (2002), Cochrane and Piazzesi (2008) and Joslin, Singleton, and Zhu (2011).

the entire yield curve, relative to BCFF forecasts, particularly for long-maturity bonds.

Not only does $\mathcal{RA}$ disagree with the BCFF consensus about the future path of bond yields, but there is substantial disagreement within the BCFF panel of forecasters.[3] Once investor disagreement is "aggregated" in bond markets, it might well affect the distribution of bond prices through investors' stochastic discount factors. To explore this possibility empirically, we extend $\mathcal{RA}$'s conditioning information set to include measures of disagreement on expected future bond yields based on the BCFF data (hereafter H_t). Learning rules conditioned on the joint history of yield information and forecaster disagreement $(\mathcal{P}_t, H_t)$ do indeed give more accurate real-time forecasts of yields than those conditioned on $\mathcal{P}$ alone.

Accommodating learning conditioned on disagreement materially affects measured risk premiums. Moreover, $\mathcal{RA}$'s constant-gain (Bayesian) learning rule gives systematically more accurate real-time forecasts (lower RMSE's) than the consensus BCFF forecasts and forecasts from random-walk models for individual yields. The outperformance with respect to BCFF is particularly evident following NBER recessions,[4] when the median BCFF forecaster repeatedly predicted that long-bond yields would rise much faster than they actually rose.

There are two potential sources of the forecasting power of disagreement in $\mathcal{RA}$'s learning rule: (i) a direct effect on future PC's through the law of motion of the state vector $(\mathcal{P}_t, H_t)$, and (ii) an indirect effect through Bayesian updates of the parameters governing the covariances of the current PC's with their future values. We show that the indirect effects are the dominant ones, and they are especially large following NBER recessions. When the U.S. economy is emerging from a recession, knowledge of the extent of disagreement among professionals is informative about how today's yield curve will impact its future shape.

Might the predictive power of H_t in $\mathcal{RA}$'s learning rule arise because yield disagreement is proxying for other macro factors? We find that the forecasting power of H_t is distinct from that of inflation, output growth, or aggregate productivity. Nor is the predictive power of H_t spanned by forward looking expectations of inflation and output, or by disagreement about future macro conditions computed from BCFF data.[5] In fact, within our learning rules, introducing macro factors or beliefs about these factors has only weak real-time predictive power for future yields, whereas conditioning on yield disagreement substantially lowers out-of-sample mean-squared forecast errors.

Our analysis of learning in bond markets complements several related studies. There

[3]For complementary evidence on how disagreement among market professionals about macroeconomic conditions impacts the future paths of bond yields, see, for examples, Dovern, Fritsche, and Slacalek (2012) and Andrade, Crump, Eusepi, and Moench (2014).

[4]Cieslak (2017) explores in depth the properties of consensus professional forecasts of the federal funds rate and finds that these short-rate expectations are similarly inaccurate during economic downturns.

[5]Buraschi and Whelan (2016) and Ehling, Gallmeyer, Heyerdahl-Larsen, and Illeditsch (2016) show that disagreement among professional forecasters on future inflation and real growth have predictive power.

is a large literature incorporating survey information directly into *DTSMs*. Extending the frameworks of Kim and Orphanides (2012) and Chun (2011), Piazzesi, Salomao, and Schneider (2013) model survey forecasts as subjective views that are distinct from those of the econometrician. In these models, the median forecaster has full knowledge of risk-factor dynamics and her forecasts are spanned by the low-order PCs of bond yields. However, empirically, a large percentage of the variation in median BCFF forecasts is *not* spanned by the risk factors $\mathcal{P}$. We extend these frameworks by: (i) introducing learning about parameters; (ii) allowing for belief heterogeneity to affect market prices of factor risks; and (iii) not tying $\mathcal{RA}$'s objective forecasts from her *DTSM* to the consensus BCFF forecasts. Together, these features materially improve the accuracy of $\mathcal{RA}$'s fitted risk premiums and real-time forecasts relative to the median professional forecaster.

Agents in the bond-pricing models of David (2008), Xiong and Yan (2009), and Buraschi and Whelan (2016) optimally filter for an unknown state (e.g., aggregate output), while "agreeing to disagree" about the *known* values of the parameters governing the state process. Equilibrium bond prices depend on the pairwise relative beliefs across all agent types, thereby giving rise to a potentially high-dimensional factor space.[6] Yet the low-order PCs account for the vast majority of the cross-sectional variation in bond yields. Therefore, we instead follow Joslin, Singleton, and Zhu (2011) (JSZ) and represent $\mathcal{P}$ in terms of yield PCs, which market participants can reasonably be assumed to observe without error (Joslin, Le, and Singleton (2013)). This motivates a second key difference with our setting: in our model the agent is learning about the unknown parameters governing a directly observed state vector.[7]

Section 2 discusses the importance of learning and disagreement in simple theoretical settings widely used in the literature. Additionally, it provides descriptive empirical evidence (without the structure of a *DTSM*) on the importance of learning and dispersion of beliefs in forming expectations on yields and excess returns. Section 3 sets forth our formal learning problem in the context of a dynamic term structure model. Our formal learning rules are then implemented empirically in Section 4 and Section 5. The extent to which conditioning on belief dispersion H proxies for other forms of information about the macroeconomy is explored in Section 6. Concluding remarks are in Section 7.

For comparability with the vast majority of macro-finance *DTSMs*, we explore the impact

[6]Similarly, the setting of Barillas and Nimark (2014) gives rise to a "forecasting the forecast of others" problem (Townsend (1983), Singleton (1987)) which, in turn, leads to an infinite-dimensional set of higher-order beliefs affecting bond prices.

[7]Collin-Dufresne, Johannes, and Lochstoer (2016) study equity risk premiums implied by a representative-agent, consumption-based model in which there is learning about parameters. We instead focus on a reduced-form SDF to ensure high accuracy in pricing of the entire yield curve (Dai and Singleton (2000), Duffee (2002)), while exploring learning about the objective distribution of the (in our case bond-relevant) state of the economy. Additionally, the Bayesian learner's SDF in our model explicitly recognizes that investors are heterogeneous in their beliefs and that this heterogeneity may be a source of priced risk.

of learning on risk premiums within a conditionally Gaussian *DTSM*. This is natural given that risk premiums are determined by the conditional first moments of yields. Of course investors may update views about factor volatilities and, in fact, this appears to be the case with $\mathcal{RA}$ within her conditionally Gaussian framework. Both the descriptive evidence presented in Section 2 and the analysis of an extended *DTSM* with stochastic volatility in Appendix H suggest that our core findings about risk premiums in the presence of learning are robust to the presence of stochastic volatility.

2 Motivating an Impact of Learning on Risk Premiums

Before characterizing learning in the context of an arbitrage-free *DTSM*, it will be instructive to highlight several robust implications of learning and investor disagreement within a simple equilibrium setting. We first show that learning on the part of a representative marginal agent can lead to risk premiums that are substantially different than those that an econometrician would compute under the assumption that the marginal agent knows the parameters governing her consumption process. This setting is then extended to allow for disagreement among market participants. The latter analysis provides motivation for conditioning on measures of disagreement when modeling risk premiums, and the descriptive evidence presented subsequently supports such conditioning.

2.1 Learning in the Bond Market

Suppose that a representative marginal agent $\mathcal{RA}$ has preferences described by a constant relative risk aversion utility function: $u(C_t) = \delta \frac{C_t^{1-\alpha}}{1-\alpha}$, with consumption growth $dC_t/C_t = dc_t$ being the only source of uncertainty in this economy. The evolution of c_t is governed by

$$dc_t = \mu_{ct} + \sigma_c dW_t^c, \tag{1}$$

where dW_t^c is a Brownian motion, $\mu_{ct} = \kappa(\theta_c - c_t)$, and κ captures mean reversion.

If $\mathcal{RA}$ knows the parameters determining the dynamics of c_t, then the instantaneous short rate is given by

$$r_t = -\log(\delta) - \alpha(1+\alpha)\frac{\sigma_c^2}{2} + \alpha\mu_{ct}.$$

The logarithm of the price of a *real* zero-coupon bond with maturity of n periods is an affine function of c_t (Vasicek (1977)),[8]

$$p_t^n = a_n\big(\theta_c, \kappa\big) + b_n\big(\kappa\big)c_t. \tag{2}$$

Finally, the expected value of the instantaneous excess return for this bond is

$$E[xr_{t+dt}^n|\kappa] = \alpha b_n\big(\kappa\big)\sigma_c^2 \equiv \alpha\sigma_{nc}. \tag{3}$$

Risk compensation is determined by the coefficient of relative risk aversion and the covariance σ_{nc} between shocks to c_t and bond prices.

Risk premiums for bonds are time-invariant in this simple model. However, the model still provides an informative framework for exploring the implications of learning. Suppose that $\mathcal{RA}$ does not know κ and that she has the Normal prior belief $\tilde{\kappa}_0 \sim \mathcal{N}\big(\hat{\kappa}_0, \sigma_{\kappa 0}\big)$. Over time she learns about κ from the past history of c_t based on the dynamic updating rules:

$$\begin{aligned} dc_t &= \tilde{\kappa}_t\big(\theta_c - c_t\big) + \sigma_c d\widetilde{W}_t^c \;,\; \tilde{\kappa}_t \sim \mathcal{N}\big(\hat{\kappa}_t, \sigma_{\kappa t}^2\big), \\ d\hat{\kappa}_t &= \frac{\sigma_{\kappa t}^2(\theta_c - c_t)}{\sigma_c^2 + \sigma_{\kappa t}^2\big(\theta_c - c_t\big)^2}\big(dc_t - \hat{\kappa}_t(\theta_c - c_t)\big), \\ d\sigma_{\kappa t}^2 &= -\frac{\sigma_{\kappa t}^4(\theta_c - c_t)^2}{\sigma_c^2 + \sigma_{\kappa t}^2\big(\theta_c - c_t\big)^2}. \end{aligned}$$

To illustrate the impact of learning on risk premiums in a simplified setting we suppose that $\mathcal{RA}$ ignores parameter uncertainty when pricing bonds. This assumption, referred to as "anticipated utility" in equilibrium models with Bayesian learning (see, e.g., Kreps (1998) and Cogley and Sargent (2008)), is often made for tractability, particularly in high dimensional models.[9] Under this convention, risk compensation at time t is given by

$$E[xr_{t+dt}^n|\hat{\kappa}_t] = \alpha b_n\big(\hat{\kappa}_t\big)\sigma_c^2 = \alpha\sigma_{nc}(\hat{\kappa}_t). \tag{4}$$

[8]Here $b_n\big(\kappa\big)$ solves the ordinary differential equation $b_n^{'} = \alpha\kappa - \kappa b_n + \frac{1}{2}b_n^2\sigma_c^2$ and $a_n\big(\theta_c, \kappa\big)$ solves $a_n^{'} = -\log(\delta) + \alpha\big(1+\alpha\big)\frac{\sigma_c^2}{2} - \alpha\kappa\theta_g + b_n\kappa\theta_c + \frac{1}{2}b_n^2\sigma_c^2$. For ease of notation, we explicitly highlight only the dependence of a_n and b_n on the parameters determining μ_{ct}.

[9]If $\mathcal{RA}$ factors uncertainty on $\tilde{\kappa}_t$ into her calculations, she would demand risk compensation equal to:

$$E[xr_{t+dt}^n|\hat{\kappa}_t, \sigma_{\kappa t}] = \int E[xr_{t+1}^n|\tilde{\kappa}_t]f(\tilde{\kappa}_t)d\tilde{\kappa}_t,$$

where $f(\cdot)$ stands for the Normal density function. Among recent studies of learning and the pricing of equities, Johannes, Lochstoer, and Mou (2016) also adopt the assumption of anticipated utility, whereas Collin-Dufresne, Johannes, and Lochstoer (2016) examine the impact of parameter uncertainty and learning in a more tractable low-dimensional Markov setting.

In this setting there is a direct link from learning about κ_t to the dynamics of risk premiums. An econometrician who assumes that κ is known to $\mathcal{RA}$ and estimates this parameter using a full sample up to date T would set κ equal to $\hat{\kappa}_T$. In contrast, $\mathcal{RA}$ would compute real-time updates κ_t based on her learning rule, with $\hat{\kappa}_T$ generally not equal to $\hat{\kappa}_t$, for $t < T$. Introducing learning about the parameters of an inflation process would add an additional channel through which learning could affect risk premiums on *nominal* bonds.

Motivated by the structure of (4) and the goal of more market-relevant insight into the effects of learning on risk premiums on *nominal* bonds, we extend this representation of xr^n to allow for state-dependence. A typical representation of risk premiums in affine *DTSMs* for nominal bonds has xr^n depending linearly on the first three *PC*s of yields according to

$$xr^n_{t+h} = \alpha_n + \mathcal{B}_{n\mathcal{P}}\mathcal{P}_t + \sigma_v v_{t+h}, \tag{5}$$

where $\mathcal{B}_{n\mathcal{P}}$ is a row vector of coefficients and v_{t+h} is a mean zero residual. Suppose that $\mathcal{RA}$ updates $\mathcal{B}_{n\mathcal{P}}$ in real time using recursive least-squares.[10] For this simplified setting, $\mathcal{RA}$'s time t estimate $\hat{\mathcal{B}}_{nt}$ is given by:

$$\begin{aligned} \hat{\mathcal{B}}_{nt} &= \hat{\mathcal{B}}_{n,t-1} + R_t^{-1}X_t' \left(xr^n_{t+h} - \hat{\alpha}_{n,t-1} - \hat{\mathcal{B}}_{n,\mathcal{P},t-1}\mathcal{P}_t \right), \\ R_t &= R_{t-1} + X_t' X_t, \end{aligned} \tag{6}$$

where $\mathcal{B}_n = \left[\alpha \vdots \mathcal{B}_{n\mathcal{P}}\right]$ and $X_t \equiv [1, \mathcal{P}_t']$.

Figure 1 shows the difference in fitted expected excess returns from (5) over a one-quarter horizon ($h = 0.25y$) for a ten-year zero-coupon US Treasury bond, full-sample estimates minus those based on recursive least-squares learning.[11] The shaded areas correspond to NBER recessions. Even this simple learning scheme gives rise to risk premiums that are notably different from their full-sample counterparts, as much as 4% annualized. (The differences exceed 5% for the $h = 1y$ horizon.) Moreover, the mechanical convergence of the recursive least-squares estimates to the full-sample estimates is quite slow.

2.2 Accommodating Dispersion of Beliefs

Multiple institutional traders participate in US Treasury markets, and survey evidence suggests that these professionals substantially disagree on the future path of macroeconomic

[10]In Section 3.2 we show formally that this is a Bayesian learning rule when $\mathcal{RA}$ views $\mathcal{B}_{n\mathcal{P}}$ as unknown and fixed over time. Within the *DTSM* setting, $\mathcal{RA}$ will be following a more general rule that nests this special case.

[11]We use yields on zero-coupon bonds with maturities of 6 months and 1, 2, 3, 5, 7, and 10 years calculated from coupon-bond yields as reported in the CRSP database using the Fama-Bliss methodology for the sample period June, 1961 through December, 2015.

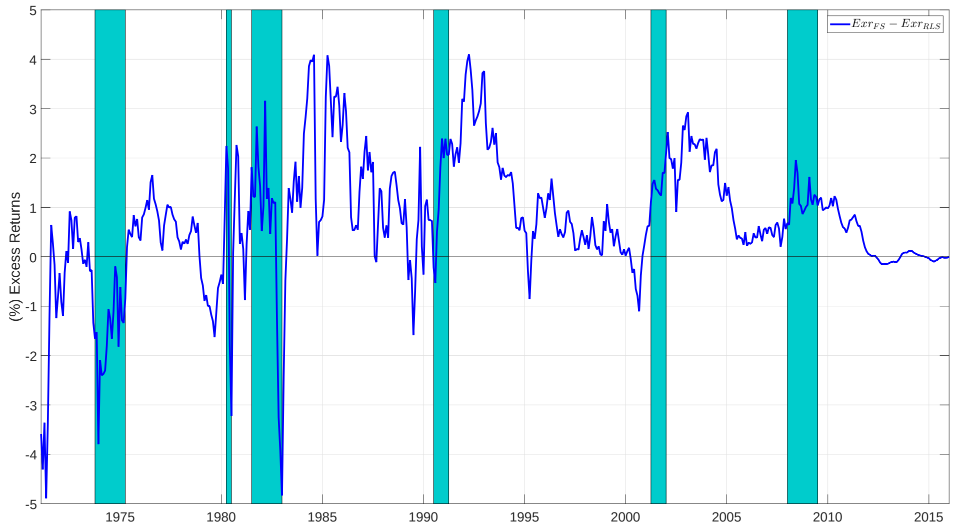

Figure 1: Differences in one-quarter ahead expected excess returns for a 10-year zero-coupon bond implied by the full-sample estimates less those from the recursive least-squares learning scheme. The sample is monthly for the period from January 1972 to December 2015.

fundamentals and bond yields. Recent studies of equilibrium models in which heterogeneous agents "agree to disagree" (e.g., Xiong and Yan (2009) and Buraschi and Whelan (2016)) show theoretically how this disagreement can affect equilibrium pricing of bonds.[12]

To illustrate how disagreement and learning interact in affecting risk premiums, suppose that the economy has the same features as in the previous section except that bonds are now priced by two agents, a and b, who both have constant relative risk aversion utility with relative risk aversion coefficient α. For simplicity, suppose that these agents agree to disagree about the value of the mean-reversion coefficient κ of the state c_t, so that their subjective dynamics of consumption are respectively:

$$dc_t^a = \hat{\mu}_{ct}^a + \sigma_c dW_t^{ac},$$
$$dc_t^b = \hat{\mu}_{ct}^b + \sigma_c dW_t^{bc},$$

where $\hat{\mu}_{ct}^a = \hat{\kappa}_t^a \big(\theta_c - c_t\big)$ and $\hat{\mu}_{ct}^b = \hat{\kappa}_t^b \big(\theta_c - c_t\big)$. The (scaled) disagreement $\Psi_t = \frac{\hat{\mu}_{ct}^a - \hat{\mu}_{ct}^b}{\sigma_c}$ determines the change of measure between the subjective physical probability measures used by a and b, $\eta_t = \frac{d\mathbb{P}_t^a}{d\mathbb{P}_t^b}$. Basak (2005) shows in this setting that with complete markets the

[12] David (2008) develops a difference-of-opinion model with similar features for understanding risk premiums in the equity markets.

consumption shares for the two agents, ω^a and ω^b, depend on the history of disagreement η_t, and that the short rate takes the form,

$$r_t = \delta - \frac{1}{2}\alpha(1+\alpha)\sigma_c^2 + \alpha\widehat{\mu}_{ct} + \\ + \alpha\omega^a(\eta_t)(\hat{\mu}_{gt}^a - \widehat{\mu}_{ct}) + \alpha\omega^b(\eta_t)(\hat{\mu}_{ct}^b - \widehat{\mu}_{ct}) + \frac{\alpha-1}{2\alpha}\omega^a(\eta_t)\omega^b(\eta_t)\Psi_t^2. \tag{7}$$

The implied risk premium on a real bond of maturity n is (Buraschi and Whelan (2016)):

$$E[xr_{t+dt}^n|\kappa_t^a, \kappa_t^b, \Psi_t, \eta_t] = \alpha\bar{\sigma}_{nc}(\Psi_t, \eta_t) + \bar{\sigma}_{nc}(\Psi_t, \eta_t)\sum_{i=\{a,b\}}\omega^i(\eta_t)\frac{\hat{\mu}_{ct}^i - \mu_{ct}}{\sigma_c}, \tag{8}$$

where $\bar{\sigma}_{nc}(\Psi_t, \eta_t)$ is the covariance between bond excess returns and consumption shocks and μ_{ct} is the drift of c_t under the objective physical measure. Disagreement indirectly affects risk premiums through the covariance between bond prices and shocks to c (recall from (3) that the entire premium is $\alpha\sigma_{nc}$ in the single-agent economy), and it directly enters through the weight on the average of the individual agents' "expectation biases."

While this stylized model motivates our subsequent conditioning of $\mathcal{RA}$'s learning rule on measures of disagreement, our econometric analysis does not literally build upon this construction. Allowing for the diversity of investors in actual bond markets in this class of difference-of-opinion models would give rise to a very high-dimensional state vector that includes $(\eta^{(a,b)}, \Psi_t^{(a,b)})$ for many pairs of investors. This is counter to the extensive evidence that the covariance structure of yields is well described by a low-dimensional set of risk factors. Additionally, we endow $\mathcal{RA}$ with an arbitrage-free *DTSM*-based learning rule, instead of one derived from a consumption-based equilibrium model. Relative to reduced-form *DTSMs*, consumption-based models typically have large pricing errors. Our goal is to explore the impact of learning on bond-market risk premiums in a setting where bonds are priced accurately.

Investor disagreement is introduced in a parsimonious way by having $\mathcal{RA}$ condition her learning rule on a summary measure of the point-in-time cross-investor dispersion of beliefs about future bond yields. Atmaz and Basak (2017) provide a formal underpinning for using this "sufficient statistic" for the impact of multiple-investor disagreements on bond prices. They show that, in their economy with a *continuum* of investors who hold different beliefs about the future payoff on a common stock, equilibrium prices are driven by the cross-sectional *average* belief bias and *cross-investor dispersion* of beliefs.

Empirical counterparts to the constructs "consensus beliefs" and "investor disagreement" are constructed using the BCFF survey of forecasts of yields over the period from January, 1985 through December, 2015, with the start date determined by data availability. The

Horizon	0.25y	0.5y	0.75y	1y
$ID(\mathcal{P}_1)$	48.18%	59.93%	60.82%	58.46%
$ID(\mathcal{P}_2)$	19.63%	24.36%	25.09%	33.25%
$ID(\mathcal{P}_3)$	26.10%	31.36%	33.73%	38.98%
$ID(y^{2y})$	51.18%	58.84%	57.48%	55.11%
$ID(y^{7y})$	41.33%	52.85%	57.22%	56.84%

Table 1: R^2's from the projections of inter-quantile differences in BCFF forecasts onto $\mathcal{P}$, over forecast horizons of one through four quarters. The sample period is January, 1985 through December 2015.

survey forecasts are for 6-months U.S. Treasury bill yield and par yields on coupon bonds with maturities of 1, 2, 3, 5, 7, and 10 years. The survey is run each month, and is typically released at the beginning of the following month (usually the first business day), based on information collected over a two-day period (usually between the 20th and the 26th of the month). To facilitate comparisons of forecasts from our *DTSMs* with those by the BCFF professionals, we use the survey-implied forecasts of zero-coupon bond yields computed by Le and Singleton (2012).[13]

That the cross-sectional dispersion in beliefs about future bond yields does in fact have predictive power for future yields is documented by extending (5) to

$$xr^n_{t+h} = \alpha_n + \mathcal{B}_{n\mathcal{P}}\mathcal{P}_t + \mathcal{B}_{nH}H_t + \sigma_{nw}w_{t+h}, \tag{9}$$

where $h \geq 1$ is the forecast horizon and H_t measures investor disagreement about future yields.[14] We have omitted the median "consensus" beliefs of the BCFF professionals about future $\mathcal{P}$ from (9) (even though a weighted average belief appears in (8)) because, after controlling for $\mathcal{P}_t$, measured median beliefs have negligable predictive power for excess returns. Indeed, the first PC of the consensus BCFF yield forecasts[15] is largely spanned by the yield PC's $\mathcal{P}_t$: linear projections of the first PC of BCFF forecasts onto $\mathcal{P}_t$ gives R^2's of 99.5% for one-quarter ahead forecasts and over 98% for one-year forecasts.

The components of H_t are constructed as the inter-decile ranges of the professional yield

[13]Whereas forecasting zero-coupon yields in an affine *DTSM* is a linear forecasting problem (see below), par yields are nonlinear functions of zero-coupon yields. We avoid this complexity by interpolating the forecasts of par yields to obtain approximate forecasts of zero yields.

[14]Buraschi and Whelan (2016) and Andrade, Crump, Eusepi, and Moench (2014), among others, present evidence that disagreement about future output growth and inflation have predictive power for yields. Our focus on dispersion in forecasts of future yields is motivated by the evidence that the priced factors in bond markets are spanned by the yields themselves (the yield PCs have a low-dimensional factor structure). Notably, in Section 6 we show that $\mathcal{RA}$'s forecasts are more accurate when she conditions on disagreement about future yields than on disagreement about future output growth or inflation.

[15]For fixed horizon j, we construct the first PC of the median forecasts of y^n_{t+j} by the BCFF professionals across different maturities n.

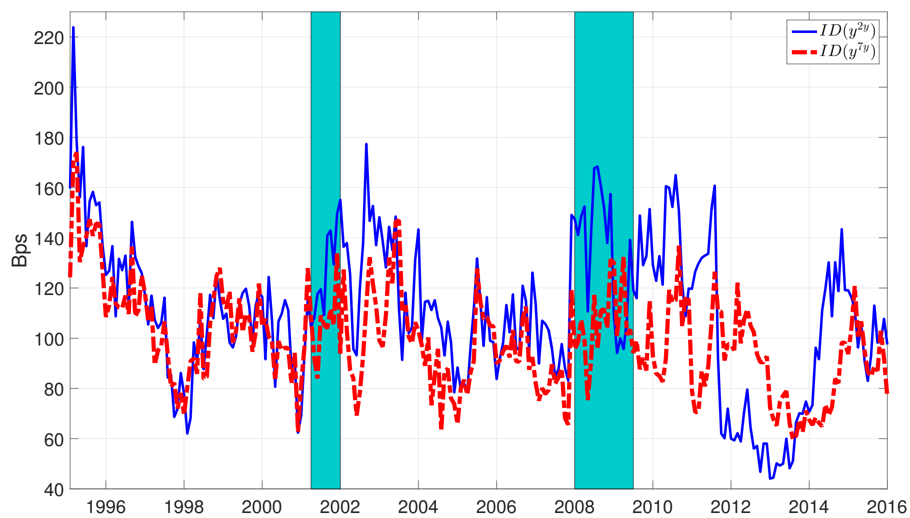

Figure 2: Historical measures of dispersions in professional forecasts one-year ahead for the two- and seven-year bond yields, $ID(y^{2y})$ and $ID(y^{7y})$.

forecasts from the BCFF survey, which begins in January 1985.[16] The differences between the ninetieth and tenth percentiles of the cross-sectional distribution of forecasts over horizon j are denoted by $ID_{jt}(y^m)$ for yields and $ID_{jt}(PCi)$ for PCs.[17] In Table 1 we report the R^2's from the projections of these dispersion measures onto the risk factors $\mathcal{P}_t$ for forecast horizons of one through four quarters. Over 50% of the dispersion in beliefs about the level of the yield curve, $ID(\mathcal{P}_1)$, is spanned by the first three PCs of Treasury yields. Thus, disagreement among the professional forecasters is effectively a priced risk in Treasury markets. There is also considerable leeway for H_t to have incremental predictive power in (9).

For our subsequent analysis of learning we set $H_t' = \left[ID_t(y^{2y}), ID_t(y^{7y})\right]$ with the horizon of the professional forecasts always set to one year (so we drop the subscript $j = 1y$). Figure 2 shows that disagreement is counter-cyclical as it tends to rise during and shortly after NBER recessions. Moreover, $ID(y^{2y})$ tends to be higher than $ID(y^{7y})$ and the gap between them $(ID(y^{2y}) - ID(y^{7y}))$ is relatively large following the two recessions in our sample. The years 2012-13 are exceptional for the persistently low level of $ID(y^{2y})$.

As background for our formal analysis of learning, we examine the one-quarter forecast

[16] Our analysis of learning is qualitatively robust to measuring dispersion in beliefs as the cross-sectional (point-in-time) volatility of professional forecasts (Patton and Timmerman (2010)) or the cross-sectional mean-absolute-deviation in forecasts (Buraschi and Whelan (2016)), and our measure is similar to that used by Andrade, Crump, Eusepi, and Moench (2014).

[17] In each month we check how many forecasters have published a forecast for the desired yield and predictive horizon. Out of the total 117 forecasters, we usually find approximately 45 forecasts.

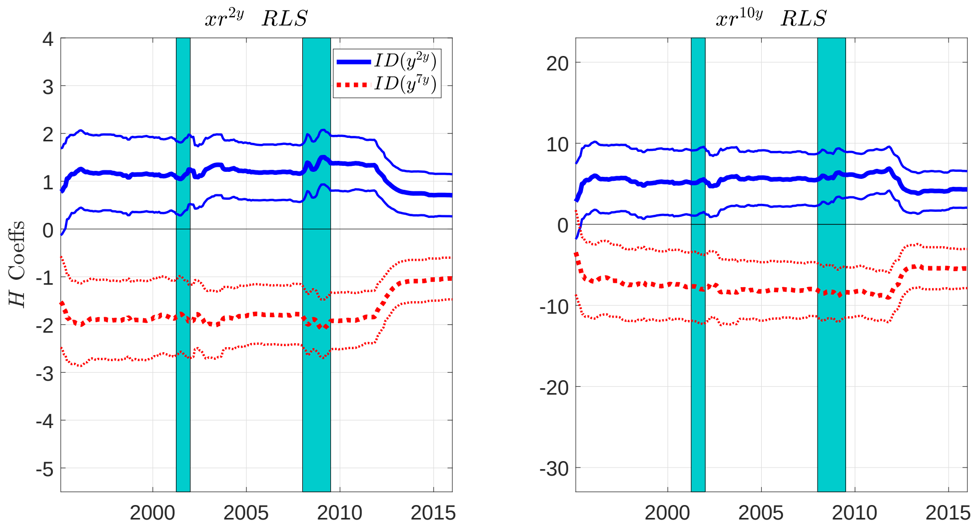

Figure 3: Evolution of the coefficients on $ID(y^{2y})$ (solid lines) and $ID(y^{7y})$ (dotted lines) based on recursive least squares estimates of (9) with 95% confidence bands. The left (right) panel is for the 2 (10) year bond. The sample is from January 1995 through December 2014.

horizon within the recursive least-squares learning setting from Section 2.1.[18] Figure 3 reports estimates of the $\mathcal{B}_{nH}$ for excess returns on bonds with maturities of 2 and 10 years. The coefficients on $ID(y^{2y})$ ($ID(y^{7y})$) are positive (negative) over the entire sample. In both cases the coefficients are statistically significant at conventional levels, supporting our premise that investor disagreement is incrementally informative about future yields.

3 Learning with a Dynamic Term Structure Model

The background in Section 2 informs the following specification of our learning problem. Consider the space of future risky payoffs generated by portfolios of U.S. Treasury bonds with weights that are functions of the history $Z_1^t \equiv (Z_t, Z_{t-1}, \ldots, Z_1)$ of a state vector Z_t.[19] In the absence of arbitrage opportunities and under weak regularity properties of the portfolio payoff space, there exists a stochastic discount factor (SDF) $\mathcal{M}_{t+1}^{\mathcal{B}}$, and an associated equivalent martingale measure $\mathbb{Q}$ that prices these Treasury portfolio payoffs (Dalang, Morton, and Willinger (1990)). $\mathcal{M}^{\mathcal{B}}$ can be generically represented in terms of the risk-neutral and historical

[18] Statistical issues related to the choice of horizon for the prediction regressions (9) are discussed in depth in Section 3.3 for our *DTSM*-based learning rules.

[19] For any variable X, the notation X_i^k, $k > i$, is short-hand for $(X_k, X_{k-1}, \ldots, X_i)$.

($\mathbb{P}$) conditional distributions of the priced risk factors $\mathcal{P}_t$ in Treasury markets:

$$\mathcal{M}^{\mathcal{B}}(\mathcal{P}_{t+1}, Z_1^t) = e^{-r_t} \times \frac{f^{\mathbb{Q}}(\mathcal{P}_{t+1}|\mathcal{P}_t)}{f^{\mathbb{P}}(\mathcal{P}_{t+1}|Z_1^t)}. \tag{10}$$

Bond prices are recovered by discounting Treasury coupons by the appropriate multiples of this *bond-market specific* $\mathcal{M}^{\mathcal{B}}(\mathcal{P}_{t+s+1}, Z_1^{t+s})$ under the objective measure $\mathbb{P}$.

We endow our "econometrician" $\mathcal{RA}$ with $\mathcal{M}^{\mathcal{B}}$ and assume that it is the pricing kernel she uses to compute risk premiums in Treasury markets *in real time*. As such, $\mathcal{M}^{\mathcal{B}}$ can be interpreted as the reduced-form SDF that our forward looking econometrician uses to learn about bond-yield dynamics. In addition, if there happened to be a marginal agent that priced bonds using $\mathcal{M}^{\mathcal{B}}$ (an active arbitrageur, for example) then we are also effectively modeling this agent's SDF under the presumption that her trading did not affect bond prices. For ease of reference, we subsequently refer to $\mathcal{M}^{\mathcal{B}}$ as $\mathcal{RA}$'s bond-market SDF.[20]

There is substantial evidence that bond yields follow a low-dimensional factor structure and, in fact, such structure underlies the pricing and risk management systems of primary dealers. Accordingly, we assume that $\mathcal{P}$ is comprised of the observed first three PCs of bond yields (there is no latent state that $\mathcal{RA}$ needs to infer in order to price Treasury bonds). In constructing $\mathcal{RA}$'s learning problem we first argue, on both conceptual and empirical grounds, that $\mathcal{RA}$ is likely to know the $\mathbb{Q}$ distribution of $\mathcal{P}$, the numerator of $\mathcal{M}^{\mathcal{B}}$. This, together with the fact that $\mathcal{RA}$ observes $\mathcal{P}_t$ at date t, implies that the central learning problem for $\mathcal{RA}$ in Treasury markets is about the data-generating process for $\mathcal{P}$. We subsequently specify a specific functional form for $f^{\mathbb{P}}(\mathcal{P}_t|Z_1^{t-1})$ in (10) that formalizes our assumptions about $\mathcal{RA}$'s learning rule underlying her perceptions of risk premiums.

3.1 Risk-Neutral Pricing of Bonds

Absent arbitrage opportunities, and under regularity, market participants can reverse engineer the risk-neutral distribution $\mathbb{Q}$ from the prices of traded bonds. This distribution will not in general be unique, unless agents live in a dynamically complete economy. Therefore, just as in prior studies of arbitrage-free *DTSMs*, our parametric specification of the $\mathbb{Q}$ distribution $f^{\mathbb{Q}}(\mathcal{P}_t|\mathcal{P}_{t-1})$ is presumed to represent an econometrically identified member of the family of admissible distributions.

Concretely, as in a $\mathbb{Q}$-affine Gaussian *DTSM*, we assume that $\mathcal{RA}$ believes the one-period

[20] The potential connections between $\mathcal{M}^{\mathcal{B}}$ and the SDF of a representative agent, should one exist, are discussed subsequently in Section 5.

riskless rate r_t follows the factor structure

$$r_t = \rho_0 + \rho_{\mathcal{P}} \mathcal{P}_t, \tag{11}$$

with $\mathcal{P}_t$ following the autonomous Gaussian $\mathbb{Q}$ process

$$\mathcal{P}_{t+1} = K^{\mathbb{Q}}_{0\mathcal{P}} + K^{\mathbb{Q}}_{\mathcal{P}\mathcal{P}} \mathcal{P}_t + \Sigma^{1/2}_{\mathcal{P}\mathcal{P}} e^{\mathbb{Q}}_{\mathcal{P},t+1}. \tag{12}$$

This leads her to set the price D^m_t of a zero-coupon bond issued at date t and maturing at date $t+m$ using the standard no-arbitrage formula

$$D^m_t = E^{\mathbb{Q}}_t \left[\prod_{u=0}^{m-1} \exp(-r_{t+u}) | \Theta^{\mathbb{Q}}, \mathcal{P}_t \right]. \tag{13}$$

For the econometric identification of $\Theta^{\mathbb{Q}}$, we follow JSZ and normalize ρ_0, $\rho_{\mathcal{P}}$, $K^{\mathbb{Q}}_{0\mathcal{P}}$, and $K^{\mathbb{Q}}_{\mathcal{P}\mathcal{P}}$ as known functions of $\Theta^{\mathbb{Q}} \equiv \left(k^{\mathbb{Q}}_{\infty}, \lambda^{\mathbb{Q}}, \Sigma_{\mathcal{P}\mathcal{P}}\right)$, with $k^{\mathbb{Q}}_{\infty}$ a scalar,[21] $\Sigma_{\mathcal{P}\mathcal{P}}$ the upper $K \times K$ block of Σ_Z, and $\lambda^{\mathbb{Q}}$ the K-vector of eigenvalues of $K^{\mathbb{Q}}_{\mathcal{P}\mathcal{P}}$ (see JSZ and Appendix B).

This JSZ normalization reveals clearly why it is reasonable to presume that all investors (including $\mathcal{RA}$) effectively know and agree on key elements of $\Theta^{\mathbb{Q}}$. Yields take the form

$$y^m_t = A_m(k^{\mathbb{Q}}_{\infty}, \lambda^{\mathbb{Q}}, \Sigma_{\mathcal{P}\mathcal{P}}) + B_m(\lambda^{\mathbb{Q}}) \mathcal{P}_t. \tag{14}$$

Therefore, at each date t, the $B_m(\lambda^{\mathbb{Q}})$ (and hence $\lambda^{\mathbb{Q}}$) can be estimated very precisely from the cross-section of Treasury yields;[22] agents will effectively know $\lambda^{\mathbb{Q}}$ from the date t cross-section y_t. The maturity-specific intercepts A_m depend also on $k^{\mathbb{Q}}_{\infty}$ and $\Sigma_{\mathcal{P}\mathcal{P}}$. However, the impact of $\Sigma_{\mathcal{P}\mathcal{P}}$ on A_m is through a convexity adjustment that is typically very small (see Appendix A). Therefore, knowledge of $\lambda^{\mathbb{Q}}$ and a tight prior on $k^{\mathbb{Q}}_{\infty}$ (also estimable from the cross-section of yields) imply that agents effectively know the A_m's as well.

Some insight into the empirical plausibility of the assumption that market participants know the parameters governing the drift of the $\mathbb{Q}$ distribution of $\mathcal{P}_t$ is revealed by the forecasts of individual BCFF professionals. If all of the BCFF professionals believe that yields follow (14), then the yield forecasts for horizon h ordered by deciles, $y^h_{t,o_1} < ... < y^h_{t,o_{10}}$, must satisfy

$$\hat{y}^{mh}_{t,o_k} = \bar{A}_m + \bar{B}_m \hat{\mathcal{P}}^h_{t,o_k} + e^{mh}_{t,o_k}, \tag{15}$$

[21] When $\mathcal{P}$ follows a stationary process under $\mathbb{Q}$, $k^{\mathbb{Q}}_{\infty}$ is proportional to the risk-neutral long-run mean of r. We adopt this more robust normalization, since the shape of the yield curve may call for the largest eigenvalue $\lambda^{\mathbb{Q}}_1$ to be very close to or even larger than unity. See JSZ for details.

[22] This is why the factor loadings B_m are reliably recovered from contemporaneous correlations among bond yields y^m_t and $\mathcal{P}_t$ (Duffee (2011)). It also explains why, holding $(K, N, Z, \mathcal{P})$ fixed, estimates of $\lambda^{\mathbb{Q}}$ in *DTSMs* without learning are nearly invariant across specifications of the $\mathbb{P}$ distribution of Z.

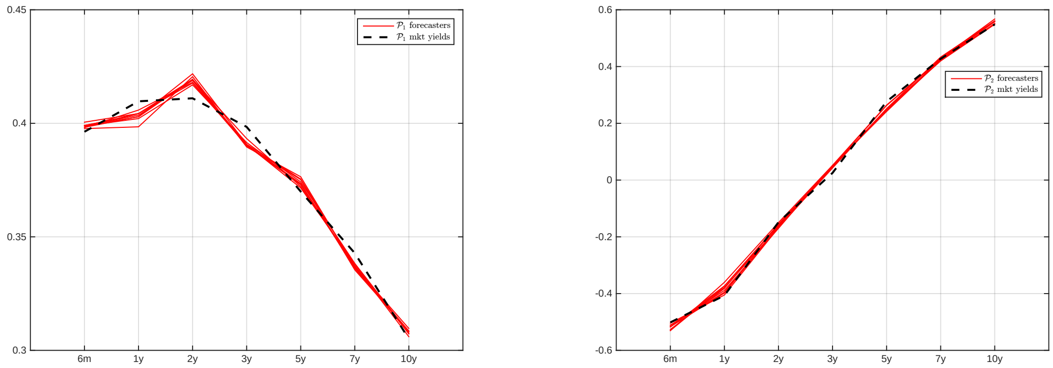

(a) Loadings on $PC1$

(b) Loadings on $PC2$

Figure 4: Time-series projections of the decile-ordered one-year ahead yields forecasts by BCFF forecasters onto their forecasts of $PC1$ and $PC2$.

where y^h_{t,o_1} is the forecast of the professional falling at the tenth percentile, y^h_{t,o_2} is the twentieth percentile, and so on up to the ninetieth percentile (we focus on order statistics, because the individual forecasters change over our sample). Holding h fixed, the loadings should be common across the ordered professionals. Figure 4 displays the full-sample estimates of these loadings by decile for $PC1$ and $PC2$ and h equal to one year (solid lines), along with the corresponding sample least-squares estimates of the loadings based on (14). The loadings are in fact remarkably similar across forecaster deciles, and the professionals' values correspond very closely to their sample counterparts. This is the case even though there are large differences in the forecasts of future yields across deciles (substantial disagreement).

Proceeding under the assumption that $\Theta^{\mathbb{Q}}$ is known to $\mathcal{RA}$, her SDF becomes

$$\mathcal{M}^{\mathcal{B}}(\Theta^{\mathbb{Q}}, \mathcal{P}_{t+1}, Z_1^t) = e^{-r_t} \times \frac{f^{\mathbb{Q}}(\mathcal{P}_{t+1}|\mathcal{P}_t; \Theta^{\mathbb{Q}})}{f^{\mathbb{P}}(\mathcal{P}_{t+1}|Z_1^t)}, \tag{16}$$

with the numerator derived from (12) evaluated at $\Theta^{\mathbb{Q}}$.

3.2 $\mathcal{RA}$'s Learning Problem

The learning problem faced by $\mathcal{RA}$ is that of inferring the parameters governing the state-dependence of $f^{\mathbb{P}}(\mathcal{P}_{t+1}|Z_1^t)$ from historical time-series data. We explore the pricing implications of the belief that $Z_t' = (\mathcal{P}_t', H_t')$ follows the historical process

$$\begin{bmatrix} \mathcal{P}_{t+1} \\ H_{t+1} \end{bmatrix} = \begin{bmatrix} K^{\mathbb{P}}_{\mathcal{P}0,t} \\ K^{\mathbb{P}}_{H0,t} \end{bmatrix} + \begin{bmatrix} K^{\mathbb{P}}_{\mathcal{P}\mathcal{P},t} & K^{\mathbb{P}}_{\mathcal{P}H,t} \\ K^{\mathbb{P}}_{H\mathcal{P},t} & K^{\mathbb{P}}_{HH,t} \end{bmatrix} \begin{bmatrix} \mathcal{P}_t \\ H_t \end{bmatrix} + \Sigma_Z^{1/2} \begin{bmatrix} e^{\mathbb{P}}_{\mathcal{P},t+1} \\ e^{\mathbb{P}}_{H,t+1} \end{bmatrix}, \tag{17}$$

where $\Theta_t^{\mathbb{P}}$, the vectorized $(K^{\mathbb{P}}_{0t}, K^{\mathbb{P}}_{Zt})$ governing the conditional mean of Z_t, is time varying and unknown. The shocks $e^{\mathbb{P}}_{\mathcal{P},t+1}$ and $e^{\mathbb{P}}_{H,t+1}$ are jointly Gaussian. The portfolios $\mathcal{P}_t$ are assumed to be priced perfectly by (13). The learning environment is completed with the following pricing equation for $\mathcal{O}_t$, the higher-order fourth through seventh PCs:

$$\mathcal{O}_t = A_{\mathcal{O}}\left(\Theta^{\mathbb{Q}}\right) + B_{\mathcal{O}}\left(\Theta^{\mathbb{Q}}\right)\mathcal{P}_t + \varepsilon_{\mathcal{O},t}, \tag{18}$$

where $(\mathcal{P}_t, \mathcal{O}_t)$ fully spans y_t. The errors $\varepsilon_{\mathcal{O},t}$ are assumed to be *iid* $Normal(0, \Sigma_{\mathcal{O}})$, with $\Sigma_{\mathcal{O}}$ diagonal (consistent with its sample counterpart from a regression of $\mathcal{O}_t$ on $\mathcal{P}_t$).

Though the discussion in Section 3.1 suggests that λ^Q can be reliably assumed to be known, trading desks recalibrate their yield curve models to new data on a regular basis. This recalibration of the hedge ratios $B_m(\lambda^{\mathbb{Q}})$ is perhaps premised on the view that (14) is an approximation to the true bond pricing function. With this possibility in mind, for our subsequent econometric analysis, we allow $\mathcal{RA}$ to update $\Theta^{\mathbb{Q}}$ monthly as new data becomes available, using the model-implied likelihood. We assume that $\mathcal{RA}$ updates $\lambda^{\mathbb{Q}}$ each month using the likelihood function (22), even though at each point in time she prices bonds treating $\Theta^{\mathbb{Q}}$ as known and fixed. In our setting, this is analogous to the widely adopted approximation to Bayesian learning of endowing agents with "anticipated utility" (see, e.g., Kreps (1998) and Cogley and Sargent (2008)). Strikingly, even with this flexibility, $\mathcal{RA}$ treats $(\lambda^{\mathbb{Q}}, k^{\mathbb{Q}}_\infty)$ as known and *virtually fixed* over our entire sample period (see Section 4).

Additionally, we assume that parameter variation per se is not a source of priced risk in bond markets.[23] Within this learning environment $\mathcal{RA}$'s SDF $\mathcal{M}^B$ takes the form:

$$\mathcal{M}(\Theta^{\mathbb{Q}}, \mathcal{P}_{t+1}, Z_1^t) = e^{\{-r_t - \frac{1}{2}\log|\Gamma_t| - \frac{1}{2}\hat{\Lambda}'_{\mathcal{P}t}\Gamma_t^{-1}\hat{\Lambda}_{\mathcal{P}t} - \hat{\Lambda}'_{\mathcal{P}t}\Gamma_t^{-1}\varepsilon^{\mathbb{P}}_{t+1} + \frac{1}{2}(\varepsilon^{\mathbb{P}}_{t+1})'(I - \Gamma_t^{-1})\varepsilon^{\mathbb{P}}_{t+1}\}}, \tag{19}$$

$$\Gamma_t = \Omega^{-1/2}_{\mathcal{P}\mathcal{P},t}\Sigma_{\mathcal{P}\mathcal{P}}(\Omega^{-1/2}_{\mathcal{P}\mathcal{P},t})',$$

$$\Omega^{1/2}_{\mathcal{P}\mathcal{P},t}\hat{\Lambda}_{\mathcal{P}t} = \hat{\Lambda}_{0t}(\Theta^{\mathbb{Q}}, \hat{\Theta}^{\mathbb{P}}_t) + \hat{\Lambda}_{1t}(\Theta^{\mathbb{Q}}, \hat{\Theta}^{\mathbb{P}}_t)Z_t, \tag{20}$$

where the market price of risk $\hat{\Lambda}_{\mathcal{P}t}$ depends on the posterior mean $\hat{\Theta}^{\mathbb{P}}_t$ and, therefore, implicitly

[23] The absence of compensation for parameter risk is fairly standard in the literature on pricing with Bayesian learning, and it greatly simplifies what is already a challenging modeling problem. Collin-Dufresne, Johannes, and Lochstoer (2016) explore implications of priced parameter uncertainty in a single-agent, consumption based setting with a much lower dimensional state space than what is required to reliably price bonds. Accommodating priced parameter uncertainty within a higher dimensional $DTSM$ is an interesting topic for further research.

on the entire history Z_1^t (Appendix D). The form of $\Lambda_{\mathcal{P}t}$ is familiar from Duffee (2002)'s model without learning, but importantly here the weights are state-dependent owing to learning.

Given $\Theta^{\mathbb{Q}}$ and the specification in (17), the physical distribution of Z_t is determined by the evolution of the unknown and drifting parameters $\Theta_t^{\mathbb{P}}$. Below we show that if innovations to $\Theta_t^{\mathbb{P}}$ are normally distributed, then the posterior density $f(\Theta_{t+1}^{\mathbb{P}}|Z_1^t, \mathcal{O}_1^t)$ also follows a normal distribution. Furthermore, the conditional $\mathbb{P}$-distribution of Z_{t+1} given Z_1^t is:

$$f^{\mathbb{P}}(Z_{t+1}|Z_1^t) = Normal\left(\hat{K}_{0t}^{\mathbb{P}} + \hat{K}_{Zt}^{\mathbb{P}} Z_t, \Omega_t\right), \tag{21}$$

where $(\hat{K}_{0t}^{\mathbb{P}}, \hat{K}_{Zt}^{\mathbb{P}})$ denotes $\mathcal{RA}$'s posterior mean of $\Theta_t^{\mathbb{P}}$ and her one-period-ahead forecast covariance matrix Ω_t depends on Σ_Z and the uncertainty regarding $\Theta_t^{\mathbb{P}}$. This $\mathbb{P}$-process for Z_t is nonlinear (in particular, non-affine) and we compute $\mathcal{RA}$'s optimal forecasts accordingly.

At date t a Bayesian $\mathcal{RA}$, faced with new observations $(Z_t, \mathcal{O}_t)$ and the past history $(Z_1^{t-1}, \mathcal{O}_1^{t-1})$, evaluates an (approximate) likelihood function by integrating out the uncertainty about $\Theta_t^{\mathbb{P}}$ using her posterior distribution. Thus, with $(\Theta^{\mathbb{Q}}, \Sigma_{\mathcal{O}})$ known,

$$\begin{aligned} f(Z_1^t, \mathcal{O}_1^t) \;\; = \;\; & \prod_{s=1}^{t} f(\mathcal{O}_s|Z_1^s, \mathcal{O}_1^{s-1}; \Theta^{\mathbb{Q}}, \Sigma_{\mathcal{O}}) \times \\ & \int f(Z_s|Z_1^{s-1}, \mathcal{O}_1^{s-1}, \Theta_{s-1}^{\mathbb{P}}; \Sigma_Z) f(\Theta_{s-1}^{\mathbb{P}}|Z_1^{s-1}, \mathcal{O}_1^{s-1}) d(\Theta_{s-1}^{\mathbb{P}}). \end{aligned} \tag{22}$$

To allow for constraints on the market prices of risk, we partition $\Theta_t^{\mathbb{P}}$ as $(\psi^r, \psi_t^{\mathbb{P}})$, where $\psi_t^{\mathbb{P}}$ is the vectorized set of free parameters and ψ^r is the vectorized set of parameters that are fixed conditional on $\Theta^{\mathbb{Q}}$. Letting ι_r and ι_f denote the matrices that select the columns of $(I \otimes [1, Z_{t-1}'])$ corresponding to the restricted and free parameters, and collecting the known terms in (17) into $\mathcal{Y}_t = Z_t - \left(I \otimes [1, Z_{t-1}']\right) \iota_r \psi^r$, we rewrite the state equation as

$$\mathcal{Y}_{t+1} = \mathcal{X}_t \psi_t^{\mathbb{P}} + \Sigma_Z^{1/2} e_{Z,t+1}^{\mathbb{P}}, \tag{23}$$

where $\mathcal{X}_t = (I \otimes [1, Z_t'])\, \iota_f$.

To accommodate the possibility of permanent structural change in the underlying economic environment, we assume that $\psi_t^{\mathbb{P}}$ evolves according to

$$\psi_t^{\mathbb{P}} = \psi_{t-1}^{\mathbb{P}} + Q_{t-1}^{1/2} \eta_t, \qquad \eta_t \overset{iid}{\sim} Normal(0, I), \tag{24}$$

where Q_{t-1} denotes the (possibly) time-varying covariance matrix of η_t, with η_t independent of all past and future $e_{Zt}^{\mathbb{P}}$. $\mathcal{RA}$ knows that $\psi_t^{\mathbb{P}}$ follows (24), but she does not observe the realized $\psi_t^{\mathbb{P}}$. Her Bayesian learning rule filters for $\psi_t^{\mathbb{P}}$ conditional on $(\Theta^{\mathbb{Q}}, \psi^r)$. Adopting a Gaussian prior on $\psi_0^{\mathbb{P}}$ leads to a posterior distribution for $\psi_t^{\mathbb{P}}$ that is also Gaussian, $\psi_t^{\mathbb{P}}|Z_1^t \sim Normal(\hat{\psi}_t^{\mathbb{P}}, P_t)$. In Appendix C we show that her posterior mean follows the recursion

$$\hat{\psi}_t^{\mathbb{P}} = \hat{\psi}_{t-1}^{\mathbb{P}} + R_t^{-1}\mathcal{X}_{t-1}'\Sigma_Z^{-1}(\mathcal{Y}_t - \mathcal{X}_{t-1}\hat{\psi}_{t-1}^{\mathbb{P}}), \tag{25}$$

which depends on the posterior variance P_t through $R_t^{-1} \equiv P_t - Q_t$, with R_t satisfying

$$R_t = \left(I - P_{t-1}^{-1}Q_{t-2}\right)R_{t-1} + \mathcal{X}_{t-1}'\Sigma_Z^{-1}\mathcal{X}_{t-1}. \tag{26}$$

This rule has a revealing interpretation within the class of *adaptive least-squares estimators* (*ALS*) of $\psi_t^{\mathbb{P}}$. We say that $\widehat{\psi}_t^{\mathbb{P}}$ is an *ALS* estimator if there exists a sequence of scalars $\gamma_t > 0$ such that $\widehat{\psi}_t^{\mathbb{P}}$ can be expressed recursively as

$$\begin{aligned} \widehat{\psi}_t^{\mathbb{P}} &= \widehat{\psi}_{t-1}^{\mathbb{P}} + R_t^{-1}\mathcal{X}_{t-1}'\Sigma_Z^{-1}(\mathcal{Y}_t - \mathcal{X}_{t-1}\widehat{\psi}_{t-1}^{\mathbb{P}}), & (27)\\ R_t &= \gamma_{t-1}R_{t-1} + \mathcal{X}_{t-1}'\Sigma_Z^{-1}\mathcal{X}_{t-1}. & (28) \end{aligned}$$

It follows immediately from (25) - (26) that the posterior mean in the Kalman filter used by $\mathcal{RA}$ to update $\psi_t^{\mathbb{P}}$ can be represented as a generalized *ALS* estimator. Moreover, (26) reveals three special cases where the filtering underlying Bayesian learning reduces to an actual *ALS* estimator (that is, (26) reduces to (28)):[24]

$\mathcal{B}\downarrow$ALS: Setting $P_{t-1}^{-1}Q_{t-2} = (1-\delta_{t-1}) \cdot I$[25] for some sequence of scalars $0 < \delta_t \leq 1$, $\hat{\psi}_t$ becomes an *ALS* estimator of $\psi^{\mathbb{P}}$ with $\gamma_t = \delta_t$.

$\mathcal{B}\downarrow$CGLS: Specializing further by setting $\delta_t = \delta$ to a constant leads to $\hat{\psi}_t$ being a *constant gain least-squares (CGLS)* estimator of $\psi^{\mathbb{P}}$ with $\gamma = \delta$.

$\mathcal{B}\downarrow$RLS: If the constant $\delta = 1$, then $\hat{\psi}_t$ is the *recursive least-squares (RLS)* estimator of $\psi^{\mathbb{P}}$.

Among the insights that emerge from this construction is that a Bayesian agent whose learning rule specializes to the *RLS* estimator is not adaptive in the following potentially important sense. With $\gamma = 1$ we have $Q_t = 0$, so an agent following a *RLS* rule is learning about an unknown value of $\psi^{\mathbb{P}}$ that is presumed to be fixed over time. Consequently, sudden changes in market conditions that result in sharp movements in recent values of Z may have an imperceptible effect on $\widehat{\psi}_t^{\mathbb{P}}$ as updated by $\mathcal{RA}$. Indeed, in environments where the *ML* estimator converges to a constant for large T, an *RLS*-based $\mathcal{RA}$ will be virtually non-adaptive on $\widehat{\psi}^{\mathbb{P}}$ to new information after a long training period.

[24]See McCulloch (2007), and the references therein, for discussions of similar issues in a setting of univariate y_t and econometrically exogenous x_t.

[25]This condition can be obtained by recursively setting $Q_{t-1} = \frac{1}{\delta_t}(P_{t-1} - P_{t-1}x_{t-1}'\Omega_{t-1}^{-1}x_{t-1}P_{t-1})$.

A more adaptive rule that responds to changes in the structure of the economy (owing say to changes in government policies) is obtained by giving less weight to values of Z far in the past. Such down-weighting arises naturally when $\mathcal{RA}$'s learning specializes to Case $\mathcal{B}\!\downarrow$**CGLS**. The constant-gain coefficient γ determines the "half-life" of the weight on past data. This follows from the observation that, conditional on $\Theta^{\mathbb{Q}}$, the first-order conditions to the likelihood function implied by Bayesian learning with *CGLS* updating (Appendix C) are identical to those of a likelihood with terms of the form $\gamma^t \epsilon_{Zt}^{\mathbb{P}\prime} \Sigma_Z^{-1} \epsilon_{Zt}^{\mathbb{P}}$.[26]

Expressions (23) and (24) are the measurement and transition equations in a Gaussian linear filter over the unknown parameters. Therefore, the distribution of $\mathcal{Y}_{t+1}$ conditional on Z_1^t is distributed $f^{\mathbb{P}}(\mathcal{Y}_{t+1}|Z_1^t) = Normal\left(\mathcal{X}_t \hat{\psi}_t^{\mathbb{P}}, \Omega_t\right)$, with the one-step ahead forecast variance determined inductively by $\Omega_t = \mathcal{X}_t P_t \mathcal{X}_t' + \Sigma_Z$. The term $\mathcal{X}_t P_t \mathcal{X}_t'$ captures the uncertainty related to the unknown $\Theta^{\mathbb{P}}$, while the second term is the innovation variance of the state Z_t.

Throughout this construction the direct dependence of Ω_t on Σ_Z is a consequence of $\mathcal{RA}$ treating Σ_Z as known, not as an object to be learned. This is an admittedly strong assumption as, empirically, we will see that $\mathcal{RA}$'s learning rule shows sizable revisions in Σ_Z. Though revisions in $\Sigma_{\mathcal{PP}}$ through learning would be largely inconsequential for pricing (convexity effects are small), they could be material for how $\mathcal{RA}$ updates beliefs about $\Theta^{\mathbb{P}}$. In Appendix H we show that our core findings are robust to the introduction of learning about the structure of the conditional covariances of Z in a model with time-varying second moments, which is described in detail in Appendix G.

3.3 Empirical Learning Rules

As a benchmark case, we estimate a three-factor *DTSM* in which $\mathcal{RA}$ follows a constant-gain learning rule (*CGLS*) with conditioning on past information on $\mathcal{P}$ alone. We call this framework rule $\ell_{CG}(\mathcal{P})$. The parameters of the pricing distribution are normalized as in JSZ and, owing to learning, the coefficients on $\mathcal{P}_t$ in (20) that determine the market prices of risk $\hat{\Lambda}$ depend on the entire history $\mathcal{P}_1^t$. We also consider rule $\ell(\mathcal{P})$, corresponding to the special case of *RLS* learning, with $\gamma = 1$.

In Section 4 we explore the properties of rule $\ell(\mathcal{P})$ initialized using ML estimates for the "training" period June 1961 through January 1972. Then every month, up through December 2014, as new data becomes available, $\mathcal{RA}$ updates her posterior and the associated forecasts of future $\mathcal{P}$. Moving through the sample, *DTSM*-based rules impose the JSZ normalizations based on current-month information about yields and updated weights determining the first three *PC*s from the sample covariance matrix of yields. This is the longest sample period we

[26]The latter is the likelihood function of a naive learner who simply re-estimates the likelihood function of a fixed-parameter model every period using the latest data and with down weighting by γ^t.

use in our analysis. Thus, we highlight the rule estimated over this period with the superscript L for "L"ong sample, $\ell^L(\mathcal{P})$.

The *DTSM*-based rules are quite highly parametrized. For parsimony, which is relevant for the subsequent out-of-sample assessments, the parameters governing the market prices of the risks $\mathcal{P}$ are set to zero if their p-value *during the training period* is larger than 0.4. Since $K^{\mathbb{Q}}_{\mathcal{P}\mathcal{P}}$ is presumed known by $\mathcal{RA}$, these constraints on $\Lambda_{\mathcal{P}}$ effectively transfer *a priori* knowledge of $\lambda^{\mathbb{Q}}$ to (some) knowledge about $K^{\mathbb{P}}_{Z}$. All constraints on $\Lambda_{\mathcal{P}}$ selected during the training period are maintained throughout the remainder of the sample period.[27]

For the case of constant-gain learning ($\gamma < 1$), we set $\gamma = 0.99$. Appendix E offers two complementary perspectives on this choice. First, if we allow $\mathcal{RA}$ to adjust γ over time, then she selects fitted γ_t's that remain quite close to 0.99. Second, her choice is evidently ex post optimal. In fact, searching over fixed γ's to minimize $\mathcal{RA}$'s out-of-sample forecast accuracy leads to a value of γ that is approximately 0.99 over the period January 1995 through December 2014.

From the fitted *DTSM* at date t, an h-period ahead forecast of Z is given by

$$\hat{Z}_{t+h} = \widehat{K}^{\mathbb{P}}_{0t} + \left(\widehat{K}^{\mathbb{P}}_{Zt}\right)\widehat{K}^{\mathbb{P}}_{0t} + ... + \left(\widehat{K}^{\mathbb{P}}_{Zt}\right)^{h-1}\widehat{K}^{\mathbb{P}}_{0t} + \left(\widehat{K}^{\mathbb{P}}_{Zt}\right)^{h} Z_t. \tag{29}$$

This leads directly to the h-period ahead forecasts of yields:

$$\hat{y}^m_{t+h} = A_m\left(\widetilde{K}^{\mathbb{Q}}_0, \widetilde{K}^{\mathbb{Q}}_{\mathcal{P}\mathcal{P}}, \widetilde{\Sigma}_{\mathcal{P}\mathcal{P}}\right) + B_m\left(\widetilde{K}^{\mathbb{Q}}_{\mathcal{P}\mathcal{P}}\right)\hat{\mathcal{P}}_{t+h}, \tag{30}$$

where the tildes indicate maximum likelihood estimators as of the forecast date.

In the literatures on forecasting with Gaussian *DTSMs* and vector autoregressions the choice of the horizon h has not been without controversy, especially as h extends out a year or longer, owing to potential small-sample biases.[28] We emphasize that our assessments of forecast accuracy are based on *out-of-sample* fit. Moreover, $\mathcal{RA}$'s *nonlinear* forecasting rules have the flexibility to uncover much richer forms of predictive power of Z_t for future excess returns than in standard affine *DTSMs*. With adaptive learning about $\widehat{K}^{\mathbb{P}}_{Zt}$, $\mathcal{RA}$ may change her weights on components of Z_t so that auxiliary (non-bond market) information in Z_t shows

[27]We wondered whether adjusting the constraints in real time would improve out-of-sample forecasts. Interestingly, for the rules we examine, such real-time updating leads to a *deterioration* in the quality of forecasts, by a substantial degree. We found that this was true for a variety of training periods. Evidently, real-time adjustments induce a form of over-fitting that compromises forecast accuracy.

[28]Stambaugh (1999) shows that the lack of strict econometric exogeneity in predictive regressions can lead to significant small-sample biases in estimated coefficients, especially with highly persistent variables. More recently, Bauer and Hamilton (2018) argue that there is a tendency for an upward bias in estimated R^2's in excess return regressions in studies of bond-market returns owing to a "standard error bias," and this bias is potentially amplified when studying long-horizon forecasts using overlapping data.

	const	$PC1$	$PC2$	$PC3$
$PC1$	0.0031	-0.0331	-0.1130	$\mathcal{C}(0)$
p-value	0.0071	0.0398	0.0042	-
$PC2$	-0.00004	0.0217	$\mathcal{C}(0)$	-0.5040
p-value	0.2440	0.0786	-	0.0205
$PC3$	0.0005	$\mathcal{C}(0)$	$\mathcal{C}(0)$	-0.3043
p-value	0.0002	-	-	0.0000

Table 2: Estimated parameters of the market prices of risk (20), along with their probability values, for rule $\ell^L_{CG}(\mathcal{P})$ trained from June 1961 through January 1972. $\mathcal{C}(0)$ are entries constrained to zero.

substantial forecasting power under some economic/market conditions, and virtually none in other periods.

In evaluating $\mathcal{RA}$'s out-of-sample forecasting accuracy we reference two alternative rules. The first, $\ell(BCFF)$, is the forecast rule implicitly used by the median professional forecaster as surveyed by the Blue Chip Financial Forecasts. We also examine the simple yield-based rule that has each zero yield following a random walk, rule $\ell(RW)$. BCFF forecasts are averages over calendar quarters and cover horizons out to five quarters ahead. For example, in January, 1999, the two-quarter ahead forecast for a specific variable will be equal to its average value between April and June. For comparability across all forecast rules, we compute similar quarterly averages for rules $\ell^L(\mathcal{P})$ and $\ell(RW)$.

4 Learning from Information in the Yield Curve: Rule $\ell^L_{CG}(\mathcal{P})$

The estimated parameters of the market prices of risk (MPR) for rule $\ell^L_{CG}(\mathcal{P})$ are displayed in Table 2, along with their probability values, with $\mathcal{C}(0)$ denoting the parameters set to zero in $\mathcal{RA}$'s learning rule. The "level" of the yield curve ($PC1$) is a major driver of the MPR's for $PC1$ and $PC2$, but not $PC3$. For the chosen training period, $PC2$ does not have a significant economic impact on its own MPR, though it does impact the MPR of $PC1$.

The real-time estimates of $\lambda^{\mathbb{Q}}$ from this learning scheme are displayed in Figure 5 (the patterns for $\ell^L(\mathcal{P})$ are nearly identical). $\mathcal{RA}$ holds $\lambda^{\mathbb{Q}}$ virtually fixed over the entire sample, consistent with the premise of her Bayesian rule for learning about $\Theta^{\mathbb{P}}_t$. There is some drift in the second eigenvalue $\lambda^{\mathbb{Q}}_2$. However, repeating our learning exercise with the full vector $\lambda^{\mathbb{Q}}$ fixed from the initial training period onward has a very small effect on the quantitative properties of the rule-implied prices or forecasts.[29]

[29]Reassurance that this near constant $\lambda^{\mathbb{Q}}$ is not a mechanical implication of learning about $f^{\mathbb{P}}(\mathcal{P}_{t+1}|Z^t_1)$

Figure 5: Estimates from model $\ell^L_{CG}(\mathcal{P})$ (with constraints on the market prices of risk determined in December 1971) of the eigenvalues $\lambda^{\mathbb{Q}}$ ($\lambda^{\mathbb{P}}$) of the feedback matrix $K_1^{\mathbb{Q}}$ ($K_1^{\mathbb{P}}$) governing the persistence in $\mathcal{P}$. The estimates at date t are based on the historical data up to observation t, over the period January 1972 through December 2014.

Pursuing this insight, if $\lambda^{\mathbb{Q}}$ is known and fixed over time, then so are the loadings B_m on $\mathcal{P}$ in the affine pricing expression (14). Combining this with the fact that $\mathcal{P}_t$ is measured with negligible error, the state-dependent components of bond yields that emerge from (14) with learning take the same form $B_m\left(\lambda^{\mathbb{Q}}\right)\mathcal{P}_t$, just as in a *DTSM* without learning. Furthermore, agents will use fixed "hedge ratios" over time to manage the risks of their bond portfolios.

This finding is especially striking in relation to how $\mathcal{RA}$ updates the historical eigenvalues $\lambda^{\mathbb{P}}$. These adjustments are relatively much larger for all three low-order *PC*s. This implies that $\mathcal{RA}$'s views about the objective feedback matrix $K^{\mathbb{P}}_{Z,t}$, and thereby risk premiums (see

under no-arbitrage restrictions is provided by running reduced-form, expanding-window regressions of the principal components on individual bond yields:

$$y^n_t = a^n_t + B^n_t\mathcal{P}_t + u^n_t.$$

The loadings B^n_t remain quite stable for yields across the maturity spectrum. Similar results are obtained using constant gains least squares with $\gamma = 0.99$. The estimates of the weights that define $\mathcal{P}$ are also stable over time.

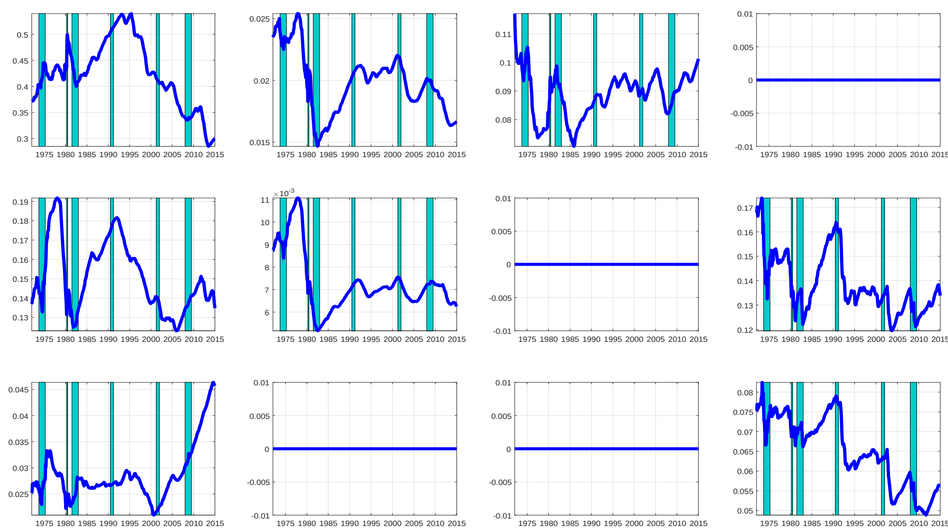

Figure 6: Posterior standard deviations of the elements of $[K_{0t}^{\mathbb{P}}, K_{1t}^{\mathbb{P}}]$ implied by the constant-gain $\ell_{CG}^{L}(\mathcal{P})$ learning rules.

below), are changing substantially over time. The largest changes in the eigenvalues $\lambda_t^{\mathbb{P}}$ of $K_Z^{\mathbb{P}}$ (Figure 5) occur during the Fed experiment in the early 1980's (when $\lambda^{\mathbb{Q}}$ remained remarkably stable). Even outside of this turbulent period there is substantial drift in $\lambda_{1t}^{\mathbb{P}}$ and $\lambda_{2t}^{\mathbb{P}}$. The estimate of $\lambda_t^{\mathbb{P}}$ for $\ell_{CG}^{L}(\mathcal{P})$ shows more temporal variation than the estimate for $\ell^{L}(\mathcal{P})$ owing to the former's down-weighting of historical data.

An interesting perspective on $\mathcal{RA}$'s confidence in forecasts based on $\mathcal{M}^{\mathcal{B}}$ comes from examining the standard deviations of her posterior distributions of the non-zero entries in $[K_{0t}^{\mathbb{P}}, K_{\mathcal{PP},t}^{\mathbb{P}}]$ implied by rules $\ell_{CG}^{L}(\mathcal{P})$ and $\ell^{L}(\mathcal{P})$. The latter rule presumes– quite likely falsely– that $\Theta^{\mathbb{P}}$ is constant, so an *RLS*-based $\mathcal{RA}$ becomes increasingly confident about its value about over time. In contrast, we can see in Figure 6 that under $\ell_{CG}^{L}(\mathcal{P})$, confidence in the fitted components of $K_{\mathcal{PP},t}^{\mathbb{P}}$ associated with $(PC1, PC2)$ tends to improve during the collapse of inflation in the late 1970s but then remains flat or slightly decline over the remainder of our sample period. This is a reflection of perceived risks regarding structural changes.

The relative accuracies of the rule-based forecasts, which depend primarily on $\mathcal{RA}$'s choice of $\widehat{\Theta}_t^{\mathbb{P}}$, are summarized by the root-mean-squared errors (RMSE's) displayed in Table 3 for both quarterly and annual horizons. Below each RMSE are Diebold and Mariano (1995) (D-M) statistics for assessing whether two RMSE's are statistically the same, calculated as

RMSE's (in basis points) for Quarterly Horizon

Rule	6m	1Y	2Y	3Y	5Y	7Y	10Y
$\ell(RW)$	30.7 (−3.81) [−]	32.6 (−2.98) [−]	35.2 (−3.98) [−]	36.1 (−5.06) [−]	36.4 (−4.89) [−]	36.0 (−3.96) [−]	32.5 (−3.01) [−]
$\ell(BCFF)$	38.6 (−) [3.81]	37.6 (−) [2.98]	43.8 (−) [3.98]	50.7 (−) [5.06]	44.6 (−) [4.89]	44.1 (−) [3.96]	40.0 (−) [3.01]
$\ell^L(\mathcal{P})$	29.7 (−3.35) [−0.82]	31.3 (−2.61) [−0.86]	36.6 (−3.31) [3.05]	37.01 (−4.97) [2.05]	37.8 (−4.28) [1.77]	37.03 (−3.62) [1.86]	34.1 (−2.59) [2.50]
$\ell^L_{CG}(\mathcal{P})$	31.1 (−3.52) [0.44]	32.0 (−2.67) [−0.49]	37.5 (−3.54) [2.74]	37.7 (−5.30) [1.90]	38.0 (−4.21) [1.63]	36.5 (−4.02) [0.72]	33.8 (−2.81) [1.37]

RMSE's (in basis points) for Annual Horizon

Rule	6m	1Y	2Y	3Y	5Y	7Y	10Y
$\ell(RW)$	118.8 (−1.00) [−]	115.3 (−0.83) [−]	103.3 (−1.90) [−]	94.1 (−2.65) [−]	84.9 (−2.82) [−]	78.8 (−2.75) [−]	70.8 (−2.70) [−]
$\ell(BCFF)$	128.8 (−) [1.00]	123.9 (−) [0.83]	122.1 (−) [1.90]	122.5 (−) [2.65]	105.9 (−) [2.82]	100.6 (−) [2.75]	88.1 (−) [2.70]
$\ell^L(\mathcal{P})$	115.5 (−1.22) [−0.58]	112.5 (−1.19) [−0.56]	107.4 (−1.77) [1.41]	100.0 (−2.45) [2.10]	91.8 (−2.10) [2.12]	84.7 (−1.99) [1.84]	78.1 (−1.31) [2.05]
$\ell^L_{CG}(\mathcal{P})$	125.0 (−0.64) [0.90]	122.2 (−0.25) [0.96]	117.8 (−0.67) [1.85]	109.0 (−1.59) [2.00]	96.3 (−1.20) [1.59]	85.5 (−1.59) [1.02]	76.9 (−1.16) [0.91]

Table 3: RMSE's for forecasts from January, 1995 to December, 2014. The D-M statistics for the differences between the *DTSM*- and BCFF-implied (*DTSM*- and *RW*-implied) forecasts are given in parentheses (brackets).

extended by Harvey, Leybourne, and Newbold (1997).[30] The first (in parentheses) tests against $\ell(BCFF)$, and the second (in brackets) tests against $\ell(RW)$. Both *DTSM*-based rules are statistically more accurate than the median professional $\ell(BCFF)$, while having comparable accuracy to $\ell(RW)$. The outperformance of $(\ell^L(\mathcal{P}), \ell^L_{CG}(\mathcal{P}))$ over $\ell(BCFF)$ is seen across the entire maturity spectrum, which is notable given that $\mathcal{RA}$ is conditioning only

[30]Consider two sequences of forecast errors e_{1t} and e_{2t}, $t = \{1, 2, ..., T\}$, define $d_t \equiv e_{1t}^2 - e_{2t}^2$, and let

$$\hat{\mu}_d = \frac{1}{T}\sum_{t=1}^{T} d_t \quad \text{and} \quad \hat{V}_d = \sum_{t=1}^{T} (d_t - \hat{\mu}_d)^2 + 2\sum_{j=1}^{h} k(j/h) \sum_{t=1}^{T-j} (d_t - \hat{\mu}_d)(d_{t+j} - \hat{\mu}_d),$$

where $k(.)$ is a Bartlett kernel that down-weights past lags to ensure that the variance of the difference in mean squared errors stays positive. The number of lags h is set to three for the one-quarter ahead forecasts and to twelve for the four-quarters ahead forecasts. Then the D-M statistic is equal to $\sqrt{T}\hat{\mu}_d/\hat{V}_d^{1/2}$.

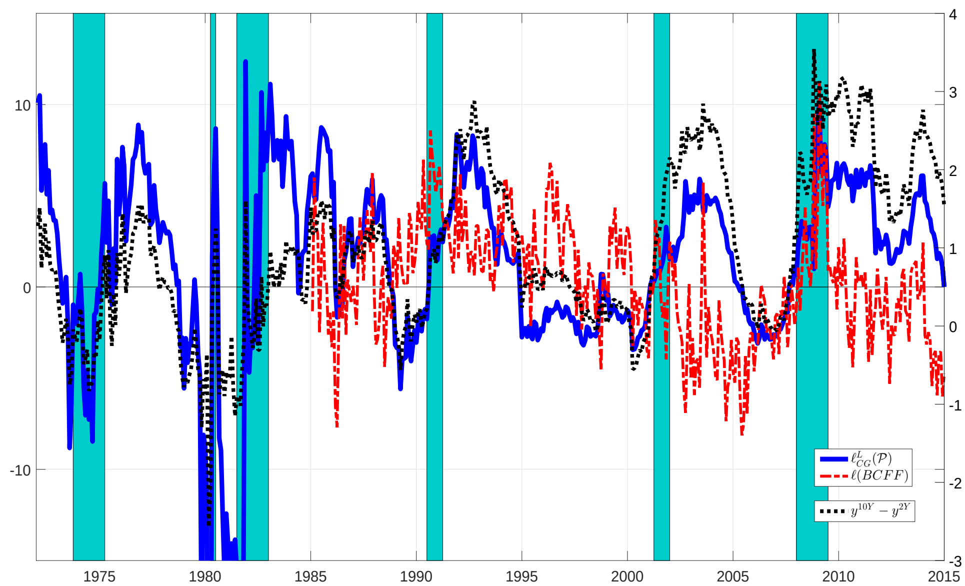

Figure 7: Average expected excess returns over holding periods of ten, eleven and twelve months for the ten-year bond based on $\ell^L_{CG}(\mathcal{P})$ and $\ell(BCFF)$ (left axis) and the slope of the Treasury curve measured as $y^{10y} - y^{2y}$ (right axis), January, 1972 to December, 2014.

on ex ante information about $\mathcal{P}_t$ whereas professionals have access to a broader information set.

Figure 7 displays the *expected* excess returns for a one-year holding periods on ten-year bonds implied by $\ell(BCFF)$ and $\ell^L_{CG}(\mathcal{P})$. The premiums implied by these rules show major differences after every NBER recession in our sample. Key to understanding these differences is the strong positive correlation between the risk premium on ten-year bonds and the steepness of the yield curve. Precisely when the Treasury curve is relatively steep, the median $BCFF$ forecaster believes that risk compensation is much lower than what is implied by our *DTSM*-based learning rule. This finding complements (and is distinct from) Rudebusch and Williams (2009)'s finding that the slope of the yield curve gives more reliable forecasts of recessions than the one-year ahead recession probabilities from the Survey of Professional Forecasters. Our focus is on risk compensation post-recessions as the U.S. economy recovers.

The risk premiums implied by $\ell^L_{CG}(\mathcal{P})$ are much more highly correlated with $(y^{10y}_t - y^2_t)$ than are the premiums from $\ell(BCFF)$. The differences between the expected excess returns from $\ell^L_{CG}(\mathcal{P})$ and $\ell(BCFF)$ are attributable primarily to differences in forecasts of the ten-year yield coming out of recessions. For example, following the low (recession) levels of y^{10y} from

Rule	RMSE's (in basis points) for Annual Horizon						
	6m	1Y	2Y	3Y	5Y	7Y	10Y
$\ell(RW)$	118.8	115.3	103.3	94.1	84.9	78.8	70.8
	(−1.00)	(−0.83)	(−1.90)	(−2.65)	(−2.82)	(−2.75)	(−2.69)
	[−]	[−]	[−]	[−]	[−]	[−]	[−]
$\ell(BCFF)$	128.8	123.9	122.1	122.5	105.9	100.6	88.1
	(−)	(−)	(−)	(−)	(−)	(−)	(−)
	[1.00]	[0.83]	[1.90]	[2.65]	[2.82]	[2.75]	[2.69]
$\ell_{CG}(\mathcal{P})$	108.9	105.7	98.7	90.9	83.0	77.0	70.8
	(−1.60)	(−1.68)	(−2.28)	(−2.93)	(−2.98)	(−3.10)	(−2.74)
	[−1.34]	[−1.27]	[−0.79]	[−0.55]	[−0.33]	[−0.38]	[0.01]
$\ell_{CG}(\mathcal{P}, H)$	109.0	104.9	95.9	86.7	76.9	69.9	63.4
	(−1.42)	(−1.51)	(−2.20)	(−2.94)	(−3.36)	(−3.88)	(−3.75)
	[−1.30]	[−1.41]	[−1.35]	[−1.26]	[−1.33]	[−1.66]	[−1.85]

Table 4: RMSE's for one-year ahead forecasts, January 1995 to December 2014. The D-M statistics for the differences between the *DTSM*- and BCFF-implied (*DTSM*- and *RW*-implied) forecasts are given in parentheses (brackets).

late 2002 until 2004 the $BCFF$ forecasters expected a much more rapid rise in y^{10y} than did $\mathcal{RA}$'s more accurate rule $\ell^L_{CG}(\mathcal{P})$. Put differently, the widespread advice to reduce long-term bond positions as the US economy emerged from recent recessions, while consistent with the subjective beliefs of the median $BCFF$ forecaster, was in fact poor advice relative to the *ex ante* signal from $\ell^L_{CG}(\mathcal{P})$ and (with hindsight) the actual performance of bonds. Moreover, using median $BCFF$ forecasts of long-term bond yields to calibrate empirical learning rules would likely lead to distorted measures of required risk compensations.[31]

5 Learning From Disagreement

The extended rule $\ell_{CG}(\mathcal{P}, H)$ fits directly into our learning framework (same priced risk factors $\mathcal{P}_t$ and SDF (16)) by expanding Z_t' to $(\mathcal{P}_t', H_t')$. Owing to the nonlinear updating rules upon which $\mathcal{RA}$ bases her forecasts of future yields, the inclusion of H_t in Z_t may materially change $\mathcal{RA}$'s forecasts of expected excess returns.[32]

Conditioning on H_t leads to a substantial improvement in forecast accuracy relative to the *DTSM*-based rules that condition only on $\mathcal{P}$ (Table 4). For the $\mathcal{B}$↓**CGLS** learner this pickup in accuracy occurs across the maturity spectrum with a slight tendency for larger gains at the long end of the Treasury curve. Moreover, $\ell_{CG}(\mathcal{P}, H)$ outperforms rule $\ell(RW)$ across the maturity spectrum (most significantly for long-maturity bonds). Again, $\ell(BCFF)$ is the least accurate rule. As reassurance that the superior performance of $\ell_{CG}(\mathcal{P}, H)$ is not an artifact of

[31] Gains in forecast performance may come from using information embedded in survey forecasts of short-term rates and, indeed, Altavilla, Giacomini, and Ragusa (2014) present evidence consistent with this view.

[32] Z could be augmented to also include macroeconomic information, but we defer this extension until Section 6, and focus for now on characterizing learning conditional on belief heterogeneity.

Rule	RMSE's by Bond Maturity 6m	1Y	2Y	3Y	5Y	7Y	10Y
	January, 1995 – December, 2000						
$\ell(RW)$	128	130	119	108	100	93	84
$\ell(BCFF)$	136	131	125	113	104	95	85
$\ell_{CG}(\mathcal{P})$	113	114	106	96	90	86	82
$\ell_{CG}(\mathcal{P}, H)$	111	112	101	89	81	76	71
	January, 2001 – December, 2007						
$\ell(RW)$	154	144	127	114	90	71	56
$\ell(BCFF)$	157	149	142	136	110	95	77
$\ell_{CG}(\mathcal{P})$	142	134	124	111	90	74	58
$\ell_{CG}(\mathcal{P}, H)$	143	134	121	106	83	65	48
	January, 2008 – December, 2014						
$\ell(RW)$	51	52	47	49	64	72	72
$\ell(BCFF)$	82	84	95	115	102	110	100
$\ell_{CG}(\mathcal{P})$	53	51	54	59	68	71	72
$\ell_{CG}(\mathcal{P}, H)$	55	52	53	58	66	69	69

Table 5: RMSE's in basis points for one-year-ahead forecasts of individual bond yields over the indicated sample periods.

small-sample bias, we note that the corresponding D-M statistics have even lower probability values at the quarterly horizon. As with rule $\ell_{CG}(\mathcal{P})$, the outperformance of $\ell_{CG}(\mathcal{P}, H)$ over rule $\ell(BCFF)$ is especially large following recessions.

Table 5 provides a more nuanced view of the relative one-year forecast accuracies across sub-periods. The outperformance of $\ell_{CG}(\mathcal{P}, H)$ relative to both $\ell(RW)$ and $\ell(BCFF)$ was especially large during the early 2000's leading up to the global financial crisis, the most challenging subperiod to forecast shorter term Treasury yields. The poor relative performance of $\ell(BCFF)$ in forecasting the two- to three-year segment of the Treasury curve is interesting in light of the findings of Fleming and Remolona (1999) and Piazzesi (2005) that this segment of the yield curve shows the largest responses to surprise macroeconomic announcements.

The portion of our sample covering the global financial crisis was the easiest subperiod for forecasting Treasury yields. With very short-term rates pegged essentially at zero, $\ell(RW)$ was the best performing rule out to the five-year maturity. $\ell_{CG}(\mathcal{P}, H)$ and $\ell(RW)$ showed nearly identical accuracy for long-maturity bonds. $\mathcal{RA}$'s DTSM-based learning rules do not directly incorporate a zero lower bound for the Federal Reserve's policy rate (see, e.g., Kim and Singleton (2012) and Christensen and Rudebusch (2015)). Nevertheless, the adaptive nature of her rule gives a partial adjustment for the zero-rate policy. Indeed, shortening our

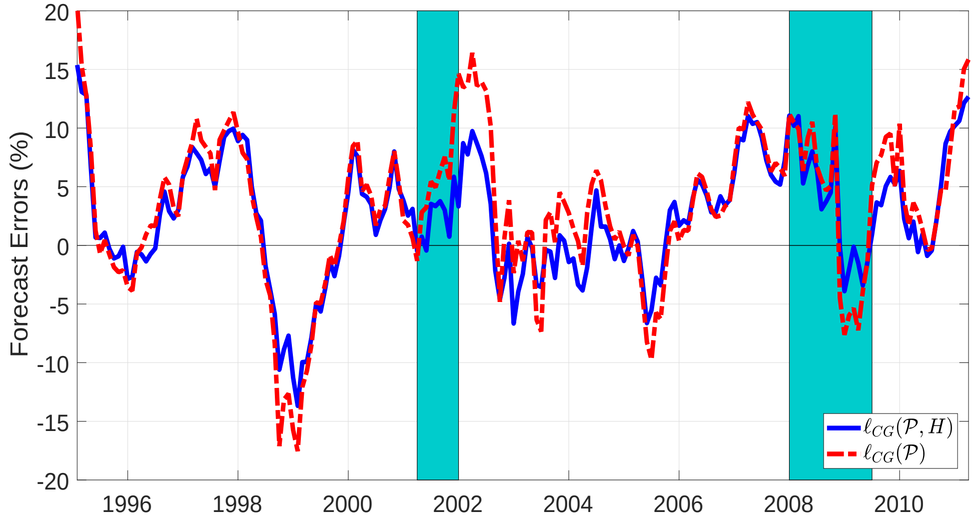

Figure 8: Errors (in percent per annum) from forecasting the realized excess returns on the ten-year bond over a one-year horizon based on rules $\ell_{CG}(\mathcal{P}, H)$ (solid) and $\ell_{CG}(\mathcal{P})$ (dashed).

evaluation window to the period from 2008 to 2011 leads to outperformance of $\ell_{CG}(\mathcal{P}, H)$ over $\ell(RW)$. Only after several years of short rates near a zero lower bound does rule $\ell(RW)$ slightly outperform $\ell_{CG}(\mathcal{P}, H)$ over the intermediate segment of the Treasury curve.

In which economic environments does conditioning on H improve $\mathcal{RA}$'s forecast accuracy over the simpler rule $\ell_{CG}(\mathcal{P})$? From Figure 8 it is seen that H_t is particularly informative about risk premiums– $\ell_{CG}(\mathcal{P}, H)$ gives smaller forecast errors than $\ell_{CG}(\mathcal{P})$– during major turning points (peaks and troughs) in $\mathcal{RA}$'s subjective risk premiums. The differences are particularly large during the Asian financial crisis of 1998, following the burst of the "dot-com bubble," and during the recent financial crisis. That $\ell_{CG}(\mathcal{P}, H)$ can manifest such selected periods of relevance for H is possible owing to its nonlinear updating of the feedback matrix in (17). Since $\ell_{CG}(\mathcal{P}, H)$ is well approximated by a recursively updated constant-gain least-squares rule, it is not unreasonable to argue that market participants could have implemented in real time an approximate version of $\mathcal{RA}$'s Bayesian learning rule.

6 Macroeconomic Information, Beliefs and Disagreement

Within-sample analyses of fixed-parameter vector-autoregressions provide substantial evidence that macroeconomic fundamentals (e.g., Ludvigson and Ng (2010) and Joslin, Priebsch, and Singleton (2014) (JPS)) and disagreement among professional forecasters about future inflation and real growth (e.g., Buraschi and Whelan (2016) and Ehling, Gallmeyer, Heyerdahl-Larsen, and Illeditsch (2016)) have predictive power for excess returns in bond markets. This leads us naturally to inquire as to whether our evidence on the predictive power of H arises as a consequence of our omission of macro conditioning variables. That is, might the role of H be simply that it is a stand-in for business cycle information?

An increasingly recognized issue when studying the predictive power of macro-variables for future yields is that official macro-time series are continuously updated after their original release date. To address this issue, we construct measures of inflation (INF) and real economic activity (REA) that market participants would have known in real time. We use data from the Archival Federal Reserve Economic Data (ALFRED) database, which reports the original releases of macroeconomic series. Letting $x_{s|t}$ denote an economic statistic indexed to time s and available at time $t \geq s$, and recognizing that most economic statistics are released with a one-month delay, an investor at time t can typically condition on

$$x_{t_0-1|t}, x_{t_0|t}, ..., x_{t-2|t}, x_{t-1|t},$$

where t_0 indicates the start of the training sample. Importantly, this is the fully updated series through time t, and not the series as it was released in real time.[33] INF is the twelve-month log difference of the Consumer price index for all urban consumers that is available at the time of estimation. REA is the three-month moving average of the first principal component of six series related to real economic activity.[34]

Using the $BCFF$ panel of forecasters, we construct consensus (median) forecasts of one-year inflation $Cons(INF)$ and real GDP growth $Cons(RGDP)$ from the monthly cross-sections of forecasters.[35] Similarly, disagreement about one year ahead inflation $ID(INF)$ and growth $ID(RGDP)$ are measured as the inter-decile ranges of the cross-sections of forecasts. Both INF and REA are negatively correlated with forecasters' disagreement about future macroeconomic variables, $ID(INF)$ and $ID(RGDP)$. Just as with disagreement about future yields, disagreement about the macroeconomy increases during weak economic times, as inflation and real economic activity decline.

[33] Prior studies using original release data have not always updated their series through time t as we do (e.g., Ghysels, Horan, and Moench (2014)). Such studies are using stale data relative to what market participants knew at the time they constructed their forecasts.

[34] The series are the difference in the logarithm of Industrial production index (INDPRO), the difference in the logarithm of total nonfarm payroll (PAYEMS), the difference of the civilian unemployment rate (UNRATE), the difference of the logarithm of "All employees: Durable goods" (DMANEMP), the difference of the logarithm of "All employees: Manufacturing" (MANEMP), and the difference of the logarithm of "All employees: NonDurable goods" (NDMANEMP). The first PC is smoothed similarly to the Chicago Fed National Activity Index.

[35] We compute one-year-ahead expected inflation and real GDP growth for each forecaster as the average of the one, two, three and four quarter ahead forecasts.

	Dep. Variable: ID(2y)			Dep. Variable: ID(7y)-ID(2y)		
Const.	0.0137	−0.0024	−0.0070	−0.0001	0.0011	0.0020
	[11.3905]	[−0.9475]	[−3.5984]	[−0.1831]	[0.6936]	[1.3320]
REA	−0.0002	−0.0007	−0.0003	0.0002	0.0003	0.0002
	[−0.6188]	[−2.7495]	[−1.3915]	[1.0369]	[1.4204]	[1.0436]
INF	−0.0380	−0.0120	0.0056	−0.0330	−0.0340	−0.0386
	[−0.8705]	[−0.3056]	[0.2125]	[−1.2323]	[−1.2129]	[−1.4753]
Cons(RGDP)	–	0.1654	0.1786	–	−0.0693	−0.0745
	[–]	[2.2722]	[3.4220]	[–]	[−1.9647]	[−2.2288]
Cons(INF)	–	0.3867	0.2564	–	0.0245	0.0489
	[–]	[9.0096]	[6.8984]	[–]	[1.1187]	[1.7259]
ID(RGDP)	–	–	0.2345	–	–	−0.0811
	[–]	[–]	[4.0740]	[–]	[–]	[−1.9215]
ID(INF)	–	–	0.3433	–	–	−0.0192
	[–]	[–]	[4.2397]	[–]	[–]	[−0.3253]
R-square	2.44%	52.70%	68.16%	1.96%	7.32%	10.19%

Table 6: Regressions of yields disagreement on macroeconomic information. Sample from January 1985 through December 2014.

Table 6 reports the contemporaneous projection of the level $ID(y^{2y})$ and slope $ID(y^{7y}) - ID(y^{2y})$ of disagreement onto current (INF, REA) and beliefs about future $(INF, RGDP)$. Nearly seventy percent of the variation in $ID(y^{2y})$ is accounted for by variation in consensus forecasts $(Cons(INF), Cons(RGDP))$ and disagreement $(ID(INF), ID(RGDP))$. Yet the pattern of coefficients in Figure 3 for the excess return on the ten-year bond shows that weights on $ID(y^{2y})$ and $ID(y^{7y})$ are approximately equal in opposite signs, suggesting that disagreement about the yield curve has predictive power for risk premiums primarily through the slope $ID(y^{2y}) - ID(y^{7r})$. It is striking that this slope is virtually orthogonal to (INF, REA) and $(ID(INF), ID(RGDP))$, and shows only weak correlations with $(Cons(INF), Cons(RGDP))$. This provides an initial hint that the predictive power of H in $\mathcal{RA}$'s learning rule is not because it is a proxy for macroeconomic information.

To formally evaluate the contribution of macro variables to $\mathcal{RA}$'s learning we expand the conditioning information in the dynamic learning framework of Section 3, again setting the gain coefficient γ equal to 0.99. The out-of-sample RMSEs of forecasts of excess returns for the one-year holding period are reported in Table 7. The macro learning rule $\ell_{CG}(\mathcal{P}, REA, INF)$ performs comparably to or actually *underperforms* the simpler nested rule $\ell_{CG}(\mathcal{P})$. The only subperiod during which $\ell_{CG}(\mathcal{P}, REA, INF)$ outperforms $\ell_{CG}(\mathcal{P})$ is leading up to the recent crisis (Part B), and this relative accuracy is attained only out to the three-year maturity.[36]

[36]These findings suggest that the full-sample analysis of JPS likely overstates the real-time predictive power of output growth and inflation for risk premiums in bonds markets. Though based on a very different form of evidence, our findings are consistent with the concerns of Bauer and Hamilton (2018). Consistent with JPS, there is some evidence of predictive power, particularly for REA, during the first part of the 2000's.

	2Y	3Y	5Y	7Y	10Y
Part A: January, 1995 – December, 2014					
$\ell_{CG}(\mathcal{P})$	1.10%	1.97%	3.39%	4.76%	6.65%
$\ell_{CG}(\mathcal{P}, H)$	1.09%	1.92%	3.17%	4.36%	5.96%
$\ell_{CG}(\mathcal{P}, REA)$	1.07%	1.96%	3.50%	5.02%	7.22%
$\ell_{CG}(\mathcal{P}, REA, INF)$	1.07%	1.97%	3.51%	5.04%	7.24%
$\ell_{CG}(\mathcal{P}, H, REA)$	1.07%	1.92%	3.32%	4.68%	6.60%
$\ell_{CG}(\mathcal{P}, ID(RGDP), ID(INF))$	1.20%	2.16%	3.71%	5.23%	7.14%
Part B: January, 2001 – December, 2007					
$\ell_{CG}(\mathcal{P})$	1.37%	2.44%	3.84%	4.95%	5.72%
$\ell_{CG}(\mathcal{P}, H)$	1.36%	2.39%	3.60%	4.47%	4.79%
$\ell_{CG}(\mathcal{P}, REA)$	1.22%	2.29%	4.02%	5.72%	7.71%
$\ell_{CG}(\mathcal{P}, REA, INF)$	1.23%	2.32%	4.09%	5.84%	7.92%
$\ell_{CG}(\mathcal{P}, H, REA)$	1.21%	2.21%	3.69%	5.08%	6.53%
$\ell_{CG}(\mathcal{P}, ID(RGDP), ID(INF))$	1.48%	2.65%	4.21%	5.48%	6.70%

Table 7: RMSEs for average expected excess returns over holding periods of ten, eleven and twelve months, based on learning rules with different choices of conditioning information.

By contrast, conditioning on H in rule $\ell_{CG}(\mathcal{P}, H)$ reduces the RMSE of the ten-year bond return by 0.7% relative to rule $\ell_{CG}(\mathcal{P})$ (Part A, a 10.5% reduction of the RMSE), and by 1.3% relative to rule $\ell_{CG}(\mathcal{P}, REA)$ (an 18.5% reduction of the RMSE). Rule $\ell_{CG}(\mathcal{P}, H, REA)$ delivers RMSEs that are very close to those obtained from $\ell_{CG}(\mathcal{P}, H)$. Thus, the source of forecasting power of H for the long end of the curve is distinct from the macroeconomic information (INF, REA). Moreover, combining $\mathcal{P}$ with beliefs on future macroeconomic information $(Cons(RGDP), Cons(INF))$ and $(ID(RGDP), ID(INF))$ leads to sizable *deteriorations* in forecasting accuracy relative to rule $\ell_{CG}(\mathcal{P}, H)$, *across the entire maturity spectrum*.

The relative outperformance of rule $\ell_{CG}(\mathcal{P}, H)$ over $\ell_{CG}(\mathcal{P}, REA)$ can be seen from Figure 9, which displays the average expected excess returns for the ten-year bonds against the corresponding realized returns. The outperformance of $\ell_{CG}(\mathcal{P}, H)$ is at times large, especially during and immediately after NBER recessions. The primary exceptions, when $\ell_{CG}(\mathcal{P}, REA)$ outperforms, are during portions of the post-crisis period, from 2012 through 2015. This incremental role of REA over and above H can also be seen from rule $\ell_{CG}(\mathcal{P}, H, REA)$.

6.1 Direct and Indirect Contributions of Disagreement

Within our learning models there are two channels through which conditioning information not spanned by the yield curve can impact expected excess returns. The first is the direct

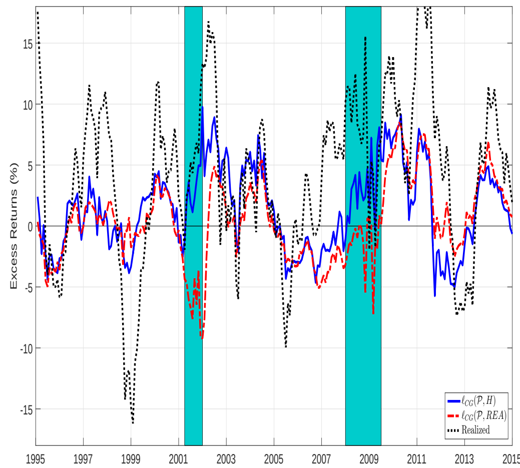

(a) Ten-Year ZCB: $\ell_{CG}(\mathcal{P}, H)$, $\ell_{CG}(\mathcal{P}, REA)$

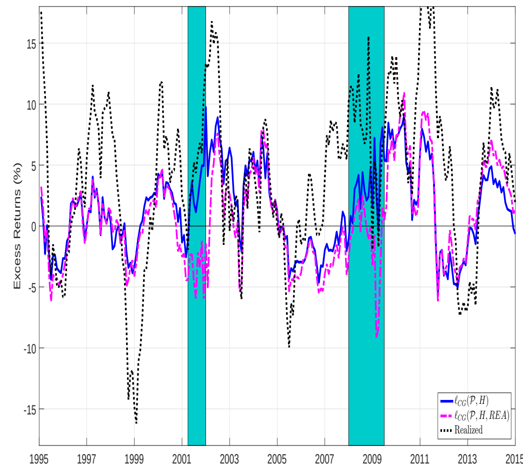

(b) Ten-Year ZCB: rules $\ell_{CG}(\mathcal{P}, H)$, $\ell_{CG}(\mathcal{P}, H, REA)$

Figure 9: Average expected excess returns over holding periods of ten, eleven and twelve months for ten-year bonds based on $\ell_{CG}(\mathcal{P}, H)$, $\ell_{CG}(\mathcal{P}, REA)$ and $\ell_{CG}(\mathcal{P}, H, REA)$, overlaid with the realized returns.

effect that this information has on forecasts of future PCs as components of Z_t in (29). The second is the indirect effect on how $\mathcal{RA}$ updates the parameters $(\widehat{K}^{\mathbb{P}}_{0t}, \widehat{K}^{\mathbb{P}}_{Zt})$ by conditioning on information beyond $\mathcal{P}$ as part of the learning process. We can frame these effects in terms of risk premiums, and within our favorite model with conditioning information $Z = [\mathcal{P}, H]$. Let $er^{n,1}_t(\mathcal{P}, H)$ $(er^{n,1}_t(\mathcal{P}))$ denote $\mathcal{RA}$'s expected excess return from holding a bond with maturity n for one year under rule $\ell_{CG}(\mathcal{P}, H)$ $(\ell_{CG}(\mathcal{P}))$. Using (29) and the expression for realized excess returns in terms of $\hat{y}^n_{t+h}$, we can write:

$$er^{n,1}_t(\mathcal{P}, H) = \hat{a}^{\mathcal{P},H}_{n,t} + \hat{b}^{\mathcal{P},H}_{n,t}\mathcal{P}_t + \hat{c}_{n,t}H_t, \tag{31}$$

$$er^{n,1}_t(\mathcal{P}) = \hat{a}^{\mathcal{P}}_{n,t} + \hat{b}^{\mathcal{P}}_{n,t}\mathcal{P}_t. \tag{32}$$

We separate the direct and indirect effects of H in (31) by estimating a rolling constant-gain least-squares projection of H_t onto $\mathcal{P}_t$,

$$H_t = \hat{\alpha}_t + \hat{\beta}_t\mathcal{P}_t + u_t, \tag{33}$$

and then construct the pseudo risk premium

$$er^{n,1}_t(\mathcal{P}, 0) = (\hat{a}^{\mathcal{P},H}_{n,t} + \hat{c}_{n,t}\hat{\alpha}_t) + (\hat{b}^{\mathcal{P},H}_{n,t} + \hat{c}_{n,t}\hat{\beta}_t)\mathcal{P}_t \equiv \hat{a}^{\mathcal{P},0}_{n,t} + \hat{b}^{\mathcal{P},0}_{n,t}\mathcal{P}_t. \tag{34}$$

	$\sigma^2[er^{10,1}]$	$\sigma^2[er^{10,1}(\mathcal{P},0)]$	$\rho(er^{10,1}, er^{10,1}(\mathcal{P},0))$	$\frac{\sigma^2[er^{10,1}-er^{10,1}(\mathcal{P},0)]}{\sigma^2[er^{10,1}]}$
		January, 1995 – December, 2014		
$\ell_{CG}(\mathcal{P},H)$	0.12%	0.09%	94.67%	10.83%
$\ell_{CG}(\mathcal{P},REA)$	0.11%	0.07%	83.23%	31.10%
$\ell_{CG}(\mathcal{P},ID(RGDP),ID(INF))$	0.16%	0.17%	96.03%	7.99%
$\ell_{CG}(\mathcal{P},Cons(RGDP),Cons(INF))$	0.16%	0.08%	77.47%	40.32%
		January, 2001 – December, 2007		
$\ell_{CG}(\mathcal{P},H)$	0.14%	0.12%	95.53%	8.77%
$\ell_{CG}(\mathcal{P},REA)$	0.11%	0.04%	71.48%	49.71%
$\ell_{CG}(\mathcal{P},ID(RGDP),ID(INF))$	0.15%	0.16%	96.54%	7.48%
$\ell_{CG}(\mathcal{P},Cons(RGDP),Cons(INF))$	0.09%	0.07%	77.98%	40.34%
		January, 2008 – December, 2014		
$\ell_{CG}(\mathcal{P},H)$	0.17%	0.11%	95.94%	10.62%
$\ell_{CG}(\mathcal{P},REA)$	0.13%	0.08%	81.53%	33.53%
$\ell_{CG}(\mathcal{P},ID(RGDP),ID(INF))$	0.11%	0.09%	86.69%	25.15%
$\ell_{CG}(\mathcal{P},Cons(RGDP),Cons(INF))$	0.36%	0.12%	88.72%	31.30%

Table 8: Statistics for risk premia and pseudo-risk premia (see (34)) from various learning rules. All statistics refer to the average expected excess returns over holding periods of ten, eleven and twelve months for the ten-year bonds.

The difference between $er_t^{n,1}(\mathcal{P},0)$ and $er_t^{n,1}(\mathcal{P})$ arises entirely from the effect that H has on the updating of the weights on $\mathcal{P}_t$ when learning is conditioned on the full information generated by $(\mathcal{P},H)$[37]. The same decomposition can be constructed for the "macro-rules" discussed in the previous section.

Figure 10, Panel (a) displays the one-year-ahead expected excess returns (average of twelve, eleven and ten months ahead) for the ten-year bond from rule $\ell_{CG}(\mathcal{P},H)$ against the corresponding pseudo risk premia $er^{10,1}(\mathcal{P},0)$. These premia track each other very closely, and in Table 8 we show that their sample correlation is close to 95%. This suggests that the direct effect of H on expected excess returns is small (its effect is mostly indirect through coefficient updating). This finding also holds for the subsamples reported in Table 8. Further, the ratio of the variance of the difference between the risk premia and pseudo-risk premia to the variance of the risk premia (last column) is always below 11%, again consistent with the hypothesis that the direct effect of H on $\mathcal{RA}$'s expected excess returns is unsubstantial.

Precisely how are these indirect effects of H manifested in $\mathcal{RA}$'s risk premiums? Figure 11 shows the difference $[\hat{b}^{\mathcal{P},0}_{10y,t} - \hat{b}^{\mathcal{P}}_{10y,t}]$ of the weight on $PC1$ (the "level" factor) weighted by the sample standard deviation of $PC1$. This captures the relative sensitivity of the pseudo excess return $er_t^{n,1}(\mathcal{P},0)$ implied by rule $\ell_{CG}(\mathcal{P},H)$ and the excess return $er_t^{n,1}(\mathcal{P})$ from rule $\ell_{CG}(\mathcal{P})$ to a one standard deviation shock in $PC1$. Any differences are entirely the consequence of

[37] While it is true that (33) is fit outside of our *DTSM*, the $(\hat{\alpha}_t, \hat{\beta}_t)$ that we recover using monthly data would be literally identical to those recovered within a *DTSM* without constraints on the market prices of risk. This is an immediate implication of the propositions in JSZ. Therefore, we believe we are obtaining a reliable picture of the impact of H on the loadings on $\mathcal{P}$ in the expression for $er_t^{n,1}(\mathcal{P},0)$.

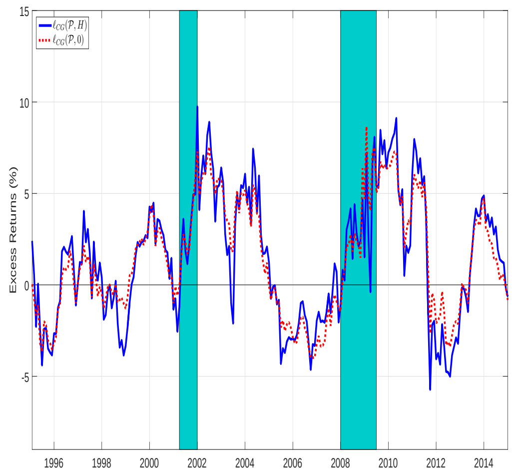

(a) $\ell_{CG}(\mathcal{P}, H)$ vs. $\ell_{CG}(\mathcal{P}, 0)$

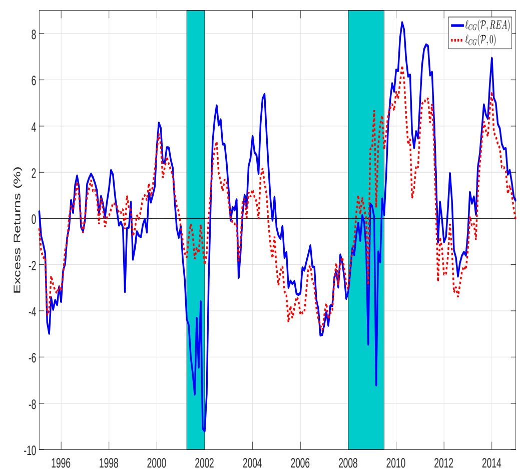

(b) $\ell_{CG}(\mathcal{P}, REA)$ vs. $\ell_{CG}(\mathcal{P}, 0)$

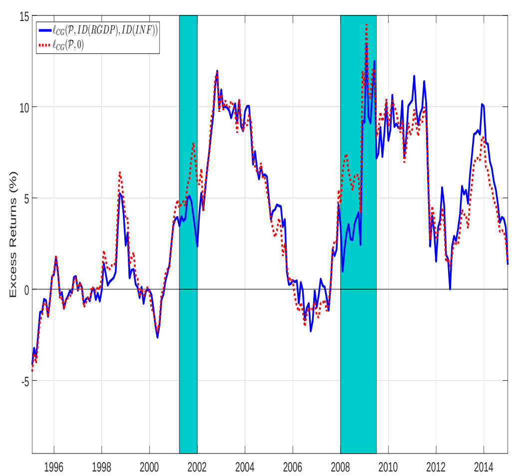

(c) $\ell_{CG}(\mathcal{P}, Cons(RGDP), Cons(INF))$ vs. $\ell_{CG}(\mathcal{P}, 0)$

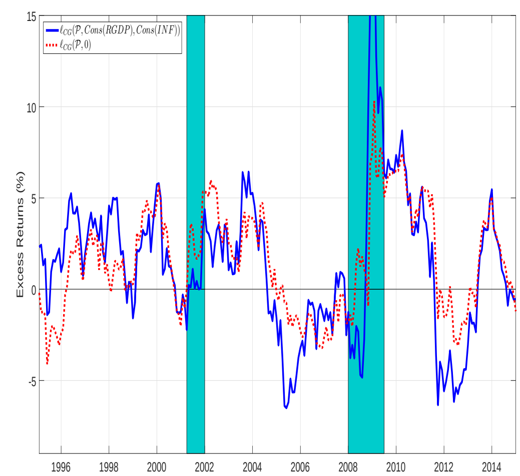

(d) $\ell_{CG}(\mathcal{P}, ID(RGDP), ID(INF))$ vs. $\ell_{CG}(\mathcal{P}, 0)$

Figure 10: Average expected excess returns over holding periods of ten, eleven and twelve months for the ten-year bonds. Comparison between risk premia generated by the learning rule, and pseudo risk premia computed according to (34).

the indirect effect of H on $\mathcal{RA}$'s subjective risk premia. The pattern of negative differences implies that excess returns tend to be *less* responsive to shocks to the level of interest rates under $\ell_{CG}(\mathcal{P}, H)$ than under $\ell_{CG}(\mathcal{P})$ *during and following recessions*. These periods are also when relative disagreement $(ID(y^{2y}) - ID(y^{7y}))$ is large. Taken together, the results

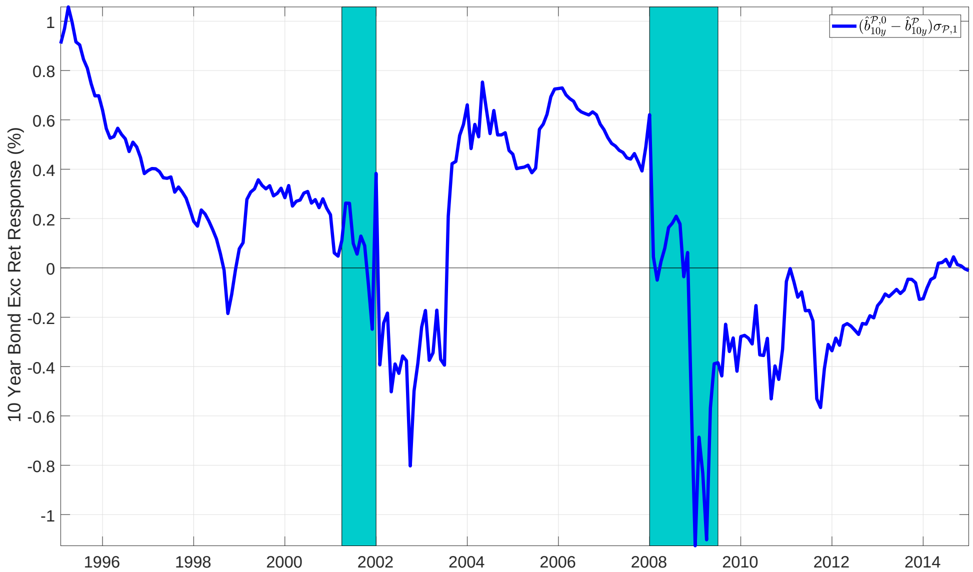

Figure 11: Relative sensitivity of the excess returns of the 10-year bond in rules $\ell_{CG}(\mathcal{P}, 0)$ and $\ell_{CG}(\mathcal{P})$ to a one standard-deviation shock to $PC1$: the difference $[\hat{b}^{P,0}_{t,10y} - \hat{b}^{P}_{t,10y}]_1$ (estimated every month from (32) and (34)) of the weight on $PC1$ scaled by the unconditional standard deviation of the PC1, computed over the sample from January 1985 through December 2014.

suggest that the higher trading volume naturally associated with heightened disagreement reduces the pressure on prices following shocks to the level of interest rates. In contrast, when disagreement is relatively low, H induces greater price sensitivity to shocks to $PC1$ under $\ell_{CG}(\mathcal{P}, H)$ compared to $\ell_{CG}(\mathcal{P})$.[38]

6.2 Direct and Indirect Contributions of Macro Information

Interestingly, the direct effects are much larger in the macro rules. Panel (b) of Figure 10 shows substantial differences for rule $\ell_{CG}(\mathcal{P}, REA)$, especially during and immediately following recessions. Table 8 shows that the correlation between risk premia and pseudo-risk premia is only 83% over the entire sample, and the last column suggests that much of the variation in $er^{10,1}(\mathcal{P}, REA)$ arises from the direct effects of REA on risk premia. Rule $\ell_{CG}(\mathcal{P}, REA)$ delivers its best performance over the sample from January 2001 through December 2007 when the correlation between the risk premia is only 77.47%, and the variance ratio is close to 50%.

[38] Malmendier, Pouzo, and Vanasco (2017) develop a theory of experiential effects in financial markets that induces a similar pattern for how disagreement affects price-pressure effects. Their disagreement is across generations, and our measures are across active professionals in the market.

Similar observations apply for $\ell_{CG}(\mathcal{P}, Cons(CPI), Cons(RGDP))$ (Panel (d) of Figure 10). Thus, both real-time macroeconomic information and consensus beliefs have a substantial direct effect on expected excess returns.

The corresponding results for rule $\ell_{CG}(\mathcal{P}, ID(CPI), ID(RGDP))$ (Panel (c) of Figure 10) show that the risk and pseudo-risk premia track each other closely, and their correlation is over 96% (Table 8). This rule delivers its best performance in the latest subsample. During this subperiod, the correlation between risk premia and pseudo-risk premia is only 86.69%, and the variance ratio in the last column is larger than 25%. Thus, again, there is a moderate direct effect of $ID(CPI)$ and $ID(RGDP)$).

7 Concluding Remarks

Three notable patterns emerge from our analysis: (i) $\mathcal{RA}$ effectively treats the parameters governing the risk-neutral distribution of the pricing factors $\mathcal{P}$ as known and constant over time (they were held virtually constant over the past thirty years); (ii) given this finding, a constrained version of the optimal Bayesian learning rule specializes to constant-gain, least-squares learning; and (iii) implementing the rule $\ell_{CG}(\mathcal{P}, H)$ in real time gives rise to forecasts of future bond yields and risk premiums that substantially outperform the analogous rule $\ell_{CG}(\mathcal{P})$ based on yield curve information alone, the learning rule implicitly followed by the median BCFF professional forecaster, and learning rules that condition on real-time information about the macroeconomy. These outperformances are especially large following recessions, when disagreement about future two-year Treasury yields is high.

Since our learning rules are inherently reduced-form, we cannot reach definitive conclusions regarding the economic mechanisms through which the dispersion measures H gain their predictive power. Nevertheless, several intriguing patterns emerge that are suggestive of fruitful directions for theoretical modeling. Under the premise that there is heterogeneity of views in the U.S. Treasury markets, the most likely source of disagreement is regarding the future paths of bond yields and not about the connection between the current state of the economy and current yields. Financial institutions have long recognized that the cross-section of bond yields is well described by the low-order PCs which are readily observable. Indeed, most use pricing and risk management systems that presume that current economic conditions are fully reflected in the PCs.

A notable feature of rule $\ell_{CG}(\mathcal{P}, H)$ is that, to a substantial degree, measures of dispersion in beliefs about future bond yields (H) affect risk premiums through the weights that $\mathcal{RA}$ assigns to the current PCs when forecasting future yields. $\mathcal{RA}$ finds it optimal to adjust the predictive content of $\mathcal{P}_t$ about future $\mathcal{P}_{t+s}$ depending on the degree of disagreement in the market, and this is especially the case after NBER recessions. During periods of heightened disagreement, $\mathcal{RA}$'s risk premiums are relatively *less* sensitive to shocks to the level of interest rates than during periods when disagreement is relatively low. Additionally, the effects of H on forecasts persist over several months. This leads us to doubt that the primary impact of H is through high-frequency episodes of flight-to-quality.

H does not seem to be a proxy for omitted macroeconomic information. Indeed, learning rules that include macro information frequently perform worse than the rule $\ell_{CG}(\mathcal{P})$, which uses only the shape of the yield curve as conditioning information.

Pursuing this further, we computed the correlations between $ID_t^{slp} \equiv 1/2(ID_t^{7y} - ID_t^{2y}))$ and the measures of "economic policy uncertainty" constructed by Baker, Bloom, and Davis. We included both their overall index and twelve of their subcategories of policy uncertainty[39]. In all cases the correlations were very small. This is further evidence that the predictive power of H is at most weakly tied to uncertainty about the macroeconomy as conventionally measured.

[39] This data was downloaded from www.policyuncertainty.com.

Appendices

A Bond Pricing in *GTSMs*

Suppose that bonds are priced under the presumption that $\Theta^{\mathbb{Q}}$ is fixed.

$$\begin{aligned} r_s &= \rho_0 + \rho_1 \mathcal{P}_s, \\ \mathcal{P}_{s+1} &= K^{\mathbb{Q}}_{0\mathcal{P}} + K^{\mathbb{Q}}_{\mathcal{P}\mathcal{P}} \mathcal{P}_s + \Sigma^{1/2}_{\mathcal{P}\mathcal{P}} e^{\mathbb{Q}}_{\mathcal{P},s+1}, \end{aligned}$$

The price of a zero coupon bond is then given by:

$$D^m_s = e^{\mathcal{A}_m + \mathcal{B}_m \mathcal{P}_s},$$

Where m denotes the maturity of the bond. We can calculate $\mathcal{A}_m$ and $\mathcal{B}_m$ by solving the first order difference equation:

$$\begin{aligned} \mathcal{A}_{m+1} - \mathcal{A}_m &= \left(K^Q_0\right)' \mathcal{B}_m + \frac{1}{2}\mathcal{B}'_m \Sigma_{\mathcal{P}} \mathcal{B}_m - \rho_0, \\ \mathcal{B}_{m+1} - \mathcal{B}_m &= \left(K^Q_1 - I\right)' \mathcal{B}_m - \rho_1. \end{aligned}$$

With initial conditions $\mathcal{A}_0 = 0$ and $\mathcal{B}_0 = 0$. The loadings for the corresponding bond yields are $A_m = -\mathcal{A}_m/m$ and $B_m = -\mathcal{B}_m/m$.

B The Canonical Model

Under the assumption that $\Theta^{\mathbb{Q}}$ is fixed, we can proceed by adopting the computationally convenient Joslin, Singleton, and Zhu (2011) normalization scheme. Specifically, let X_s denote a set of latent risk factors with

$$\begin{aligned} r_s &= 1' X_s, \\ X_{s+1} &= \begin{bmatrix} k^{\mathbb{Q}}_{\infty} \\ 0 \\ 0 \end{bmatrix} + J(\lambda^{\mathbb{Q}}) X_s + \Sigma^{1/2}_{XX} e^{\mathbb{Q}}_{X,s+1}, \end{aligned}$$

with bond loadings:

$$y^m_s = A_{X,m}(k^{\mathbb{Q}}_{\infty}, \lambda^{\mathbb{Q}}, \Sigma_{XX}) + B_{X,m}\left(\lambda^{\mathbb{Q}}\right) X_s.$$

Under this normalizations Joslin, Singleton, and Zhu (2011) show that there exits a unique rotation of X_s so that the factors are the first three principal components of bond yields:

$$\mathcal{P}_s = v(k_\infty^{\mathbb{Q}}, \lambda^{\mathbb{Q}}, \Sigma_{XX}, W) + L\left(\lambda^{\mathbb{Q}}, W\right) X_s,$$

where $v = W[A_{X,m_1}, ..., A_{X,m_J}]$ and $L = W[B_{X,m_1}, ..., B_{X,m_J}]$, and W denote the principal component weights. The bond loadings A_m and B_m in the expression

$$y_s^m = A_m(\Theta^{\mathbb{Q}}) + B_m(\lambda^{\mathbb{Q}})\mathcal{P}_s,$$

are fully determined by the parameters $\Theta^{\mathbb{Q}} = (k_\infty^{\mathbb{Q}}, \lambda^{\mathbb{Q}}, \Sigma_{\mathcal{PP}})$ and the principal component weights W. Furthermore, the parameters $K_{0\mathcal{P}}^{\mathbb{Q}}$, $K_{\mathcal{PP}}^{\mathbb{Q}}$, ρ_0 , ρ_1, and Σ_{XX} are all transformations of the elements in $\Theta^{\mathbb{Q}}$:

$$\begin{aligned}
K_{0\mathcal{P}}^{\mathbb{Q}} &= L\begin{bmatrix} k_\infty^{\mathbb{Q}} \\ 0 \\ 0 \end{bmatrix} - LJ(\lambda^{\mathbb{Q}})L^{-1}v, \\
K_{\mathcal{PP}}^{\mathbb{Q}} &= LJ(\lambda^{\mathbb{Q}})L^{-1}, \\
\rho_0 &= -1'L^{-1}v, \\
\rho_1 &= (L^{-1})'1, \\
\Sigma_{\mathcal{PP}} &= L\Sigma_{XX}L^{-1}.
\end{aligned}$$

See Joslin, Singleton, and Zhu (2011) for details and proofs.

C Log likelihood function

We begin by noting that when $(\Theta^{\mathbb{Q}}, \Sigma_{\mathcal{O}})$ and Σ_Z are presumed to be constant, equation (22) implies that we can decompose the log likelihood function into a $\mathbb{P}$ and $\mathbb{Q}$ part

$$-2\log L = -2\log L^{\mathbb{Q}}(\Theta^{\mathbb{Q}}, \Sigma_{\mathcal{PP}}, \Sigma_{\mathcal{O}}) - 2\log L^{\mathbb{P}}(\Theta_t^{\mathbb{P}}, \Sigma_Z, Q_t).$$

$\log L^{\mathbb{Q}}$ denotes the part of the likelihood function associated with pricing errors and $\log L^{\mathbb{P}}$ the likelihood function of the dynamic evolution of Z_t,

$$Z_{t+1} = K_{Z,0t}^{\mathbb{P}} + K_{Z,1t}^{\mathbb{P}} Z_t + \Sigma_Z^{1/2} e_{Z,t+1}^{\mathbb{P}}, \tag{35}$$

where $Z_t' = (\mathcal{P}_t', H_t')'$ and $\Theta_t^{\mathbb{P}} = [K_{Z,0t}^{\mathbb{P}}, K_{Z,1t}^{\mathbb{P}}]$ denotes the drifting parameters. We assume $\Theta_t^{\mathbb{P}}$ can be partitioned as $(\psi^r, \psi_t^{\mathbb{P}})$, where $\psi_t^{\mathbb{P}}$ is the vectorized set of free parameters and ψ^r is the vectorized set of parameters that are fixed conditional on $\Theta^{\mathbb{Q}}$. The unrestricted parameters, $\psi_t^{\mathbb{P}}$, evolve according to a random walk

$$\psi_t^{\mathbb{P}} = \psi_{t-1}^{\mathbb{P}} + Q_{t-1}^{1/2}\eta_t \qquad \eta_t \overset{iid}{\sim} N(0, I), \tag{36}$$

with stochastic covariance matrix Q_{t-1}. By moving terms that involve known parameters and observable states to the left hand side we can rewrite equation (35) into

$$\mathcal{Y}_t = \mathcal{X}_{t-1}\psi_{t-1}^{\mathbb{P}} + \Sigma_Z^{1/2} e_{Z,t}^{\mathbb{P}}, \tag{37}$$

where

$$\begin{aligned} \mathcal{Y}_t &= Z_t - \left(I \otimes [1, Z'_{t-1}]\right)\iota_r\psi^r, \\ \mathcal{X}_t &= \left(I \otimes [1, Z'_t]\right)\iota_f, \end{aligned}$$

with ι_r and ι_f denoting the matrices that select the columns of $(I \otimes [1, Z'_{t-1}])$ corresponding to the restricted and free parameters respectively. With normally distributed innovations to the latent parameter states (36) (the transition equation) and to the factor dynamics (37) (the measurement equation) we have a well-defined linear Kalman filter.[40] Conditional on $(\Theta^{\mathbb{Q}}, \Sigma)$ the solution to the Kalman filter is given by recursively updating the posterior mean $\hat{\psi}_t^{\mathbb{P}} = \mathbb{E}^{\mathbb{P}}(\psi_t^{\mathbb{P}}|Z_1^t)$, posterior variance $P_t = \mathbb{V}^{\mathbb{P}}(\psi_t^{\mathbb{P}}|Z_1^t)$, and forecast variance $\Omega_t = \mathbb{V}^{\mathbb{P}}(Z_{t+1}|Z_1^t)$ according to:

$$\begin{aligned} \hat{\psi}_t^{\mathbb{P}} &= \hat{\psi}_{t-1}^{\mathbb{P}} + P_{t-1}\mathcal{X}'_{t-1}\Omega_{t-1}^{-1}(\mathcal{Y}_t - \mathcal{X}_{t-1}\hat{\psi}_{t-1}^{\mathbb{P}}), && (38) \\ P_t &= P_{t-1} + Q_{t-1} - P_{t-1}\mathcal{X}'_{t-1}\Omega_{t-1}^{-1}\mathcal{X}_{t-1}P_{t-1}, && (39) \\ \Omega_{t-1} &= \mathcal{X}_{t-1}P_{t-1}\mathcal{X}'_{t-1} + \Sigma_Z, && (40) \end{aligned}$$

with $\mathbb{P}$ log likelihood function given by

$$\begin{aligned} -2\log L^{\mathbb{P}} &= (t-1)N\log(2\pi) + \sum_{s=2}^{t}\log|\Omega_{s-1}| \\ &+ \frac{1}{2}\sum_{s=2}^{t}(\mathcal{Y}_s - \mathcal{X}_{s-1}\hat{\psi}_{s-1})'\Omega_{s-1}^{-1}(\mathcal{Y}_s - \mathcal{X}_{s-1}\hat{\psi}_{s-1}). \end{aligned} \tag{41}$$

[40]Note that the latent states in the filtering problem are the parameters and not the factors.

Reworking equation (38) gives[41]

$$\hat{\psi}_t^{\mathbb{P}} = \hat{\psi}_{t-1}^{\mathbb{P}} + (P_t - Q_{t-1})\mathcal{X}_{t-1}'\Sigma_Z^{-1}\left(\mathcal{Y}_t - \mathcal{X}_{t-1}\hat{\psi}_{t-1}^{\mathbb{P}}\right). \tag{42}$$

Letting $R_t = (P_t - Q_{t-1})^{-1}$, (42) reduces to the first equation in the definition of an adaptive least squares estimator (see (27)). Equation (39) can then be rewritten as[42]

$$\begin{aligned}(P_t - Q_{t-1})^{-1} &= P_{t-1}^{-1} + \mathcal{X}_{t-1}'\Sigma_Z^{-1}\mathcal{X}_{t-1} \\ &= (I - P_{t-1}^{-1}Q_{t-2})(P_{t-1} - Q_{t-2})^{-1} + \mathcal{X}_{t-1}'\Sigma_Z^{-1}\mathcal{X}_{t-1},\end{aligned} \tag{43}$$

which reduces to (28) if $Q_{t-2} = (1-\gamma_{t-1})P_{t-1}$, for a sequence of scalars $0 < \gamma_t \leq 1$. Using (39) it follows that this condition is satisfied by choosing

$$Q_{t-1} = \frac{1-\gamma_t}{\gamma_t}\left(P_{t-1} - P_{t-1}\mathcal{X}_{t-1}'\Omega_{t-1}^{-1}\mathcal{X}_{t-1}P_{t-1}\right).$$

From this expression it also follows that Q_{t-1} is measurable with respect to Z_1^{t-1} as long as γ_t is measurable. We can summarize the preceding calculations as:

$$R_t\hat{\psi}_t^{\mathbb{P}} = \gamma_{t-1}R_{t-1}\hat{\psi}_{t-1}^{\mathbb{P}} + \mathcal{X}_{t-1}'\Sigma_Z^{-1}\mathcal{Y}_t, \tag{44}$$

$$R_t = \gamma_{t-1}R_{t-1} + \mathcal{X}_{t-1}'\Sigma_Z^{-1}\mathcal{X}_{t-1}, \tag{45}$$

$$\hat{\psi}_t^{\mathbb{P}} = R_t^{-1}R_t\hat{\psi}_t^{\mathbb{P}}, \tag{46}$$

$$P_t = \frac{1}{\gamma_t}R_t^{-1}, \tag{47}$$

$$\Omega_{t-1} = \mathcal{X}_{t-1}P_{t-1}\mathcal{X}_{t-1}' + \Sigma_Z, \tag{48}$$

with log likelihood function given by (41). The constant gain estimator corresponds to the special case where $\gamma_t = \gamma$ for all t.

D Pricing Kernel

The pricing kernel can be expressed as

$$\mathcal{M}_{t,t+1} = e^{-r_t} \times \frac{f_{t,t+1}^{\mathbb{Q}}(\mathcal{P}_{t+1})}{f_{t,t+1}^{\mathbb{P}}(\mathcal{P}_{t+1})}.$$

[41] Substitute (40) into (39) and the resulting equation into (38).

[42] This expression is obtained by substituting (40) into (39), plugging the resulting equation back into (39), and multiplying by $(P_t - Q_{-1})^{-1}$ from the left and P_{t-1}^{-1} from the right.

Since the distributions are conditionally normal under both measures, they have equal support. Then, $\mathcal{M}_{t,t+1}$ defines a strictly positive pricing kernel. We can rewrite the conditional distributions as

$$\begin{aligned} f^{\mathbb{P}}_{t,t+1} &= N(\hat{K}^{\mathbb{P}}_{\mathcal{P}0,t} + [\hat{K}^{\mathbb{P}}_{\mathcal{PP},t}, \hat{K}^{\mathbb{P}}_{\mathcal{P}H,t}]Z_t, \Omega_{\mathcal{PP},t}) = N(\hat{\mu}^{\mathbb{P}}_t, \Omega_{\mathcal{PP},t}), \\ f^{\mathbb{Q}}_{t,t+1} &= N(K^{\mathbb{Q}}_0 + K^{\mathbb{Q}}_1\mathcal{P}_t, \Sigma_{\mathcal{PP}}) = N(\mu^{\mathbb{Q}}_t, \Sigma_{\mathcal{PP}}), \end{aligned}$$

where $(\hat{K}^{\mathbb{P}}_{\mathcal{P}0,t}, [\hat{K}^{\mathbb{P}}_{\mathcal{PP},t}, \hat{K}^{\mathbb{P}}_{\mathcal{P}H,t}])$ denote the posterior means of the latent parameters states, and $\Omega_{\mathcal{PP},t}$ the upper left 3×3 entries of the conditional covariance matrix Ω_t given in equation (48). We can reduce this expression as follows (we use the notation c_t to terms that are $\mathcal{F}_t$ measurable but not of direct interest):

$$\begin{aligned} \log \mathcal{M}_{t,t+1} + r_t &= c_t + \frac{1}{2}(\mathcal{P}_{t+1} - \hat{\mu}^{\mathbb{P}}_t)'\Omega^{-1}_{\mathcal{PP},t}(\mathcal{P}_{t+1} - \hat{\mu}^{\mathbb{P}}_t) - \frac{1}{2}(\mathcal{P}_{t+1} - \mu^{\mathbb{Q}}_t)'\Sigma^{-1}_{\mathcal{PP}}(\mathcal{P}_{t+1} - \mu^{\mathbb{Q}}_t) \\ &= c'_t - \left(\Omega^{-1}_{\mathcal{PP},t}\hat{\mu}^{\mathbb{P}}_t - \Sigma^{-1}_{\mathcal{PP}}\mu^{\mathbb{Q}}_t\right)' \mathcal{P}_{t+1} + \frac{1}{2}\mathcal{P}'_{t+1}(\Omega^{-1}_{\mathcal{PP},t} - \Sigma^{-1}_{\mathcal{PP}})\mathcal{P}_{t+1} \\ &= c''_t - \Lambda'_{\mathcal{P}t}\Gamma^{-1}_t\varepsilon^{\mathbb{P}}_{t+1} + \frac{1}{2}(\varepsilon^{\mathbb{P}}_{t+1})'\left(I - \Gamma^{-1}_t\right)\varepsilon^{\mathbb{P}}_{t+1}, \end{aligned}$$

where

$$\begin{aligned} \Lambda_{\mathcal{P}t} &= \Omega^{-1/2}_{\mathcal{PP},t}(\hat{\mu}^{\mathbb{P}}_t - \mu^{\mathbb{Q}}_t) \\ \Gamma_t &= \Omega^{-1/2}_{\mathcal{PP},t}\Sigma_{\mathcal{PP}}(\Omega^{-1/2}_{\mathcal{PP},t})' \\ c''_t &= -\frac{1}{2}\log|\Gamma_t| - \frac{1}{2}\Lambda'_{\mathcal{P}t}\Gamma^{-1}_t\Lambda_{\mathcal{P}t} \end{aligned}$$

Thus the stochastic discount factor resembles a stochastic discount factor under full information, though with the parameters determining the market price of risks replaced by their posterior means, and with an additional stochastic convexity term and matrix Γ_t representing the change of conditional covariance matrix from $\mathbb{P}$ to $\mathbb{Q}$.

To show that $\Lambda_{\mathcal{P}t}$ is naturally interpreted as the market prices of risk in our learning setting, consider an asset with log total-return spanned by the factors $\mathcal{P}_t$: $r^a_t = \alpha + \beta'\mathcal{P}_t$ and satisfying $\mathbb{E}_t\left[e^{r^a_t+1}\mathcal{M}_{t,t+1}\right] = 1$. Using the fact that $\mathbb{E}[e^{\theta'\varepsilon + \frac{1}{2}\varepsilon'(I-\Gamma^{-1})\varepsilon}] = e^{\frac{1}{2}\theta'\Gamma\theta + \frac{1}{2}\log|\Gamma|}$, for $\varepsilon \sim N(0, I)$, the left-hand side of the last expression can be rewritten as

$$\begin{aligned} &\exp\{\alpha + \beta'\hat{\mu}^{\mathbb{P}}_t - r_t + c''_t + \frac{1}{2}(\beta'\Omega^{1/2}_t - \Lambda'_{\mathcal{P}t}\Gamma^{-1}_t)\Gamma_t(\beta'\Omega^{1/2}_t - \Lambda'_{\mathcal{P}t}\Gamma^{-1}_t)' + \frac{1}{2}\log|\Gamma_t|\} \\ &\quad = \exp\{\mathbb{E}_t[r^a_{t+1}] - r_t + \frac{1}{2}\beta'\Omega_t\beta - \beta'\Omega^{1/2}_t\Lambda_{\mathcal{P}t}\}. \end{aligned}$$

This leads to

$$\mathbb{E}_t[r^a_{t+1}] - r_t + \frac{1}{2}\mathbb{V}[r^a_{t+1}] = \beta' \Omega_t^{1/2} \Lambda_{\mathcal{P}t};$$

the expected log excess return equals the quantity of risk times the market price of risk (after adjusting for a convexity term).

E Selecting the Constant Gain Coefficient γ

Figure 12 shows RMSEs, based on the benchmark learning rule $\ell^L_{CG}(\mathcal{P})$, for the first principal component of US Treasury yields over the sample from January 1995 through December 2014. The minimal RMSE for one-year-ahead forecasts (Panel (b)) is achieved for $\gamma = 0.99$. The result is similar when for the one-quarter-ahead forecasts (Panel (a)).

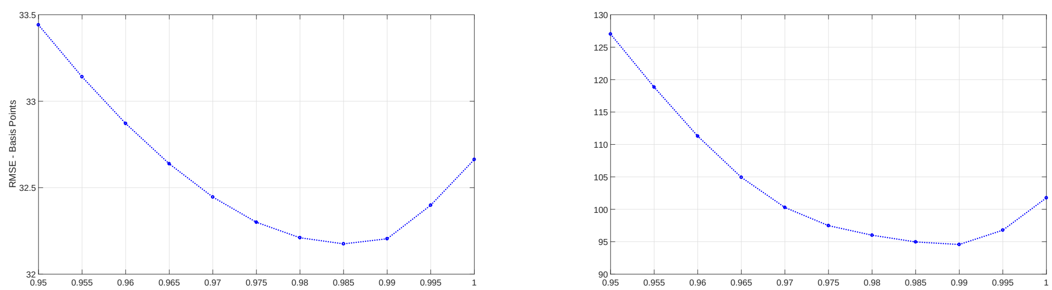

(a) Quarterly Forecasts

(b) Annual Forecasts

Figure 12: RMSE for one-quarter ahead and one-year ahead forecasts of $PC1$ of bond yields, over the sample from January 1995 through December 2014. The RMSEs are reported for different values of the constant gains coefficient γ.

This finding led us to wonder whether $\mathcal{RA}$ would find it optimal (from the perspective of accurate forecasting) to adjust her constant-gain parameter γ over time. Intuitively, a real-time learner might lower γ (give more weight to recent data) when recent forecast errors are large (i.e., when there is strong evidence for a structural change). To ascertain whether $\mathcal{RA}$'s forecasts are more accurate with a state-dependent γ, we estimated the learning model on an equally spaced grid of down-weight parameters $\{0.95, 0.955, ..., 0.995, 1\}$. The three- and twelve-month ahead forecasts for each value of γ in the grid are recorded for each month from January 1995 through December 2014. Using a look-back period of ten years we select the "optimal" γ to minimize the RMSE of the first PC. For example, in January 1995 we chose γ to minimize the out-of-sample RMSE between January 1985 (post training period) and January 1995 (using forecasts made between January 1985 and January 1995). We then step forward one month at a time, fixing the ten-year look-back window, and repeat this exercise[43]. Figure 13 shows the evolution of the dynamically updated γ parameters (γ_t) over the sample from January 1995 through December 2014. For one-year-ahead forecasts, the constant gains coefficient has a value of 0.99 for most of the sample (Panel (b)), and γ_t always remains between 0.98 and 1. The estimated γ_t is more volatile for one-quarter-ahead forecasts (Panel (a)), though γ_t takes values close to 0.99 for most of the sample.

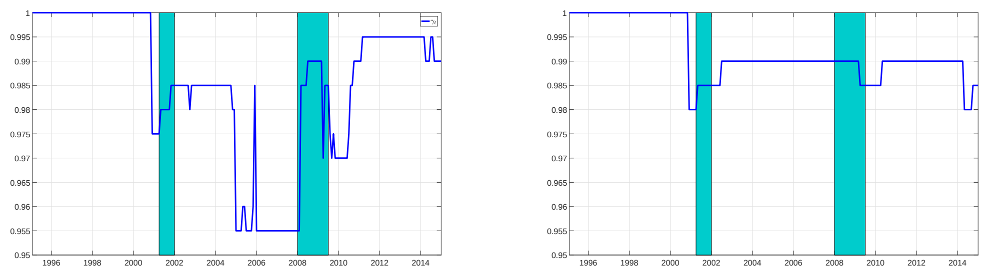

(a) Quarterly Forecasts

(b) Annual Forecasts

Figure 13: Real-time varying constant-gain parameters γ that minimize one quarter and one year ahead RMSE's of the first PC over the previous 10 years.

F Constraints on Λ_t for Rules $\ell_{CG}(\mathcal{P})$ and $\ell_{CG}(\mathcal{P}, H)$

We use the training sample to reduce the dimension of $\Theta^{\mathbb{P}}$. For models evaluated out of sample between January 1995 and March 2011, the training sample consists of the prior 10 years from January 1985 through December 1994. Our dimension reduction strategy is based on restricting the physical measure towards the risk neutral. First we estimate a model without restrictions imposed, and then we inspect the statistic significance of each of the parameters in $PmQ_t = \left(\begin{smallmatrix} \hat{K}^{\mathbb{P}}_{\mathcal{P}0,t} - \hat{K}^{\mathbb{Q}}_{\mathcal{P}0} & \hat{K}^{\mathbb{P}}_{\mathcal{P}\mathcal{P},t} - \hat{K}^{\mathbb{Q}}_{\mathcal{P}\mathcal{P}} & \hat{K}^{\mathbb{P}}_{\mathcal{P}H,t} \\ \hat{K}^{\mathbb{P}}_{H0,t} & \hat{K}^{\mathbb{P}}_{H\mathcal{P},t} & \hat{K}^{\mathbb{P}}_{HH,t} \end{smallmatrix}\right)$. If the p-value, induced by the posterior variance, at the end of the training sample is above 0.1, the corresponding coefficient in $K^{\mathbb{P}}_{Z,t}$ is concentrated out such that the corresponding entry in PmQ_t is zero. There are only two exceptions to this rule. First, we deem that the coefficient of the lagged second principal component in the second principal component equation plays an important role in capturing the persistence of the second PC. Thus, we leave it unrestricted even when the p-value is above 0.1. Second, we choose to restrict the market price of risk of the third principal component to be equal to zero. This is in line with what is found by Joslin, Priebsch,

[43]This is in the spirit of the adaptive step-size algorithm proposed by Kostyshyna (2012) that draws upon the engineering literature to adjust the gain parameter based on past forecast errors.

and Singleton (2014), and consistent with the idea that the third principal component is a spread portfolio that hedges away US Treasury bonds risks. In the data, we find that most of the coefficients in the equation of the third principal component are not significant, with the exception of the coefficient for the second principal component, which is borderline significant with 0.9 confidence. Table 9 displays the restrictions imposed on the autoregressive feedback matrix for rule $\ell_{CG}(\mathcal{P}, H)$. Similarly, Table 10 reports the restrictions for rule $\ell_{CG}(\mathcal{P})$.

	Λ_{0t}	Λ_{1t}				
		PC_1	PC_2	PC_3	$ID(y^{2y})$	$ID(y^{7y})$
PC_1	*	*	*	0	*	*
PC_2	*	0	*	*	0	0
PC_3	0	0	0	0	0	0
$ID(y^{2y})$	0	*	*	*	*	*
$ID(y^{7y})$	0	*	0	0	0	*

Table 9: Restrictions applied in rule $\ell_{CG}(\mathcal{P}, H)$ to the parameters in PmQ_t.

	const	PC_1	PC_2	PC_3
PC_1	*	*	*	0
PC_2	*	0	*	*
PC_3	0	0	0	0

Table 10: Restrictions applied in rule $\ell_{CG}(\mathcal{P})$ to the parameters in PmQ_t.

G Stochastic Volatility Model

Suppose that there exist a 3-dimensional state-variable, consisting of a univariate volatility factor V_t, and 2 conditionally Gaussian factors X_t. Following our specification of the Gaussian models with learning, we assume that the parameters governing the risk neutral measure are known and constant. Joslin and Le (2014) show that an econometrically exactly identified specification is given by

$$\begin{aligned} V_{t+1}|V_t &\sim CAR(\rho^{\mathbb{Q}}, c^{\mathbb{Q}}, v^{\mathbb{Q}}), \\ X_{t+1} &= K^{\mathbb{Q}}_{XV} V_t + J(\lambda^{\mathbb{Q}}) X_t + \sqrt{\Sigma_0 + \Sigma_1 V_t} \cdot \varepsilon^{\mathbb{Q}}_t, \\ r_t &= r^{\mathbb{Q}}_{\infty} + \rho_V V_t + 1' X_t, \end{aligned}$$

where CAR is short for the compound autoregressive gamma process. The CAR process has a conditional Laplace transform that is exponentially affine and first and second moments given by

$$\begin{aligned}
\log \mathbb{E}^{\mathbb{Q}}(e^{uV_{t+1}}|V_t) &= -v^{\mathbb{Q}} \log(1 - uc^{\mathbb{Q}}) + \frac{\rho^{\mathbb{Q}} u}{1 - uc^{\mathbb{Q}}} V_t, \\
\mathbb{E}_t^{\mathbb{Q}}(V_{t+1}|V_t) &= v^{\mathbb{Q}} c^{\mathbb{Q}} + \rho^{\mathbb{Q}} V_t, \\
\mathbb{V}_t^{\mathbb{Q}}(V_{t+1}|V_t) &= v^{\mathbb{Q}^2} c^{\mathbb{Q}} + 2\rho^{\mathbb{Q}} V_t.
\end{aligned}$$

The innovation to the non-volatility factors, $\varepsilon_{t+1}^{\mathbb{Q}}$, is assumed to be normally distributed and independent of V_{t+1}. It follows that zero coupon bond prices are exponentially affine, $D_t^n = e^{A_n + B_{n,V} V_t + B_{n,X} X_t}$, with loadings that satisfy the recursions

$$\begin{aligned}
A_{n+1} &= A_n + \frac{1}{2} B'_{n,X} \Sigma_0 B_{n,X} - v^{\mathbb{Q}} \log\left(1 - c^{\mathbb{Q}} B_{n,V}\right) - r_{\infty}^{\mathbb{Q}}, \\
B_{n+1,X} &= J(\lambda^{\mathbb{Q}})' B_{n,X} - 1, \\
B_{n+1,V} &= B'_{n,X} K_{XV} + \frac{1}{2} B'_{n,X} \Sigma_1 B_{n,X} + \frac{\rho^{\mathbb{Q}} B_{n,V}}{1 - c^{\mathbb{Q}} B_{n,V}} - \rho_V.
\end{aligned}$$

Under the physical measure we assume that parameters that govern the dynamics of the volatility factor are known and constant, while the parameters that govern the conditional Gaussian factors are drifting and unknown

$$V_{t+1}|V_t \sim CAR(\rho^{\mathbb{P}}, c^{\mathbb{P}}, v^{\mathbb{P}}), \tag{49}$$

$$X_{t+1} = K_{X0,t}^{\mathbb{P}} + K_{XV,t}^{\mathbb{P}} V_t + K_{XX,t}^{\mathbb{P}} X_t + \sqrt{\Sigma_{0X} + \Sigma_{1X} V_t} \cdot \varepsilon_t^{\mathbb{P}}. \tag{50}$$

As yields are affine in the state-variables,

$$y_t = A(\Theta^{\mathbb{Q}}, \Sigma_{0X}, \Sigma_{1X}) + B_V(\Theta^{\mathbb{Q}}, \Sigma_{0X}, \Sigma_{1X}) V_t + B_X(\lambda^{\mathbb{Q}}) X_t,$$

the principal components $\mathcal{P}$ are also affine in the state, since $\mathcal{P}_t = W y_t$. This in turn implies that V_t can be written as an affine function of f_t:

$$V_t = \alpha(\Theta^{\mathbb{Q}}, \Sigma_{0X}, \Sigma_{1X}) + \beta(\Theta^{\mathbb{Q}}, \Sigma_{0X}, \Sigma_{1X})' \mathcal{P}_t.$$

Joslin and Le (2014) show that we can rewrite and reparameterize equation (50) with

$$\mathcal{P}_{t+1}^{2:3} - W^{2:3} B_V V_{t+1} = \tilde{K}_{\mathcal{P}0,t}^{\mathbb{P}} + \tilde{K}_{\mathcal{P}V,t}^{\mathbb{P}} V_t + \tilde{K}_{\mathcal{P}\mathcal{P},t}^{\mathbb{P}} \mathcal{P}_t^{2:3} + \sqrt{\tilde{\Sigma}_{0\mathcal{P}} + \tilde{\Sigma}_{1\mathcal{P}} V_t} \cdot \varepsilon_t^{\mathbb{P}}, \tag{51}$$

where the superscripts 2 : 3 refer to the second and third PCs, and the tilde is used to indicate that these are parameters governing the dynamics of $(V_t, (\mathcal{P}_t^{2:3})')'$ (and not $\mathcal{P}_t$). Therefore, the model's parameters can be decomposed into constant and known $\mathbb{Q}$-parameters $(r_\infty^{\mathbb{Q}}, \rho_V, \rho^{\mathbb{Q}}, c^{\mathbb{Q}}, v^{\mathbb{Q}}, K_{XV}^{\mathbb{Q}}, \lambda^{\mathbb{Q}})$, constant and known covariance matrices $(\tilde{\Sigma}_{\mathcal{P}0}, \tilde{\Sigma}_{1\mathcal{P}})$, constant and known $\mathbb{P}$-parameters $(\rho^{\mathbb{P}}, c^{\mathbb{P}}, v^{\mathbb{P}})$, and unknown drifting $\mathbb{P}$ parameters $(\tilde{K}_{\mathcal{P}0,t}^{\mathbb{P}}, \tilde{K}_{\mathcal{P}V,t}^{\mathbb{P}}, \tilde{K}_{\mathcal{P}\mathcal{P},t}^{\mathbb{P}})$. We impose that $c^{\mathbb{P}} = c^{\mathbb{Q}}$ and $v^{\mathbb{P}} = v^{\mathbb{Q}}$. These two conditions guarantee diffusion invariance of V_t, and that the market prices of risks are non-exploding in the continuous time limit (see Joslin and Le (2014)). From equations (49) - (51) it is seen that the conditional first and second moments of the principal principal components are given by

$$\mathbb{E}_t^{\mathbb{P}}(\mathcal{P}_{t+1}) = K_{0,t}^{\mathbb{P}} + K_{1,t}^{\mathbb{P}}\mathcal{P}_t \qquad (52)$$

$$\mathbb{V}_t^{\mathbb{P}}(\mathcal{P}_{t+1}) = \Sigma_0^{\mathbb{P}} + \Sigma_1^{\mathbb{P}}\mathcal{P}_t \qquad (53)$$

where $(K_{0,t}^{\mathbb{P}}, K_{1,t}^{\mathbb{P}}, \Sigma_0^{\mathbb{P}}, \Sigma_1^{\mathbb{P}})$ are known functions of $(\Theta^{\mathbb{Q}}, \tilde{\Sigma}_{\mathcal{P}0}, \tilde{\Sigma}_{1\mathcal{P}}, \Theta_t^{\mathbb{P}})$ induced by rotating $(V_t, (\mathcal{P}_t^{2:3})')'$ to $\mathcal{P}_t$. Similarly to what done for the Gaussian learning model, we impose restrictions on $[K_{0,t}^{\mathbb{P}}, K_{1,t}^{\mathbb{P}}]$ based on a training sample. These restrictions can be written as $vec\left([\tilde{K}_{\mathcal{P}0,t}^{\mathbb{P}}, \tilde{K}_{\mathcal{P}V,t}^{\mathbb{P}}, \tilde{K}_{\mathcal{P}\mathcal{P},t}^{\mathbb{P}}]\right) = R\psi_t + q$, where ψ evolves according to a random walk

$$\psi_t = \psi_{t-1} + Q_{t-1}^{1/2}\eta_t.$$

A set of sufficient conditions that guarantees that the innovation co-variance matrix of ψ_t is proportional to the posterior co-variance matrix will ensure that the posterior means of ψ is given by a constant gain estimator. The proof is similar to the derivations discussed in the paper for the Gaussian learning model.

H Robustness to Time-Varying Volatility

A less constrained Bayesian (relative to $\mathcal{RA}$) would formally build updating of $\Sigma_{\mathcal{P}\mathcal{P}}$ into her learning rule. A priori, we would not expect this generalization of our learning rules to materially affect $\mathcal{RA}$'s conditional forecasts of bond yields, our primary focus for modeling risk premiums. Updating of $\Sigma_{\mathcal{P}\mathcal{P}}$ would only change the posterior conditional means indirectly through interactions with $\Theta^{\mathbb{Q}}$, passed onto the $\mathbb{P}$-feedback parameters by the restrictions on the market price of risk. In our current setting $\mathcal{RA}$ keeps the $\mathbb{Q}$ parameters $(k_\infty^{\mathbb{Q}}, \lambda^{\mathbb{Q}})$ nearly constant. Therefore, it seems unlikely that formally introducing learning about $\Sigma_{\mathcal{P}\mathcal{P}}$ would lead to large changes in the inferred posterior conditional $\mathbb{P}$-means of bond yields.

To provide further reassurance on this front, we proceed to investigate learning within a setting of stochastic volatility. Suppose there are three risk factors consisting of a univariate volatility factor V_t and a bivariate X_t that is Gaussian conditional on V_t. We adopt the following normalized just-identified representation of the state under $\mathbb{Q}$:

$$V_{t+1}|V_t \sim CAR(\rho^{\mathbb{Q}}, c^{\mathbb{Q}}, v^{\mathbb{Q}}), \tag{54}$$

$$X_{t+1} = K_V^{\mathbb{Q}} V_t + diag(\lambda^{\mathbb{Q}}) X_t + \sqrt{\Sigma_{0X} + \Sigma_{1X} V_t}\, \varepsilon_t^{\mathbb{Q}}, \tag{55}$$

$$r_t = r_\infty^{\mathbb{Q}} + \rho_V V_t + 1' X_t, \tag{56}$$

where CAR denotes a compound autoregressive gamma process (Gourieroux and Jasiak (2006)) and $\Theta^{\mathbb{Q}} \equiv (r_\infty^{\mathbb{Q}}, \rho_V, \rho^{\mathbb{Q}}, c^{\mathbb{Q}}, v^{\mathbb{Q}}, K_V^{\mathbb{Q}}, \lambda^{\mathbb{Q}})$. As before, we assume that $\mathcal{RA}$ treats $\Theta^{\mathbb{Q}}$ as constant and known, which implies that yields are given by

$$y_t = A(\Theta^{\mathbb{Q}}, \Sigma_{0X}, \Sigma_{1X}) + B_V(\Theta^{\mathbb{Q}}, \Sigma_{1X}) V_t + B_X(\Theta^{\mathbb{Q}}) X_t,$$

and the principal components are affine in (V_t, X_t) (see Appendix G for details). The market prices of risk are assumed to be such that, under $\mathbb{P}$, the state follows the process

$$V_{t+1}|V_t \sim CAR(\rho^{\mathbb{P}}, c^{\mathbb{P}}, v^{\mathbb{P}}), \tag{57}$$

$$X_{t+1} = K_{0t}^{\mathbb{P}} + K_{Vt}^{\mathbb{P}} V_t + K_{Xt}^{\mathbb{P}} X_t + \sqrt{\Sigma_{0X} + \Sigma_{1X} V_t}\, \varepsilon_{t+1}^{\mathbb{P}}, \tag{58}$$

where $\varepsilon_{t+1}^{\mathbb{P}}$ is independent of V_{t+1} and we let $\Theta_t^{\mathbb{P}} = (\rho^{\mathbb{P}}, c^{\mathbb{P}}, v^{\mathbb{P}}, K_{0t}^{\mathbb{P}}, K_{Vt}^{\mathbb{P}}, K_{Xt}^{\mathbb{P}})$. $\mathcal{RA}$ presumes that the volatility parameters $(\rho^{\mathbb{Q}}, c^{\mathbb{Q}}, v^{\mathbb{Q}}, \Sigma_{0X}, \Sigma_{1X})$ are constant, while those governing the conditional means of X_t are unknown and drifting. In Appendix G we show that the conditional first moments of the principal components are given by

$$\mathbb{E}_t^{\mathbb{Q}}(\mathcal{P}_{t+1}) = K_{0\mathcal{P}}^{\mathbb{Q}} + K_{1\mathcal{P}}^{\mathbb{Q}} \mathcal{P}_t \quad \text{and} \quad \mathbb{E}_t^{\mathbb{P}}(\mathcal{P}_{t+1}) = K_{0\mathcal{P},t}^{\mathbb{P}} + K_{1\mathcal{P},t}^{\mathbb{P}} \mathcal{P}_t,$$

where $(K_{0\mathcal{P}}^{\mathbb{Q}}, K_{1\mathcal{P}}^{\mathbb{Q}}, K_{0\mathcal{P},t}^{\mathbb{P}}, K_{1\mathcal{P},t}^{\mathbb{P}})$ are known functions of $(\Theta^{\mathbb{Q}}, \Sigma_{X0}, \Sigma_{1X}, \Theta_t^{\mathbb{P}})$ from the rotation of $(V_t, X_t')'$ to $\mathcal{P}_t$. As before, a subset of the parameters in $[K_{0\mathcal{P}t}^{\mathbb{P}}, K_{1\mathcal{P}t}^{\mathbb{P}}]$ is constrained based on the training sample.

Figure 14 plots the eigenvalues of the feedback matrices $K_{1\mathcal{P}}^{\mathbb{Q}}$ and $K_{1\mathcal{P},t}^{\mathbb{P}}$ from the perspective of $\mathcal{RA}$'s real-time learning rule in the presence of V_t and conditioning only on the history of the PCs. The eigenvalues of $K_{1\mathcal{P}}^{\mathbb{Q}}$ are $(\rho^{\mathbb{Q}}, \lambda^{\mathbb{Q}})$ and the eigenvalues of $K_{1\mathcal{P},t}^{\mathbb{P}}$ are $(\rho^{\mathbb{P}}, eig(K_{Xt}^{\mathbb{P}}))$.[44] Relaxing the assumption of constant conditional volatility does not alter our prior finding that the $\mathbb{Q}$ eigenvalues are nearly constant over the entire sample period. The variation in the

[44] The feedback matrices in the conditional first moments of $\mathcal{P}_t$ and $(V_t, X_t')'$ will have equal eigenvalues, as $\mathcal{P}_t$ is an affine function of $(V_t, X_t')'$.

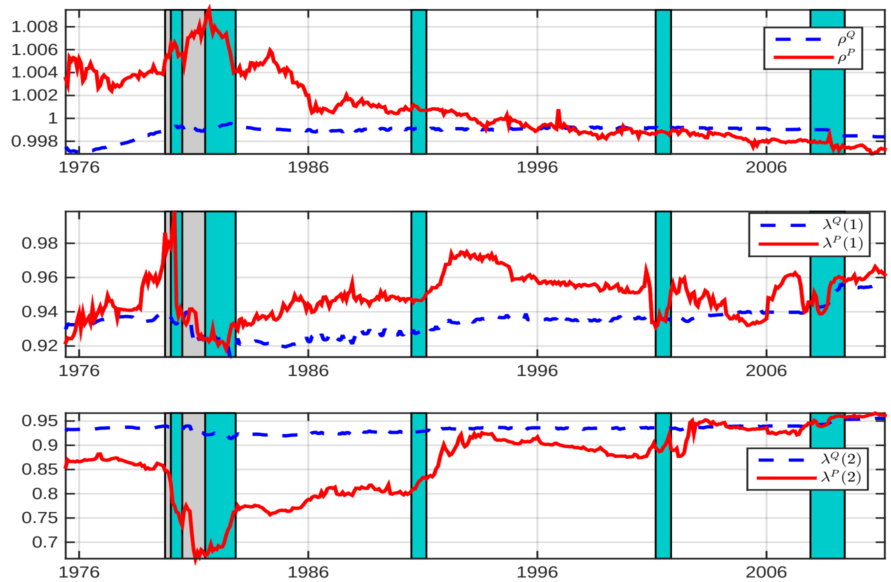

Figure 14: Estimates from model $\ell_{CG}^{1,3}(\mathcal{P})$ of the eigenvalues of the feedback matrix $K_{1\mathcal{P}}^{\mathbb{Q}}$ ($K_{1\mathcal{P},t}^{\mathbb{P}}$). The eigenvalues of $K_{1\mathcal{P}}^{\mathbb{Q}}$ are ($\rho^{\mathbb{Q}}, \lambda^{\mathbb{Q}}$) and the eigenvalues of $K_{1\mathcal{P},t}^{\mathbb{P}}$ are ($\rho^{\mathbb{P}}, eig(K_{Xt}^{\mathbb{P}})$).

eigenvalues of $K_{1\mathcal{P},t}^{\mathbb{P}}$ reflects the substantial variation in the market prices of risk.

Figure 15 offers an interesting perspective on the degree to which the learning rule $\ell_{CG}^{L}(\mathcal{P})$ (that presumes constant $\Sigma_{\mathcal{P}\mathcal{P}}$) captures the swings in the conditional covariance matrix that would be perceived by an agent learning in the presence of stochastic volatility. On the diagonal are the estimated conditional standard deviations from models both with and without stochastic volatility. Rule $\ell_{CG}^{L}(\mathcal{P})$ captures the overall evolution of the conditional standard deviations, but fails to pick up the huge increment in volatilities during the Fed experiment. Perceptions about volatility under $\ell_{CG}^{L}(\mathcal{P})$ also decay relatively slowly during the great moderation. The constant conditional correlations are updated by $\ell_{CG}^{L}(\mathcal{P})$ in a manner very similar to the learning rule for the stochastic volatility model.

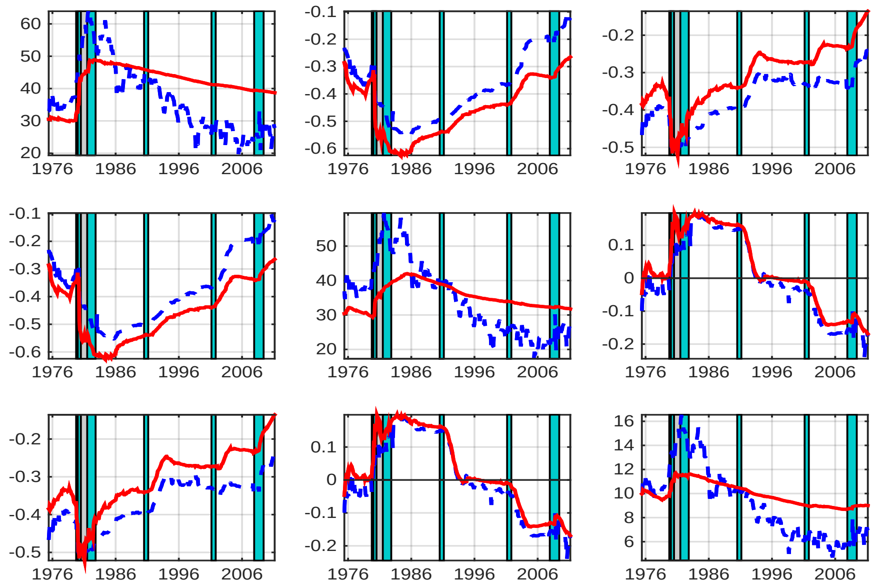

Figure 15: Summary of $\Sigma_{\mathcal{P}\mathcal{P}}$. Conditional standard deviations (main diagonal elements) and correlations (off-diagonal elements) estimates from learning models with (blue line) and without (red line) stochastic volatility. The estimates at date t are based on the historical data up to observation t, over the period July, 1975 to March, 2011.

References

Altavilla, C., R. Giacomini, and G. Ragusa, 2014, Anchoring the Yield Curve Using Survey Expectations, Working paper, Working Paper No 1632, European Central Bank.

Andrade, P., R. Crump, S. Eusepi, and E. Moench, 2014, Fundamental Disagreement, Working paper, Federal Reserve Bank of New York Staff Report No. 655.

Atmaz, A., and S. Basak, 2017, Belief Dispersion in the Stock Market, *Journal of Finance*.

Barillas, F., and K. Nimark, 2014, Speculation, Risk Premia and Expectations in the Yield Curve, Working paper, Cornell University.

Basak, S., 2005, Asset Pricing with Heterogeneous Beliefs, *Journal of Banking and Finance* 29, 2849–2881.

Bauer, M., and J. Hamilton, 2018, Robust Bond Risk Premia, *Review of Financial Studies* 31, 399–448.

Buraschi, A., and P. Whelan, 2016, Speculation, Hedging, and Interest Rates, Working paper, Imperial College.

Christensen, J., and G. Rudebusch, 2015, Estimating Shadow-Rate Term Structure Models with Near-Zero Yields, *Journal of Financial Econometrics* 13, 226–259.

Chun, A., 2011, Expectations, Bond Yields and Monetary Policy, *Review of Financial Studies* 24, 208–247.

Cieslak, A., 2017, Short-rate Expectations and Unexpected Returns in Treasury Bonds, Working paper, forthcoming, *Review of Financial Studies*.

Cochrane, J., and M. Piazzesi, 2002, The Fed and Interest Rates: A High-Frequency Identification, *American Economic Review P&P* 92, 90–95.

Cochrane, J., and M. Piazzesi, 2008, The Fed and Interest Rates: A High-Frequency Identification, *Working Paper*.

Cogley, T., and T. Sargent, 2008, Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision-Making, *International Economic Review* 49, 185–221.

Collin-Dufresne, P., M. Johannes, and L. Lochstoer, 2016, Parameter Learning in General Equilibrium: The Asset Pricing Implications, *American Economic Review* 106, 664 – 698.

Dai, Q., and K. Singleton, 2000, Specification Analysis of Affine Term Structure Models, *Journal of Finance* 55, 1943–1978.

Dalang, R., A. Morton, and W. Willinger, 1990, Equivalent Martingale Measures and No-arbitrage in Stochastic Securities Market Models, *Stochastics Stochastics Rep. 29* 2, 185–201.

David, A., 2008, Heterogeneous Beliefs, Speculation, and the Equity Premium, *Journal of Finance* 63, 41–83.

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

Dovern, J., U. Fritsche, and J. Slacalek, 2012, Disagreement Among Forecasters in G7 Countries, *Review of Economics and Statistics* 94, 1081–1096.

Duffee, G., 2002, Term Premia and Interest Rate Forecasts in Affine Models, *Journal of Finance* 57, 405–443.

Duffee, G., 2011, Forecasting with the Term Structure: the Role of No-Arbitrage, Working paper, Johns Hopkins University.

Ehling, P., M. Gallmeyer, C. Heyerdahl-Larsen, and P. Illeditsch, 2016, Disagreement about Inflation and the Yield Curve, Working paper, University of Virginia.

Fleming, M., and E. Remolona, 1999, Price Formation and Liquidity in the US Treasury Market: The Response to Public Information, *Journal of Finance* 54, 1901–1915.

Ghysels, E., C. Horan, and E. Moench, 2014, Forecasting though the Rear-View Mirror: Data Revisions and Bond Return Predictability, Working paper, FRB of New York Staff Report No. 581.

Gourieroux, C., and J. Jasiak, 2006, Autoregressive Gamma Processes, *Journal of Forecasting* 25, 129–152.

Harvey, D., S. Leybourne, and P. Newbold, 1997, Testing the Equality of Prediction Mean-Squared Errors, *International Journal of Forecasting* 13, 281–291.

Johannes, M., L. Lochstoer, and Y. Mou, 2016, Learning about Consumption Dynamics, *Journal of Finance* 72, 551–600.

Joslin, S., and A. Le, 2014, Interest Rate Volatility and No Arbitrage Affine Term Structure Models, Working paper, Pennsylvania State University.

Joslin, S., A. Le, and K. Singleton, 2013, Why Gaussian Macro-Finance Term Structure Models are (Nearly) Unconstrained Factor-VARs, *Journal of Financial Economics* 109, 604–622.

Joslin, S., M. Priebsch, and K. Singleton, 2014, Risk Premiums in Dynamic Term Structure Models with Unspanned Macro Risks, *Journal of Finance* LXIX, 1197–1233.

Joslin, S., K. Singleton, and H. Zhu, 2011, A New Perspective on Gaussian Dynamic Term Structure Models, *Review of Financial Studies* 24, 926–970.

Kim, D., and A. Orphanides, 2012, Term Structure Estimation with Survey Data on Interest Rate Forecasts, *Journal of Financial and Quantitative Analysis* 47, 241–272.

Kim, D., and K. Singleton, 2012, Term Structure Models and the Zero Bound: An Empirical Investigation of Japanese Yields, *Journal of Econometrics* 170, 32–49.

Kostyshyna, O., 2012, Application of an Adaptive Step-Size Algorithm in Models of Hyperinflation, *Macroeconomic Dynamics* 16, 355–375.

Kreps, D., 1998, Anticipated Utility and Dynamic Choice, in D. Jacobs, E. Kalai, and M Kamien, eds.: *Frontiers in Research in Economic Theory* (Cambridge: Cambridge University Press,).

Le, A., and K. Singleton, 2012, A Robust Analysis of the Risk-Structure of Equilibrium Term Structures of Bond Yields, Working paper, University of North Carolina.

Ludvigson, S., and S. Ng, 2010, Macro Factors in Bond Risk Premia, *Review of Financial Studies* 22, 5027–5067.

Malmendier, U., D. Pouzo, and V. Vanasco, 2017, A Theory of Experience Effects, Working paper, UC Berkeley.

McCulloch, J., 2007, The Kalman Foundations of Adaptive Least Squares, With Applications to U.S. Inflation, Working paper, Ohio State University.

Patton, A., and A. Timmerman, 2010, Why do Forecasters Disagree? Lessons from the Term Structure of Cross-Sectional Dispersion, *Journal of Monetary Economics* 57, 803–820.

Piazzesi, M., 2005, Bond Yields and the Federal Reserve, *Journal of Political Economy* 113, 311–344.

Piazzesi, M., J. Salomao, and M. Schneider, 2013, Trend and Cycle in Bond Premia, Working paper, Stanford University.

Rudebusch, G., and J. Williams, 2009, Forecasting Recessions: The Puzzle of the Enduring Power of the Yield Curve, *Journal of Business and Economic Statistics* 27, 492–503.

Singleton, K., 1987, Asset Prices in a Time Series Model with Disparately Informed, Competitive Traders, in W. Burnett, and K. Singleton, eds.: *New Approaches to Monetary Economics* (Cambridge University Press,).

Stambaugh, R., 1999, Predictive Regressions, *Journal of Financial Economics* 54, 375–421.

Townsend, R., 1983, Forecasting the Forecasts of Others, *Journal of Political Economy* 91, 546–588.

Vasicek, O., 1977, An Equilibrium Characterization of the Term Structure, *Journal of Financial Economics* 5, 177–188.

Xiong, W., and H. Yan, 2009, Heterogeneous Expectations and Bond Markets, *Review of Financial Studies.*