UNIL | Université de Lausanne
Département de finance

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

FINANCE RESEARCH SEMINAR SUPPORTED BY UNIGESTION

Learning From Disagreement in the U.S. Treasury Bond Market

Kenneth J. SINGLETON,
Stanford University

Friday, June 14, 2019, 10:30-12:00
Room 126, Extranef building at the University of Lausanne

UNIGESTION

Learning From Disagreement in the U.S. Treasury Bond Market*

Marco Giacoletti Kristoffer T. Laursen Kenneth J. Singleton †

June 6, 2019

Abstract

We study the evolution of risk premiums on US Treasury bonds from the perspective of a *real-time Bayesian learner* $\mathcal{BL}$ who conditions her beliefs on measures of disagreement among professional forecasters about the future paths of bond yields. Learning about historical yields and disagreement within a dynamic term structure model leads to substantial variation in $\mathcal{BL}$'s subjective expected excess returns on bonds. The informativeness of disagreement is shows to be distinct from the (much weaker) forecasting power of inflation and output growth. Rather, it appears to reflect policy uncertainty and, in particular, uncertainty about fiscal policy. $\mathcal{BL}$'s learning rule substantially outperforms 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 learning about changing policy, regulatory or political environments. However, most of the literature that studies risk compensation in bond markets has followed and Dai and Singleton (2000) and Duffee (2002) in focusing on *retrospective within-sample* measures of risk premiums, and in presuming that the parameters governing the evolution of the risk factors in the economy are fixed and known by investors.

That learning can materially impact measured risk premiums is illustrated in Figure 1.[1] For a one-quarter holding period on a ten-year Treasury zero-coupon bond, the within-sample, least-squares forecasts of (annualized) excess returns and the forecasts from a real-time, recursive least-squares learning model differ by as much as 4%; analogous differences for the one-year horizon exceed 5%.[2] Moreover, the mechanical end-of-sample convergence of the recursive to the full-sample least-squares estimates is quite slow.

In this paper we explore in depth the structure and complexity of *ex ante* risk premiums in US Treasury markets through the lens of a Bayesian econometrician– referenced as $\mathcal{BL}$– who is learning in real-time about the distribution of future bond yields. As an "outside observer" of bond markets, $\mathcal{BL}$ approaches her learning problem with several insights in hand, including: (i) bond yields are well described by a low-dimensional factor model with the principal components (PCs) of yields as factors; (ii) the cross-sectional covariance structure of yields in an arbitrage-free market essentially reveals the factor loadings to investors; and and (iii) "market prices" of the factor risks vary over the business cycle.

Most distinctively, relative to extant empirical learning models in bond markets, $\mathcal{BL}$ believes that the observed differences of opinion among market participants about future yields influence the actual (objective) conditional distributions of realized yields. She learns from disagreement by conditioning on the cross-sectional dispersion of beliefs across market participants when forecasting future realized excess returns in bond markets. Atmaz and Basak (2017) show that, in markets where there are a large number of participants (as in Treasury markets), this measure of disagreement may serve as sufficient statistic for the impact of disagreement on equilibrium asset prices.

[1] This exercise reinforces prior evidence on learning and yield forecasts. See, e.g., Laubach, Tetlow, and Williams (2007), Dewachter and Lyrio (2008), and Gargano, Pettenuzzo, and Timmerman (2018).

[2] For the learning model, forecasts are the fitted values from recursive least-squares projections of the realized excess return $xr^n_{t+0.25y}$ for an n-period bond over a one-quarter horizon onto the first three principal components of bond yields $\mathcal{P}_t$:

$$xr^n_{t+0.25y} = \alpha_{n,t} + \mathcal{B}_{n\mathcal{P},t}\mathcal{P}_t + \sigma_v v_{t+h}.$$

Figure 1 compares these "expanding-window" estimates to their full-sample counterparts.

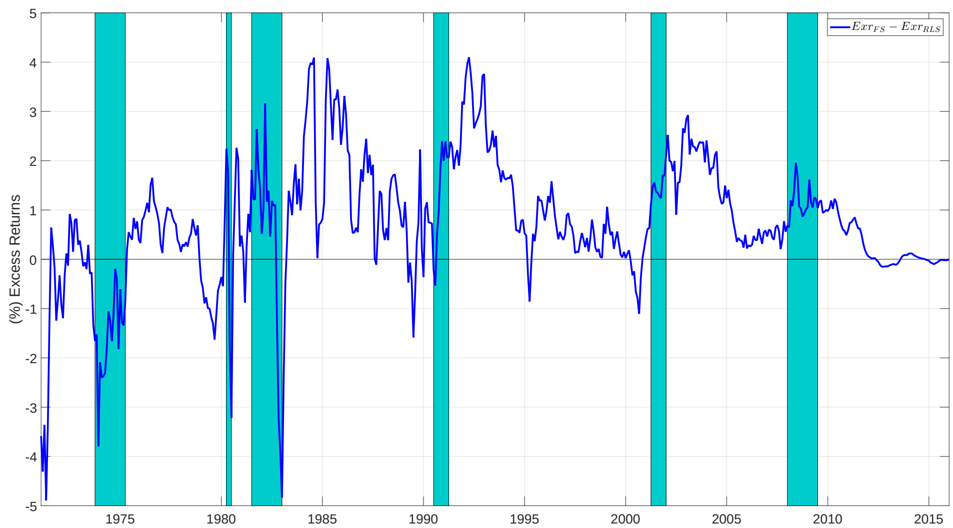

Figure 1: The three-month expected excess returns for a 10-year zero-coupon bond implied by the full-sample estimates minus those from a recursive least-squares learning scheme. The sample is monthly from January 1972 to December 2014. NBER recessions are shaded.

$\mathcal{BL}$'s learning rule is based on a Gaussian Dynamic Term Structure Model ($DTSM$),[3] under the presumption that the parameters of the historical distribution of yields may shift with economic conditions, and she adopts a Bayesian rule for updating her subjective beliefs about these parameters. Cross-sectional disagreement about yields potentially impacts $\mathcal{BL}$'s learning rule through two channels: her market prices of factor risks may be depend directly on measured disagreement; and it may inform in real time how she updates the parameters governing the time-series dynamics of the yield PCs.

Empirically, the fitted market prices of risk from $\mathcal{BL}$'s learning rule differ systematically from those implied by the full-sample analysis of the fixed-parameter version of this $DTSM$. In particular, $\mathcal{BL}$'s learning rule generates risk premiums that are more sensitive to shocks to the level of the yield curve, especially at turning points of business cycles. Moreover, building on feature (ii), we find the striking result that $\mathcal{BL}$ holds the factor loadings ("hedge ratios") nearly constant over the past twenty-five years. In contrast, $\mathcal{BL}$ makes substantial revisions to the parameters governing the conditional distribution of the PCs of bond yields (i.e., her

[3] A Gaussian framework offers substantial flexibility in modeling the conditional first moments of bond yields which determine risk premiums. Now, in fact, $\mathcal{BL}$ updates her estimates of the conditional covariance matrix of the yield PCs monthly as new information arrives (see below). Both our descriptive evidence and analysis of an extended $DTSM$ with stochastic volatility in Appendix E suggest that our central findings about risk premiums under $\mathcal{BL}$'s learning scheme are robust to the presence of stochastic volatility.

views on the future paths of yields).

Investor disagreement– proxied by the cross-sectional dispersion in forecasts from the Blue Chip Financial Forecasters (BCFF) panel– does indeed have substantial *real-time* predictive power for future yields, over and above the history of bond yields. $\mathcal{BL}$'s forecasts are substantially more accurate than the consensus BCFF forecasts of yields for both one- and four-quarter ahead horizons, especially coming out of recessions.[4] As such, $\mathcal{BL}$'s *ex ante* term premiums are notably different than those implied by a learning rule that conditions only on bond yields or those implied by the BCFF consensus beliefs.

Moreover, once $\mathcal{BL}$ conditions on disagreement, many of the macroeconomic variables that previous researchers have found to have predictive power for expected excess returns (e.g., Ludvigson and Ng (2010), and Joslin, Priebsch, and Singleton (2014), JPS) are largely rendered redundant.[5] Digging deeper, we provide evidence that the predictive power of yield disagreement within $\mathcal{BL}$'s learning rule has roots in policy uncertainty and, in particular, to uncertainty about future government spending. In Section 6 we discuss ways these findings might inform future development of equilibrium models in which market participants disagree about the conditional distributions of bond yields.

2 Disagreement in US Treasury Markets and Bond Yields

From various theoretical perspectives, a role for disagreement in the determination of equilibrium bond prices has been well established.[6] From $\mathcal{BL}$'s perspective, the relevant risk factors are the known yield PC's. For given her presumed factor structure under (i) above the *theoretical* PC's span any underlying macro factors that are relevant for pricing bonds. Her learning problem is that of determining the values of the parameters governing the evolution of these PC's over time. Prior to formalizing this learning problem, it is instructive to examine some of the descriptive statistics of the joint distribution of yields and measured disagreement.

Throughout this analysis we focus on the U.S. Treasury zero-coupon bond yields of maturities 6 months, and 1, 2, 3, 5, 7, and 10 years. For a matching set of bond maturities, we construct empirical counterparts to disagreement using the BCFF survey of yield forecasts over the period from January, 1985 through December, 2015, with the start date determined

[4]Thus, disciplining $\mathcal{BL}$'s learning rule with survey information would likely lead to a material deterioration in forecast accuracy. Kim and Orphanides (2012), Chun (2011), and Piazzesi, Salomao, and Schneider (2013) use survey data directly in the estimation of *DTSMs* in which the median forecaster has full knowledge of risk-factor dynamics and her forecasts are spanned by the low-order PCs of bond yields. We show subsequently that roughly 50% of the variation in consensus BCFF forecasts is not spanned by the PCs of bond yields.

[5]These earlier studies did not explore real-time learning; they were full-sample analyses.

[6]One illustrative set of models are those of David (2008), Xiong and Yan (2009), and Buraschi and Whelan (2016) in which the pairwise relative beliefs across heterogeneous agent types impact bond prices as "agree to disagree" about the values of an unobserved risk factor that influences aggregate consumption growth.

	Panel (a): Bond Yields			
	0.25y	0.5y	0.75y	1y
$ID(y^{6m})$	60.31%	69.54%	66.76%	64.23%
$ID(y^{2y})$	56.37%	61.83%	59.84%	57.41%
$ID(y^{5y})$	43.39%	54.32%	55.54%	54.66%
$ID(y^{7y})$	45.58%	55.76%	58.31%	59.13%
$ID(y^{10y})$	41.21%	54.03%	55.64%	56.49%
	Panel (b): Bond Yields, $ID(CPI)$, $ID(RGDP)$			
	0.25y	0.5y	0.75y	1y
$ID(y^{6m})$	66.52%	76.04%	76.36%	74.13%
$ID(y^{2y})$	60.42%	69.17%	72.01%	72.00%
$ID(y^{5y})$	50.60%	65.10%	71.54%	71.31%
$ID(y^{7y})$	53.57%	66.30%	73.64%	74.63%
$ID(y^{10y})$	51.38%	66.29%	72.26%	72.05%
$ID(y^{10y}) - ID(y^{2y})$	23.12%	18.14%	17.49%	13.98%

Table 1: R^2's from the projections of inter-quantile differences in BCFF forecasts onto all yields (maturities of 6 months, and 2, 5, 7 and 10 years), and onto all yields and disagreement about inflation and real output growth over forecast horizons of one through four quarters. The sample period is January, 1985 through December 2014.

by data availability. 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). Both the zero yields and their corresponding survey-implied forecasts are computed as in Le and Singleton (2012).[7]

Disagreement is constructed as the inter-decile range of the professional yield forecasts from the BCFF survey, which begin in January 1985.[8] The differences between the ninetieth and tenth percentiles of the cross-sectional distribution of forecasts over horizon j and bond maturity m is denoted by $ID_{jt}(y^m)$.[9] In Panel (a) of Table 1 we report the R^2's from the projections of these dispersion measures onto bond yields for forecast horizons of one through four quarters. Yield disagreement is only partially spanned by the yield curve, with R^2's ranging from 40% to 69% across forecast horizons and bond maturities.

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. We focus on disagreement about future yields, because the priced factors in bond markets are spanned by the low-dimensional yield PCs. Moreover,

[7]By interpolating the forecasts of par yields to obtain approximate forecasts of zero yields, we simplify our analysis to one of forecasting zero-coupon yields in an affine *DTSM*.

[8]Our results are 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).

[9]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.

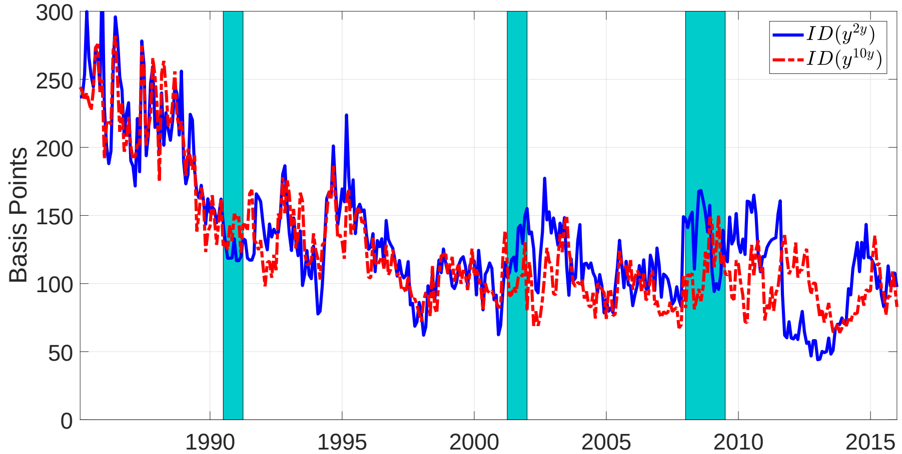

Figure 2: Historical measures of dispersions in professional forecasts one-year ahead for the two- and seven-year bond yields, $ID(y^{2y})$ and $ID(y^{10y})$.

disagreement about yields is not spanned by disagreement about inflation and output growth. Using the $BCFF$ panel of forecasters, we construct forecasts of one-year ahead inflation and real GDP growth (again as inter-decile ranges).[10] Panel (b) of Table 1 shows that R^2's in the projections of yield disagreement on yields and macroeconomic disagreement are less than 75%. Equally notable, when the dependent variable is the slope of disagreement, $ID(y^{10y}) - ID(y^{2y})$, the R^2 is only 14% for the four quarters horizon. In subsequent sections we show that yield disagreement is more strongly predictive of future excess bond returns than inflation, output growth, and disagreement about these macro variables.

A large part of its variation over time of our measures of disagreement can be explained their low-order PCs: the first PC of the covariance matrix of the $ID_{jt}(y^m)$, across j and m, explains 97% of the variation. The second PC is a "slope of disagreement" factor. Therefore, in our subsequent analysis of learning we summarize information on bond yield disagreement at time t using $H_t' = \left[ID_t(y^{2y}), ID_t(y^{10y})\right]$ over the horizon of one year (so we drop the subscript $j = 1y$). Figure 2 shows that these measures of disagreement are counter-cyclical as they tends to rise during and shortly after NBER recessions. Moreover, $ID(y^{2y})$ tends to be higher than $ID(y^{10y})$ and the gap between them $(ID(y^{2y}) - ID(y^{10y}))$ is relatively large following the two recessions in our sample. The years 2012-13 are exceptional for the persistently low level of $ID(y^{2y})$. The cyclical patterns of H for different choices of the professional forecasters' horizon are qualitatively similar.

[10] 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.

3 Formalizing $\mathcal{BL}$'s Learning Rule

Figure 1 provides compelling evidence that learning shapes beliefs about expected excess returns. Our goal is to explore how real-time learning, enriched by conditioning on disagreement among professionals, affects the market prices of factor risks, insights that real-time predictive regressions alone do not reveal. Additionally, we are interested in how $\mathcal{BL}$ updates the parameters of the historical and pricing distributions as new information arrives, and the role of disagreement in this updating process. To identify these properties of $\mathcal{BL}$'s learning rule, we endow her with a learning model based on a Gaussian *DTSM*.

In a standard Gaussian *DTSM* the one-period riskless rate r_t has an affine factor representation, $r_t = \rho_0 + \rho_{\mathcal{P}}\mathcal{P}_t,$. Consistent with the well-documented covariance structure of yields, the risk factors $\mathcal{P}_t$ are taken to be the first three *PC*s of bond yields. Under the pricing distribution $\mathbb{Q}$, $\mathcal{P}$ is assumed to follow the autonomous Gaussian process

$$\mathcal{P}_{t+1} = K_{0\mathcal{P}}^{\mathbb{Q}} + K_{\mathcal{P}\mathcal{P}}^{\mathbb{Q}}\mathcal{P}_t \, + \, \Sigma_{\mathcal{P}\mathcal{P}}^{1/2} e_{\mathcal{P},t+1}^{\mathbb{Q}}. \tag{1}$$

The arbitrage-free price D_t^m of a zero-coupon bond issued at date t and maturing at date $t+m$ is then determined as

$$D_t^m = E_t^{\mathbb{Q}}\left[\prod\nolimits_{u=0}^{m-1} \exp(-r_{t+u})|\Theta^{\mathbb{Q}}, \mathcal{P}_t\right], \tag{2}$$

where $\Theta^{\mathbb{Q}}$ is the parameter vector governing the $\mathbb{Q}$ distribution of $\mathcal{P}$. Econometric identification of $\Theta^{\mathbb{Q}}$ can be achieved (see Joslin, Singleton, and Zhu (2011), JSZ) by normalizing ρ_0, $\rho_{\mathcal{P}}$, $K_{0\mathcal{P}}^{\mathbb{Q}}$, and $K_{\mathcal{P}\mathcal{P}}^{\mathbb{Q}}$ to be known functions of $\left(k_\infty^{\mathbb{Q}}, \lambda^{\mathbb{Q}}, \Sigma_{\mathcal{P}\mathcal{P}}\right)$, with $k_\infty^{\mathbb{Q}}$ a scalar and $\lambda^{\mathbb{Q}}$ the eigenvalues of $K_{\mathcal{P}\mathcal{P}}^{\mathbb{Q}}$.[11] Zero-coupon bond yields in this model are affine functions of $\mathcal{P}$:

$$y_t^m = A_m(k_\infty^{\mathbb{Q}}, \lambda^{\mathbb{Q}}, \Sigma_{\mathcal{P}\mathcal{P}}) + B_m(\lambda^{\mathbb{Q}})\mathcal{P}_t. \tag{3}$$

Professional traders typically adopt a pricing model that shares the factor structure of (3). This industry practice presumes knowledge of the loadings (A_m, B_m), for $\Theta^{\mathbb{Q}}$ is being treated as both *known* and *fixed* over the interval $[t, t+m]$ when computing D_t^m.[12] By following this practice, $\mathcal{BL}$ can precisely calibrate $\lambda^{\mathbb{Q}}$ from the loadings $B_m(\lambda^{\mathbb{Q}})$ and the cross-maturity

[11] When $\mathcal{P}$ follows a stationary process under $\mathbb{Q}$, $k_\infty^{\mathbb{Q}}$ 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 $\lambda_1^{\mathbb{Q}}$ to be very close to or even larger than unity. See JSZ for details.

[12] This is reminiscent of the assumption of "anticipated utility" that is often made in equilibrium representative-agent models with Bayesian learning (see, e.g., Kreps (1998) and Cogley and Sargent (2008)) for tractability, particularly in high dimensional models. Among recent studies of learning and the pricing of equities, Johannes, Lochstoer, and Mou (2016) also adopt the assumption of anticipated utility. Importantly, here we are discussing the martingale pricing measure and there is no presumption of shared beliefs across market participants.

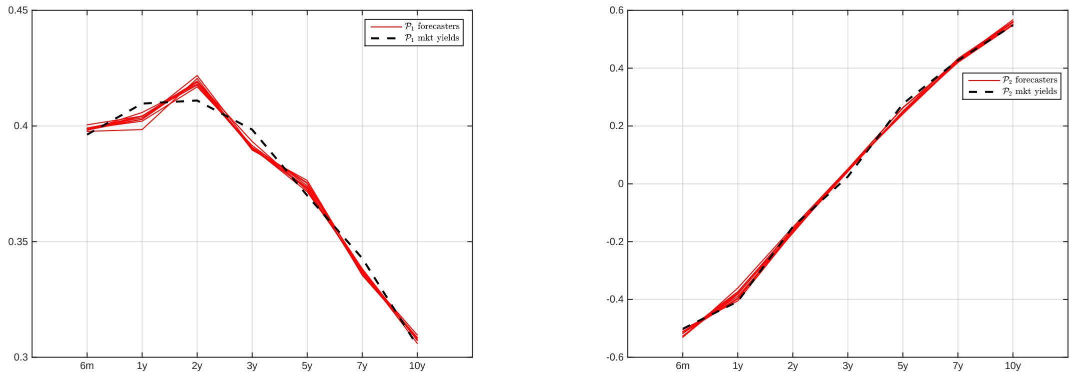

(a) Loadings on PC1

(b) Loadings on PC2

Figure 3: Time-series projections of the decile-ordered one-year ahead forecasts by BCFF forecasters of yields onto their forecasts of $PC1$ and $PC2$.

factor structure of y_t^m and $\mathcal{P}_t$ (see also Duffee (2011)). The intercepts A_m depend in addition on $k_\infty^{\mathbb{Q}}$ and $\Sigma_{\mathcal{PP}}$. However, the impact of $\Sigma_{\mathcal{PP}}$ on A_m is through a convexity adjustment that is typically very small. Therefore, knowledge of $\lambda^{\mathbb{Q}}$ and a tight prior on $k_\infty^{\mathbb{Q}}$ (also estimable from the cross-section of yields) imply that $\mathcal{BL}$ effectively knows the A_m's as well.

Now if the factor loadings are known to $\mathcal{BL}$, then they are also knowable by other bond market participants, so one might wonder whether professionals do in fact share the same views about the B_m. Fortunately, we can shed light on this using the monthly BCFF yield forecasts. Consider the yield forecasts for horizon h ordered by deciles in the BCFF panel, $y_{t,o_1}^h < ... < y_{t,o_{10}}^h$, where y_{t,o_1}^h is the forecast of the professional falling at the tenth percentile, y_{t,o_2}^h is the forecast falling at the twentieth percentile, and so on up to the ninetieth percentile (we focus on order statistics, because the individual forecasters change over our sample). If the BCFF professionals share common views on the loadings, then (for each horizon h) we should find identical loadings across deciles in the projections

$$\hat{y}_{t,o_k}^{mh} = \bar{A}_m + \bar{B}_m \hat{\mathcal{P}}_{t,o_k}^h + e_{t,o_k}^{mh}, \tag{4}$$

where $\hat{\mathcal{P}}_{t,o_k}^h$ are the BCFF decile forecasts of the PCs. Figure 3 displays the full-sample estimates (for $h = 1y$) of these loadings on $PC1$ and $PC2$ for the decile expectations (solid lines) and for the sample yields (dashed lines). The loadings are remarkably similar across forecaster deciles, and they are all close to the sample counterparts. *This is the case even though there are large differences in the forecasts of future yields across deciles (Figure 2).*

Taken together, these observations suggest that $\mathcal{BL}$'s central learning problem in Treasury markets is about the physical dynamics of the risk factors $f^{\mathbb{P}}(\mathcal{P}_t|Z_1^{t-1})$, where $Z_1^{t-1} \equiv \{Z_1, Z_2, ..., Z_{t-1}\}$. Two considerations motivate the conditioning of $f^{\mathbb{P}}(\mathcal{P}_t|Z_1^{t-1})$ on a richer information set than the history of $\mathcal{P}_t$. First, the market prices of the factor risks will generally depend on more information than bond yields alone (see JPS for discussion of this point in a non-learning setting). Second, as $\mathcal{BL}$ learns, her time-varying posterior means of the unknown parameters of $f^{\mathbb{P}}(\mathcal{P}_t|Z_1^{t-1})$ will change with the arrival of new information Z. Given our focus on the roles of disagreement H_t in $\mathcal{BL}$'s learning rule, for most of our analysis $Z_t' = (\mathcal{P}_t', H_t')$. However, to shed broader light on the nature of the incremental forecasting power of H_t we subsequently expand Z_t to include various macro variables as well.

Within the Gaussian *DTSM* framework, $\mathcal{BL}$ assumes that Z_t follows the process

$$Z_{t+1} = K_{0t}^{\mathbb{P}} + K_{Zt}^{\mathbb{P}} Z_t \, + \, \Sigma_Z^{1/2} e_{Z,t+1}^{\mathbb{P}}, \quad e_{Z,t+1}^{\mathbb{P}} \; N(0, \Sigma_Z), \tag{5}$$

with the elements of $K_{0t}^{\mathbb{P}}$ and $K_{Zt}^{\mathbb{P}}$ unknown and potentially time varying. We define $\Theta_t^{\mathbb{P}}$ as the vectorized $(K_{0t}^{\mathbb{P}}, K_t^{\mathbb{P}})$. The portfolios $\mathcal{P}_t$ are assumed to be priced perfectly by (2),[13] while the higher-order PCs $\mathcal{O}_t$ are priced with errors:

$$\mathcal{O}_t = A_{\mathcal{O}} \left(\Theta^{\mathbb{Q}}\right) + B_{\mathcal{O}} \left(\Theta^{\mathbb{Q}}\right) \mathcal{P}_t + \varepsilon_{\mathcal{O},t}, \tag{6}$$

where $(\mathcal{P}_t, \mathcal{O}_t)$ fully spans y_t. The pricing 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 $\mathcal{P}_t$).

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 (5) into $\mathcal{Y}_t = Z_t - \left(I \otimes [1, Z_{t-1}']\right) \iota_r \psi^r$, we rewrite the state equation as

$$\mathcal{Y}_{t+1} = \mathcal{X}_t \psi_t^{\mathbb{P}} + \Sigma_Z^{1/2} e_{Z,t+1}^{\mathbb{P}}, \tag{7}$$

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

In recognition of the possibility of permanent structural changes in the underlying economic environment, $\mathcal{BL}$ assumes that $\psi_t^{\mathbb{P}}$ evolves according to

$$\psi_t^{\mathbb{P}} = \psi_{t-1}^{\mathbb{P}} + Q_{t-1}^{1/2} \eta_t, \qquad \eta_t \overset{iid}{\sim} Normal(0, I), \tag{8}$$

[13] In a similar Gaussian setting without learning, Joslin, Le, and Singleton (2013) show that model-implied risk premiums and forecasts of future yields are nearly identical from a model in which $\mathcal{P}_t$ is priced perfectly and a model with all of the bonds (and hence $\mathcal{P}$) are priced with errors.

where Q_{t-1} denotes the (possibly) time-varying covariance matrix of η_t, with η_t independent of all past and future $e^{\mathbb{P}}_{Zt}$. $\mathcal{BL}$'s Bayesian learning rule filters for the unknown $\psi^{\mathbb{P}}_t$ conditional on her knowledge of $(\Theta^{\mathbb{Q}}, \psi^r)$. Adopting a Gaussian prior on $\psi^{\mathbb{P}}_0$ leads to a posterior distribution for $\psi^{\mathbb{P}}_t$ that is also Gaussian, $\psi^{\mathbb{P}}_t | Z^t_1 \sim Normal(\hat{\psi}^{\mathbb{P}}_t, P_t)$. In Appendix A we show that her posterior mean follows the recursion

$$\hat{\psi}^{\mathbb{P}}_t = \hat{\psi}^{\mathbb{P}}_{t-1} + R^{-1}_t \mathcal{X}'_{t-1} \Sigma^{-1}_Z (\mathcal{Y}_t - \mathcal{X}_{t-1} \hat{\psi}^{\mathbb{P}}_{t-1}), \tag{9}$$

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

$$R_t = \left(I - P^{-1}_{t-1} Q_{t-2}\right) R_{t-1} + \mathcal{X}'_{t-1} \Sigma^{-1}_Z \mathcal{X}_{t-1}. \tag{10}$$

This rule has a revealing interpretation within the class of adaptive least-squares estimators of $\psi^{\mathbb{P}}_t$. We subsequently focus on the following two special cases:[14]

$\mathcal{B}\!\downarrow$CGLS: Setting $P^{-1}_{t-1} Q_{t-2} = (1-\delta) \cdot I$[15] for a constant scalar $0 < \delta \leq 1$, $\hat{\psi}_t$ becomes the *constant gain least-squares (CGLS)* estimator of $\psi^{\mathbb{P}}$ with $\gamma = \delta$.

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

Note that a Bayesian agent whose learning rule specializes to the *RLS* estimator is not adaptive in the following important sense. With $\gamma = 1$ we have $Q_t = 0$, so an agent following a *RLS* rule is learning about an unknown $\psi^{\mathbb{P}}$ 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 $\widehat{\psi^{\mathbb{P}}_t}$ as updated by $\mathcal{BL}$. Indeed, in environments where the ML estimator converges to a constant for large T, an *RLS*-based $\mathcal{BL}$ will be virtually non-adaptive on $\widehat{\psi^{\mathbb{P}}}$ to new information after a long training period.

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 $\mathcal{BL}$'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 A) are identical to those of a likelihood with terms of the form $\gamma^t \epsilon^{\mathbb{P}\prime}_{Zt} \Sigma^{-1}_Z \epsilon^{\mathbb{P}}_{Zt}$.[16]

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

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

[16]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.

The density of $\mathcal{Y}_{t+1}$ conditional on Z_1^t is $f^{\mathbb{P}}(\mathcal{Y}_{t+1}|Z_1^t) = Normal\left(\mathcal{X}_t\hat{\psi}_t^{\mathbb{P}}, \Omega_t\right)$, 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. Out-of-sample forecasts of yields and excess returns are computed directly from $\mathcal{BL}$'s learning rule. Since yields are spanned by $\mathcal{P}$, the first step is to compute the out-of-sample forecasts of the state vector Z, which includes $\mathcal{P}$. The h-period ahead forecast of Z from the fitted $DTSM$ at date t, is given by

$$\hat{Z}_{t+h} = \widehat{K}_{0t}^{\mathbb{P}} + \left(\widehat{K}_{Zt}^{\mathbb{P}}\right)\widehat{K}_{0t}^{\mathbb{P}} + ... + \left(\widehat{K}_{Zt}^{\mathbb{P}}\right)^{h-1}\widehat{K}_{0t}^{\mathbb{P}} + \left(\widehat{K}_{Zt}^{\mathbb{P}}\right)^{h} Z_t. \tag{11}$$

This leads directly[17] to the h-period ahead forecasts of yields:

$$\hat{y}_{t+h}^m = A_m\left(K_0^{\mathbb{Q}}, K_{\mathcal{PP}}^{\mathbb{Q}}, \Sigma_{\mathcal{PP}}\right) + B_m\left(K_{\mathcal{PP}}^{\mathbb{Q}}\right)\hat{\mathcal{P}}_{t+h}. \tag{12}$$

Importantly, (11) and (12) reveal that *$\mathcal{BL}$'s learning rule is not an affine model*; $\mathcal{BL}$'s optimal forecasts and expected excess returns are nonlinear functions of the history of Z_t. This is also revealed by $\mathcal{BL}$'s stochastic discount factor $\mathcal{M}^B$:

$$\begin{aligned}
\mathcal{M}(\Theta^{\mathbb{Q}}, \mathcal{P}_{t+1}, Z_1^t) &= e^{\{-r_t - \frac{1}{2}\log|\Gamma_t| - \frac{1}{2}\hat{\Lambda}_{\mathcal{P}t}'\Gamma_t^{-1}\hat{\Lambda}_{\mathcal{P}t} - \hat{\Lambda}_{\mathcal{P}t}'\Gamma_t^{-1}\varepsilon_{t+1}^{\mathbb{P}} + \frac{1}{2}(\varepsilon_{t+1}^{\mathbb{P}})'(I-\Gamma_t^{-1})\varepsilon_{t+1}^{\mathbb{P}}\}}, \quad (13)\\
\Gamma_t &= \Omega_{\mathcal{PP},t}^{-1/2}\Sigma_{\mathcal{PP}}(\Omega_{\mathcal{PP},t}^{-1/2})',\\
\Omega_{\mathcal{PP},t}^{1/2}\hat{\Lambda}_{\mathcal{P}t} &= \hat{\Lambda}_{0t}(\Theta^{\mathbb{Q}}, \hat{\Theta}_t^{\mathbb{P}}) + \hat{\Lambda}_{1t}(\Theta^{\mathbb{Q}}, \hat{\Theta}_t^{\mathbb{P}})Z_t, \quad (14)
\end{aligned}$$

where the market prices of risk $\hat{\Lambda}_{\mathcal{P}t}$ depend on the posterior mean $\hat{\Theta}_t^{\mathbb{P}}$ and, therefore, implicitly on the entire history Z_1^t (Appendix B). 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.

At date t a Bayesian $\mathcal{BL}$, 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,

$$\begin{aligned}
f(Z_1^t, \mathcal{O}_1^t) &= \prod_{s=1}^{t} f(\mathcal{O}_s|Z_1^s, \mathcal{O}_1^{s-1}; \Theta^{\mathbb{Q}}, \Sigma_{\mathcal{O}}) \times \\
&\quad \int f(Z_s|Z_1^{s-1}, \mathcal{O}_1^{s-1}, \Theta_{s-1}^{\mathbb{P}}; \Sigma_Z) f(\Theta_{s-1}^{\mathbb{P}}|Z_1^{s-1}, \mathcal{O}_1^{s-1}) d(\Theta_{s-1}^{\mathbb{P}}). \quad (15)
\end{aligned}$$

For robustness– and consistent with practice by sophisticated market participants– we extend

[17] Notice that in computing the forecast we are neglecting the effects of parameter uncertainty, but rather we are assuming that $\mathcal{BL}$ takes the physical dynamics of Z as known when forecasting. We have found that accounting for parameter uncertainty does not materially change yields forecasts.

$\mathcal{BL}$'s learning rule to have her update $\Theta^{\mathbb{Q}}$ monthly using (15) as new data becomes available. This allows us to assess whether the cross-maturity structure of bond yields calls for revisions of the factor loadings as $\mathcal{BL}$ learns over time. Strikingly, $\mathcal{BL}$'s estimates of $\lambda^{\mathbb{Q}}$ change very little over the twenty-five years that she is learning in our sample (Section 4.1). This supports the premise that $\mathcal{BL}$ prices bonds as if $\lambda^{\mathbb{Q}}$ is known.

The constraints ϕ^r on $(K_{0t}^{\mathbb{P}}, K_{Zt}^{\mathbb{P}})$ in (5) mitigate overfitting of the forecasts of yields relative to unconstrained linear regressions. As we explain in Section 4.1 and Appendix D, $\mathcal{BL}$ selects constraints during a model "training period" that predates her learning. As such, she is truly computing out-of-sample forecasts using a real-time learning rule.

Throughout this construction the direct dependence of Ω_t on Σ_Z is a consequence of $\mathcal{BL}$ treating Σ_Z as known, and not as an object to be learned. This is an admittedly strong assumption as, empirically, $\mathcal{BL}$'s learning rule shows 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 she updates her beliefs about $\Theta^{\mathbb{P}}$. In Appendices E and F 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.

4 Learning and Risks Premiums: Empirical Evidence

Equipped with this parametric learning rule for $\mathcal{BL}$, we next explore empirically the properties of her subjective risk premiums, and of the roles of disagreement in shaping her views about future excess returns. In setting out on this analysis it is instructive to focus initially on learning from information in the yield curve alone. This serves as a benchmark for evaluating the *incremental* role of disagreement, and it allows examination of learning over a much longer time period (owing to relatively limited historical BCFF data).

4.1 Learning From the Yield Curve: Rule $\ell_{CG}^{L}(\mathcal{P})$

Consider a three-factor *DTSM* in which $\mathcal{BL}$ follows a constant-gain learning rule with conditioning on past information on $\mathcal{P}$ alone, rule $\ell_{CG}(\mathcal{P})$. $\mathcal{BL}$'s rule is trained over the period from June 1961 through January 1972, and it is during this period that constraints are set on the market prices of factor risks ($\Lambda_{\mathcal{P}t}$ in (14)). Every month thereafter, through December 2014, $\mathcal{BL}$ updates her posterior $\widehat{\Theta}_t$ as new data becomes available.

The gain parameter γ is set to 0.99. Appendix C offers two complementary perspectives on this choice. First, if we allow $\mathcal{BL}$ to adjust γ over time based on realized out-of-sample forecast accuracy, for both one- and four-quarter ahead horizons, she selects fitted γ_t's close to 0.99. This is especially true over the period January 1995 through December 2014, which will be the focus of most of our subsequent analysis. Second, setting γ equal to 0.99 is ex-post optimal. In fact, searching over fixed γ's to minimize out-of-sample forecast errors at the one-quarter ahead and one-year-ahead horizon over the period from January 1995 through December 2014 leads to an "optimal" value of γ that is approximately 0.99.

The real-time estimates of $\lambda^{\mathbb{Q}}$ from rule $\ell^L_{CG}(\mathcal{P})$ are displayed in Figure 4. Notably, $\mathcal{BL}$ holds $\lambda^{\mathbb{Q}}$ virtually fixed over the entire sample, consistent with the premise of her Bayesian rule for learning about $\Theta^{\mathbb{P}}_t$. (They are virtually identical for the RLS rule $\gamma = 1$). Though there is mild drift in the second eigenvalue $\lambda^{\mathbb{Q}}_2$, 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. This is reflected in the loadings $B(\lambda^{\mathbb{Q}})$ (see (3)) of the first two PC's of the yields $(y^{2y}, y^{5y}, y^{10y})$ as they are close to constant over the entire sample (Figure 5).[18] Thus, $\mathcal{BL}$ uses nearly fixed "hedge ratios" in managing bond portfolio risks.

This finding is especially striking in relation to how $\mathcal{BL}$ updates the historical eigenvalues $\lambda^{\mathbb{P}}$ (see Figure 4). Her views about the objective feedback matrix $K^{\mathbb{P}}_{\mathcal{PP},t}$ change substantially over time, hence so do the implied responses of risk premiums to innovations in $\mathcal{P}$. For all of the periods of large changes in $\lambda^{\mathbb{P}}_t$, $\lambda^{\mathbb{Q}}$ remains remarkably stable. This is particularly evident for $\lambda^{\mathbb{P}}_1$ during the Fed experiment in the early 1980's.

The *DTSM*-based rules are potentially highly parametrized. For parsimony, and forecast precision in the out-of-sample assessments, the parameters governing the market prices of $\mathcal{P}$ risks (MPR's) are set to zero if their p-values *during the training period* are larger than 0.4. Since $K^{\mathbb{Q}}_{\mathcal{PP}}$ is presumed known by $\mathcal{BL}$, these constraints on $\Lambda_{\mathcal{P}}$ effectively transfer *a priori* knowledge of $\lambda^{\mathbb{Q}}$ to (some) knowledge about $K^{\mathbb{P}}_{\mathcal{PP},t}$. All constraints on $\Lambda_{\mathcal{P}}$ selected during the training period are maintained throughout the remainder of the sample period.[19] See Appendix D for details on the constraints imposed.

Figure 6 compares annualized MPR's for $\mathcal{P}_1$ and $\mathcal{P}_2$, for $\ell^L_{CG}(\mathcal{P})$ against those implied by the fixed-parameter, full-sample estimates. Real-time learning induces substantial differencea in measured MPR's compared to results without learning. Over the sample from January 1985 through December 2014, the root-mean-squared differences are 23 basis points for $\mathcal{P}_1$ and

[18]Further reassurance that the near constant $\lambda^{\mathbb{Q}}$ is not a mechanical implication of the no-arbitrage term structure framework is provided by running the reduced-form, expanding-window regressions

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

The loadings B^n_t remain quite stable for yields across the maturity spectrum. Importantly, the estimates of the weights that define $\mathcal{P}$ are also stable over time.

[19]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.

Figure 4: Estimates from model $\ell^L_{CG}(\mathcal{P})$ of the eigenvalues $\lambda^{\mathbb{Q}}$ ($\lambda^{\mathbb{P}}$) of the feedback matrices $K^{\mathbb{Q}}_{\mathcal{P}\mathcal{P},t}$ ($K^{\mathbb{P}}_{\mathcal{P}\mathcal{P},t}$) 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.

133 basis points for $\mathcal{P}_2$. For $\mathcal{P}_1$ the effects of learning are particularly large around business cycle turning points. There are larger and more persistent effects of learning for $\mathcal{P}_2$, especially after 1995. The slope of the yield curve is unquestionably a priced risk for $\mathcal{BL}$.

The expected excess returns over a one-year holding period implied by $\mathcal{BL}$'s learning rule $\ell^L_{CG}(\mathcal{P})$ differ substantially from the corresponding risk premiums implicit in the learning rule followed by the median BCFF forecaster, $\ell(BCFF)$ (Figure 7).[20] Key to understanding these differences– which are particularly large following NBER recessions– is the strong positive correlation between $\mathcal{BL}$'s 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 $\mathcal{BL}$'s learning rule. For example, following the low (recession) levels of y^{10y} from late 2002 until 2004 the $BCFF$ forecasters

[20]BCFF forecasts are averages over calendar quarters and cover horizons out to five quarters ahead. For example, in January, 1999, the one-quarter ahead forecast for a specific variable will be equal to its average value between February and April. For comparability across all forecast rules, we compute similar quarterly averages for rule $\ell^L_{CG}(\mathcal{P})$ and, indeed, all subsequent rules explored in this paper.

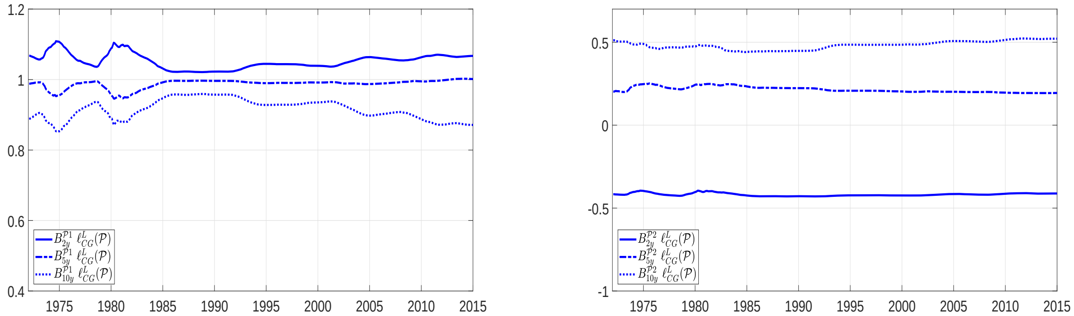

Figure 5: Loadings of $\mathcal{P}_1$ and $\mathcal{P}_2$ on bond yields of maturity 2, 5 and 10 years and for model $\ell^L_{CG}\mathcal{P}$, based on (12) over the sample January 1972 through December 2014.

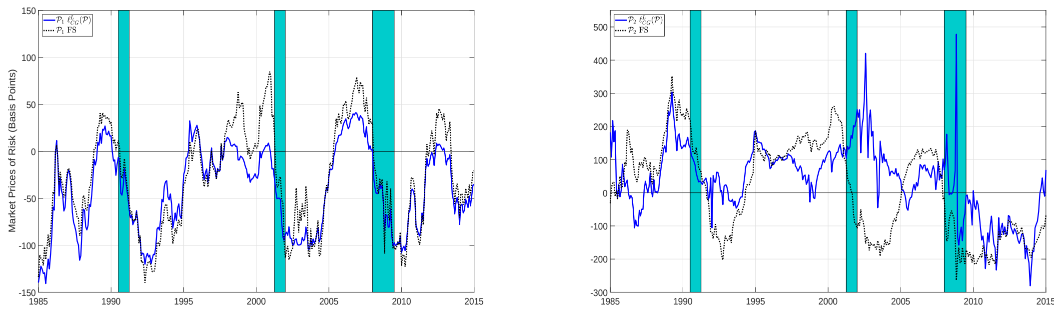

Figure 6: One-month ahead market prices of risk ($\Lambda_{\mathcal{P},t}$) for $\mathcal{P}_1$ and $\mathcal{P}_2$. We compare $\ell^L_{CG}(\mathcal{P})$ against full sample (FS) estimates over the period June 1961 through December 2014.

expected a much more rapid rise in y^{10y} than did $\mathcal{BL}$'s more accurate rule $\ell^L_{CG}(\mathcal{P})$. Thus, 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 the benefit of hindsight) the actual performance of bonds.[21]

4.2 Learning from Disagreement: Rule $\ell_{CG}(\mathcal{P}, H)$

To accommodate dependence of expected excess returns on both information spanned by the current yield curve and disagreement about the paths of future yields, we expand the

[21] This finding is complementary to (and 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. It suggests that using median $BCFF$ forecasts of long-term bond yields to calibrate empirical learning rules would lead to distorted measures of required risk compensations. 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.

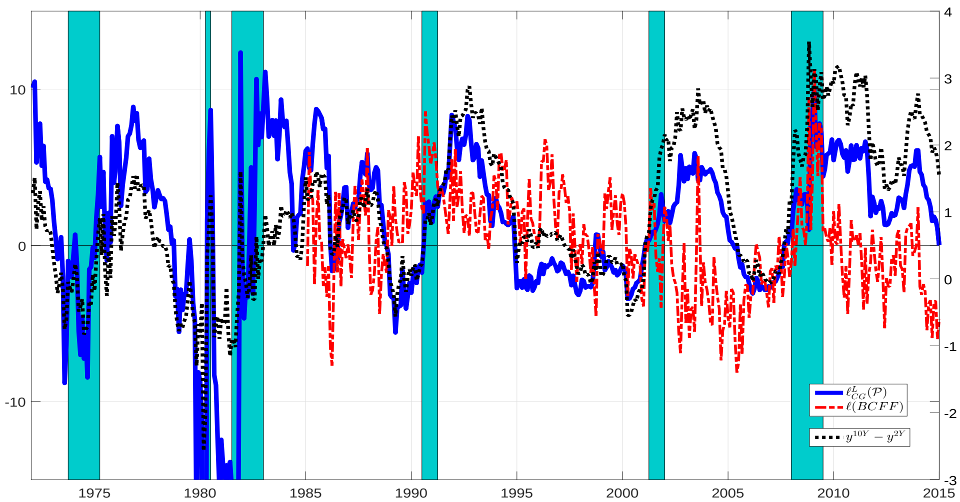

Figure 7: Average expected excess returns over holding periods of ten, eleven and twelve months for the ten-year bond based on $\ell^L_{CG}(\mathcal{P})$ and $\ell(BCFF)$ (left axis) and the slope of the Treasury curve measured as $y^{10y} - y^{2y}$ (right axis), January, 1972 to December, 2014.

conditioning information vector Z'_t to $(\mathcal{P}'_t, H'_t)$. This rule– referred to as $\ell_{CG}(\mathcal{P}, H)$– fits directly into the framework of Section 3. Since data for H_t is available only from January 1985, rule $\ell_{CG}(\mathcal{P}, H)$ is trained from January 1985 through January 1995, and then $\mathcal{BL}$ proceeds forward in time using a monthly expanding window up to December 2014. For comparison we also construct forecasts from the simple yield-based rule that has each zero yield following a random walk, rule $\ell(RW)$.

The relative accuracies of these rule-based forecasts, which depend primarily on $\widehat{\Theta}^{\mathbb{P}}_t$, can be assessed from the RMSE's displayed in Table 2. Below each RMSE are Diebold and Mariano (1995) (D-M) statistics for assessing whether two RMSE's are statistically the same, calculated as extended by Harvey, Leybourne, and Newbold (1997). Conditioning on H_t leads to a substantial improvement in forecast accuracy relative to the *DTSM*-based rules that conditions only on $\mathcal{P}$. This pickup in accuracy occurs across the maturity spectrum, with larger gains at the long end of the Treasury yield curve. Moreover, $\ell_{CG}(\mathcal{P}, H)$ outperforms rule $\ell(RW)$ across the maturity spectrum (most significantly for long-maturity bonds). The outperformance of $\ell_{CG}(\mathcal{P}, H)$ relative to the least accurate rule $\ell(BCFF)$ is statistically significant across the entire yield curve for the one-quarter horizons.[22]

[22]In the literature 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 (see, e.g., Stambaugh (1999)). In a term structure context, Bauer and Hamilton (2018) argue that there

	RMSE's (in basis points) for Quarterly Horizon						
Rule	6m	1Y	2Y	3Y	5Y	7Y	10Y
$\ell(RW)$	30.7	32.6	35.2	36.1	36.4	36.0	32.5
	(−3.81)	(−2.98)	(−3.98)	(−5.06)	(−4.89)	(−3.96)	(−3.01)
	[−]	[−]	[−]	[−]	[−]	[−]	[−]
$\ell(BCFF)$	38.6	37.6	43.8	50.7	44.6	44.1	40.0
	(−)	(−)	(−)	(−)	(−)	(−)	(−)
	[]	[]	[]	[]	[]	[]	[]
$\ell_{CG}(\mathcal{P})$	28.9	31.1	35.8	36.1	37.1	36.6	33.4
	(−3.47)	(−2.60)	(−3.55)	(−5.22)	(−4.53)	(−3.75)	(−2.74)
	[−1.45]	[−0.86]	[1.31]	[−0.06]	[0.71]	[0.85]	[1.65]
$\ell_{CG}(\mathcal{P}, H)$	28.8	29.9	34.2	34.7	36	35.3	31.6
	(−3.38)	(−2.86)	(−3.73)	(−5.16)	(−4.39)	(−4.00)	(−3.10)
	[−1.45]	[−1.45]	[−1.42]	[−1.60]	[−0.33]	[−0.67]	[−1.36]

	RMSE's (in basis points) for Annual Horizon						
Rule	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]
$\ell_{CG}(\mathcal{P}, H)$	109.4	105.0	95.8	86.5	77.0	70.2	64.0
	(−1.38)	(−1.48)	(−2.17)	(−2.91)	(−3.28)	(−3.72)	(−3.49)
	[−1.20]	[−1.35]	[−1.44]	[−1.36]	[−1.41]	[−1.69]	[−1.75]

Table 2: RMSE's for one-quarter and 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).

The outperformance of $\ell_{CG}(\mathcal{P}, H)$ relative to both $\ell(RW)$ and $\ell(BCFF)$ was especially large in the early 2000's (Table 3). The portion of our sample covering the global financial crisis was the easiest subperiod for forecasting Treasury yields. With short-term rates pegged essentially at zero, $\ell(RW)$ was the best performing rule out to the five-year maturity. $\ell_{CG}(\mathcal{P}, H)$ outperformed $\ell(RW)$ for long-maturity bonds, even though the 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)). Moreover, shortening our evaluation window to the period from 2008 to 2011 leads to outperformance of $\ell_{CG}(\mathcal{P}, H)$ over $\ell(RW)$ across all yields. 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

is a tendency for an upward bias in estimated R^2's in excess return regressions owing to a "standard error bias," and this bias is potentially amplified when studying long-horizon forecasts using overlapping data. We emphasize that our assessments of forecast accuracy are based on *out-of-sample* fit. Moreover, our findings on out-performance are consistent across one-quarter- and one-year-ahead horizons.

	RMSE's by Bond Maturity						
Rule	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	81
$\ell_{CG}(\mathcal{P}, H)$	113	113	103	92	84	80	76
	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	135	124	111	90	74	58
$\ell_{CG}(\mathcal{P}, H)$	140	132	119	104	81	64	47
	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
$\ell_{CG}(\mathcal{P}, H)$	60	56	55	58	65	67	67

Table 3: RMSE's in basis points for one-year-ahead forecasts of individual bond yields over the indicated sample periods.

of the Treasury curve.

How does conditioning on H change $\mathcal{BL}$'s perceived risk premiums? Figure 8 shows one-month annualized market prices of risk for the first two PC's of the yield curve. Introducing H in the conditioning information substantially increases the volatility of risk premiums. Over the sample from January 1995 through December 2014 the sample standard deviation of the estimated MPR for the first PC (second PC) is 0.69% (1.24%) for model $\ell_{CG}(\mathcal{P})$, while it is almost 2.5% (3.0%) for rule $\ell_{CG}(\mathcal{P}, H)$. With this increased volatility of the MPR's from rule $\ell_{CG}(\mathcal{P}, H)$ comes more accurate out-of-sample forecasts of future yields (Table 3).

Letting $xr_{t+h}^{n,h}$ denote the realized excess return on an n-maturity bond over an h-period holding period, the left panel of Figure 9 displays $\mathcal{BL}$'s forecast $exr_{1y,t}^{10y}(\mathcal{P}, H)$ of $exr_{t+1y}^{10y,1y}$ from rule $\ell_{CG}(\mathcal{P}, H)$ and $er_{1y,t}^{10y}(\mathcal{P})$ from rule $\ell_{CG}(\mathcal{P})$. The right panel shows the difference in RMSEs across these two learning rules from forecasting the realized excess returns $xr_{t+1y}^{10y,1y}$. The greater forecast accuracy of rule $\ell_{CG}(\mathcal{P}, H)$ relative to rule $\ell_{CG}(\mathcal{P})$ is particularly large during and immediately after the recessions of 2001 and 2008. For most of the sample, the difference in RMSEs is close to 80 bps and drops to 65 bps only after 2013. This forecasting power of disagreement is more pronounced for long-term bonds, as the corresponding differences in the RMSEs from rules $\ell_{CG}(\mathcal{P})$ and $\ell_{CG}(\mathcal{P}, H)$ for the two-year Treasury bond are less than 5

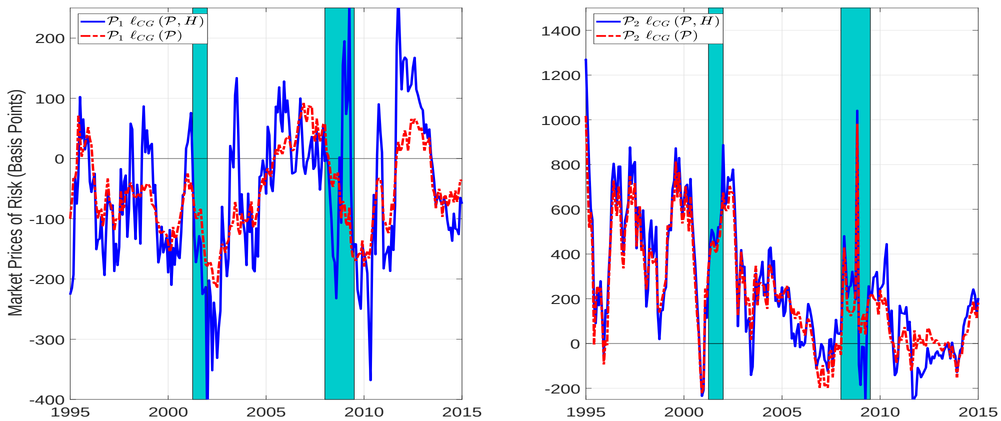

Figure 8: Market prices of risk for the first and second PCs of the yield curve over a one-month holding period, based on rules $\ell_{CG}(\mathcal{P})$ and $\ell_{CG}(\mathcal{P}, H)$.

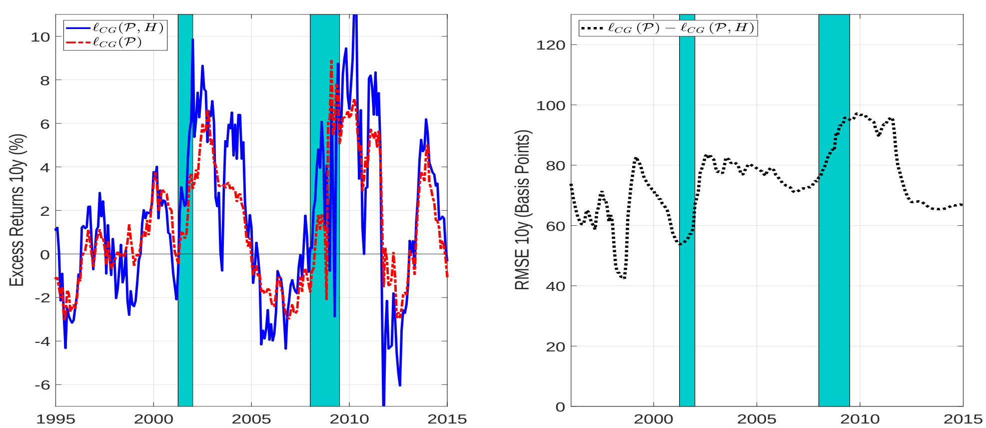

Figure 9: The left panel shows expected excess returns for a 10-year bond over a one-year horizon, estimated using rules $\ell_{CG}(\mathcal{P})$ and $\ell_{CG}(\mathcal{P}, H)$ from January 1995 through December 2015. The right panel shows the difference in RMSEs of forecasted of excess returns between rules $\ell_{CG}(\mathcal{P})$ and $\ell_{CG}(\mathcal{P}, H)$. The RMSEs are computed using an expanding window from December 1995 through December 2014.

basis point post 2000.

Within $\mathcal{BL}$'s learning rule there are two channels through which disagreement (H) can affect expected excess returns. The first is the direct effect that It has on forecasts of future PCs as components of Z_t in (11). The second is the indirect effect of $\mathcal{BL}$'s updating of the parameters $(\widehat{K}^{\mathbb{P}}_{0t}, \widehat{K}^{\mathbb{P}}_{Zt})$ as part of the learning process conditioned on $(\mathcal{P}, H)$. We can frame these effects in terms of risk premiums. Using (11) and the expression for realized excess returns $xr^{n,h}_{t+h}$ in terms of y^n_{t+h}, we can write:

$$exr^n_{1y,t}(\mathcal{P},H) = \hat{a}^{\mathcal{P},H}_{n,t} + \hat{b}^{\mathcal{P},H}_{n,t}\mathcal{P}_t + \hat{c}_{n,t}H_t, \quad (16)$$

$$exr^n_{1y,t}(\mathcal{P}) = \hat{a}^{\mathcal{P}}_{n,t} + \hat{b}^{\mathcal{P}}_{n,t}\mathcal{P}_t. \quad (17)$$

We separate the direct and indirect effects of H in (16) by estimating a rolling constant-gain least-squares projection of H_t onto $\mathcal{P}_t$,

$$H_t = \hat{\alpha}_t + \hat{\beta}_t\mathcal{P}_t + u_t, \quad (18)$$

and then we construct the pseudo expected excess return

$$exr^n_{1y,t}(\mathcal{P},0) \equiv (\hat{a}^{\mathcal{P},H}_{n,t} + \hat{c}_{n,t}\hat{\alpha}_t) + (\hat{b}^{\mathcal{P},H}_{n,t} + \hat{c}_{n,t}\hat{\beta}_t)\mathcal{P}_t \equiv \hat{a}^{\mathcal{P},0}_{n,t} + \hat{b}^{\mathcal{P},0}_{n,t}\mathcal{P}_t. \quad (19)$$

The difference between (19) and (17) 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 $(\mathcal{P}_t, H_t)$.[23]

There are large differences between $exr^{10y}_{1y,t}(\mathcal{P},H)$ and $exr^{10y}_{1y,t}(\mathcal{P},0)$, especially around cyclical turning points (left panel of Figure 10). This is indicative of H having large *direct* effects during major transitions to rising or falling expected excess returns from $\mathcal{BL}$'s perspective. The differences in the root-mean-square forecasting errors for rules $\ell_{CG}(\mathcal{P})$ and $\ell_{CG}(\mathcal{P},H)$ are above 80 bps for most of the sample (right panel), whereas the differences induced by $exr^{10y}_{1y,t}(\mathcal{P})$ and $exr^{10y}_{1y,t}(\mathcal{P},0)$ are always below 40 bps. Thus, following the dot-com bust in 2000, just under half of the increased forecast accuracy of $\ell_{CG}(\mathcal{P},H)$ arises through the *indirect* effects of H on the updating of the parameters by $\mathcal{BL}$'s learning rule; the *direct* predictive power of H for future $\mathcal{P}$ that is orthogonal to $\mathcal{P}_t$ accounts for the other half plus/minus.

5 Macroeconomic Information and Disagreement About Yields

Within-sample analyses in the absence of learning show that macroeconomic fundamentals have predictive power for excess returns in bond markets (e.g., Ludvigson and Ng (2010) and JPS). 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?

[23]While it is true that (18) 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 $exr^n_{1y,t}(\mathcal{P},0)$.

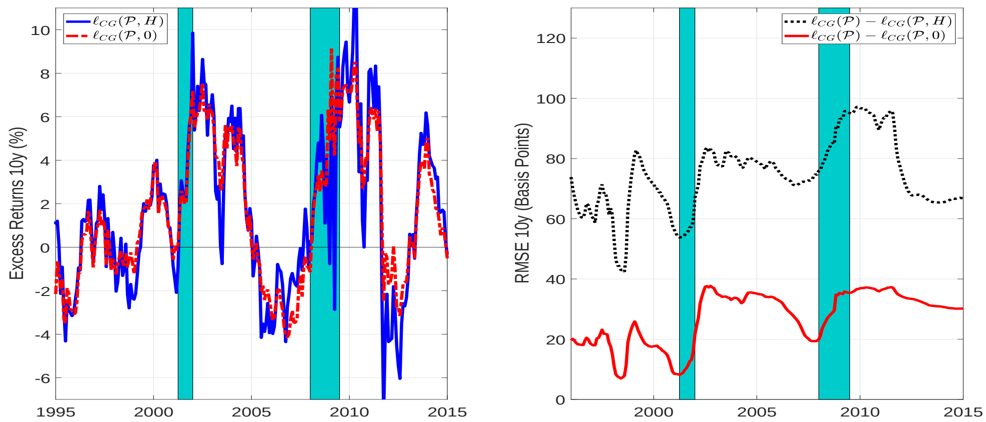

Figure 10: The left panel shows estimated using rule $\ell_{CG}(\mathcal{P}, H)$ and rule $\ell_{CG}(\mathcal{P}, 0)$ over the sample from January 1995 through December 2015. The panel on the right shows the difference in expected returns RMSE between rule $\ell_{CG}(\mathcal{P}, H)$ and rule $\ell_{CG}(\mathcal{P}, 0)$. RMSEs are calculated using the same methodology as in figure 9.

5.1 Inflation, Output Growth, and Real-Time Expected Excess Returns

An important issue when studying the predictive power of macro-variables for future yields is that official macro time-series are regularly updated after their original release date. Thus, the current releases of macro time-series are not the same as the ones available to investors in real time. 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

$$x_{t_0-1|t}, x_{t_0|t}, \ldots, x_{t-2|t}, x_{t-1|t},$$

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.[24] *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

[24]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.

	2Y	3Y	5Y	7Y	10Y
Part A: January, 1995 – December, 2014					
$\ell_{CG}(\mathcal{P})$	1.06%	1.97%	3.44%	4.82%	6.61%
$\ell_{CG}(\mathcal{P}, H)$	1.05%	1.92%	3.22%	4.43%	5.94%
$\ell_{CG}(\mathcal{P}, REA)$	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%
$\ell_{CG}(\mathcal{P}, H, REA)$	1.03%	1.93%	3.39%	4.77%	6.53%
$\ell_{CG}(\mathcal{P}, H, REA, INF)$	1.23%	2.31%	3.89%	5.33%	7.07%
$\ell_{CG}(\mathcal{P}, ID(RGDP), ID(INF))$	1.20%	2.16%	3.71%	5.23%	7.14%
$\ell_{CG}(\mathcal{P}, Cons(RGDP), Cons(INF))$	1.22%	2.21%	3.85%	5.31%	7.17%
Part B: January, 2001 – December, 2007					
$\ell_{CG}(\mathcal{P})$	1.35%	2.48%	3.94%	5.10%	5.76%
$\ell_{CG}(\mathcal{P}, H)$	1.32%	2.38%	3.62%	4.51%	4.76%
$\ell_{CG}(\mathcal{P}, REA)$	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%
$\ell_{CG}(\mathcal{P}, H, REA)$	1.18%	2.22%	3.73%	5.14%	6.45%

Table 4: 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.

of six series related to real economic activity.[25]

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.[26] 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 in the BCFF panel. 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.

To formally evaluate the contribution of macro variables to $\mathcal{BL}$'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. Over the full sample the macro learning rules $\ell_{CG}(\mathcal{P}, REA)$ and $\ell_{CG}(\mathcal{P}, REA, INF)$ perform comparably to $\ell_{CG}(\mathcal{P})$ for maturities up to 5 years, and

[25] 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.

[26] 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.

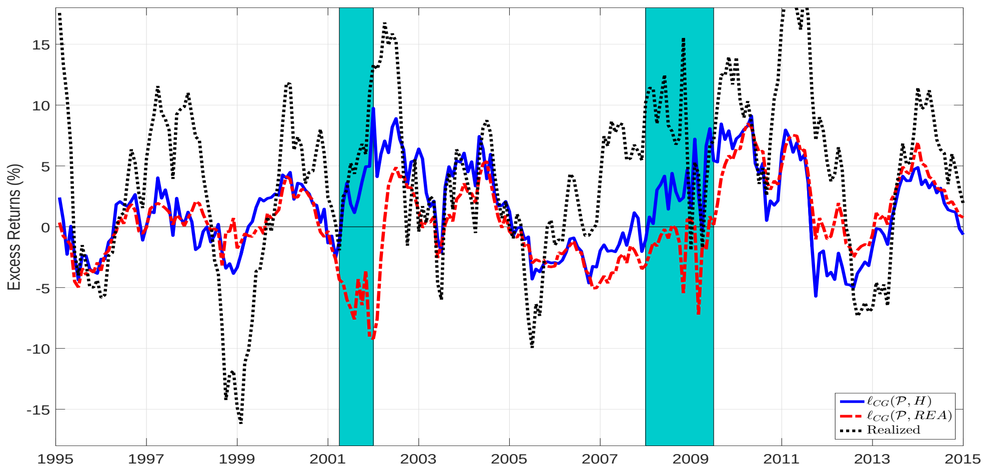

Figure 11: Average expected excess returns over holding periods of 10, 11, and 12 months for ten-year bonds based on $\ell_{CG}(\mathcal{P}, H)$ and $\ell_{CG}(\mathcal{P}, REA)$, overlaid with the realized returns.

underperform for longer maturities (Table 4). Leading up to the global financial crisis– 2001 to 2007– $\ell_{CG}(\mathcal{P}, REA)$ substantially underperforms $\ell_{CG}(\mathcal{P})$ for longer maturity bonds.[27]

By contrast, rule $\ell_{CG}(\mathcal{P}, H)$ reduces the RMSE of the 10-year bond return by 0.7% relative to rule $\ell_{CG}(\mathcal{P})$ (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). The relative outperformance of rule $\ell_{CG}(\mathcal{P}, H)$ over $\ell_{CG}(\mathcal{P}, REA)$ can be seen graphically from Figure 11, where expected excess returns for the ten-year bonds are plotted against the realized excess returns. The outperformance of $\ell_{CG}(\mathcal{P}, H)$ is at times large, especially during and immediately after NBER recessions. The primary exception when $\ell_{CG}(\mathcal{P}, REA)$ outperforms is during portions of the post-crisis period of 2011-12. Conditioning on both disagreement and current real economic activity (rule $\ell_{CG}(\mathcal{P}, H, REA)$) delivers RMSEs that are very close to those of $\ell_{CG}(\mathcal{P}, H)$.

Combining $\mathcal{P}$ with information about the beliefs of professionals about the future macroeconomy (either $(Cons(RGDP)_t, Cons(INF)_t)$ or $(ID(RGDP)_t, ID(INF)_t)$ leads to sizable *deteriorations* in forecasting accuracy relative to rule $\ell_{CG}(\mathcal{P}, H)$, *across the entire maturity spectrum*. A similar deterioration in fit arises under rule $\ell_{CG}(\mathcal{P}, H, REA, INF)$. Thus, the forecasting power of H for the long end of the curve is largely distinct from that of information about inflation and real output growth.

[27]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. Yet, consistent with JPS, there is evidence of some predictive power for REA, particularly during the first part of the 2000's.

5.2 The Information Content of Yield Disagreement: Policy Uncertainty?

Having found that disagreement is not proxying for information about output or inflation, we turn next to the possibility that the predictive power of H is related to uncertainty about macroeconomic policy. Letting EPU denote one of the indices of Economic Policy Uncertainty developed by Bloom, Backer, and Davis (2016), we construct RLS estimates of the projections

$$xr^n_{t+h} = \gamma^n_{0,t} + \gamma^n_{\mathcal{P},t}\mathcal{P}_t + \gamma^n_{\widehat{EPU},t}\widehat{EPU}_t + e_{x,t+h} \quad (20)$$

$$EPU_{t+h} = \gamma^{EPU}_{0,t} + \gamma^{EPU}_{\mathcal{P},t}\mathcal{P}_{t+h} + \gamma^{EPU}_{H,t}H_{t+h} + e_{u,t+h} \equiv \widehat{EPU}_{t+h} + e_{u,t+h} \quad (21)$$

where xr^n_{t+h} is the realized excess return of a zero coupon bond with maturity n over the period from t to $t+h$. Since these projections condition on $\mathcal{P}_t$, $\gamma^n_{\widehat{EPU},t}$ effectively reveals the extent to which H_t is informative about future excess returns, over and above its correlation with $\mathcal{P}_t$, through its association with economic policy uncertainty.

Our prior is that disagreement about future yields will, to varying degrees over time, reflect a variety of sources of risks for bond markets. We focus on two measures of policy uncertainty: EPU_{news} that is constructed from textual analysis of major newspapers based on the terms 'uncertainty' or 'uncertain' and 'economic' or 'economy,' along with one or more of the following terms: 'congress', 'legislation', 'white house', 'regulation', 'federal reserve', or 'deficit;" and EPU_{fed} that measures the more specific uncertainty about federal and local government spending.[28]

Figure 12 shows real-time recursive least squares estimates of one-quarter ahead ($h = 3$) expected excess returns for a 10 year bond based on (20), and compares them with expected excess returns from real-time conditioning only on $\mathcal{P}$ and on $(\mathcal{P}, H)$. Notice, first of all, that rules $\ell_{CG}(\mathcal{P}, H)$ and $\ell_{CG}(\mathcal{P})$ generate very different paths for expected excess returns. Additionally, the excess returns from the rule $\ell_{CG}(\mathcal{P}, \widehat{EPU}_{news})$ track the returns from $\ell_{CG}(\mathcal{P})$ much more closely than they track returns from $\ell_{CG}(\mathcal{P}, H)$, especially during the late 1990's and from the onset of the global financial crisis onward. Thus, this broad measure of economic policy uncertainty is evidently not the key source of the predictive power for H.

On the other hand, and quite strikingly, from mid 1998 through the remainder of our sample, $\mathcal{BL}$'s expected excess returns and those from rule $\ell(\mathcal{P}, \widehat{EPU}_{fed})$ are nearly identical. At least for this extended period of time, it appears that the predictive power of H in $\mathcal{BL}$'s learning rule is associated with an impact of fiscal policy uncertainty on bond yields. The

[28] These are ailable at http://www.policyuncertainty.com. We also examined measures of uncertainty related to tax and monetary policies and found them to be only weekly correlated with yield disagreement. Also, replacing H_t with $H^{M\prime}_t = (ID(CPI), ID(RGDP))$ in (21) leads to a substantial deterioration in explanatory power in these projections, which reinforces our earlier finding that disagreement about inflation and output is not the source of the predictive power of H_t for excess bond returns.

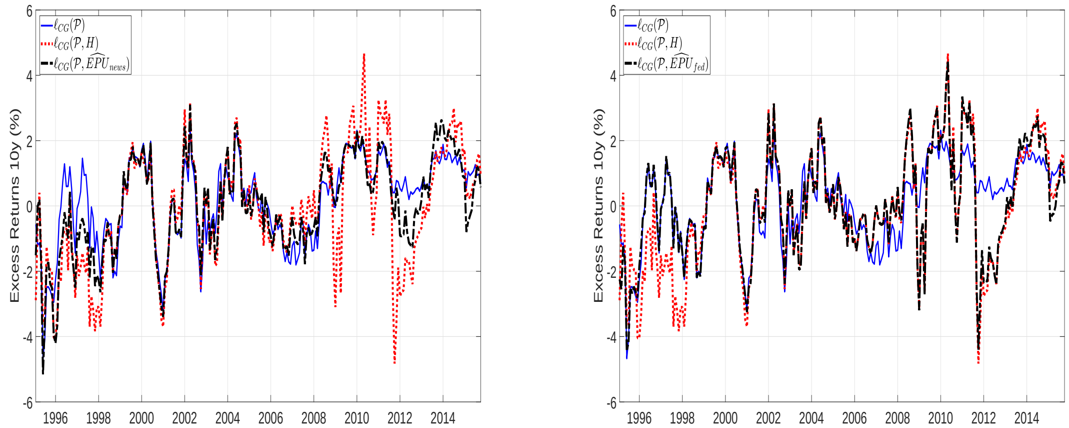

Figure 12: Recursive-real time estimates of one quarter-ahead 10-year bond excess returns conditioning on $\widehat{EPU}_{news}$ (left panel) and $\widehat{EPU}_{fed}$ (right panel).

late 1990's are again an exceptional period. This was the period of fiscal conservatism under the Clinton administration with a policy focus on deficit reduction. This focus may well have weakened the association between fiscal policy uncertainty and professionals' disagreement about the future course of interest rates.

To be clear, we are not saying that uncertainty about fiscal policy drives (spans) the observed dispersion of yield forecasts by professionals. On the contrary, $EPU_{fed,t}$ has limited explanatory power for H_t, and this is especially true when projecting the slope of the term structure of disagreement onto $EPU_{fed,t}$. What are findings suggest is that the component of $EPU_{fed,t}$ that is spanned by H_t has roughly the same correlation with future bond yields as H_t. Put differently, it is the component of the slope of the disagreement curve correlated with $EPU_{fed,t}$ that has strong predictive power for future bond yields in $\mathcal{BL}$'s learning rule.

6 Concluding Remarks

Appendices

A Log likelihood function

We begin by noting that when $(\Theta^{\mathbb{Q}}, \Sigma_{\mathcal{O}})$ and Σ_Z are presumed to be constant, equation (15) 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,

$$Z_{t+1} = K_{Z,0t}^{\mathbb{P}} + K_{Z,1t}^{\mathbb{P}} Z_t + \Sigma_Z^{1/2} e_{Z,t+1}^{\mathbb{P}}, \tag{22}$$

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

$$\psi_t^{\mathbb{P}} = \psi_{t-1}^{\mathbb{P}} + Q_{t-1}^{1/2}\eta_t \qquad \eta_t \overset{iid}{\sim} N(0, I), \tag{23}$$

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 (22) into

$$\mathcal{Y}_t = \mathcal{X}_{t-1}\psi_{t-1}^{\mathbb{P}} + \Sigma_Z^{1/2} e_{Z,t}^{\mathbb{P}}, \tag{24}$$

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 (23) (the transition equation) and to the factor dynamics (24) (the measurement equation) we have a well-defined linear Kalman filter.[29] 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 $\Omega_t = \mathbb{V}^{\mathbb{P}}(Z_{t+1} | Z_1^t)$

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

according to:

$$\begin{aligned}\hat{\psi}_t^{\mathbb{P}} &= \hat{\psi}_{t-1}^{\mathbb{P}} + P_{t-1}\mathcal{X}_{t-1}'\Omega_{t-1}^{-1}(\mathcal{Y}_t - \mathcal{X}_{t-1}\hat{\psi}_{t-1}^{\mathbb{P}}), &(25)\\ P_t &= P_{t-1} + Q_{t-1} - P_{t-1}\mathcal{X}_{t-1}'\Omega_{t-1}^{-1}\mathcal{X}_{t-1}P_{t-1}, &(26)\\ \Omega_{t-1} &= \mathcal{X}_{t-1}P_{t-1}\mathcal{X}_{t-1}' + \Sigma_Z, &(27)\end{aligned}$$

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

$$\begin{aligned}-2\log L^{\mathbb{P}} &= (t-1)N\log(2\pi) + \sum_{s=2}^{t}\log|\Omega_{s-1}| &(28)\\ &+ \frac{1}{2}\sum_{s=2}^{t}(\mathcal{Y}_s - \mathcal{X}_{s-1}\hat{\psi}_{s-1})'\Omega_{s-1}^{-1}(\mathcal{Y}_s - \mathcal{X}_{s-1}\hat{\psi}_{s-1}).\end{aligned}$$

Reworking equation (25) gives[30]

$$\hat{\psi}_t^{\mathbb{P}} = \hat{\psi}_{t-1}^{\mathbb{P}} + (P_t - Q_{t-1})\mathcal{X}_{t-1}'\Sigma_Z^{-1}\left(\mathcal{Y}_t - \mathcal{X}_{t-1}\hat{\psi}_{t-1}^{\mathbb{P}}\right). \quad (29)$$

Letting $R_t = (P_t - Q_{t-1})^{-1}$, (29) reduces to the first equation in the definition of an adaptive least squares estimator (see (**??**)). Equation (26) can then be rewritten as[31]

$$\begin{aligned}(P_t - Q_{t-1})^{-1} &= P_{t-1}^{-1} + \mathcal{X}_{t-1}'\Sigma_Z^{-1}\mathcal{X}_{t-1} &(30)\\ &= (I - P_{t-1}^{-1}Q_{t-2})(P_{t-1} - Q_{t-2})^{-1} + \mathcal{X}_{t-1}'\Sigma_Z^{-1}\mathcal{X}_{t-1},\end{aligned}$$

which reduces to (**??**) if $Q_{t-2} = (1-\gamma_{t-1})P_{t-1}$, for a sequence of scalars $0 < \gamma_t \leq 1$. Using (26) 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

[30]Substitute (27) into (26) and the resulting equation into (25).

[31]This expression is obtained by substituting (27) into (26), plugging the resulting equation back into (26), and multiplying by $(P_t - Q_{-1})^{-1}$ from the left and P_{t-1}^{-1} from the right.

as γ_t is measurable. We can summarize the preceding calculations as:

$$\begin{aligned}
R_t\hat{\psi}_t^{\mathbb{P}} &= \gamma_{t-1}R_{t-1}\hat{\psi}_{t-1}^{\mathbb{P}} + \mathcal{X}_{t-1}'\Sigma_Z^{-1}\mathcal{Y}_t, && (31)\\
R_t &= \gamma_{t-1}R_{t-1} + \mathcal{X}_{t-1}'\Sigma_Z^{-1}\mathcal{X}_{t-1}, && (32)\\
\hat{\psi}_t^{\mathbb{P}} &= R_t^{-1}R_t\hat{\psi}_t^{\mathbb{P}}, && (33)\\
P_t &= \frac{1}{\gamma_t}R_t^{-1}, && (34)\\
\Omega_{t-1} &= \mathcal{X}_{t-1}P_{t-1}\mathcal{X}_{t-1}' + \Sigma_Z, && (35)
\end{aligned}$$

with log likelihood function given by (28). The constant gain estimator corresponds to the special case where $\gamma_t = \gamma$ for all t.

B Pricing Kernel

The pricing kernel can be expressed as

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

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

$$\begin{aligned}
f_{t,t+1}^{\mathbb{P}} &= N(\hat{K}_{\mathcal{P}0,t}^{\mathbb{P}} + [\hat{K}_{\mathcal{P}\mathcal{P},t}^{\mathbb{P}}, \hat{K}_{\mathcal{P}H,t}^{\mathbb{P}}]Z_t, \Omega_{\mathcal{P}\mathcal{P},t}) = N(\hat{\mu}_t^{\mathbb{P}}, \Omega_{\mathcal{P}\mathcal{P},t}),\\
f_{t,t+1}^{\mathbb{Q}} &= N(K_0^{\mathbb{Q}} + K_1^{\mathbb{Q}}\mathcal{P}_t, \Sigma_{\mathcal{P}\mathcal{P}}) = N(\mu_t^{\mathbb{Q}}, \Sigma_{\mathcal{P}\mathcal{P}}),
\end{aligned}$$

where $(\hat{K}_{\mathcal{P}0,t}^{\mathbb{P}}, [\hat{K}_{\mathcal{P}\mathcal{P},t}^{\mathbb{P}}, \hat{K}_{\mathcal{P}H,t}^{\mathbb{P}}])$ denote the posterior means of the latent parameters states, and $\Omega_{\mathcal{P}\mathcal{P},t}$ the upper left 3×3 entries of the conditional covariance matrix Ω_t given in equation (35). 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):

$$\begin{aligned}
\log \mathcal{M}_{t,t+1} + r_t &= c_t + \frac{1}{2}(\mathcal{P}_{t+1} - \hat{\mu}_t^{\mathbb{P}})'\Omega_{\mathcal{P}\mathcal{P},t}^{-1}(\mathcal{P}_{t+1} - \hat{\mu}_t^{\mathbb{P}}) - \frac{1}{2}(\mathcal{P}_{t+1} - \mu_t^{\mathbb{Q}})'\Sigma_{\mathcal{P}\mathcal{P}}^{-1}(\mathcal{P}_{t+1} - \mu_t^{\mathbb{Q}})\\
&= c_t' - \left(\Omega_{\mathcal{P}\mathcal{P},t}^{-1}\hat{\mu}_t^{\mathbb{P}} - \Sigma_{\mathcal{P}\mathcal{P}}^{-1}\mu_t^{\mathbb{Q}}\right)'\mathcal{P}_{t+1} + \frac{1}{2}\mathcal{P}_{t+1}'(\Omega_{\mathcal{P}\mathcal{P},t}^{-1} - \Sigma_{\mathcal{P}\mathcal{P}}^{-1})\mathcal{P}_{t+1}\\
&= c_t'' - \Lambda_{\mathcal{P}t}'\Gamma_t^{-1}\varepsilon_{t+1}^{\mathbb{P}} + \frac{1}{2}(\varepsilon_{t+1}^{\mathbb{P}})'\left(I - \Gamma_t^{-1}\right)\varepsilon_{t+1}^{\mathbb{P}},
\end{aligned}$$

where

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

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 $\mathcal{P}_t$: $r_t^a = \alpha + \beta'\mathcal{P}_t$ and satisfying $\mathbb{E}_t\left[e^{r_t^a+1}\mathcal{M}_{t,t+1}\right] = 1$. 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

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

This leads to

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

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

C Selecting the Constant Gain Coefficient γ

Which value of the constant gain coefficient γ would be selected by an econometrician using the model $\ell_{CG}^L(\mathcal{P})$, developed in section 4.1, to forecast future yields? To answer this question, we let the econometrician select the value of γ in real time. We estimate the model on an equally spaced grid of constant gain parameters: $\{0.95, 0.955, ..., 0.995, 1\}$. We collect the three- and twelve-month ahead forecasts for each value of γ and for each month from January 1982 through December 2014. In each month, we select the "optimal" γ, as the value that minimizes the RMSE of out-of-sample forecasts for the first PC over the previous ten years. For example, in January 1982 we choose γ to minimize the out-of-sample RMSE between

January 1971 (post training period) and January 1981[32]. **??** shows the evolution of the dynamically updated γ parameters (γ_t) over the sample from January 1982 through December 2014. For one-year-ahead forecasts, the constant gains coefficient takes a value of 0.99 for most of the sample (Panel (b)), and especially in the period from the late 1990s to the end of the sample, which is the focus of most of the analysis in the paper. The γ_t is more volatile when minimizing in real-time the one-quarter-ahead RMSE of forecasts (Panel (a)), though γ_t takes values close to 0.99 for large past of the sample, and especially in the latter years.

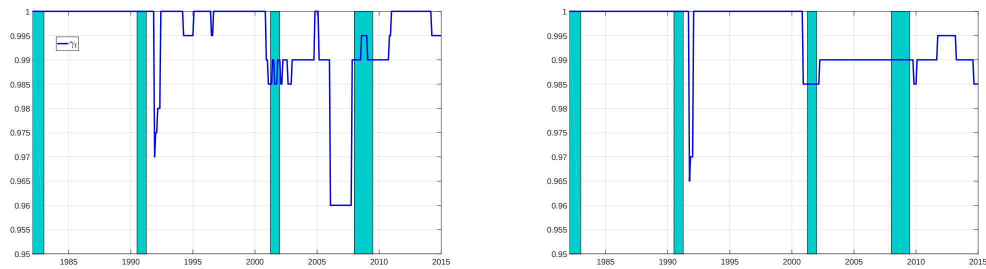

(a) Quarterly Forecasts

(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.

Figure 14 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, and for different values of γ. The minimal RMSE for one-year-ahead forecasts (Panel (b)) is achieved for $\gamma = 0.99$. For the one-quarter-ahead forecasts (Panel (a)), the smallest RMSE is delivered by values of γ equal to 0.99 and 0.995.

Overall, the choice of $\gamma = 0.99$ seems to appear optimal (in the sense of minimizing forecast errors) or close to optimal both ex-ante and ex-post.

D 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

[32]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.

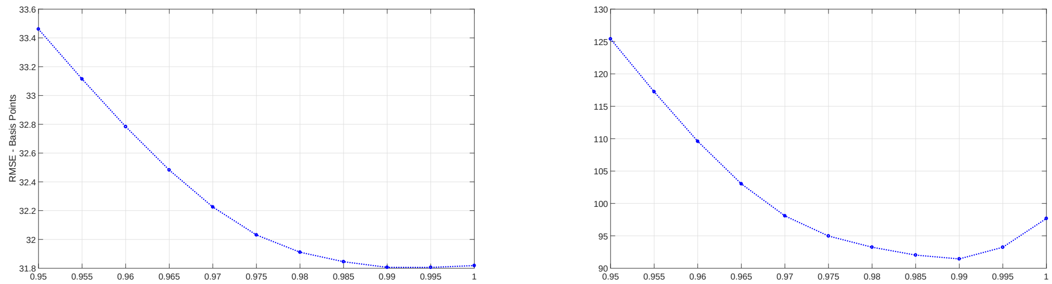

(a) Quarterly Forecasts

(b) Annual Forecasts

Figure 14: 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.

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, 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 5 displays the restrictions imposed on the autoregressive feedback matrix for rule $\ell_{CG}(\mathcal{P}, H)$. Similarly, Table 6 reports the restrictions for rule $\ell_{CG}(\mathcal{P})$.

E 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

	Λ_{0t}	Λ_{1t}				
		PC_1	PC_2	PC_3	$ID(y^{2y})$	$ID(y^{10y})$
PC_1	*	*	*	0	*	*
PC_2	*	0	*	*	0	0
PC_3	0	0	0	0	0	0
$ID(y^{2y})$	0	*	*	*	*	*
$ID(y^{10y})$	0	*	0	0	0	*

Table 5: 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 6: Restrictions applied in rule $\ell_{CG}(\mathcal{P})$ to the parameters in PmQ_t.

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

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

The innovation to the non-volatility factors, $\varepsilon^{\mathbb{Q}}_{t+1}$, 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

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

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

$$V_{t+1}|V_t \sim CAR(\rho^{\mathbb{P}}, c^{\mathbb{P}}, v^{\mathbb{P}}), \tag{36}$$

$$X_{t+1} = K^{\mathbb{P}}_{X0,t} + K^{\mathbb{P}}_{XV,t}V_t + K^{\mathbb{P}}_{XX,t}X_t + \sqrt{\Sigma_{0X} + \Sigma_{1X}V_t} \cdot \varepsilon^{\mathbb{P}}_t. \tag{37}$$

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 $\mathcal{P}_t = Wy_t$. 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 (37) with

$$\mathcal{P}^{2:3}_{t+1} - W^{2:3}B_V V_{t+1} = \tilde{K}^{\mathbb{P}}_{\mathcal{P}0,t} + \tilde{K}^{\mathbb{P}}_{\mathcal{P}V,t}V_t + \tilde{K}^{\mathbb{P}}_{\mathcal{P}\mathcal{P},t}\mathcal{P}^{2:3}_t + \sqrt{\tilde{\Sigma}_{0\mathcal{P}} + \tilde{\Sigma}_{1\mathcal{P}}V_t} \cdot \varepsilon^{\mathbb{P}}_t, \tag{38}$$

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 $(V_t, (\mathcal{P}^{2:3}_t)')'$ (and not $\mathcal{P}_t$). Therefore, the model's parameters can be decomposed into constant and known $\mathbb{Q}$-parameters $(r^{\mathbb{Q}}_{\infty}, \rho_V, \rho^{\mathbb{Q}}, c^{\mathbb{Q}}, v^{\mathbb{Q}}, K^{\mathbb{Q}}_{XV}, \lambda^{\mathbb{Q}})$, 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 $(\tilde{K}^{\mathbb{P}}_{\mathcal{P}0,t}, \tilde{K}^{\mathbb{P}}_{\mathcal{P}V,t}, \tilde{K}^{\mathbb{P}}_{\mathcal{P}\mathcal{P},t})$. 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 (36) - (38) it is seen that the conditional first and second moments of the principal principal components are given by

$$\mathbb{E}_t^{\mathbb{P}}(\mathcal{P}_{t+1}) = K_{0,t}^{\mathbb{P}} + K_{1,t}^{\mathbb{P}}\mathcal{P}_t \tag{39}$$

$$\mathbb{V}_t^{\mathbb{P}}(\mathcal{P}_{t+1}) = \Sigma_0^{\mathbb{P}} + \Sigma_1^{\mathbb{P}}\mathcal{P}_t \tag{40}$$

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.

F Robustness to Time-Varying Volatility

A less constrained Bayesian (relative to $\mathcal{BL}$) 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 $\mathcal{BL}$'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 $\mathcal{BL}$ 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 $\mathbb{Q}$:

$$V_{t+1}|V_t \sim CAR(\rho^{\mathbb{Q}}, c^{\mathbb{Q}}, v^{\mathbb{Q}}), \tag{41}$$

$$X_{t+1} = K_V^{\mathbb{Q}}V_t + diag(\lambda^{\mathbb{Q}})X_t + \sqrt{\Sigma_{0X} + \Sigma_{1X}V_t}\,\varepsilon_t^{\mathbb{Q}}, \tag{42}$$

$$r_t = r_\infty^{\mathbb{Q}} + \rho_V V_t + 1'X_t, \tag{43}$$

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 $\mathcal{BL}$ 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 E for details). The market prices of risk are assumed to be such that, under $\mathbb{P}$, the state follows the process

$$V_{t+1}|V_t \sim CAR(\rho^{\mathbb{P}}, c^{\mathbb{P}}, v^{\mathbb{P}}), \tag{44}$$

$$X_{t+1} = K_{0t}^{\mathbb{P}} + K_{Vt}^{\mathbb{P}}V_t + K_{Xt}^{\mathbb{P}}X_t + \sqrt{\Sigma_{0X} + \Sigma_{1X}V_t}\,\varepsilon_{t+1}^{\mathbb{P}}, \tag{45}$$

where $\varepsilon_{t+1}^{\mathbb{P}}$ is independent of V_{t+1} and we let $\Theta_t^{\mathbb{P}} = (\rho^{\mathbb{P}}, c^{\mathbb{P}}, v^{\mathbb{P}}, K_{0t}^{\mathbb{P}}, K_{Vt}^{\mathbb{P}}, K_{Xt}^{\mathbb{P}})$. $\mathcal{BL}$ 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 E 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 15 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 $\mathcal{BL}$'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 $(\rho^{\mathbb{P}}, eig(K_{Xt}^{\mathbb{P}}))$.[33] 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 eigenvalues of $K_{1\mathcal{P},t}^{\mathbb{P}}$ reflects the substantial variation in the market prices of risk.

Figure 16 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

[33]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 $(V_t, X_t')'$.

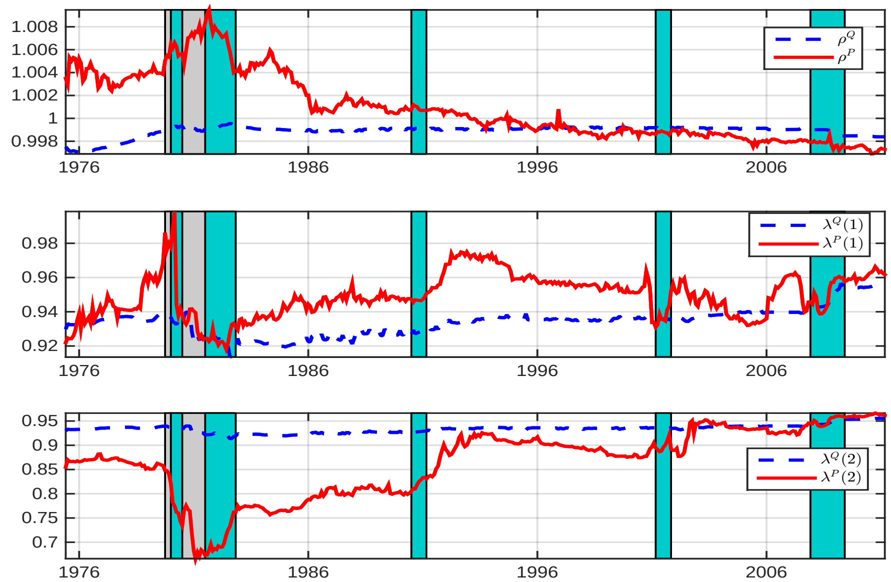

Figure 15: Estimates from model $\ell_{CG}^{1,3}(\mathcal{P})$ of the eigenvalues of the feedback matrix $K_{1\mathcal{P}}^{\mathbb{Q}}$ ($K_{1\mathcal{P},t}^{\mathbb{P}}$). The eigenvalues of $K_{1\mathcal{P}}^{\mathbb{Q}}$ are ($\rho^{\mathbb{Q}}, \lambda^{\mathbb{Q}}$) and the eigenvalues of $K_{1\mathcal{P},t}^{\mathbb{P}}$ are ($\rho^{\mathbb{P}}, eig(K_{Xt}^{\mathbb{P}})$).

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.

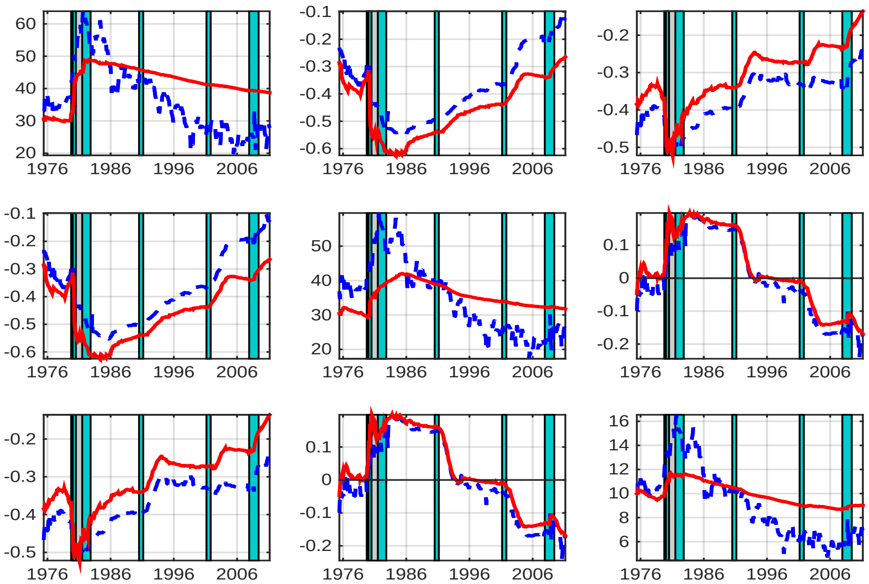

Figure 16: Summary of $\Sigma_{\mathcal{P}\mathcal{P}}$. Conditional standard deviations (main diagonal elements) and correlations (off-diagonal elements) estimates from learning models with (blue line) and without (red line) stochastic volatility. The estimates at date t are based on the historical 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*.

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

Bloom, N., S. Backer, and S. Davis, 2016, Measuring Economic Policy Uncertainty, *Quarterly Journal of Economics*.

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.

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.

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

Dewachter, H., and M. Lyrio, 2008, Learning, Macroeconomic Dynamics and the Term Structure of Interest Rates, in *Asset Prices and Monetary Policy* (National Bureau of Economic Research,).

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

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.

Gargano, A., D. Pettenuzzo, and A. Timmerman, 2018, Bond Return Predictability: Economic Value and Links to the Macroeconomy, Working paper, forthcoming, *Management Science*.

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,).

Laubach, T., R. Tetlow, and J. Williams, 2007, Learning and the Role of Macroeconomic Factors in the Term Structure of Interest Rates, Working paper, Board of Governors of the Federal Reserve System.

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.

Litterman, R., and J. Scheinkman, 1991, Common Factors Affecting Bond Returns, *Journal of Fixed Income* 1, 54–61.

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

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., 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.

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

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