LEARNING ABOUT BOND PRICES*

Pau Belda[1], Albert Marcet[2], and Kenneth Singleton[3]

[1]Bank of England
[2]Centre de Recerca en Economia Internacional (CREI), ICREA, BSE and UPF
[3]Stanford University

November 2025
Incomplete Manuscript. Please do not circulate

Abstract

We document that survey expectations systematically underestimate how the yield curve slope predicts future excess bond returns, violating Rational Expectations. We rationalize this fact via a model where investors have imperfect knowledge about the equilibrium pricing function and learn about it. Learning imparts inertia in bond prices, which naturally generates well-known features of US bond markets — high excess return volatility, a flat yield volatility term structure, and Expectations Hypothesis violations. Methodologically, we extend the Internal Rationality framework — where agents optimize given their subjective beliefs — to multiple assets, which requires specifying a joint perceived distribution of bond prices.

1 Introduction

Understanding bond yield dynamics remains one of the most important challenges in asset pricing. While existing models have made substantial progress, simultaneously

*Marcet acknowledges funding from the Axa Foundation, the Excellence Program of Banco de España, the European Research Council under grants FP/2007-2013 n. 324048 (APMPAL) and Horizon2020 GA n. 788547 (APMPAL-HET), AGAUR (Generalitat de Catalunya), and the Spanish Ministry of Economy and Competitiveness through the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S). The views expressed here are solely of the authors and do not represent, in any way, those of the Bank of England or any of its committees.

accounting for multiple empirical regularities remains challenging: the high volatility of long-term yields and excess returns, the failure of the Expectations Hypothesis, and the behavior of yield curve forecasts. We propose that these empirical features emerge naturally and parsimoniously once we recognize that investors have to learn about bond pricing itself.

This paper develops a model where investors are Internally Rational (IR): they optimize given their understanding of how markets work, but this understanding—while sophisticated—remains imperfect. The key insight is that bond prices are inherently self-referential. Long bond prices depend on expectations of future bond prices, creating a feedback loop where beliefs about yield trends become self-fulfilling. Even though this equilibrium concept is a small deviation from the standard paradigm the behavior of bond prices can be very different from Rational Expectations (RE). Even though the model is extremely simple it improves the match of several important empirical observations relative to other alternatives in the literature.

The basic idea hinges on the fact that long bond prices depend on investor's expectation about tomorrow's long bond price. When investors learn about the underlying trend in bond prices: a high *expected* slope drives long bond prices up, causing *actual* slope to be high, therefore reinforcing agents' belief, driving future *expectations* of the slope up and future slopes up and so on. This mechanism, where high expectations about bond yields reinforce actually high bond yields can potentially cause large deviations from RE and generate fluctuations of their own due to the self-referential aspect of learning about asset prices.

A key ingredient of a model under IR is specifying agents' understanding of prices. The aim is to formulate these price beliefs in such a way that actual price behavior is not too different from agents' expectation. A modelling difficulty relative to stock price applications of IR (see, for example, Adam, Marcet and Nicolini (2016), Adam, Marcet and Beutel (2017), Winkler (2020), Belda (2025)) is that many assets (bonds of different maturities) are traded at any given time, therefore, we need to specify the *joint* distribution of bonds of all maturities.[1] Therefore we need to specify the *joint* distribution of bond prices that agents have in mind. For this, we model investors' perceived bond prices as a perceived yield curve.

Even though the agents' beliefs are specified by the modeller most of the work

[1]In the mentioned stock price papers under IR what matters for today's stock price is the price for the same asset next period. With bonds today's bond price of maturity n depends on the expected price of a a different asset, namely a bond of maturity $n-1$.

will be using a very simple equilibrium model where the only driving state variables are inflation and short rates determined by policy. The rest of the model is completely standard, with homogeneous agents, time-separable utility, no transaction costs except for limits to short bond holdings, etc.

One aim of the paper is to match some basic second moments of yields an excess long bond reeturns, including variances and correlations across bond maturities and across time. We also study other issues addressed in the literature such as the failure of the Expectations Hypothesis and the flat term structure of yield volatilities. A large literature tries to understand many of these features of yields using equilibrium models under the RE or under learning about fundamentals. This literature has contributed much to our understanding of how to use equilibrium models, a brief and unfair summary is that there are still many puzzles out there that are difficult to explain jointly by one single model with few parameters, for a more detailed account see section 2.

Furthermore, we aim at explaining survey behavior. Some recent work by Giacoletti, Larsen and Singleton (2020) (GLS) and Singleton (2021) focuses on the behavior of survey expectations about bond yields as reported in the BCFF. We focus here on "RE regressions" reporting how the current slope explains both actual and expected bond returns. These regressions can be used to test RE but also to characterize the behavior of surveys (see Adam, Marcet and Beutel (2017)). We find that survey expectations systematically overestimate the effect of the current slope on future yields relative to objective expectations. This suggests that it is important to incorporate deviations from RE in understanding bond prices and/or that a good model will have to explain this bias observed in surveys.

The IR literature requires that even though investors' perceptions are not perfect (not RE) agents' perceptions should not be easy to reject. In our model, the equilibrium term structure will be approximately like the one agents expect: agents' perceived slope will actually be, approximately, the correct slope, even if their belief is far from the slope that would occur under RE. Furthermore, just because agents believe that the level and slope factors have a trend that moves slowly over time they will learn about the level and this will cause the actual level and slope to be time-varying. Thus optimal behavior given agents' model of bond prices generates a yield curve with long frequency movements similar to the perceived ones. This feature of the model will play a key role in explaining all the observations mentioned above, including the RE regressions.

We show through a series of analytic and simulation results that the model can

account for some basic features of the data that have been very difficult to explain with equilibrium models. In particular, we can show that the perceived underlying slope is a state variable of the model that imparts additional volatility to yields, more so for longer maturities. This explains the higher volatility of long yields, the very high volatility of excess returns, the failure of expectations hypothesis and the RE regressions.

Under IR agents are fully rational in the sense that they maximize their dynamic utility and they use efficiently all the information available, even though their view about prices differs from the actual one generated by the model and this perceived price process entails some learning about price behavior. The profession often takes the view that if investors are rational asset prices should equal the discounted sum of future payoffs. This questions our claim that agents who learn about prices can be consistent with utility optimization. We show that our model is consistent once the investors' state space is appropriately defined by including asset prices asin Adam and Marcet (2011). More specific to bond prices may be the view that optimality conditions imply a factor model of bond prices as in the affine term structure literature, with a certain loadings recursion as in Duffie and Kan (1996), therefore price beliefs should incorporate this factor structure. We will show how these restrictions are compatible with the price beliefs of our agents.

Although much progress has been made in the field of understanding how to model expectations deviating from RE, much of the profession sees this as some sort of arbitrary exercise. According to this view there many ways to deviate from RE therefore there are infinite degrees of freedom once you deviate from RE and one can match anything you want. This is however not the necessarily case. First because we use an extremely simple model with very few fundamental free parameters. Second because we require that agents' beliefs are somehow validated by the behavior of the model, so that with finite observations it would be impossible to reject the model. Third because we use survey data as guidance as to what is a reasonable model of bond price expectations.

Section 2 reviews the literature in more detail. Section 3 documents facts about realized and expected yields. Section 4 introduces the model. Section 5 discusses the foundations of Internal Rationality. Section 6 computes the equilibrium, derives analytical results, and shows how the model with learning can qualitatively account for the empirical observations.

2 Literature Review

Some papers address observations on the first and second moments of yields and the failure of the Expectations Hypothesis (EH) by introducing time-varying risk premia, including Dai and Singleton (2002), Wachter (2006), Gabaix (2008), and Bansal and Shaliastovich (2013). Various RE models can match some aspects of second moments, but they do not reproduce the EH and return-forecasting regressions we emphasize. The excess volatility of long-bond excess holding-period returns documented since Shiller (1979) has been rarely addressed; Giglio and Kelly (2018) argue that standard yield-curve models will not be able to generate this volatility. At the same time, a growing empirical literature based on survey data questions RE and documents systematic forecast errors and information rigidities, e.g. Coibion and Gorodnichenko (2015), Adam, Marcet and Beutel (2017), and more recently Farmer, Nakamura and Steinsson (2023), who show that even professional forecasters learn only slowly about low-frequency components of the data-generating process.

Modern theoretical and empirical work on the term structure is largely based on dynamic factor models (DFTS). Gaussian affine term-structure models are widely used by Duffee (2002), Dai and Singleton (2002), Joslin, Priebsch and Singleton (2014), and Bauer and Rudebusch (2020). In equilibrium RE or Bayesian/RE theories, the principal components of the yield curve should span the relevant macroeconomic state variables. In the data, however, no standard macro variable is well spanned by the yield PCs, and macro variables often help predict excess long-bond returns over and above yield factors. Our framework is well suited to address these issues because price beliefs themselves become a state variable that affects yield dynamics and excess-return predictability.

Bayesian/RE: A large literature assumes that agents learn about fundamentals, while taking the pricing function as known. For bond markets, this means that investors may be Bayesian learners about inflation or other macro shocks, but they share a common, correct mapping from fundamentals into bond prices. We refer to these as Bayesian/RE models: beliefs are updated using Bayes' rule about fundamentals, but the relation between fundamentals and prices satisfies RE. Examples include Xiong and Yan (2010), Buraschi and Whelan (2022), Buraschi and Jiltsov (2022), and Ehling, Gallmeyer, Heyerdahl-Larsen and Illeditsch (2018). These models match some aspects of the data, but the asymmetric treatment of learning—about fundamentals but not about prices—leaves open the question of how investors might learn about the pricing function itself. It is not obvious that such models can re-

produce the role of the slope in the EH-type regressions we document; further work along these lines is needed.

Adaptive learning and bond yields: A different strand uses adaptive learning to study bond prices. This literature shows that learning is crucial for understanding yield-curve behavior and can help explain EH failures and excess volatility. Sinha (2016) develops an adaptive-learning DSGE model that accounts for regression anomalies and part of the volatility puzzle. Related contributions include Piazzesi and Schneider (2006), Laubach, Tetlow and Williams (2007), and Ellison and Tischbirek (2021), who focus on how learning about inflation, short rates, and macro conditions shapes long yields and term premia. Our approach provides microfoundations for this literature by modeling price beliefs directly, rather than only beliefs about fundamentals, and by jointly addressing a broader set of bond-market facts, including survey expectations.

Bond survey expectations: A recent literature studies survey expectations of bond yields through the lens of factor models with learning. Giacoletti, Laursen and Singleton (2020) and Singleton's (2021) AFA Presidential Address (GLS) show that, in many environments, consensus survey expectations of bond yields resemble forecasts from principal-components models of the yield curve, but that simple PC-based forecasts systematically outperform surveys when the economy exits recessions. Farmer, Nakamura and Steinsson (2023) provide complementary evidence that survey forecasts of interest rates display persistent biases and slow learning about long-run components.

This project: We aim to improve the match with bond-market data under *Internal Rationality*. Investors do not know the true pricing function, but they hold a *consistent system of beliefs* about bond prices and behave in a fully rational way given those beliefs. Beliefs about bond prices are not RE—they do not coincide with the true pricing function—but they are close to the data-generating process and difficult to reject with available observations. At the same time, they reproduce key features of expectations surveys. Because we require investors to use a price model that is both disciplined and hard to refute, our framework can be viewed as a relatively minor departure from the standard RE paradigm that nevertheless has large implications for yield dynamics, excess-return volatility, and the behavior of survey expectations.

3 Empirical observations

The data set contains nominal monthly yields backed out from the Reuters data set. Yields are annualized and expressed in basis points. The data on expectations is from the BCFF survey. As usual in the literature on expectations it interprets the median forecast as the consensus forecast, representative of the overall market opinion, we do not study disagreement in this version of the paper. The original forecasts for calendar quarter were converted to monthly forecasts by interpolation. We have used standard procedures found in the literature for all these interpolations.

We use the sample period 1983.12-2019.12 for all data because this is the only period for which surveys are available and we prefer to exclude the covid recession and post-covid period. The observations about yields are robust to longer sample periods. Although the data is monthly all maturities n are yearly. Also, we use yearly leads and lags in correlations. This is because as we shorten the frequency all the data tends to have a correlation close to one, the yearly correlation shows a larger variety of correlations depending on the asset and regression. We also use yearly leads in the left hand side of various regressions we consider.

3.1 Bonds first and second moments, Expectations Hypothesis

We first describe some key properties of the data on bond yields that have been addressed by a number of papers in the literature, including the first and second moments of yields and excess returns and the failure of the Expectations Hypothesis (EH).

The results are in Table 1

Table 1: *Yields First and Second Moments*

	1y	2y	3y	5y	7y	10y
$\widehat{Y^n}$	388.83	417.23	441.39	481.19	510.89	541.69
$\widehat{\sigma}_{Y^n}$	289.15	294.34	291.41	281.62	272.50	258.74
$\widehat{corr}(Y_t^n, Y_{t-1y}^n)$	0.89	0.90	0.91	0.93	0.93	0.93

We find the celebrated term premium, for this sample it is of 152 points. The numbers also demostrate the flat Term structure of Yield Volatilities, as emphasized for example by Xiong and Yan (2010) (XY) most RE models with homogeneous

agents and a reasonable serial correlation of expected inflation show volatilities that decrease much more quickly with maturities. The serial correlation is very high, slightly higher for longer maturities, the subindex $t-1y$ refers to 12 months before t.

Although much more attention has been devoted in the literature to the excess volatility of stock price returns, the excess volatility of excess long bond returns has been known since Shiller (1979). More specifically, define excess bond returns as the difference between the one year holding period return (in logs) and the short yield:

$$\begin{aligned} xr^n_{t+1y} &= \log Q^{n-1}_{t+1y} - \log Q^n_t - Y^1_t \\ &= -(n-1)Y^{n-1}_{t+1y} + nY^n_t - Y^1_t. \end{aligned} \tag{1}$$

As expected our data set confirms that the volatility of excess returns is much larger than the volatility of yields and it grows very fast with maturity, see Table 2.

Table 2: *Volatility of excess bond returns*

maturity n	2y	3y	5y	7y	10y
$\widehat{\sigma}_{xr^n}$	136.60	256.20	455.74	636.45	880.10

Another celebrated feature of the data is the failure of the Expectations hypothesis (EH), namely that in a regression $Y^{n-1}_{t+1y} - Y^n_t = \alpha^n + \beta^n \frac{Y^n_t - Y^1_t}{n-1} + U^n_t$ the β's are negative, more negative for higher maturities. This is in contrast with the so-called Expectations Hypothesis (EH) which implies $\beta^n = 1$.[2]

Not surprisingly our sample confirms this observation, see Table 3.

Table 3: *Expectations Hypothesis (traditional version)*

	2y	3y	5y	7y	10y
$\widehat{\beta^n}$	-0.31	-0.50	-0.99	-1.28	-1.67
t-stat	-0.43	-0.60	-1.21	-1.56	-2.16

This formulation of the EH is the traditional one, but we will focus for the rest of the paper on an equivalent formulation of the same hypothesis requiring that in

[2]XY address this observation and the flat term structure of volatilities in a model of learning and disagreement. They also study a feature of the data described by Cochrane Piazzesi (2005) where a combination of forward rates helps predict excess returns. We do not focus on the latter as it can be shown that this feature of the data is not robust to different sample periods. See appendix for details.

the following regression

$$xr^n_{t+1y} = \alpha^{xr}_n + b^{xr}_n \frac{Y^n_t - Y^1_t}{n-1} + U^n_t \tag{2}$$

$b^{xr}_n \approx 0$. The interest on this formulation stems, in part, from the fact that the left hand side of the regression shows the marginal profits from shifting funds away from short bonds and into long bonds. Results are in Table 4, The t-stats are for the hypothesis that $b^{xr}_n = 0$. As expected the EH is rejected, the b^{xr}_n's are significantly

Table 4: *Expectations Hypothesis (xr version)*

	2y	3y	6y	8y	11y
$\widehat{b^{xr}_n}$	0.61	2.54	10.69	15.87	24.17
t-stat	0.52	1.14	2.40	2.87	3.54
R^2	0.32	1.59	6.00	7.88	10.41

positive, more so for longer maturities. We consider these coefficients no so much out of a genuine interest in the EH but as a property of the data that a good model should be able to match.

The reader may be puzzled by that fact that the last three columns of Table 4 now report statistics for maturities that are shifted one year forward. The reason for this will become apparent when we discuss Table 5 below.

3.2 RE regressions

The previous empirical results have been noted and analysed before. We now report some new results on the behavior of survey expectations.

Following Adam, Marcet and Beutel (2017) we compare regressions that explain actual realisations versus the same regression explaining survey expectations. This will serve as a test of RE but, more importantly, as a feature of the data that we aim to explain.

Let $\mathcal{E}_t(xr^n_{t+1y})$ denote observed survey expectations at t about xr_{t+1y}. We consider a regression analogous to (2) but with survey expectations in the left hand side:

$$\mathcal{E}_t(xr^n_{t+1y}) \quad = \quad \alpha^{\mathcal{E}}_n + b^{\mathcal{E}}_n \frac{Y^n_t - Y^1_t}{n-1} + U^{n,\mathcal{E}}_t \tag{3}$$

where, as before, $1y$ means 12 months. Under RE $b^{xr}_n = b^{\mathcal{E}}_n$, even if surveys have

measurement error, so testing this hypothesis is a test of RE.[3] The results are in Table 5

Table 5: *RE Regression (EH xr version)*

	2y	3y	6y	8y	11y
$\widehat{b}_n^{xr}$	0.61	2.54	10.69	15.87	24.17
$\widehat{b}_n^{\mathcal{E}}$	0.16	0.80	2.50	1.27	3.69
t-stat $b_n^{xr} = b_n^{\mathcal{E}}$	0.37	0.76	1.75	2.47	2.70

Note that the t-stat rejects RE for long maturities. Furthermore we find that the regressions show a systematic bias: the coefficients for actual excess returns are always larger $\widehat{b}_n^{xr} > \widehat{b}_n^{\mathcal{E}}$. It seems, therefore, that investors systematically underestimate the importance of the slope $\frac{Y_t^n - Y_t^1}{n-1}$ in forecasting excess returns.

The reader may note that in Tables 4 and 5 the headings of the last three columns are for maturities shifted forward one year relative to Tables 1 and 2. This is because the BCFF survey provides information about the investors' opinion on future yields of maturities 1, 2, 3, 5, 7, 10. Since Y_{t+1y}^{n-1} determines xr_{t+1y}^n (see formula (1)) it is clear that the survey provides direct information about $\mathcal{E}_t(xr_{t+1y}^n)$ for n=2, 3,4, 6, 8, 11. Using these maturities in Table 5 (we skip $n = 4$ for brevity) is convenient because we can use the survey information directly and avoid interpolating the observations to maturities not explicitely available in the survey. Thus we also used these maturities in Table 4 since they are the ones needed for the RE test of equality of coefficients.

Table 5 shows the results when the slope $\frac{Y_t^n - Y_t^1}{n-1}$ is in the right hand side. This is for comparison with the EH regression in (2). However analogous RE tests with equality of coefficients can be formulated for any other regressor in place of the slope $\frac{Y_t^n - Y_t^1}{n-1}$, as long as the regressor is in the investors' information set at time t, RE implies the coefficients should appear for these regressors, whether xr_{t+1y}^n or $\mathcal{E}_t(xr_{t+1y}^n)$ is in the left side of the regression.

We now report various other versions of RE regressions where the regressors capture the slope in different ways. The first version introduces the second principal

[3] See Adam, Marcet and Beutel (2017) for details. This is similar to the very influential approach of Coibion and Gorodnichenko (2015). They consider regressions where the survey future forecast error is in the left hand side, which is equivalent with testing $b_n^{xr} = b_n^{\mathcal{E}}$, we stick to the considered formulation because we are interested in matching both b_n^{xr} and $b_n^{\mathcal{E}}$, not only the difference. CG also work with regressions where the right hand side contains the update of the forecast, or or the current forecast error. Those regressions are sensitive to measurement error while our approach is not.

component of a PC decomposition of the yields. This PC (labelled PC2) is thought to capture the underlying slope of the yield curve, so we use it to replace the slope in the right hand side of (2) and (3). The second version uses the investors' best estimate of the underlying slope $\widehat{\mathbf{s}}_{t-3}$ according to the model we set below (see equation (26)), this is essentially the geometric average of past slopes for the 10-year bond, in this case using a discount factor of $1-g=0.5$.[4]

To allow for a more flexible parameterization of how different past slopes enter, the third version uses $\widehat{\mathbf{s}}_t, \widehat{\mathbf{s}}_{t-3}$.

The results are shown in Table 6. They confirm and even empasize those of Table 5: surveys always underestimate the effect of the slope on future excess returns. The difference in coefficients is larger and more significant for longer maturities. All the versions have much more significant rejections of RE than Table 5

Table 6: *RE Regression alternative regressors*

regressor		2y	3y	6y	8y	11y
PC_t^2	$\widehat{b}_n^{xr}$	0.06	0.13	0.43	0.61	0.89
	$\widehat{b}_n^{\mathcal{E}}$	0.02	0.05	0.14	0.13	0.19
	t-stat $b_n^{xr}=b_n^{\mathcal{E}}$	0.85	1.11	1.98	2.57	2.87
$\widehat{\mathbf{s}}_{t-3}$	$\widehat{b}_n^{xr}$	1.04	2.16	9.14	13.90	22.67
	$\widehat{b}_n^{\mathcal{E}}$	0.41	0.57	0.28	-2.28	-2.86
	t-stat $b_n^{xr}=b_n^{\mathcal{E}}$	0.77	0.95	2.12	2.79	3.08
$\widehat{\mathbf{s}}_t, \widehat{\mathbf{s}}_{t-3}$	p-value $b_n^{xr}=b_n^{\mathcal{E}}$	0.70	0.59	0.09	0.01	0.00

We are interested in these results not only because RE is rejected but because they give information about how to build a "better" model: a deviation from RE should be consistent with the above results.

Recall that, as is common in the literature on expectations, we have interpreted the median expectation as representative of the whole market. This is a reasonable choice in our setup since we have a homogeneous agent model, a more complete model would incorporate disagreement. Obviously the rejection of RE depends on our choice, it is possible that the median is not representative of the expectation actually influencing the bond prices at each month. For example, the papers on Bayesian/RE and disagreement imply that the relevant opinion is a weighted aver-

[4]To emphasize, a difference between this version of the regression and the one in the previous table is that in Table 5 a different slope $\frac{Y_t^n-Y_t^1}{n-1}$ is in the right hand side for each n, while in the current Table the right hand side contains a geometrically weighted sum of past $\frac{Y_t^{10}-Y_t^1}{10}$, same right hand side variable for all n.

age of all investors (see for example Xiong and Yan (2010) (XY), Buraschi, A., and A. Jiltsov (2006), Buraschi and Whelan (2022)). Singleton (2021) discusses how one would use the panel of surveys for each survey respondent to obtain a different summary of the market opinion. However, whatever model of disagreement is considered will have an implication about the behavior of the median expectation, and that could be compared with the above regressions.

4 A Simple Model

The previous section summarises a list of empirical observations that have been difficult to explain jointly with equilibrium models. In addition it contributes a new puzzle showing a systematic departure of survey expectations from RE. It seems reasonable to try to address these issues with a model where agents do not have RE but, instead, they learn about bond prices. We now describe the benchmark equilibrium model under Internal Rationality that we will be considering, adapting previous work on stock prices in IR to bond prices. Rational behavior will imply that agents learn about an unobserved underlying shock to prices. We start with a simple specification which will then be simplified even further in order to derive some analytic results.

4.1 The micro fundamentals

The fundamentals are standard and very simple, the only non-standard feature will be the agents' price beliefs and, possibly, the specification of the short rate.

There are many investors ex-ante equal with utility function $E_0^{\mathcal{P}} \sum_{t=0}^{\infty} \delta^t u(C_t)$. The notation $E_t^{\mathcal{P}}$ expresses that expectations are taken with respect to investor's perceived process for prices and fundamentals, denoted $\mathcal{P}$, below we describe this in much more detail.

There are competitive markets for N nominal zero-coupon bonds of maturity $n = 1, ..., N$ with prices $Q_t = \left[Q_t^1, ..., Q_t^N\right]'$. Investors behave competitively, they maximize their utility subject to the constraint

$$P_t c_t + \sum_{n=1}^{N} Q_t^n B_t^{i,n} = W_t P_t + \sum_{n=1}^{N} Q_t^{n-1} B_{t-1}^{i,n} \tag{4}$$

$$\overline{B} \geq B_t^{i,n} \geq \underline{B} \tag{5}$$

where P is the price level, W an exogenous endowment and $Q_t^0 = 1$.

The true process for log inflation $\pi_t \equiv \log P_t - \log P_{t-1}$ follows

$$\pi_t = \mu(1-\rho) + \rho\pi_{t-1} + \eta_t^\pi \tag{6}$$

$\eta_t^\pi \sim \mathcal{N}(0, \sigma_\eta)$, iid.

The short rate $Y_t^1 = -\log(Q_t^1)$ is exogenous, given by

$$Y_t^1 = -\log(\delta) + \gamma g^C + \pi_t^e + \xi_t \tag{7}$$

A natural interpretation of ξ_t is as a combination of a policy shock and expected short run inflation. The assumption that the short rate is exogenous will cause the bond constraints (5) to be binding in equilibrium for $n = 1$. All other prices are determined in equilibrium.

Investors each period observe bond prices Q_t and the exogenous processes π_t, W_t, ξ_t.

We consider a closed economy. Since the short rate is exogenous but prices for maturities $n > 2$ are endogenous the supply of bonds has to be such that markets clear when bond constraints (5) are active only for short bonds $n = 1$. This can be formalised assuming that the government buys or sells short bonds to satisfy the demand according to whether the lower or upper limit in (5) is binding, and the government funds these short bond operations with long bond issuances or purchases guaranteeing that (5) is not binding for $n > 1$. The government uses a lump sum tax insuring that bonds stay within the limits.

4.2 Agent Beliefs

As emphasised in Adam et al (2011-2016-2017) when rational investors have beliefs about asset prices investors should contemplate a state space where asset prices are included as a coordinate separate from fundamental shocks. More precisely, in our model this means that the state space for each investor should include possible histories of Q. Formally, the state space is given by $\Omega = \Omega_Q \times \Omega_{\pi,\xi,W}$ where Ω_Q contains possible realizations of bond prices $Q = [Q^n]_{n=1}^N$, and $\Omega_{\pi,\xi,W}$ contains all possible realizations of the exogenous processes for inflation, policy shock and endowment π, ξ, W.

This is a natural assumption for competitive investors who take prices as given, hence they see the Q's as a separate shock that buffets their budget constraint. Essentially all the papers written in competitive equilibrium dynamic models go

around this requirement because they take for granted that the pricing function is known to investors, therefore prices do not represent a separate shock in the minds of these investors. As advanced in section 2 the assumption that agents know the pricing function is common to RE models, Bayesian/RE, robust beliefs, etc. If investors ignore the pricing function as in our paper Ω_Q should be a separate coordinate of the state space Ω. We give more details in section 5.

We assume that investors are rational given a belief system $\mathcal{P}$ that puts joint probabilities on sequences of (Q_t, π_t, ξ_t, W_t). This choice disciplines belief formation, updating algorithms of expectations, etc, these will be a consequence of rational behavior and the assumed $\mathcal{P}$.

We maintain the following assumptions about investors' beliefs throughout the paper:

Assumptions on the belief system $\mathcal{P}$

1. full knowledge of the true process for π, ξ, W.

2. To specify the joint distribution of the bond prices Q we assume that investors believe that bond yields $Y_t^n = -\log Q_t^n/n$ are given by the following "perceived yield curve"

$$Y_t^n = \mathbf{a}_t^n + \mathbf{s}_t n + \phi^n \pi_t^e + u_t^n \tag{8}$$

   u^n is a time-maturity specific error and $\mathbf{a}_t^n, \phi^n$ follows recursions

$$(\mathbf{a}_t^n, \phi^n) \quad = \quad F(\mathbf{a}_{t-1}^n, \phi^{n-1}, n) \tag{9}$$

   and $\mathbf{a}_t^n$ are $\mathbf{s}_t$ unobserved.

   This only partly specifies the agents' beliefs, in various parts of the paper we will complete these beliefs with perceived distributions of the factors and the recursion F.

We think of 1 as a simplifying assumption. As we point out in section 2 the Bayesian/RE literature has developped considerably and it has increased our knowledge of models where agents learn about the exogenous processes when they know the pricing function, including many applications where agents are heterogeneous in their beliefs. Instead we focus on learning about bond prices, therefore assuming knowledge of the exogenous processes is a simplification.

Assumption 2 says that investors believe that there is a linear underlying yield curve with time-varying level $\mathbf{a}_t$ and slope $\mathbf{s}_t$. In addition investors understand that

expected inflation plays a role in the determination of yields.

In words, this perceived yield curve is a special case of the yield curve considered in the literature with three factors: an intercept $\mathbf{a}_t$, a slope $\mathbf{s}_t$ and inflation being the third factor. Note that the slope and inflation are assumed to have well known loadings namely n and ϕ^n. Agents will learn about the value of the unobserved $\mathbf{a}_t, \mathbf{s}_t$ by filtering the observed yields. This is a natural way to specify the joint distribution of bond prices, all investors should know that the bond price Q_t^n is closely related to the bond price Q_t^{n-1}. As we will see the RE solution is a special case of this formulation, this will allow us to claim a sense in which these beliefs are a small deviation from RE.

4.3 Equilibrium under Internal Rationality

To summarize, we look for bond prices Q such that if investors solve

$$\max_{\left\{\left[B_t^{i,n}\right]_{n=1}^{N}, C_t\right\}_{t=0}^{\infty}} E_0^{\mathcal{P}} \sum_{t=0}^{\infty} \delta^t \; u(C_t) \tag{10}$$

subject to the budget constraint (4) and bond limits taking as given the beliefs specified in the *Assumptions about Beliefs*, and if the supply for bonds is as described in the last paragraph of section 4.1, then bond markets at all maturities clear in all periods.

It is clear, therefore, that an equilibrium condition will be

$$Q_t^n = \delta \; E_t^{\mathcal{P}} \left(\frac{u'(C_{t+1})}{u'(C_t)} \frac{P_t}{P_{t+1}} Q_{t+1}^{n-1} \right) \text{ for all } 1 < n \leq N \tag{11}$$

5 Internal Rationality Foundations

Much of modern finance theory is predicated on the premise that asset prices are determined by the "true" discounted value of future asset payments, and that this follows whenever investors are rational. Therefore we have found many researchers who frown upon the idea in the previous section that rational investors can hold price beliefs that are not determined by the true value of asset payments.

In this section we explain that rational investors are compatible with price beliefs as in the perceived yield curve. We first adapt ideas already stated in the literature for stock prices, so we offer a sketchy explanation in the main text. In addition we

discuss a key issue that pertains specifically to bond pricing. A reader who is not too bothered by these theoretical considerations can simply skip this whole section.

5.1 State Space under IR

As usual in IR models, the operator $E_t^{\mathcal{P}}$ denotes the agent's expectations in a probability space, with a state space Ω and a subjective probability measure $\mathcal{P}$ defined over Ω.[5]

As mentioned before the underlying probability space Ω under IR is given by

$$\Omega \equiv \Omega_Q \times \Omega_{\pi,\xi,W} \tag{12}$$

that is it contains a separate coordinate for all possible sequences of bond prices in addition to exogenous variables. The probability measure $\mathcal{P}$ is a model primitive and agents take it as given. This measure may imply that agents optimally learn about unobserved shocks or unobserved parameters. In this paper $\mathcal{P}$ satisfies assumptions stated in section 4.2, a few additional assumptions below complete the specification of the beliefs. As usual, the set of all possible histories up to period t is denoted $\Omega^t = \Omega_Q^t \times \Omega_{\pi,\xi,W}^t$ with a typical element $\omega^t = (Q, \pi, \xi, W)^t$

Investors choose consumption and bond holdings in period t contingent on the observed history ω^t, i.e. the solution to (10) is a function

$$(C_t, B_t) : \Omega^t \to R^{N+1} \tag{13}$$

for all t, $B_t = [B_t^1, ..., B_t^N]$.

That rational investors condition their decisions on the history of exogenous variables and bond price realizations is a natural setup in a model of competitive behavior: since investors see prices as a stochastic variable beyond their control, and since prices influence their budget constraint it would be irrational not to condition investor choices on the observed Q^t.

The expected utility associated with any such contingent consumption choice is well defined and it can be written in terms of primitives as

$$E_0^{\mathcal{P}} \sum_{t=0}^{\infty} \delta^t \; u(C_t) = \int_{\Omega} \sum_{t=0}^{\infty} \delta^t \; u(C_t(\omega^t)) \; d\mathcal{P}(\omega). \tag{14}$$

[5] For a complete specification we need to define a probability space as a triplet $(\Omega, \mathcal{S}, \mathcal{P})$ where $\mathcal{S}$ is a sigma-algebra in Ω. This sigma-algebra plays no significant role in the paper so we only mention it here for completeness, we take the sigma-algebra of all Borel subsets of Ω

and similarly for any other expression involving the $E_t^{\mathcal{P}}$ operator.

The above setup is not standard. Given the fundamentals specified above most competitive equilibrium dynamic models in economics would specify a probability space $\Omega \equiv \Omega_{\pi,\xi,W}$. This is because in the literature on RE, including the literature on "rational bubbles", also in Bayesian/RE, as well as many behavioral economics papers and the robust control literature it is assumed that competitive agents know the correct pricing function $Q_t^{Bay/RE} : \Omega_\pi^t \to R_+^N$ that holds in equilibrium.[6] This is actually a special case of the above beliefs $\mathcal{P}$, with the property that under this assumption the Q coordinate in (12) is redundant, justifying the use of $\Omega \equiv \Omega_{\pi,\xi,W}$ in the literatures just mentioned.

5.2 Rationality of price beliefs, adapting previous arguments

A "standard view" (SV) among academic economists is that knowledge of the pricing function $Q_t^{Bay/RE}$ is a consequence of having rational investors. Under this view, rational investors will be be able to infer the pricing function simply from their own optimality conditions.

Let us be explicit about how this view would apply to our simple model of bond pricing. To simplify we ignore the policy shock for the remainder of section 5, set $\xi_t = 0$ and assume that (11) holds as equality also for $n = 1$.

"Equation (11) is an optimality condition known by a rational investor, therefore a rational investor could infer from this condition that

$$Q_t^n = \delta^n \; E_t^{\mathcal{P}} \left(\frac{u'(C_{t+n})}{u'(C_t)} \frac{P_t}{P_{t+n}} \right) = \delta^n \; E_t \left(\frac{u'(W_{t+n})}{u'(W_t)} \frac{P_t}{P_{t+n}} \right) \tag{15}$$

where the first equation follows because $\mathcal{P}$ is a consistent probability measure so it satisfies the law of iterated expectations (LIE), and the second equation follows because everything inside the expectation depends on exogenous processes. Since the last E_t is a function of exogenous variables it must be the case that a rational investor knows this function, she can infer $Q_t^{Bay/RE}$ from the above equation, no separate beliefs about prices Q can be compatible with rationality. In this case Q can not be a separate coordinate in the probability space as in (12)."

So would the SV probably go.

However, in general the above SV can be discarded for two reasons

[6] It should be possible to incorporate the ideas of robust control when agents have separate beliefs about asset prices, but we do not know of such an application.

1. Dependence of sdf on future prices: In the optimality condition (11) future consumptions are a function of future Q's, denoted $C_{t+n}((Q,\pi,\xi)^{t+n})$'s, under IR. Therefore the expectation after the second equality in (15) depends on the beliefs about future Q's and not only on the beliefs about exogenous variables, so the superscript $^{\mathcal{P}}$ can not be deleted, therefore the second equality is false.

2. Alternation of marginal investors. In a general case, with many heterogeneous agents who hold different beliefs, the identity of the marginal investor may change at any point in time. Therefore the measure for which (11) holds at t is possibly different from the measure for which this holds next period, therefore LIE can not be applied. In this case the first equality of (15) would fail in a heterogeneous agent version of the model.

The first argument above adapts the arguments of Adam, Marcet and Beutel (2017), and the second argument adapts the argument in Adam and Marcet (2011), both for stock prices. We do not give a full account in the main text.

5.3 Rationality of price beliefs, affine models

Affine term structure models have become important forecasting tools, both in the research literature and in the industry. In particular, Gaussian versions are employed by Duffee (2002), Dai and Singleton (2002). This literature uses LIE on standard equilibrium conditions (often called "no-arbitrage conditions") to impose restrictions on forecasting models of bonds that, if correct in the data, should help to improve out-of-sample forecasts. In particular, this literature often considers that a set of variables X_t determining the short rate are the only factors determining all the yields. Furthermore, in the language of the bond price factors literature, the factor loadings for each maturity satisfy a certain recursion that depends on the risk neutral measure, see Duffie and Kan (1996).

Since the equilibrium condition (11) is exactly the condition used by this literature[7] This may suggest that that beliefs that depend on underlying level and slope factors $\mathbf{a}_t, \mathbf{s}_t$ as stated in (8) can not be held by internally rational agents because of two reasons:

[7]We do not use the traditional term "no-arbitrage condition" because no-arbitrage only implies this condition under complete markets, since we assume incomplete markets no-arbitrage would not be enough to impose this condition, but this condition does hold as a consequence of optimal behavior.

1. Rational agents should understand that the prices of all maturities are a function of the same shocks as the short rate (in this case (π_t^e, W_t)).

2. (11) implies that factor loadings have to satisfy a certain recursion

But it turns out that the standard argument does not apply here, let us see this in some more detail. For this purpose let us ignore for this subsection that the short rate is exogenous, let us consider the standard case when (11) also holds for $n = 1$, a similar but less standard argument would apply in the case that the short rate is exogenous, see next footnote.

Retracing the steps that usually lead to a recursion for bond prices we would first define the sdf $\mathcal{M}_{t+1} \equiv \delta \frac{u'(C_{t+1})}{u'(C_t)} \frac{P_t}{P_{t+1}}$ then (11) implies

$$Q_t^n = Q_t^1 \; E_t^{\mathcal{P},\mathbb{Q}} \left(Q_{t+1}^{n-1} \right) \tag{16}$$

where $E_t^{\mathcal{P},\mathbb{Q}}$ is taken under the risk neutral measure applied to the investors' perceived expectations. This equation is correct and it normally implies 1. and 2. above. [8]

But a closer look reveals that this equation holds for a risk-neutral conditional density that involves bond prices, namely,

$$\begin{aligned} f^{\mathcal{P},\mathbb{Q}}(Q_{t+1}, \pi_{t+1}^e, W_{t+1} | (Q, \pi^e, W)^t) &= (Q_t^1)^{-1} f^{\mathcal{P}}(Q_{t+1}, \pi_{t+1}^e, W_{t+1} | (Q, \pi^e, W)^t) \\ &\quad \cdot \mathcal{M}_{t+1}((Q, \pi^e, W)^{t+1}). \end{aligned}$$

Since the sdf is defined on the space Ω, that is $\mathcal{M}_{t+1}((Q, \pi^e, W)^{t+1})$ it means that argument 1. above is not correct: even if the short rate depends only on exogenous variables π^e, W the sdf is a function of prices, therefore Q_t^2 petentially depends on any variable determining the distribution of, say, Q_{t+1}^{10}.[9]

Furthermore, since consumption is chosen optimally by investors who take as given their belief system, the standard recursion on the loadings will only confirm that the optimality condition (11) is compatible with (8), for an optimal response when the bond price is taken as an exogenous process will be compatible with the assumed (in the investors' mind) price process. Therefore any restriction on prices

[8] In the case that the short rate is exogenous one could obtain an equivalent recursion from $n = 2$ onwards based on the equality $Q_t^n = \frac{Q_t^2}{E_t^{\mathcal{P},\mathbb{Q}}(Q_{t+1}^1)} E_t^{\mathcal{P},\mathbb{Q}} \left(Q_{t+1}^{n-1} \right)$.

[9] A similar argument holds for the Bayesian/RE literature on disagreement, where some average of the distribution of beliefs is an additional factor, see Xiong and Yan (2010) (XY), Buraschi, A., and A. Jiltsov (2006).

imposed using (16) can not contradict the price process already incorporated in the belief system $\mathcal{P}$.

The SV relies on a sort of circular argument: if investors are assumed to know the pricing function then the pricing function can be derived from the investors' optimality condition. This does not mean that optimising behavior implies knowledge of the pricing function.

This shows that the role of the sdf in finance depends on the model under consideration. If agents learn about prices then the sdf pricing all assets $\mathcal{M}_{t+1}((Q, \pi^e, W)^{t+1})$ lives in the investors' state space that includes prices and this means that variables determining future prices according to $\mathcal{P}$ are actual factors. On the other hand the general idea that a few factors may summarize the behavior of bonds at all maturities is validated. The loading recursions used in much of the literature can be seen as an approximation to (16).

6 Learning Equilibrium, An even Simpler Model

In the interest of obtaining simple algebraic solutions we now simplify even further the model of section 4. This should help to clarify the role of learning about the yield curve in generating bond price volatility, why we match the RE regressions and to understand in what sense our model is a deviation from full RE.

6.1 Some further simplifications

For this purpose we introduce **six simplifications** to section 4

1st simplification: assume risk neutral utility function $u(C) = C$.

2d simplification: inflation has a unit root $\rho = 1$. Therefore

$$\pi_t^e \equiv E_t^{\mathcal{P}}(\pi_{t+1}) = E_t(\pi_{t+1}) = \pi_t$$

3d simplification: we consider a log-linear approximation to (11), that is we consider equilibrium yields that under IR satisfy

$$\begin{aligned} nY_t^n &= -\log\delta + E_t^{\mathcal{P}}(\pi_{t+1} + (n-1)Y_{t+1}^{n-1}) && (17) \\ &= -\log\delta + E_t(\pi_{t+1}) + (n-1)E_t^{\mathcal{P}}(Y_{t+1}^{n-1}) \text{ for } 1 < n \leq N && (18) \end{aligned}$$

where the second equation uses Assumption 1 on $\mathcal{P}$.

4th simplification: the policy shock is iid, independent of inflation and it satisfies $E(\xi_t) = 0$. This is obviously far from the observed behavior, such a shock would be highly serially correlated. It makes it harder to match part of the data but it is a convenient for now.

Given the simplifications made thus far the *RE equilibrium* for this simpler model is

$$Y_t^{RE,n} = -\log\delta + \pi_t^e \tag{19}$$

The next simplifications fully specify the investors' beliefs $\mathcal{P}$. Recall that Assumptions 1 and 2 partly specify investors' beliefs by stating the perceived yield curve (8) and stating that investors know the inflation process. We now complete these assumption by making precise assumptions on investor's perception about the perceived level and slope.

5th simplification: Investors know the true intercept and the true loading on inflation according to the RE solution (19), therefore the perceived yield curve (8) has $\mathbf{a}_t = -\log\delta$ and $\phi^n = 1$.

However, agents are uncertain about the unobserved underlying slope and they believe that

$$\mathbf{s}_t = \mathbf{s}_{t-1} + w_t \tag{20}$$

for an iid $w_t \sim N(0, \sigma_w^2)$ and perceived initial value $E_0^{\mathcal{P}} = 0$.

In other words agents know the intercept and loading of inflation, they even know the true average value of the slope, which is 0 under RE, but they are uncertain about the exact slope so they expressed this uncertainty by perceiving the the underlying slope $\mathbf{s}_t$ as an unobserved unit root process.

Under these assumptions the perceived yield curve (8) becomes

$$Y_t^n = -\log\delta + \mathbf{s}_t n + \pi_t^e + u_t^n \tag{21}$$

It is clear that (20) and (21) are a dynamic system, the first equation being the transition equation while the second is a measurement equation. A standard application of the Kalman filter implies that the best estimate of the unobserved variable $\mathbf{s}_t$ is

$$\begin{aligned} E_t^{\mathcal{P}}(\mathbf{s}_t) &\equiv \widehat{\mathbf{s}}_t = \widehat{\mathbf{s}}_{t-1} + g(s_t - \widehat{\mathbf{s}}_{t-1}) \\ s_t &\equiv E_t^{\mathcal{P}}(\mathbf{s}_t \mid Y_t^N, ..., Y_t^1) \end{aligned} \tag{22}$$

with a constant gain g and where the "time t-slope" s_t is the best forecast of $\mathbf{s}_t$ given only information at t. Generically s_t involves the optimal estimate of the underlying slope (under the investors' belief system) using only all the yields available at t and no past information. This would involve some regression on the equation (8) across yields n for given t with a constant and n in the right hand side. We now make

As a consequence of all these assumptions and optimal behavior we have that IR investors form their yield expectations as

$$E_t^{\mathcal{P}}(Y_{t+1}^n) \quad = \quad -\log\delta + \widehat{\mathbf{s}}_t n + \pi_t^e \tag{23}$$

this follows from the perceived yield curve (21).

Our terminology may have been confusing thus far: following a common practice in finance, we used the term "slope" loosely, referring to various related concepts. From now on we will (try to) stick to the following nomenclature:

- We refer to "the slope" as the observed value of the realised yield curve in a given period $\frac{Y_t^n - Y_t^1}{n-1}$ for some maturity n.
- On the other hand the "underlying slope" is the shock $\mathbf{s}_t$ that investors believe is buffeting all the yields and that is unobserved.
- The "estimated slope" is the best forecast of the latter using all information up to t, namely $\widehat{\mathbf{s}}_t$.
- The "t- slope" is the best forecast of the latter using only information that is acquired at t, namely s_t.

We now make a final simplification that allows us to find easily the t-slope:

6th simplification: there is a "reference bond" of maturity m. Investors believe that the variance of the shock u_t^n is much lower for the reference bond of maturity $n = m$ and for the short rate $n = 1$ than for all other maturities.

The bond pricing literature often points to the fact that there are "reference maturities" that reflect market information better than other maturities, the most widely used is the 10-year bond. Obviously one can take for granted that information on the short bond is very accurate as it is given by central bank policy. Therefore, taking a limit where $var(u_t^n)$ goes to zero for $n = m, 1$ implies that the optimal t-slope is found by Generalized Least Squares estimate in a cross section regression of (21). This would imply that investors essentially ignore all the yields except

$n = m, 1$ when they form their t-estimated slope s_t. Combining (21) for $n = m, 1$ gives that, under investor beliefs, an optimal estimate of the t-slope is

$$s_t = \frac{Y_t^m - Y_t^1}{m-1} = \mathbf{s}_t + \frac{u_t^m - u_t^1}{m-1} \tag{24}$$

.

Equilibrium Yields in the Simple Model

Putting all this together we have that $E_t^{\mathcal{P}}(Y_{t+1}^{n-1}) = -\log\delta + \gamma g^C + \pi_t^e + (n-1)\widehat{\mathbf{s}}_t$, plugging this into (18) we have that equilibrium yields under IR satisfy

$$Y_t^n \quad = \quad -\log\delta + E_t(\pi_{t+1}) + \left(\frac{n-1}{n}\right)^2 \widehat{\mathbf{s}}_t n \tag{25}$$

6.2 Closeness of price beliefs $\mathcal{P}$ and true distribution

The reader may be comfortable with simplifications 1 to 4, since they are commonly used in scores of papers. But most readers may be less comfortable with 5th and 6th simplifications since they specify the investors' beliefs about prices investors' beliefs and these assumptions are not commonly made in standard papers. An important issue of modelling under IR is to justify such assumptions about the belief system $\mathcal{P}$. In this subsection we argue that these are natural simplifications, that they deviate slightly from the standard RE paradigm and that they are a "reasonable approximation" to the equilibrium distribution.

First, notice that by assumption agents know the process for exogenous processes so that the terms $-\log\delta + \pi_t^e$ coincide both in the perceived yield curve (21) and the equilibrium yield curve (25). Therefore the agent having this belief will eventually recognise she is correct in this belief.

Second, there is a sense in which IR beliefs are a small deviation from RE beliefs. It is clear that RE beliefs are a special case of (21) when $\sigma_w^2 \approx 0$ since, in this case $\widehat{\mathbf{s}}_t = \mathbf{s}_t = 0$ for all t, the perceived (21) equals the true yield curve (25), equal to the RE case (19). An analyst willing to consider small deviations from RE should be willing to consider our simplifications 5th 6th.

Third, notice that even if $\widehat{\mathbf{s}}_t$ varies significantly over time, when the term $\frac{n-1}{n}$ is close to 1 the perceived yield curve (21) is approximately replicated by the actual yield curve (25). Therefore for large n investors would have a hard time noting that their beliefs are not exactly correct.

Fourth, notice that just because agents think that there is an underlying time-varying slope $\mathbf{s}_t$ then their estimate $\widehat{\mathbf{s}}_t$ is actually a time-varying shock influencing the yields each period. In the language of the literature on bonds' principal components, just because agents believe that there is a factor $\mathbf{s}_t$ determining bond yields causes $\widehat{\mathbf{s}}_t$ to be a factor that actually determines bond yields.[10]

Fifth, the system converges to RE. Since the literature on convergence of adaptive learning models to RE is well known we give here a cursory account: in the language of this literature the perceived law of motion (PLM) is given by perceived yield curve (21), implying that investors believe that the best way to predict future yields is (23). Given simplification 6 the pricing of the reference bond of maturity m is the relevant one for this issue. When investors that the slope is $\widehat{\mathbf{s}}_t$ as in the PLM, the actual law of motion (ALM) given by the actual yield curve (25) says that the slope multiplying n is $\left(\frac{m-1}{m}\right)^2$. This defines the T-map that gives the true slope as a function of the perceived slope:

$$T(\widehat{\mathbf{s}}_t) \equiv \left(\frac{m-1}{m}\right)^2 \widehat{\mathbf{s}}_t$$

This T-map is the hallmark of convergence to RE, and it is well known that convergence to RE equilibrium requires $T' < 1$. This holds both for the case of OLS[11] or convergence as $g \to 0$ for the constant gain model.[12] Since the derivative is $T' = \left(\frac{m-1}{m}\right)^2 < 1$ the model would converge to RE, that is $\widehat{\mathbf{s}}_t \to 0$.

All of this suggests that agents' beliefs are unlikely to be rejected given observed yields. Although in the constant gain model and with infinite data investors might understand that their belief (21) is not correct, with a limited amount of data this will be a very difficult task. Furthermore, the presence of other shocks buffeting the determination of yields as we do have in the real world will make it even harder to discard their belief (21). We will check formally that this is the case with a version of the model that resembles the data more closely than the current simple model.

[10] This version of the model would imply that yields at, say, 3 or 4 years, are not close to beliefs. However, when we introduce a serially correlated ξ as in the data this misspecification will become hard to detect.

[11] This means that we have $1/t$ instead of a constant gain g in equation (22), see Marcet and Sargent (1989) and Evans and Honkapohja (2001) for this convergence result.

[12] See Williams (2017) for this convergence result.

6.3 Analytic Results

We now show analytically some properties of the model, demonstrating the potential to address the empirical puzzles we have mentioned above. We only simplify things further by incorporating s_t with a lag

$$\widehat{\mathbf{s}}_t = \widehat{\mathbf{s}}_{t-1} + g\left(\frac{Y_{t-1}^m - Y_{t-1}^1}{m-1} - \widehat{\mathbf{s}}_{t-1}\right) \tag{26}$$

as is standard in the adaptive learning literature, this simplifies algebra slightly and it could be removed easily.

Plugging in equilibrium yields (25) in (8) we have $s_t = \frac{Y_t^m - Y_t^1}{m-1} = \widehat{\mathbf{s}}_t \ \frac{m-1}{m} - \frac{\xi_t}{m-1}$. Plugging this in (22)

$$\widehat{\mathbf{s}}_t \ = \ \left(1 - \frac{g}{m}\right)\widehat{\mathbf{s}}_{t-1} - \frac{g\xi_{t-1}}{m-1} \tag{27}$$

hence in equilibrium the estimated slope $\widehat{\mathbf{s}}_t$ is an AR(1) process with serial correlation $1 - \frac{g}{m}$ and a shock $\frac{g\xi_{t-1}}{m-1}$.

This equation shows how learning will influence the yields: any shock that bumps up the slope $s_t = \frac{Y_t^m - Y_t^1}{m-1}$ will have an effect on current and future estimated slopes $\widehat{\mathbf{s}}_t$ and, through the ALM (25), on all current and future yields. Since for reasonable parameterizations m is relatively high and g relatively low, $1 - \frac{g}{m}$ will be high and the effect on $\widehat{\mathbf{s}}_t$ of a shock to the slope is very persistent.

This shows that the model has the potential to explain the puzzles we described: The ALM (25) implies that a shock to the slope influences more strongly longer maturities because the term $\left(\frac{n-1}{n}\right)^2$ in (25) is larger for large maturities. Therefore there is a self-referential transmission of slope movements: just because $s_t = \frac{Y_t^m - Y_t^1}{m-1}$ is high in some period, this causes the perceived underlying slope $\widehat{\mathbf{s}}_t$ to go up, driving future slope up, driving future perceived slopes is also up and so on... It is clear that this mechanism has the potential to explain the high volatility puzzles described in section 3 since high maturities experience a stronger dependence on the variability of the perceived slope.

The linear structure of this model allows us to easily compute all the relevant moments. Clearly the moments of $\widehat{\mathbf{s}}_t$ are easily computed from the AR(1) represen-

tation (27). Summarising previous derivations we have

$$
\begin{aligned}
Y_t^n &= -\log(\delta) + \widehat{\mathbf{s}}_t \, \frac{(n-1)^2}{n} + \pi_t^e && (28)\\
s_t &= \frac{Y_t^m - Y_t^1}{m-1} = \widehat{\mathbf{s}}_t \, \frac{m-1}{m} - \frac{\xi_t}{m-1} && (29)\\
xr_{t+1}^n &= \left[2n - 3 + (n-2)^2 \frac{g}{m}\right] \widehat{\mathbf{s}}_t && (30)\\
&\quad + (n-1)\eta_{t+1} + \left[\frac{(n-2)^2 g}{m-1} - 1\right] \xi_t && (31)
\end{aligned}
$$

We also need to compute how investors' expectations evolve in order to compare them with the RE regressions reported above. This derivation is interesting because it is possibly non-intuitive for someone (like us) raised in the RE tradition. Under RE you can use indistinguishably relations that hold in equilibrium and relations that hold in the agents' model, since the two distributions are equal. But in the derivation for $E_t^{\mathcal{P}}$ we need to carefully distinguish when each of the distributions has to be applied. The first equality uses information known to the investors given the definition of xr (1), the fact that $E^{\mathcal{P}}$ is a proper conditional expectation, and the agents' perceptions (23)

$$
\begin{aligned}
E_t^{\mathcal{P}}(xr_{t+1}^n) &= -(n-1)E_t^{\mathcal{P}}(Y_{t+1y}^{n-1}) + nY_t^n - Y_t^1\\
&= -(n-1)(-\log\delta + \widehat{\mathbf{s}}_t(n-1) + \pi_t^e) + nY_t^n - Y_t^1 && (32)
\end{aligned}
$$

Now that this is expressed in terms of equilibrium yields and slope we apply the modeller's knowledge of equilibrium conditions (28), the evolution of the short rate and simple algebra to obtain

$$
\begin{aligned}
E_t^{\mathcal{P}}(xr_{t+1}^n) &= -(n-1)(-\log\delta + \widehat{\mathbf{s}}_t(n-1) + \pi_t^e)\\
&\quad + n(-\log(\delta) + \widehat{\mathbf{s}}_t \, \frac{(n-1)^2}{n} + \pi_t^e) + \log\delta - \pi_t^e - \xi_t\\
&= -\xi_t && (33)
\end{aligned}
$$

With this at hand we can now explore the model's potential to explain the puzzles of Section 3.

6.3.1 Flat term structure of volatilities

From the equations in the previous subsection we have

$$\sigma^2_{\widehat{\mathbf{s}}} = \frac{g^2\sigma^2_\xi}{(m-1)^2\left(1-\left(1-\frac{g}{m}\right)^2\right)} \tag{34}$$

$$\sigma^2_{Y^n} = \sigma^2_{\pi^e} + \frac{(n-1)^4}{n^2}\sigma^2_{\widehat{\mathbf{s}}} \tag{35}$$

The above formula shows that IR imparts larger volatility to long maturities, as required. This is because, as we explained before and shown by the above formula for $\sigma^2_{Y^n}$, the estimated slope $\sigma^2_{\widehat{\mathbf{s}}}$ imparts a larger volatility on longer maturities.

The above formula shows that this simpler model overshoots relative to the data: volatilities are increasing with n instead of approximately flat as in Table 1. However, once we introduce a calibrated serial correlation of inflation in section **??**, and a serially correlated policy shock as we plan to do in the next version of the paper this should help increase the volatility of short rates versus long rates, thus flattening the volatilities in the model of learning.

6.3.2 Excess volatility of excess returns

From the above formula for excess returns we have

$$\sigma^2_{xr^n} = \left[2n-3+(n-2)^2\frac{g}{m}\right]^2\sigma^2_{\widehat{\mathbf{s}}} \tag{36}$$

$$+(n-1)^2\sigma^2_{\eta^\pi} + \left[\frac{(n-2)^2 g}{m-1}-1\right]^2\sigma^2_\xi \tag{37}$$

Note that under RE excess returns also have the term $(n-1)^2\sigma^2_{\eta^\pi}$, thus increasing the volatility of xr with n, but this increase would be insufficient to explain quantitatively the observed volatility of excess returns. The above formula shows that under IR there are additional terms contributing to the volatility, and the importance of these terms increases with n. The intuition is the same as for yields: $\sigma^2_{\widehat{\mathbf{s}}}$ has a larger influence on longer maturities, therefore its variance contributes more to longer maturities.

6.3.3 Failure of the Expectations Hypothesis

The coefficients b_n^{xr} in regression (2) would be zero under RE. We now have under IR

$$\begin{aligned} b_n^{xr} &= \frac{cov(xr_{t+1}^n, \frac{Y_t^n - Y_t^1}{n-1})}{\sigma_s^2} \\ &= \frac{\left[2n-3+(n-2)^2\frac{g}{m}\right]\frac{n-1}{n}\sigma_{\widehat{\mathbf{s}}}^2 - \left(\frac{(n-2)^2 g}{m-1} - 1\right)\frac{\sigma_\xi^2}{n-1}}{\left(\frac{n-1}{n}\right)^2 \sigma_{\widehat{\mathbf{s}}}^2 + \frac{\sigma_\xi^2}{(n-1)^2}} \end{aligned} \tag{38}$$

Clearly this coefficient is now positive and it goes to infinity as n grows, consistent with the failure of EH reported in Section 3. The reason is, again, $\widehat{\mathbf{s}}_t$ plays a more important role for longer maturities, since $\widehat{\mathbf{s}}_t$ also determines the regressor $\frac{Y_t^n - Y_t^1}{n-1}$ the covariance term grows very quickly and b_n^{xr} grows with n.

6.3.4 RE Regressions

As we saw in section 3.2 survey expectations underpredict the influence of the slope on future excess returns, leading to estimates $\hat{b}_n^{xr} - \hat{b}_n^{\mathcal{E}} > 0$ in the regressions (2)-(3), with a larger and more significant difference for longer maturities. This result is robust to many alternative ways of capturing the underlying the slope in the right hand side of the regressions.

This is because in this model agents systematically overestimate the effect that the current estimated slope has on future yields. More precisely, in the context of the simpler model, we see that (21) implies that agents' prediction of yields is

$$E_t^{\mathcal{P}}(Y_{t+1}^n) = -\log\delta + \pi_t^e + \widehat{\mathbf{s}}_t n \tag{39}$$

while in equilibrium yields follow (25) so that the actual expectation of future yields is

$$E_t(Y_{t+1}^n) = -\log\delta + \pi_t^e + \left(\frac{n-1}{n}\right)^2 \widehat{\mathbf{s}}_t n \tag{40}$$

The term $\frac{n-1}{n}$ means that investors *over*predict the effect of the slope on actual future *yields*. Since future Y's enter with a negative sign in xr_{t+1}^n the inequality is reversed for the regression coefficients: investors *under*predict the effect of the slope on actual future *excess returns*. Hence, $b_n^{xr} > b_n^{\mathcal{E}}$ generically, as in the data.

Our claim is that a robust feature of a model of learning about the slope of the yield curve is that it reproduces this feature of the data. Formally, for the current model it is easy to see that the coefficients in the RE regressions (3) satisfy

$$b_n^{xr} - b_n^{\mathcal{E}} = \frac{(n-1)cov(E_t^{\mathcal{P}}(Y_{t+1}^{n-1}) - Y_{t+1}^{n-1}, s_t)}{\sigma_s^2},$$

by the argument in the previous paragraph $cov(E_t^{\mathcal{P}}(Y_{t+1}^n) - Y_{t+1}^n, s_t) > 0$, therefore $b_n^{xr} - b_n^{\mathcal{E}} > 0$. It is also possible to show this difference is larger for large n.[13]

We reckon that $b_n^{xr} > b_n^{\mathcal{E}}$ is a generic property of a model of learning about the slope of the yield curve. One can show that this inequality will hold whenever the model is stable under learning so that the T-map has a derivative less than 1. As we mentioned, that $T' < 1$ is needed for stability under learning, $T' > 1$ would imply explossive fluctuations that could drive yields to plus or minus infinity. Therefore for any model where learning imparts more volatility but no has explossive behavior the RE regressions are likely to show $b_n^{xr} > b_n^{\mathcal{E}}$, no time to explain this in detail now...

In particular, for the special case considered here the equilibrium equations (28)-(76) imply that

$$b_n^{\mathcal{E}} = \frac{\sigma_\xi^2}{\frac{(n-1)^3}{n^2}\sigma_{\hat{s}}^2 + \frac{1}{n-1}\sigma_\xi^2}$$

Clearly $b_n^{\mathcal{E}}$ goes to zero as n grows showing that $b_n^{xr} - b_n^{\mathcal{E}}$ goes to infinity as n grows.

7 A Quantitative Model

The analytical results above are derived in a deliberately simplified model. For the quantitative exercises, we relax some of the assumptions: inflation is stationary, the policy shock is serially correlated and there is forward guidance. This section summarizes the resulting law of motion for bond yields under Rational Expectations (RE) and then introduce learning.

[13] For the simpler model we have $b_n^{\mathcal{E}} = 0$, same as under under RE (same as RE). However this property is not robust, this coefficient would be positive or negative for small deviations of the model under IR.

7.1 RE Affine Term Structure Model

Time is measured in months and n denotes maturity in months. We focus on annual maturities $n = 12, 24, 36, \ldots$ for comparison with the empirical evidence based on one–year holding periods. The relevant state variables for the RE benchmark are inflation π_t and a monetary policy shock ξ_t entering the short rate. Inflation follows an AR(1) process,

$$\pi_t \equiv \log \frac{P_t}{P_{t-12}} = (1-\rho)\mu + \rho\,\pi_{t-1} + \eta_t, \qquad \eta_t \sim \mathcal{N}(0, \sigma_\eta^2), \tag{41}$$

and the policy shock is

$$\xi_t = \rho_\xi\,\xi_{t-1} + \varepsilon_t, \qquad \varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2), \tag{42}$$

with (η_t, ε_t) i.i.d. over time.

Let Y_t^n denote the annualized yield on an n–month zero–coupon bond. The log–linearized Euler equation for bond pricing implies that, for $n > 12$,

$$\frac{n}{12}\,Y_t^n = -\ln\delta + (1-\rho^{12})\mu + \rho^{12}\pi_t + \frac{n-12}{12}\,\mathbb{E}_t\big[Y_{t+12}^{n-12}\big], \tag{43}$$

while the one–year rate (the 12–month yield) satisfies

$$Y_t^{12} = -\ln\delta + (1-\rho^{12})\mu + \rho^{12}\pi_t + \xi_t. \tag{44}$$

Here δ is the subjective discount factor, μ is the (log) mean inflation rate, and ρ^{12} is the 12–month autocorrelation of inflation implied by (41). The expectations $\mathbb{E}_t[\cdot]$ are taken under the physical measure. Equations (43)–(44) provide a standard linear recursive structure in which yields at longer maturities are determined by short rates and expectations of future yields.

Under RE we postulate an affine representation of yields in terms of the state vector (π_t, ξ_t),

$$Y_t^n = a_n + \phi_n \pi_t + \psi_n \xi_t, \qquad n \geq 12, \tag{45}$$

where a_n is the constant (or level) loading, ϕ_n the inflation loading, and ψ_n the policy–shock loading at maturity n.

To simplify notation define

$$b \equiv -\ln\delta + (1-\rho^{12})\mu, \qquad c \equiv \rho^{12}.$$

Using the inflation and policy processes (41)–(42), one–year–ahead expectations are

$$\mathbb{E}_t[\pi_{t+12}] = (1-\rho^{12})\mu + \rho^{12}\pi_t, \qquad \mathbb{E}_t[\xi_{t+12}] = \rho_\xi^{12}\,\xi_t. \tag{46}$$

Substituting the affine conjecture (45) and (46) into the Euler equation (43), and equating coefficients on $(1, \pi_t, \xi_t)$, yields the recursions (for $n > 12$)

$$a_n = \frac{12}{n}\,b + \frac{n-12}{n}\big[a_{n-12} + \phi_{n-12}(1-\rho^{12})\mu\big], \tag{47}$$

$$\phi_n = \frac{c}{n}\big[12 + (n-12)\phi_{n-12}\big], \tag{48}$$

$$\psi_n = \frac{n-12}{n}\,\rho_\xi^{12}\,\psi_{n-12}, \tag{49}$$

with initial conditions at one year

$$a_{12} = b, \qquad \phi_{12} = c, \qquad \psi_{12} = 1. \tag{50}$$

For annual maturities $n = 12m$ with $m \in \mathbb{N}$, these recursions admit closed–form solutions. The inflation loadings are

$$\phi_n = \frac{12\rho^{12}}{n}\,\frac{1-\rho^n}{1-\rho^{12}}, \qquad n = 12m, \tag{51}$$

and the policy loadings are

$$\psi_n = \frac{12}{n}\,\rho_\xi^{n-12}, \qquad n = 12m. \tag{52}$$

Finally, the constant term can be written as

$$a_n = -\ln\delta + \mu\,(1-\phi_n), \qquad n = 12m. \tag{53}$$

Inflation shocks enter the term structure "horizontally": they affect every maturity directly via $\mathbb{E}_t[\pi_{t+12}]$, so ϕ_n declines only slowly with n. By contrast, the policy shock ξ_t affects the term structure only through its impact on the short rate today, and then propagates through expectations. As a result, ψ_n decays rapidly with maturity, reflecting the fact that a one–time policy surprise has a limited effect on far–out yields in the simple RE benchmark.

7.1.1 More persistent policy effects and forward guidance

Empirically, monetary policy shocks appear to have more persistent effects on long rates than the simple structure in (52) implies. To capture this, we enrich the policy block to allow for forward guidance shocks that directly affect the expected future path of short rates.

Formally, suppose that each month the central bank delivers news about the entire future short–rate path, generating a family of shocks

$$\xi_{t,h}, \qquad h = 0, 1, 2, \ldots,$$

where $\xi_{t,h}$ is the unanticipated component of the announcement at time t about the short rate at horizon h. For tractability we assume that each $\xi_{t,h}$ follows the same AR(1) process,

$$\xi_{t,h} = \rho_\xi \, \xi_{t-1,h} + \varepsilon_{t,h}, \qquad \varepsilon_{t,h} \sim \text{i.i.d.}, \tag{54}$$

independently across (t, h).

Solving the affine RE model with this richer policy block yields

$$Y_t^n = a_n + \phi_n \pi_t + \sum_{k=0}^{m-1} \psi_{n,k} \, \xi_{t,12k}, \qquad n = 12m, \tag{55}$$

with loadings

$$\psi_{n,k} = \frac{12}{n} \, \rho_\xi^{n-12-12k}, \qquad k = 0, \ldots, m-1. \tag{56}$$

In this specification, a surprise about the policy rate several years ahead (k large) can move long yields directly, rather than only through repeated revaluation of the short rate.

The fully general structure (55) introduces many separate policy shocks. To keep the quantitative model parsimonious, we collapse these into a single effective policy shock by imposing

$$\xi_{t,h} = \xi_t \quad \text{for all } h.$$

Under this restriction the yield expression simplifies back to a single policy shock,

$$Y_t^n = a_n + \phi_n \pi_t + \bar{\psi}_n \, \xi_t, \qquad n = 12m, \tag{57}$$

where

$$\bar{\psi}_n = \frac{12}{n} \, \frac{1 - \rho_\xi^n}{1 - \rho_\xi^{12}}. \tag{58}$$

Relative to (52), the effective policy loading $\bar{\psi}_n$ decays more slowly with maturity, so that unexpected policy innovations have more persistent and quantitatively relevant effects on long yields, in line with the data.

7.2 Embedding learning in the quantitative structure

The affine RE representation (57) serves as the backbone for the quantitative implementation of our IR model. Investors are assumed to know the processes (41)–(42) for (π_t, ξ_t) and the RE loadings $(a_n, \phi_n, \bar{\psi}_n)$, but they are uncertain about a latent *underlying slope* factor s_t that shifts the term structure.

More precisely, their perceived yield curve is

$$Y_t^n = a_n + \phi_n \pi_t + \bar{\psi}_n \xi_t + \boldsymbol{s}_t n + u_t^n, \tag{59}$$

where $\boldsymbol{s}_t$ follows a random walk (as in the analytic model),

$$\boldsymbol{s}_t = \boldsymbol{s}_{t-1} + w_t, \qquad w_t \sim \mathcal{N}(0, \sigma_w^2), \tag{60}$$

and u_t^n is an idiosyncratic yield residual. Given (59)–(60), investors optimally filter s_t using the cross section of yields and update their estimate $\hat{\boldsymbol{s}}_t \equiv \mathbb{E}_t^P[\boldsymbol{s}_t]$ with a (constant) Kalman gain g, exactly as in the simple model. The perceived law of motion under IR is thus obtained by replacing $\boldsymbol{s}_t$ with $\hat{\boldsymbol{s}}_t$ in (59). The actual equilibrium yield curve is determined by substituting investors' perceived law of motion into the Euler equations (43)–(44), analogously to the derivations in Section 6 and Appendix 1.

7.3 Baseline Estimation

The parameters of the quantitative model are estimated by Simulation Method of Moments (SMM). Let

$$\theta = (g, \rho_\xi, \sigma_\xi)'$$

collect the constant–gain learning parameter, the persistence of the policy shock, and its standard deviation. All remaining parameters (such as ρ, μ, and δ) are pinned down outside the SMM step from the inflation data and the calibration described in the previous section. The idea of SMM is to choose θ so that the model reproduces a set of sample moments that summarize the joint dynamics of yields and excess returns in the data.

Denote by m^{data} a 27×1 vector stacking the empirical moments and by $m(\theta)$ the corresponding model–simulated moments. In the implementation, m^{data} contains five groups of statistics: (i) the standard deviations of annualized zero–coupon yields at maturities 1, 2, 3, 5, 7, and 10 years; (ii) the first–order autocorrelations (at a 12–month lag) of these yields; (iii) the standard deviations of one–year excess holding–period returns on 2, 3, 5, 7, and 10–year bonds; (iv) the expectations–hypothesis (EH) slope coefficients b_{xr}^n from regressions of excess returns on the corresponding term spreads for maturities 2, 3, 6, 8, and 11 years; and (v) the "RE regression" coefficients b_s^n obtained by regressing excess returns on the model counterpart of survey–based expectations for the same maturities.

For a given candidate parameter vector θ, the code simulates the model and computes the same set of moments. The SMM criterion function is a weighted quadratic form in the distance between model and data moments,

$$Q(\theta) = \big[m(\theta) - m^{\text{data}}\big]' W \big[m(\theta) - m^{\text{data}}\big],$$

where W is a positive–definite weighting matrix. In the implementation, we use a diagonal weighting matrix, with the diagonal entries consisting of the inverse of the individually Newey-West variances of the corresponding observed statistics. This criterion minimizes the sum of squared t-statistics across moments, therefore focusing more directly on matching them. We then compute t–statistics for each moment as

$$t_j(\theta) = \frac{m^{\text{data},j} - m^j(\theta)}{se_j},$$

Table **??** compares key moments of the term structure in the data to their model counterparts across maturities. The parameter triplet $(g, \rho_\xi, \sigma_\xi)$ is estimated by simulated method of moments, while (ρ, μ) and δ are pinned down using historical expected inflation and the average one-month T-bill rate. For each statistic and maturity, the last column reports a t-statistic for the difference between the data and the model, computed from Monte Carlo simulations; values close to zero indicate that the model reproduces the empirical moment within the sampling uncertainty of the simulation.

The model performs well in matching the cross-section of yield volatilities. For maturities between one and five years, the standard deviations of yields in the model lie very close to their empirical counterparts, with t-statistics well below conventional significance thresholds. At longer horizons the model somewhat overstates

Table 7: *Summary statistics by maturity: data vs. model vs. t-stats. Parameters estimated via SMM* $(g, \rho_\xi, \sigma_\xi) = (0.0995, 1.0006, 30.9473)$ *using Monte Carlo simulation with 100 repetitions. Statistics produced using 1000 repetitions.* $(\rho, \mu) = (0.9947, 288.79)$ *are estimated from the historical expected inflation;* $\delta = 0.9987$ *is calibrated to match the 1 month T-bill.*

Statistic	Maturity	Data	Model	t-stat
Standard deviation yields σ_{Y^n}	1y	280.00	273.75	0.28
	2y	284.04	274.94	0.38
	3y	280.63	277.47	0.13
	5y	269.98	285.35	-0.62
	7y	260.52	296.43	-1.42
	10y	247.59	318.00	-2.92
Serial correlation yields ρ_{Y^n}	1y	0.89	0.87	0.65
	2y	0.90	0.87	1.41
	3y	0.91	0.87	2.08
	5y	0.92	0.87	2.79
	7y	0.92	0.88	2.59
	10y	0.92	0.89	2.39
Standard deviation excess returns σ_{xr^n}	2y	127.11	104.86	1.92
	3y	239.72	213.80	1.29
	5y	521.33	433.49	1.78
	7y	676.19	655.58	0.31
	10y	895.49	994.95	-1.08
EH coefficient b^n_{xr}	2y	1.37	1.80	-0.54
	3y	3.10	4.22	-0.68
	6y	10.84	10.51	0.08
	8y	15.53	14.62	0.17
	11y	24.61	20.82	0.51
RE coefficient b^n_s	2y	0.89	0.04	4.28
	3y	1.61	0.06	4.00
	6y	3.02	0.11	2.08
	8y	1.87	0.14	0.75
	11y	3.91	0.18	0.98

yield volatility, especially at the ten-year maturity, where the simulated standard deviation is around 30 basis points higher than in the data and the associated t-statistic reaches about -3. Overall, however, the level and the mild downward slope of yield volatility with maturity are both well captured by the estimated learning model.

The autocorrelation of yields is also broadly in line with the data. Empirically, yields are highly persistent, with first–order autocorrelations between 0.89 and 0.92 across maturities. The model generates similarly high persistence, with autocorrelations between 0.87 and 0.89. The resulting t-statistics range from 0.65 at the short end to around 2.5–2.8 for longer maturities, indicating that the model slightly understates the degree of persistence but remains close in economic terms: yields remain highly persistent in the simulated economy, as in the data.

Turning to excess holding-period returns, the model reproduces both their large magnitude and their maturity profile. In the data, the standard deviation of one-year excess returns increases sharply with maturity, from roughly 130 basis points at two years to almost 900 basis points at ten years. The model generates a similar steep rise in return volatility, and the levels are close: it modestly understates volatility at intermediate maturities (two to five years) and slightly overstates it at the ten-year maturity. The associated t-statistics are all below two in absolute value, so that the high volatility of long-term bond risk premia is matched reasonably well.

In terms of predictive regressions, the central success of the model is that it generates a large wedge between the EH and RE coefficients, very much like in the data. Empirically, the EH slopes b^n_{xr} are well above one and rise steeply with maturity, while the RE coefficients b^n_s are much smaller and comparatively flat. The learning model reproduces this pattern: it delivers a strongly upward-sloping term structure for b^n_{xr} and simultaneously keeps b^n_s close to zero across maturities. Although the individual RE coefficients in the model are below their empirical counterparts—especially at short and medium maturities—the gap $b^n_{xr} - b^n_s$ is of the right order of magnitude and displays the right maturity profile. In this sense, the estimated model not only matches the key second-moment properties of yields and excess returns, but also captures the core feature of the predictability evidence: large EH coefficients coexisting with much weaker RE-style predictability.

7.4 Cross-correlation in fundamentals

As a robustness test, in this section we model the joint dynamics of monthly inflation and policy news with a VAR(1) process. Let the data–generating process be

$$\begin{pmatrix} \pi_t \\ \xi_t \end{pmatrix} = \kappa + A \begin{pmatrix} \pi_{t-1} \\ \xi_{t-1} \end{pmatrix} + \begin{pmatrix} \nu_t \\ \varepsilon_t \end{pmatrix}, \qquad \begin{pmatrix} \nu_t \\ \varepsilon_t \end{pmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{pmatrix} \sigma^2_\nu & \sigma_{\nu,\varepsilon} \\ \sigma_{\nu,\varepsilon} & \sigma^2_\varepsilon \end{pmatrix} \right), \tag{61}$$

where $\kappa = (k_0, k_1)'$ is a vector of intercepts, A is a 2×2 coefficient matrix, and the shock vector $(\nu_t, \varepsilon_t)'$ has a full covariance matrix with variances $(\sigma_\nu^2, \sigma_\varepsilon^2)$ and covariance $\sigma_{\nu,\varepsilon}$.

We estimate the nine parameters of (61) using SMM. The targeted moments summarize the joint distribution of (π_t, ξ_t) and the persistence of the one–year yield. Specifically, the moment vector includes: (i) first– and twelfth–order autocorrelations of π_t and ξ_t; (ii) contemporaneous and lagged cross–correlations between ξ_t and inflation at different lags; (iii) unconditional means and standard deviations of π_t and ξ_t; and (iv) the autocorrelation of the one–year yield at a 12–month lag, $\text{corr}(Y_t^1, Y_{t-12}^1)$. The empirical moments are computed using Newey–West standard errors with lag length $L = 12$.

For any candidate parameter vector, we simulate the VAR(1) in (61) and the implied one–year yield over T periods, where T equals the length of the empirical sample. The initial conditions (π_0, ξ_0) are set equal to the initial values of the empirical time series. We then compute the same set of statistics as in the data.

Table 8 reports the targeted statistics (Panel A) and the estimated parameters of the VAR(1) (Panel B). The model matches the empirical autocorrelations, cross–correlations, means, and standard deviations reasonably well: the t–statistics of the moment discrepancies (computed using Newey–West standard errors with $L = 12$) are generally modest in magnitude. The estimated VAR coefficients are close to unity on the diagonal ($A_{00} \approx 0.98$ and $A_{11} \approx 0.98$), confirming the high persistence of both inflation and policy news, while the off–diagonal terms capture the comovement between the two series. The shock variances and covariance imply a positive correlation between inflation and policy news innovations.

Given the VAR dynamics of (π_t, ξ_t), the only remaining free parameter of the quantitative model is the Kalman gain g that governs how strongly agents update their estimate of the latent slope factor. We estimate g in a second SMM step that focuses on yield and excess–return statistics. Hence, we follow a two-stage procedure: i) estimate the VAR parameters by SMM to match the joint moments of (π_t, ξ_t) and the one–year yield autocorrelation; ii) hold the VAR parameters fixed at their SMM estimates and estimate the Kalman gain g using a separate SMM criterion to match the remaining statistics.

Table 9 reports the estimation results. The SMM estimate of the Kalman gain is $\hat{g} = 0.0554$. The table compares the data and model moments for all targeted statistics and reports the associated t–statistics. The estimated model reproduces fairly well the persistence and volatility of yields across maturities, as well as the

Table 8: *SMM estimation of the VAR(1): data vs. model statistics (t-stats via NW, $L = 12$) and estimated VAR(1) parameters.*

Panel A: Targeted statistics			
Statistic	Data	Estimate	t-stat
$\text{corr}(\pi_t, \pi_{t-1})$	0.9953	0.9944	1.03
$\text{corr}(\pi_t, \pi_{t-12})$	0.9383	0.9285	0.74
$\text{corr}(\xi_t, \xi_{t-1})$	0.9932	0.9891	2.86
$\text{corr}(\xi_t, \xi_{t-12})$	0.8412	0.8757	-0.83
$\text{corr}(\xi_t, \pi_t)$	0.6963	0.7303	-0.45
$\text{corr}(\xi_t, \pi_{t-1})$	0.6966	0.7298	-0.44
$\text{corr}(\xi_t, \pi_{t-12})$	0.6686	0.7166	-0.59
$\text{sd}(\xi_t)$	210.5944	234.7870	-1.81
$\text{sd}(\pi_t)$	94.9148	98.6257	-0.34
$\text{E}(\xi_t)$	90.8341	97.0946	-0.18
$\text{E}(\pi_t)$	288.7949	280.0654	0.55
$\text{corr}(Y_t^1, Y_{t-12}^1)$	88.5742	91.5579	-1.05
Panel B: Estimated parameters			
Parameter		Estimate	
k_0		4.0462	
k_1		-6.2784	
A_{00}		0.9821	
A_{01}		0.0006	
A_{10}		0.0234	
A_{11}		0.9818	
σ_ν		0.0994	
σ_ε		0.3235	
$\sigma_{\nu,\varepsilon}$		0.0354	

magnitude and maturity profile of excess–return volatility. The model generates a pronounced wedge between b_{xr}^n and the RE coefficients b_s^n, which is a key feature of the data and central to our interpretation of the role of learning.

8 Conclusions

Table 9: *SMM estimation of g to match the yields statistics. $\hat{g} = 0.0554$. SMM with 1000 repetitions. VAR parameters from Table 8.*

Statistic	Maturity	Data	Model	t-stat
Standard deviation yields σ_{Y^n}	1y	280.00	296.11	-0.72
	2y	284.04	275.48	0.36
	3y	280.63	258.46	0.91
	5y	269.98	246.55	0.95
	7y	260.52	267.36	-0.27
	10y	247.59	340.35	-3.85
Serial correlation yields ρ_{Y^n}	1y	0.89	0.91	-0.96
	2y	0.90	0.92	-0.66
	3y	0.91	0.92	-0.36
	5y	0.92	0.92	-0.29
	7y	0.92	0.94	-1.01
	10y	0.92	0.95	-1.79
Standard deviation excess returns σ_{xr^n}	2y	127.11	101.27	2.23
	3y	239.72	201.37	1.91
	5y	521.33	384.68	2.77
	7y	676.19	571.83	1.58
	10y	895.49	917.47	-0.24
EH coefficient b^n_{xr}	2y	1.37	0.18	1.52
	3y	3.10	1.38	1.04
	6y	10.84	5.06	1.41
	8y	15.53	7.77	1.45
	11y	24.61	12.39	1.65
RE coefficient b^n_s	2y	0.89	0.09	4.05
	3y	1.61	0.12	3.84
	6y	3.02	0.18	2.03
	8y	1.87	0.19	0.73
	11y	3.91	0.19	0.98

References

- Adam K., A. Marcet (2011). "Internal Rationality and Asset Prices" *Journal of Economic Theory.*
- Adam, K., A. Marcet and J.P. Nicolini (2016). "Stock Market Volatility and Learning", *Journal of Finance.*
- Adam, K., A. Marcet and J. Beutel (2017). "Stock Price Booms and Expected Capital Gains", *American Economic Review.*

- Bansal, R. and Ivan Shaliastovich (2013) "A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets", The Review of Financial Studies, Volume 26, Issue 1.
- Bauer M., Glenn D. Rudebusch (2020) "Interest Rates under Falling Stars", *American Economic Review*, 110(5): 1316–1354.
- Belda, P. (2023) "Capital Gains Taxation, Learning and Bubbles." Working paper.
- Buraschi, A. and P Whelan (2022) "Speculation, sentiment, and interest rates" Management Science.
- Buraschi, A., and A. Jiltsov, 2006, Model uncertainty and option markets with heterogeneous beliefs, The Journal of Finance 61, 2841–2897.
- Chernov, M. and P. Mueller, 2012 "The term structure of inflation expectations", *Journal of Financial Economics*, 106, 367–394.
- Cochrane, J., and M. Piazzesi. (2005) "Bond Risk Premia." *American Economic Review* 95:138–60.
- Coibion, O. and Y. Gorodnichenko (2015) "Information rigidity and the expectations formation process: A simple framework and new facts", *American Economic Review* 105 (8).
- Dai, Qiang, and Kenneth J. Singleton, 2000, "Specification analysis of affine term structure models", Journal of Finance 55, 1943-1978.
- Duffee, Gregory R., 2002, "Term premia and interest rate forecasts in affine models", Journal of Finance 57, 405-443.
- Duffie, Darrell, and Rui Kan, 1996, "A yield-factor model of interest rates", Mathematical Finance 6, 379-406.
- Ehling P., M Gallmeyer, C Heyerdahl-Larsen, P OIlledistch (2018) "Disagreement about inflation and the yield curve", Journal of Financial.
- Ellison, M. and Andreas Tischbirek (2021) "Beauty Contests and the Term Structure" Journal of the European Economic Association.

- Evans, G. and S. Honkapohja, (2001) *Learning and Expectations in Macroeconomics*, Princeton University Press.
- Gabaix, X. (2008) "Variable rare disasters: A tractable theory of ten puzzles in macro-finance" American Economic Review.
- Giacoletti, M., K.T. Laursen and K. J. Singleton (2020) "Learning From Disagreement in the U.S. Treasury Bond Market" *Journal of Finance*, July 2020.
- Giglio, S and Bryan Kelly (2018) "Excess Volatility: Beyond Discount Rates" The Quarterly Journal of Economics, Volume 133, Issue 1.
- Joslin S., M. Priebsch, K.J. Singleton (2014) "Risk Premiums in Dynamic Term Structure Models with Unspanned Macro Risks" The Journal of Finance.
- Laubach T., RJ Tetlow, JC Williams (2007) "Learning and the role of macroeconomic factors in the term structure of interest rates", working paper, Federal Reserve Board.
- Marcet, A. and T. Sargent, (1989), 'Convergence of Least Squares Mechanisms in Self-Referential Linear Stochastic Models', *Journal of Economic Theory* 48, 337 - 368.
- Piazzesi, M. and Martin Schneider (2006) "Equilibrium Yield Curves" NBER Macroeconomics Annual Volume 21.
- Shiller, R. ("1979) The Volatility of Long-Term Interest Rates and Expectations Models of the Term Structure" Journal of Political Economy 87:1190–219.
- Singleton K. J., "How Much 'Rationality' Is There in Bond-Market Risk Premiums?", AFA Presidential address, *Journal of Finance* August 2021 Vol. 76 Issue 4. Video:
- Sinha, A. (2016) "Learning and the yield curve" Journal of Money, Credit and Banking, 2016.
- Wachter, J. (2006) "A consumption-based model of the term structure of interest rates", Journal of Financial Economics.
- Williams, N. (2019) "Escape Dynamics in Learning Models", *The Review of Economic Studies*, Volume 86, Issue 2.

- Winkler, F. (2020) "The Role of Learning for Asset Prices and Business Cycles", Journal of Monetary Economics.
- Xiong, W. (2013) "Bubbles Crises and Heterogeneous Beliefs", *Handbook on Systemic Risk*, Cambridge University Press.
- Xiong W., Hongjun Yan (2010) "Heterogeneous Expectations and Bond Markets" *The Review of Financial Studies*.

Appendix 1: Allowing for $\rho < 1$.

In the even simpler model of section 6 we now remove simplification 2 and we allow for any value of $|\rho| \leq 1$ in the inflation process (6).

This serves as an example of the fact that, under IR, after all assumptions have been made the equilibrium of the model follows from logical and mathematical derivations.

Recall that we maintain Assumptions 1- 2 in section 4.2 throughout the paper and that we maintain simplification 5 in section 6 saying that investors know the true intercepts and the true loadings on inflation ϕ^n. Since we are now considering $|\rho| < 1$ it is no longer true that $\phi^n = 1$ in the perceived yield curve equals the true loadings on inflation. Examining the RE equilibrium in this case it is easy to see that when $|\rho| < 1$

$$\begin{aligned} Y_t^{RE,n} &= a^n + \phi^n \pi_t && (62) \\ \phi^n &= \frac{1}{n}\sum_{i=1}^{n} \rho = \rho\left[\left(1 - \frac{1}{n}\right)\phi^{n-1} + \frac{1}{n}\right] && (63) \\ a^n &= -\log\delta + \mu(1-\phi^n) && (64) \end{aligned}$$

for $\phi^{0,RE} = 1$.

A consequence of simplification 5 is that the loadings of inflation in the perceived yield curve (8) equal those of RE: It can be checked easily that if we set the intercept and inflation loadings $\mathbf{a}_t^n = a^n$ and ϕ^n as in the last equation then the actual yield curve becomes

$$\begin{aligned} nY_t^n &= -\log\delta + E_t(\pi_{t+1}) + (n-1)(\mathbf{a}_t^{n-1} + \widehat{\mathbf{s}}_t(n-1) + \phi^{n-1}E_t(\pi_{t+1})) && (65) \\ &= -\log\delta + \mu(1-\rho) + (n-1)a^{n-1} && (66) \\ &\quad + \widehat{\mathbf{s}}_t(n-1)^2 + (n-1)\phi^{n-1}\mu(1-\rho) + [(n-1)\phi^{n-1}\rho + \rho]\pi_t) && (67) \end{aligned}$$

where the first equation comes from equation (18), the second from Assumptions 1-2. Simple algebra gives that the actual yield curve is

$$Y_t^n = a^n + \widehat{\mathbf{s}}_t(n-1)^2/n + \phi^n\pi_t \quad (68)$$

Notice that the constant a^n and the coefficient multiplying inflation ϕ^n are the same in the perceived and the actual yield curve or, equivalently, investors are correct about their beliefs on the intercept a^n and the inflation loadings ϕ^n.

Now, given the belief (62), from the agents' point of view the slope has the following form

$$\frac{Y_t^m - Y_t^1}{m-1} = \mathbf{s}_t + \frac{(\mu - \pi_t)(\rho - \phi^m) + u_t^m - u_t^1}{m-1}$$

Therefore the t-slope s_t (that is the best estimate of $\mathbf{s}_t$ given information at t and the beliefs $\mathcal{P}$) to be plugged in the Kalman filter (22) is now

$$s_t = \frac{Y_t^m - Y_t^1}{m-1} - \frac{(\mu - \pi_t)(\rho - \phi^m)}{m-1}.$$

Plugging (68) in the previous equation we have that the *actual* law of motion for s is

$$s_t = \frac{(m-1)\widehat{\mathbf{s}}_t}{m} - \frac{\xi_t}{m-1} \tag{69}$$

Therefore s_t and $\widehat{\mathbf{s}}_t$ are given by the same expression as in the unit root case (27).

Applying (68) to the formula for excess returns we have

$$\begin{aligned} xr_{t+1}^n &= \left[(n-1)^2 - (n-2)^2(1 - \frac{g}{m})\right]\widehat{\mathbf{s}}_t && (70)\\ &\quad + \phi^{n-1}(n-1)\eta_{t+1} + \left[\frac{(n-2)^2 g}{m-1} - 1\right]\xi_t && (71)\\ &= \left[2n - 3 + (n-2)^2\frac{g}{m}\right]\widehat{\mathbf{s}}_t && (72)\\ &\quad + \phi^{n-1}(n-1)\eta_{t+1} + \left[\frac{(n-2)^2 g}{m-1} - 1\right]\xi_t && (73) \end{aligned}$$

summarising all these derivations we have that in addition to this equation the actual equilibrium is given by (68), (69).

All the second moments in Table **??** are easily computed with these equations.

For the RE regression coefficients notice that now the slope $\frac{Y_t^m - Y_t^1}{m-1}$ does not equal s, therefore we need to use the actual law of motion

$$\frac{Y_t^n - Y_t^1}{n-1} = \frac{(\mu - \pi_t)(\rho - \phi^n)}{n-1} + \frac{(n-1)\widehat{\mathbf{s}}_t}{n} - \frac{\xi_t}{n-1}$$

Therefore

$$\begin{aligned} b_n^{xr} &= \frac{cov(xr_{t+1}^n, \frac{Y_t^n - Y_t^1}{n-1})}{\sigma^2_{\frac{Y_t^n - Y_t^1}{n-1}}} \\ &= \frac{\left[2n-3+(n-2)^2\frac{g}{m}\right]\frac{n-1}{n}\sigma^2_{\widehat{\mathbf{s}}} - \left(\frac{(n-2)^2 g}{m-1} - 1\right)\frac{\sigma^2_\xi}{n-1}}{\left(\frac{\rho-\phi^n}{n-1}\right)^2\sigma^2_\pi + \left(\frac{n-1}{n}\right)^2\sigma^2_{\widehat{\mathbf{s}}} + \frac{\sigma^2_\xi}{(n-1)^2}} \end{aligned} \tag{74}$$

As for the second RE regression we have, using investors' beliefs

$$\begin{aligned} E_t^{\mathcal{P}}(xr_{t+1}^n) &= -(n-1)E_t^{\mathcal{P}}(Y_{t+1y}^{n-1}) + nY_t^n - Y_t^1 \\ &= -(n-1)[-\log\delta + \mu(1-\rho\phi^{n-1}) + \rho\phi^{n-1}\pi_t + \widehat{\mathbf{s}}_t(n-1)] + nY_t^n - Y_t^1 \end{aligned} \tag{75}$$

using the actual yield curve

$$\begin{aligned} E_t^{\mathcal{P}}(xr_{t+1}^n) &= -(n-1)[-\log\delta + \mu(1-\rho\phi^{n-1}) + \rho\phi^{n-1}\pi_t + \widehat{\mathbf{s}}_t(n-1)] \\ &\quad + n(-\log(\delta) + \mu(1-\phi^n) + \phi^n\pi_t + \widehat{\mathbf{s}}_t\,\frac{(n-1)^2}{n}) + \log\delta - \mu(1-\rho) - \rho\pi_t - \xi_t \\ &= -\xi_t \end{aligned} \tag{76}$$

thus

$$b_n^{\mathcal{E}} = \frac{cov(\xi_t, \frac{Y_t^n - Y_t^1}{n-1})}{\sigma^2_{\frac{Y_t^n - Y_t^1}{n-1}}} = \frac{\sigma^2_\xi}{\frac{(\rho-\phi^n)^2}{n-1}\sigma^2_\pi + \frac{(n-1)^3}{n^2}\sigma^2_{\widehat{\mathbf{s}}} + \frac{\sigma^2_\xi}{n-1}} \tag{77}$$