BANCO CENTRAL DE RESERVA DEL PERÚ

Learning Through the Yield Curve

Marco Ortiz*

* Banco Central de Reserva del Perú.

DT. N° 2013-018
Serie de Documentos de Trabajo
Working Paper series
Diciembre 2013

Learning Through the Yield Curve *

Marco Ortiz

Banco Central de Reserva del Perú (BCRP)
and London School of Economics (LSE)

December 12, 2013

Abstract

This paper presents a model in which investors form their expectations in an adaptive way to price bonds, in the spirit of Adam, Marcet and Nicolini (2011). We follow different assumptions regarding the learning process followed by agents. In the case of finite maturity bonds, the knowledge of the pricing of the first maturity will act as an "anchor", limiting the price volatility of bonds with short maturities. As the maturity increases, the price volatility converges to that of the consol bond. Our results suggests this learning mechanism can help the model capture some of the observed empirical dynamics of yield curve.

Resumen

Este documento presenta un modelo en el cual los inversionistas forman sus expectativas sobre el precio de los bonos de forma adaptativa, en línea con el trabajo de Adam et al. (2011). Exploramos diferentes supuestos relativos al proceso de aprendizaje seguido por los agentes. En el caso de bonos con un periodo de maduración finito, el conocimiento respecto a la regla de formación de precios del bono a un periodo actúa como un 'ancla', limitando la volatilidad del precio de los bonos con periodos de maduración cortos. A medida que el periodo de maduración se incrementa, la volatilidad del precio de los bonos converge a la de los bonos perpetuos. Nuestros resultados sugieren que este mecanismo de aprendizaje puede ayudar al modelo a capturar algunos de los hechos estilizados asociados con la dinámica de la curva de rendimiento.

JEL Code: G12, D83, D84.

*The opinions expressed in this paper are not necessarily shared by the institutions with which the author is currently affiliated. I wish to thank Albert Marcet and Wouter den Haan for helpful comments and suggestions. I thank the participants of the BCRP research seminar. Correspondence to: Marco Ortiz, email: marco.ortiz@bcrp.gob.pe

1 Introduction

The adaptive learning approach proposes the modelling of agents as econometricians, who use all the available data to construct forecasting functions, revised over time as new data becomes available. Adam et al. (2011) use this approach to successfully explain a series of equity pricing puzzles. Modelling agents who learn about the behaviour of equity prices is key to their results. This paper explores whether this particular mechanism is useful for the study of bond yield dynamics.

Carceles-Poveda and Giannitsarou (2008) provide a thorough analysis of the introduction of *self-referential* learning general equilibrium framework. Self-referential learning means that agents' beliefs affect the behaviour of economic variables, which in turn affect agents' beliefs and so on. In the models presented by Carceles-Poveda and Giannitsarou (2008), agents are assumed to know the functional form of the law of motion relating the variable of interest to the state variables, but they do not know the value of the coefficients in this law of motion. The authors conclude that the effects of learning in their setup are modest, providing very little improvements over the rational expectations (RE) case. Timmermann (1996) analyzes self-referential learning when agents also learn about the exogenous dividend process. This mechanism is capable of increasing the volatility of simulated stock prices but, as in the case of Carceles-Poveda and Giannitsarou (2008), its impact is modest. The key difference between these results and the ones obtained by Adam et al. (2011) is that in the latter, agents are not endowed with knowledge of the mapping between dividends and prices. One of the critical implications of this assumption is that agents are unable to express the value of the asset as a discounted sum of payoffs. Moreover, in a related paper, Adam and Marcet (2011) show that endowing agents with this knowledge involves a large set of strong assumptions about agents' beliefs. This means that agents have to form beliefs about the law of motion of stock prices directly. These beliefs are used to forecast next period's stock price. Since this period's stock price is affected by expectations of next period's stock price, there is a very direct link between beliefs about stock prices and the actual behaviour of stock prices.

We follow Adam et al. (2011) by presenting a model in which agents learn about bond prices. Applying the idea of Adam et al. (2011) to bond prices means that learning is strictly speaking no longer self-referential, unless a consol bond is considered. That is beliefs about the price of an n-th period bond do not affect the behaviour of the price of this n-th period bond, but they do affect the behaviour of the prices of bonds with a higher maturity. We obtain three key results: First, learning about price dynamics affects the relative volatility of bonds in a heterogeneous way, as the prices of higher maturity assets become more volatile relative to the rational expectations (RE) case. Second, as maturity increases, the price volatility converges to the one implied by the consol bond case, which we treat as a benchmark since, similar to equity, this instrument has no redemption date. Finally, we perform a series of numerical exercises that suggest this mechanism can be useful for explaining the pattern of volatilities observed in the term structure.

The present document is organized in the following way. In Section 2 we present the baseline model and the proposed cases for study. In Section 3 we discuss the learning rule and contrast analytically the case of consol bonds, in which learning is *self-referential* and the case of finite-maturity bonds. In Section 4, we present numerical simulations and report findings regarding the implied behaviour of price and yield volatility in both cases. Section 5 presents our conclusions.

2 Setup

2.1 The Economy

The economy is composed by infinite-lived consumers-investors who face the following problem:

$$\max_{c_t} E_t \sum_{t=0}^{\infty} \delta^t \ln c_t, \tag{1}$$

Agents can save their resources in nominal assets. In particular, we assume they have access to a series of $\mathcal{T}$ risk-free bonds with a payoff at maturity and a series of coupons in each period, denoted by ϕ_t. $Q_t^{(\tau)}$ stands for the nominal price of the bond with τ periods to maturity in period t, while $B_{t+1}^{(\tau)}$ is the quantity of bonds. We assume $\mathcal{T}$ is an arbitrary high number. Additionally, agents have access to a consol type bond, which price we denote by Q_t^{∞}, and pays the same coupon as the finite maturity bonds.

The representative investor budget constraint will be given by:

$$P_t c_t + \sum_{\tau=1}^{\mathcal{T}} Q_t^{(\tau)} B_{t+1}^{(\tau)} + Q_t^{\infty} B_{t+1}^{\infty} \leq P_t y_t + \sum_{\tau=1}^{\mathcal{T}} \left(Q_t^{(\tau)} + \phi_t \right) B_t^{(\tau)} + (Q_t^{\infty} + \phi_t) B_t^{\infty}, \quad \forall_t.$$

The first order conditions over the bonds yield the following set of Euler equations:

$$Q_t^{(i)} = \delta E_t \left[\left(\frac{P_t}{P_{t+1}} \frac{c_t}{c_{t+1}} \right) \left(Q_{t+1}^{(i-1)} + \phi_{t+1} \right) \right], \ \forall_{i \in [1, \mathcal{T}]} \tag{2}$$

$$Q_t^{(\infty)} = \delta E_t \left[\left(\frac{P_t}{P_{t+1}} \frac{c_t}{c_{t+1}} \right) \left(Q_{t+1}^{(\infty)} + \phi_{t+1} \right) \right], \tag{3}$$

where Q_t^0 is the principal paid by the bond at maturity, in period t. The principal is given by $Q_t^{0,I} = \phi_t \bar{Q}^{0,I}$. Thus, coupons and principal will share the same growth rate.

2.2 Cases

The cases we present in this paper will follow the pricing equations given by 2 and 3, however they will differ in the assumptions we make regarding the process for inflation and the coupon payments, as well as in the way agents form their beliefs, defined by the *perceived law of motion* (PLM).

Case I: Stochastic coupon growth and learning about nominal bond price growth rates.
We begin as close as possible to Adam et al. (2011) to provide a benchmark. This exercise will show us how the introduction of a finite maturity alters the results obtained by the

aforementioned authors for the case of equity. We make the pay-off structure stochastic by defining the following process for the coupon and principal:

$$\phi_{t+1} = \gamma_\phi \phi_t \varepsilon_t, \;\; \varepsilon_t \sim N(1, \sigma_\phi^2), \;\; \gamma_\phi > 0. \tag{4}$$

Under stochastic coupons, a consol-type bond will have the same pay-off structure as equity. Since here we focus on the coupon as the main driver of these assets dynamics, consumption will be parameterized as random walk process, following Cochrane and Campbell (1999):

$$\frac{c_t}{c_{t+1}} = \mu_c + \varepsilon_t^c \tag{5}$$

where $\varepsilon_t^c \sim N(0, \sigma_c^2)$.[1] We assume as well the consumption growth covariance with the rest of the processes in the model is zero.[2] Prices are assumed constant. We follow Adam et al. (2011) by specifying agents with the following beliefs:

$$Q_{t+1}^{(i,I)} = \beta_t^i Q_t^{(i,I)} \epsilon_t^{(i)}, \; \forall_{i>1} \tag{6}$$

where $Q_{t+1}^{(i,I)}$ stands for the price of the bond with i periods to maturity, in the Case I setup, and $\epsilon_t \sim i.i.d\, N(1, \sigma_\epsilon^2)$. Note that agents will hold an individual PLM for each maturity.[3]

Case II: Stochastic inflation and learning about bond price real growth rates

The case of stochastic coupon payments provides a good approximation to the effects of different maturities for the learning mechanism in Adam et al. (2011). However, bonds exhibit a deterministic path for coupons (usually flat). For this reason we change this assumption and introduce now a deterministic growth rate for coupons.

$$\phi_{t+1} = \gamma_\phi \phi_t; \tag{7}$$

which considers the case of flat coupon-payment schedule when $\gamma_\phi = 1$. The processes for consumption remains the same as in the previous case. However, inflation will now be stochastic:

$$\frac{P_t}{P_{t+1}} = \gamma_\pi \varepsilon_t^\pi, \;\; \varepsilon_t^\pi \sim N(1, \sigma_\pi^2), \;\; \gamma_\pi \geq 0. \tag{8}$$

[1]The literature has not reached a definitive consensus on the process consumption growth follows. While some authors find that the implications for consumption growth in the model of Hall (1978) - consumption follows a random walk - cannot be completely rejected by the data, others claim that the series fits a process with high serial autocorrelation. For a discussion see Carroll et al. (2011).

[2]Introducing a covariance would complicate the analysis as the actual growth rate of prices would stop being constant. Throughout the present paper we will assume agents only care about first moments.

[3]We simply follow the proposed learning rule in Adam et al. (2011). Alternatively, we could have set a model where agents iterate 2 forward and obtain a yield curve where all maturities depend on the same factors. This would be equivalent to the learning mechanism proposed by Carceles-Poveda and Giannitsarou (2008), which yields modest results. Adam and Marcet (2011) discuss this element of arbitrariness, which is often present in learning models, and propose a microfounded framework where the *Law of iterated expectations* ceases to hold and learning about price dynamics arises as an optimal behaviour. This type of learning can be rationalized in several ways: (1) as a model where agents possess short-term buy and sell strategies; (2) as a model where agents do not know they are the marginal pricer; and (3) as a model of vanishing heterogeneity across agents. For a discussion see Adam and Marcet (2011).

We assume inflation is uncorrelated with the rest of the processes in the economy.[4] Agents will focus on learning the real growth rate of bonds. For this reason, the PLM will be given by:

$$\frac{Q_{t+1}^{(i,II)}}{P_{t+1}} = \beta_t^i \frac{Q_t^{(i,II)}}{P_t} \epsilon_t^{(i)}, \; \forall_{i>1} \tag{9}$$

In this case, we have chosen to use inflation as the process behind bond price dynamics. Assuming similar dynamics for the consumption growth rate will generate equivalent results.[5]

3 Learning

3.1 Analytic results

In this section we study analytically how the introduction of the learning mechanisms proposed in Section 2 affects the dynamics of asset prices. This is important since in order to stay as close to Adam et al. (2011), we restrict our attention to learning mechanisms that comply with two desirable properties: (1) the laws of motion under a least-squares learning rule converge to those corresponding to the rational expectations equilibrium, and (2) learning should be *reasonable*. To check whether the first condition is satisfied, we analyse the ordinary differential equations (ODE) associated with the stochastic recursive algorithm (SRA) that describes the dynamics of the processes under learning. For this purpose we construct T-maps and check if the rational expectations equilibrium constitutes an equilibrium of the system of ODE. For the second condition, we verify that the PLM presented in Section 2 are not *misspecified*, in the sense that PLMs used cannot possibly converge to a REE.[6]

We start by defining a general learning function. For the case of nominal price growth learning, beliefs will be updated following:

$$\hat{\beta}_t^{(i)} = \hat{\beta}_{t-1}^{(i)} + g_t \left[\frac{Q_{t-1}^{(i)}}{Q_{t-2}^{(i)}} - \hat{\beta}_{t-1}^{(i)} \right], \; \forall_{i \in [1,\mathcal{T}], i=\infty} \tag{10}$$

where $g_t(0) = 0$ and $g'(\cdot) > 0$. These conditions define a learning process that adjusts beliefs in the same direction as the prediction error.[7]

[4] As we show in appendix A, under rational expectations, the inflation rate and bond prices will be uncorrelated. However the covariance will not be zero *along the learning path*. For this reason we need agents who hold *linear* beliefs as a condition for convergence to the rational expectations equilibrium.

[5] In the case consumption growth is chosen as the source of dynamics agents will learn about the *risk-adjusted* price growth rate. For an example see Adam et al. (2011).

[6] It is important differentiate this from econometric misspecification. Most learning models are misspecified from an econometric viewpoint, since agents fail to recognize the *self-referential* nature of the process they estimate. See Evans and Honkapohja (2001), Ch.13.

[7] The assumption that agents use lagged information to update their beliefs is standard in the learning literature as simultaneous updating gives rise to a series of difficulties. For a discussion see Evans and Honkapohja (2001).

Since both cases share the same properties, we focus on Case I for our analytical results.[8] First we focus on the consol bond. Substituting assumptions 4, 5, together with the PLM for this case into Eq. 3, we obtain:

$$Q_t^{(\infty,I)} = \delta\mu_c\gamma_\phi\phi_t + \delta\mu_c\hat{\beta}_t^\infty Q_t^{(\infty,I)},$$

where $\hat{\beta}_t^\infty$ is the belief agents hold about β^∞. Solving for Q_t^∞ yields:

$$Q_t^{(\infty,I)} = \frac{\delta\mu_c\gamma_\phi}{1-\delta\mu_c\hat{\beta}_t^\infty}\phi_t \tag{11}$$

From this result we can derive the actual behaviour of the consol bond growth rate, which is given by:

$$\frac{Q_t^{(\infty,I)}}{Q_{t-1}^{(\infty,I)}} = \frac{1-\delta\mu_c\hat{\beta}_{t-1}^\infty}{1-\delta\mu_c\hat{\beta}_t^\infty}\frac{\phi_t}{\phi_{t-1}} \tag{12}$$

Using 4 and collecting terms we obtain:

$$\frac{Q_t^{(\infty,I)}}{Q_{t-1}^{(\infty,I)}} = \left(\gamma_\phi + \frac{\gamma_\phi\delta\mu_c\Delta\hat{\beta}_t^\infty}{1-\delta\mu_c\hat{\beta}_t^\infty}\right)\varepsilon_t = T\left(\hat{\beta}_t^\infty, \Delta\hat{\beta}_t^\infty\right)\varepsilon_t \tag{13}$$

where:

$$T\left(\beta^\infty, \Delta\hat{\beta}^\infty\right) \equiv \gamma_\phi + \frac{\gamma_\phi\delta\mu_c\Delta\hat{\beta}_\infty}{1-\delta\mu_c\hat{\beta}^\infty} \tag{14}$$

is the T-mapping, which summarizes the actual behaviour of the infinite-maturity asset growth rate for given values of β and $\Delta\beta$. Therefore, the dynamics of consol bond prices are not only defined by the beliefs agents hold on the growth rate, $\hat{\beta}$, but also by the change in these beliefs. This generates *momentum* in the consol price dynamics, which is the key element explaining the low frequency "ups and downs" observed in the data.[9] This result is obtained because learning is *self-referential*. When a positive shock to beliefs occur, it increases the observed future growth rates, which in turn, pushes up future beliefs. Notice as well that the specified PLM allows for *reasonable* learning, since provided beliefs converge, consol bond prices will follow the behaviour implied by rational expectations.[10]

For the case of finite-maturity assets, we obtain from Eq. 2:

$$Q_t^{(i,I)} = \delta\mu_c\gamma_\phi\phi_t + \delta\mu_c\hat{\beta}_t^i Q_t^{(i-1,I)}, \ \forall_{i\in[2,\mathcal{T}]} \tag{15}$$

while the one period to maturity asset price will be given by:

$$Q_t^{1,I} = \delta\mu_c\gamma_\phi\left(\phi_t + \phi_t\bar{Q}^{0,I}\right) \tag{16}$$

[8] We refer the reader to appendix B for derivations for Case II.

[9] See Adam et al. (2011).

[10] See Appendix A for derivations.

Substituting the actual growth rates for each maturity, it is possible to express the actual growth rate as:[11]

$$\frac{Q_t^{(i,I)}}{Q_{t-1}^{(i,I)}} = \left(\gamma_\phi + \frac{\gamma_\phi \Delta\hat{\beta}_t^{(i-1)} + \gamma_\phi \hat{\beta}_t^{(i-1)} \Omega_t^{(i-1)}}{\gamma_\phi \frac{\phi_{t-1}}{Q_{t-1}^{(i-1,I)}} + \hat{\beta}_t^{(i-1)}} \right) \varepsilon_t, \quad \forall_{i \in [2,\mathcal{T}]}, \tag{17}$$

where:

$$\Omega_t^{(i)} = \frac{\Delta\hat{\beta}_t^{(i-1)} + \hat{\beta}_t^{(i-1)} \Omega_t^{(i-1)}}{\gamma_\phi \frac{\phi_{t-1}}{Q_{t-1}^{(i-1,I)}} + \hat{\beta}_{t-1}^{(i-1)}}, \quad \forall_{i \in [2,\mathcal{T}]}. \tag{18}$$

and:

$$\Omega_t^{(1)} = 0, \tag{19}$$

Equation 17 represents the actual dynamics followed by the price growth rate of each maturity. We notice several interesting features. First, learning stops being *self-referential*. In this case the shocks affecting the beliefs do not have a feedback effect. An increase in the estimated growth rate for the price of maturity i, (i.e.: $\Delta\hat{\beta}_t^{(i)} > 0$), will affect prices for maturities i and higher, but, given the family of learning rules defined by 10, not maturity i. When agents believe the price of the i periods to maturity asset will increase over the next period, the price of the subsequent maturity $(i+1)$ would increase today. This affects the perceived dynamics of this maturity and consequently the price of the $i+2$ periods to maturity asset. In this sense, agents learn *through* the yield curve.

Second, we verify that these *perceived laws of motion* are well specified, as the change in the bond price would converge under least squares learning to the predicted growth rate under rational expectations for all maturities.[12] [13] If $\Delta\hat{\beta}^{(i)} = 0, \ \forall_i$, then:

$$\hat{\beta}_t^i = \gamma_\phi, \ \forall_i \tag{20}$$

The third result is related in the way in which pricing errors in previous maturities affect new ones. As in the case of rational expectations, a shock in the coupon growth rate will increase the price of all maturities. However, due to the presence of price learning, an additional effect emerges as higher maturities carry changes in expectations from lower maturities. This effect will not necessarily push the price of bonds in the same direction as the shock, since the updating of beliefs occurs using lagged information. Given the non-linearities and interaction among different maturities, we analyse these dynamics in Section 4 through numerical simulations.

Finally, it is important to stress that when agents learn the growth rate for each maturity, the growth rates they observe will actually differ across maturities, even if their priors are the correct ones.[14]

[11] See appendix B for derivations.

[12] Strictly, we need to define additional conditions for convergence. In specific, we would have to make use of the *Projection Facility*, which imposes bounds on the values of the $\beta's$. Some additional assumptions must be made on the support of ε. For a discussion see Adam et al. (2011).

[13] See Appendix A for derivations of the rational expectations equilibrium growth rates.

[14] This result is related to the *self-fulfilled* dynamics observed in other applications of learning models. For an example see Branch and Evans (2011).

3.2 Learning

In order to complete the characterization of bond price dynamics it is necessary to define how agents update their beliefs. The literature presents several alternatives. The most popular ones are: (a) *least-squares learning* (LSL) and (b) *constant-gain learning* (CGL). The former depicts agents who put the same weight on each observation. Therefore, it is a *decreasing-gain* learning mechanism. By contrast, under CGL, agents always put a higher weight on new observations relative to previous ones. As Sargent (1993) points out, the use of this type of learning may reflect agents' concerns with regime changes or a preference for adaptability.[15] In this section we consider the CGL case. Branch and Evans (2006) find ample empirical evidence supporting the use of this type of learning when modelling agents expectations.[16] Now we define the beliefs updating equations. For the case of simple growth rates, as in Case I:

$$\hat{\beta}_t^{(i)} = \hat{\beta}_{t-1}^{(i)} + \alpha \left[\frac{Q_{t-1}^{(i)}}{Q_{t-2}^{(i)}} - \hat{\beta}_{t-1}^{(i)} \right], \; \forall_{i \in [1,\mathcal{T}], i=\infty} \tag{21}$$

where α is a constant and positive gain parameter.[17] When agents learn about the price real growth rate, as in Case II, beliefs follow:

$$\hat{\beta}_t^{(i)} = \hat{\beta}_{t-1}^{(i)} + \alpha \left[\frac{P_{t-2}}{P_{t-1}} \frac{Q_{t-1}^{(i)}}{Q_{t-2}^{(i)}} - \hat{\beta}_{t-1}^{(i)} \right], \; \forall_{i \in [1,\mathcal{T}], i=\infty} \tag{22}$$

4 Numerical Exercises

4.1 Baseline Calibration

Each model involves five parameters, reported in tables 1 and 2 for cases I and II. We set a value of 0.994 for the time-preference parameter, in order to match the United States ex-post 3-month treasury bill average real return rate for the 1969-2013 period. The growth rates of consumption and inflation are taken from the *Bureau of Economic Analysis* NIPA tables, for the years 1969 to 2013. In the stochastic coupon case we set the growth rate and standard deviation of coupons to match the behaviour of consumption. Finally, the gain parameter is set at 0.001 for the case of stochastic coupons and to 0.01 for the case of learning over the real growth rate. We set this parameter to avoid falling into the projection facility upper and lower bounds, for the simulated maturities. This means that all dynamics are generated by the learning process and not by additional constraints. We remind the reader that our emphasis is not into matching any empirical moments, but simply to analyse how maturity plays a role under the proposed learning mechanism.

[15] On the other side, the presence of constant-gain learning can give rise to unexpected dynamics. For discussion see Williams (2001). See Evans and Honkapohja (2001), Ch. 14 for a general discussion of the properties of models of constant-gain learning.

[16] Even if we assume CGL, we consider PLMs that allow for convergence to the rational expectations' beliefs when LSL is followed.

[17] We make use of the standard timing assumption in order to avoid the joint determination of beliefs and observed prices. For a discussion see Evans and Honkapohja (2001), Ch. 3.

Table 1: **Baseline Calibration: Case I**

Parameter	*Value*	*Description*
δ	0.994	consumers time-preference parameter.
μ_c	0.993	inverse of consumption growth factor.
γ_ϕ	1/0.993	growth rate of coupons and principal.
σ_ϕ	0.005	standard deviation of coupon growth rate.
α	0.001	gain parameter (fixed).

Table 2: **Baseline Calibration: Case II**

Parameter	*Value*	*Description*
δ	0.994	consumers time-preference parameter.
μ_c	0.993	inverse of consumption growth factor.
γ_π	0.989	inverse of inflation factor.
σ_π	0.007	standard deviation of inflation factor.
α	0.01	gain parameter (fixed).

4.2 Results

We perform numerical simulations for both cases. First, we confirm our analytical results. As the maturity increases, the volatility of the price/coupon ratio under learning rises. Figure 1, shows the results for Case I, the model with stochastic coupon growth. We report the sample standard deviations for the learning case for both the finite maturity and consol bonds. We observe a pattern that increases with maturity, approximating the consol bond. Notice that under rational expectations the price/coupon ratio is constant.

The numerical results confirm the insights obtained from our analytical derivations. Learning affects volatilities in a heterogeneous way across maturities. There are two reasons for this: First, higher maturities carry on pricing errors from previous ones. Due to the fact that the one period to maturity bond is not subject to pricing errors, the number of terms to maturity limits the pricing error of a given asset. Second, as maturity increases, the importance of capital gains, relative to coupon income, rises. The fact that agents know perfectly how to price next period coupons, makes learning about prices more important for higher maturities. As maturity increases, this effect vanishes and the volatility of the asset starts converging to the once implied by the consol bond case.

Figure 2 shows the evolution of beliefs for three particular maturities. As we can observe, the behaviour of beliefs for the highest maturity approximates the one of the infinite-lived asset.

For Case II we observe similar dynamics. This is not surprising since, from our analytical derivations, we know that the way in which changes in beliefs affect higher maturities is analogous to the case with an stochastic coupon growth rate. Because under rational expectations prices would be fixed, we report the price standard deviation for each maturity. Figure 3 shows the results. As we observe, simulated standard deviations of prices approximate the one of the consol bond as maturity increases. Beliefs show a similar pattern as the dynamics of the beliefs for the highest maturity approximate the ones of the consol bond in each simulation. Note that only at very high maturities is the volatility of bond prices of finite-maturity bonds similar to that of the consol bond. Learning the price of a one-period bond is simple since it is only involves learning about

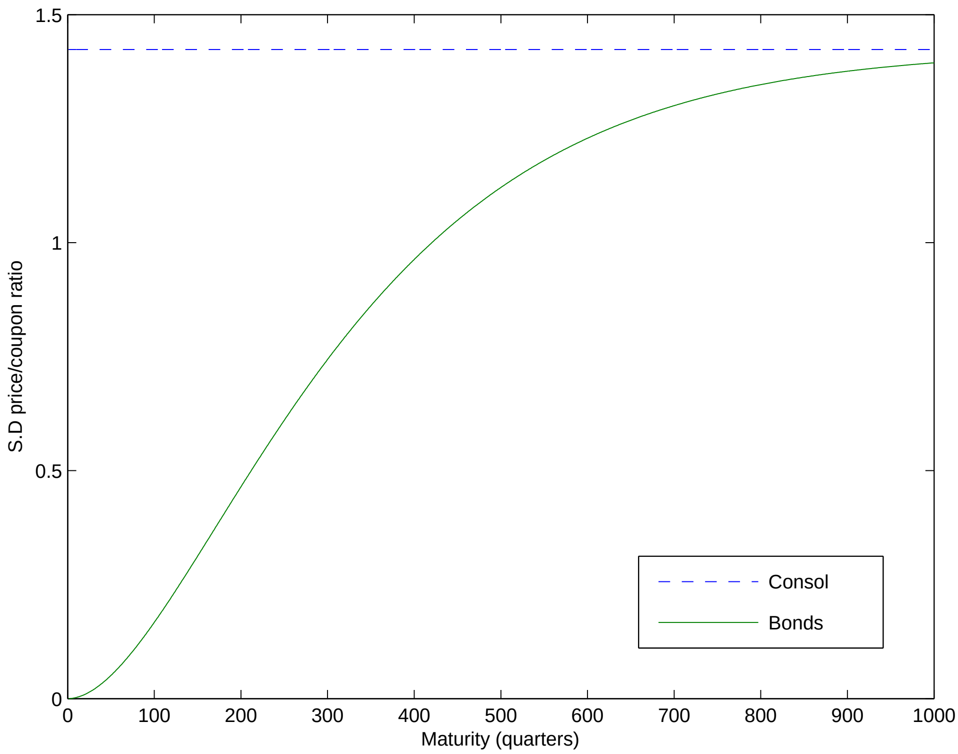

Figure 1: Case I - Std. Dev. of price/coupon ratio

Note: Figure reports the standard deviation of the price/coupon ratio of the simulated prices under learning. Simulations follow the calibration in Table 1. Sample size in each draw is 1000. We report averages of 20 draws. Simulated values never hit the *projection facility.*

an exogenous variable. Learning about a two-period bond involves learning about an exogenous variable and the law of motion of a one-period bond. As the maturity increases the learning exercise involves more endogenous variables, but this only gradually leads to higher volatility.

Finally, we use these prices to calculate the yield to maturity. We report the simulated statistics for a selection of maturities. The pattern generated is similar to the one found in the data: yields decrease slowly through the term structure. It must be noticed that here yields correspond to the ones of non-zero coupon bonds. Nonetheless, volatility arises from the behaviour of expected capital gains, while coupons are fixed. The volatility for the first maturity is zero since it is priced assuming rational expectations. Although our purpose is not to match the empirical data, the results suggests this learning mechanism can help the model capture the observed empirical dynamics.

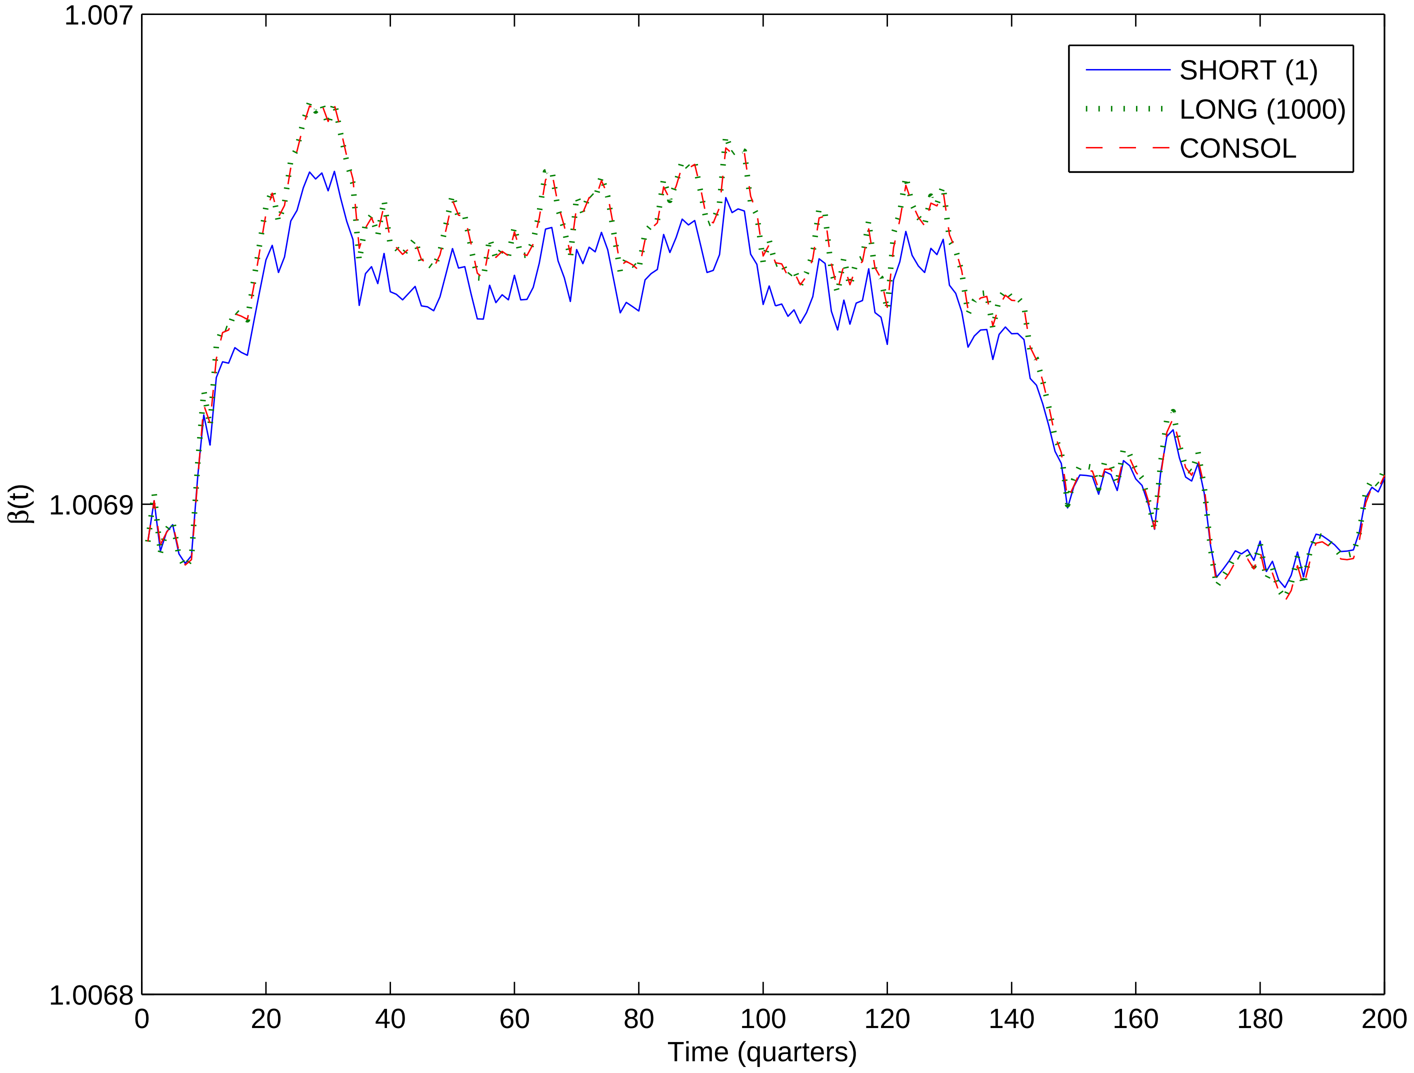

Figure 2: Case I -Evolution of beliefs

Note: Figure reports the beliefs of the price growth for each of the reported maturities under learning. Simulations follow the calibration in Table 1. Sample size of the draw is 1000. We report first 200 observations

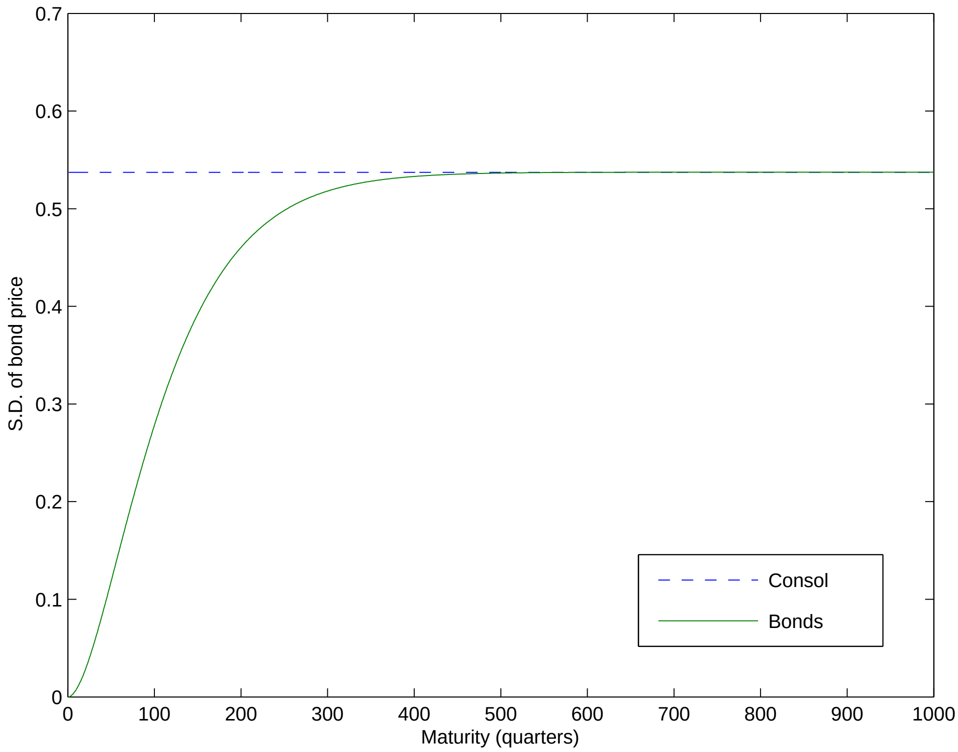

Figure 3: Case II - Standard deviation of bond prices under learning

Note: Figure reports the standard deviation of the simulated prices under learning and rational expectations. Simulations follow the calibration in Table 1. Sample size in each draw is 1000. We report first 200 observations.

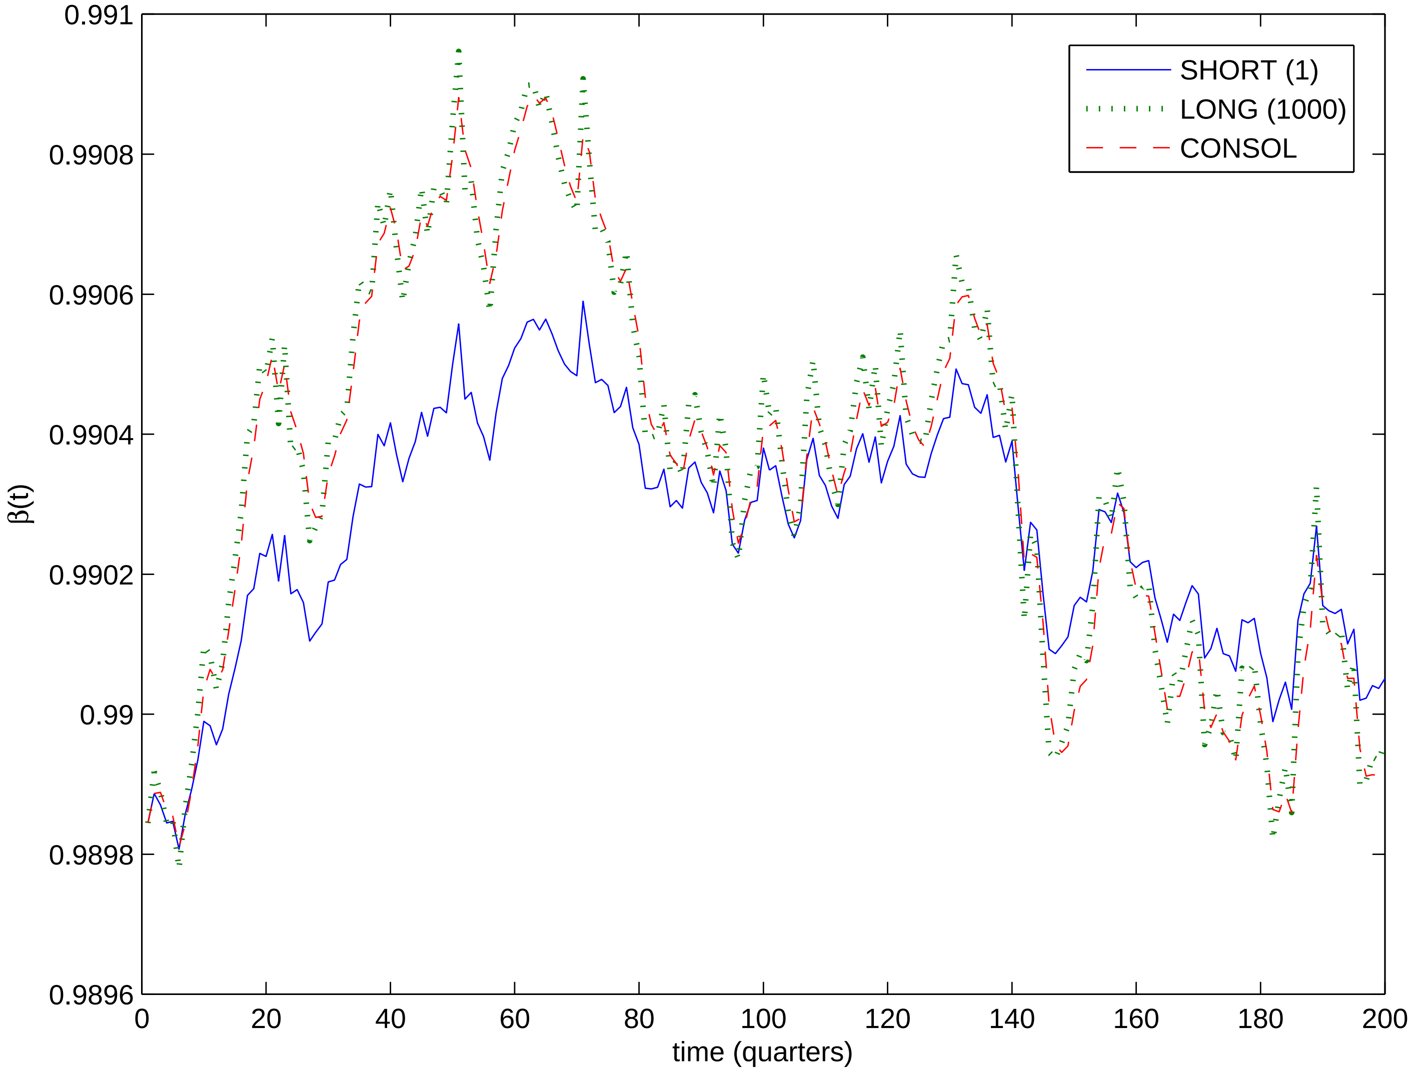

Figure 4: Case II - Evolution of beliefs

Note: Figure reports the beliefs of the price growth for each of the reported maturities under learning. Simulations follow the calibration in Table 1. Sample size of the draw is 1000. We report first 200 observations.

5 Conclusions

The present paper presents a model in which agents learn about the growth rate of bonds adaptively, following Adam et al. (2011). The proposed learning mechanism, when applied to finite-maturity assets, ceases to be *self-referential*, although, it generates interesting dynamics. First, changes in beliefs are carried over to higher maturities. This generates an amplification of shocks *through* the yield curve. Second, the impact of learning *through* the yield curve affects maturities in an heterogeneous way. Higher maturities exhibit a larger increase in their volatilities relative to the rational expectations results. Finally, the numerical results suggest that as maturity increases, the volatility of the asset converges to the one of the consol bond, which is subject to a *self-referential* learning mechanism as in Adam et al. (2011).

In addition, we present a model in which coupons are fixed. In this case the volatility will come from agents updating their beliefs about the real growth rate of bond prices. Even if under rational expectations the nominal price of these bonds would be fixed, the introduction of learning is capable of generating a slow decaying volatility pattern across maturities, similar to the one observed in the data. We confirm that the learning mechanism presented by Adam et al. (2011) for the case of equity, can help in the understanding of the behaviour of finite-maturity assets.

Even though these results are promising, we consider that the adaptive learning literature still has some open issues regarding its microfoundations. Although Adam and Marcet (2011) address this subject by proposing a model in which the PLM arises from a well-specified agent-based problem, their formulation still has problems in terms of determining the existence and uniqueness of equilibrium prices. This is one of the key topics that must be tackled by future research.

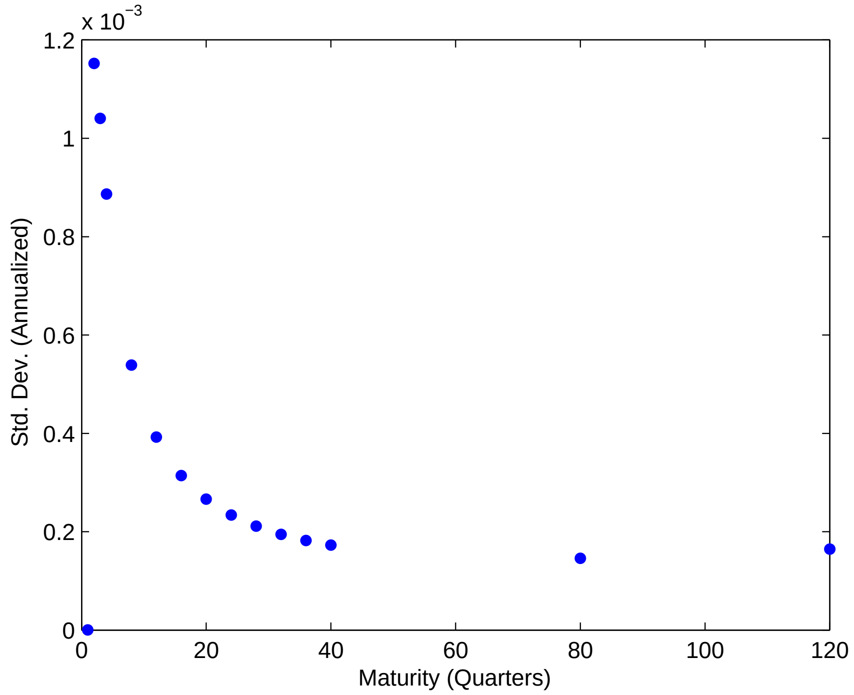

Figure 5: Case II - Yield volatility under learning

Note: Figure reports the average standard deviation of calculated yield to maturity for simulated coupon-paying bonds prices under learning. Simulations follow the calibration in Table 1. Sample size in each draw is 1000. We report averages of 20 draws.

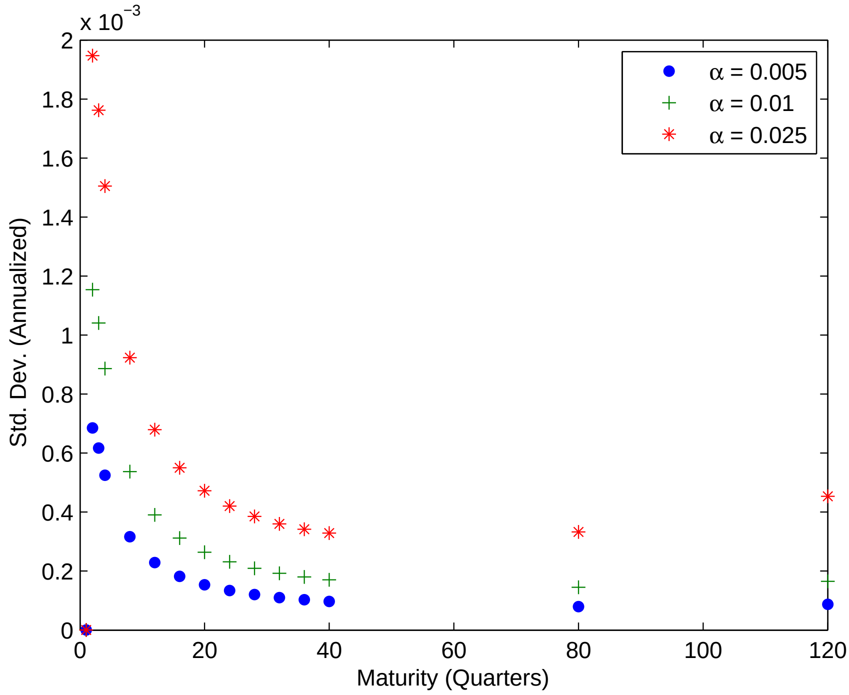

Figure 6: Case II: Sensitivity to gain parameter

Note: Figure reports the average standard deviation of calculated yield to maturity for simulated coupon-paying bonds prices under learning. Simulations follow the calibration in Table 1 except for the value of α, which stands for the constant-gain parameter. Sample size in each draw is 1000. We report averages of 20 draws.

References

Adam, K. and A. Marcet (2011, May). Internal rationality, imperfect market knowledge and asset prices. *Journal of Economic Theory 146*(3), 1224–1252.

Adam, K., A. Marcet, and J. P. Nicolini (2011, September). Stock market volatility and learning. CEP Discussion Papers dp1077, Centre for Economic Performance, LSE.

Branch, W. A. and G. W. Evans (2006, May). A simple recursive forecasting model. *Economics Letters 91*(2), 158–166.

Branch, W. A. and G. W. Evans (2011, July). Learning about risk and return: A simple model of bubbles and crashes. *American Economic Journal: Macroeconomics 3*(3), 159–91.

Carceles-Poveda, E. and C. Giannitsarou (2008, July). Asset pricing with adaptive learning. *Review of Economic Dynamics 11*(3), 629–651.

Carroll, C. D., J. Slacalek, and M. Sommer (2011, November). International evidence on sticky consumption growth. *The Review of Economics and Statistics 93*(4), 1135–1145.

Cochrane, J. H. and J. Campbell (1999, April). By force of habit: A consumption-based explanation of aggregate stock market behavior. *The Journal of Political Economy 107*(2), 205–51.

Evans, G. and S. Honkapohja (2001). *Learning and Expectations in Macroeconomics.* Frontiers of Economic Research. Princeton University Press.

Hall, R. E. (1978, December). Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. *Journal of Political Economy 86*(6), 971–87.

Sargent, T. J. (1993). *Bounded Rationality in Macroeconomics: The Arne Ryde Memorial Lectures.* Number 9780198288695 in OUP Catalogue. Oxford University Press.

Timmermann, A. (1996, October). Excess volatility and predictability of stock prices in autoregressive dividend models with learning. *Review of Economic Studies 63*(4), 523–57.

Williams, N. (2001). Escape dynamics in learning models. Technical report.

A. Rational expectations

Here we derive the rational expectations behaviour of assets for each case. For this we use equations 2 and 3 together with the assumptions for each case.

Case I: Stochastic coupon growth and learning about nominal bond price growth rates.

In this case, the euler equation for consol bonds (Eq. 3) becomes:

$$Q_t^{\{(\infty,I),RE\}} = \delta\mu_c E_t \left[\gamma_\phi \phi_t + Q_{t+1}^{\{(\infty,I),RE\}}\right] \tag{23}$$

solving forward, we obtain:[18]

$$Q_t^{\{(\infty,I),RE\}} = \frac{\delta\mu_c\gamma_\phi}{1-\delta\mu_c\gamma_\phi}\phi_t \tag{24}$$

The rational expectations growth rate is be given by:

$$\frac{Q_t^{\{(\infty,I),RE\}}}{Q_{t-1}^{\{(\infty,I),RE\}}} = \frac{\phi_t}{\phi_{t-1}} = \gamma_\phi \varepsilon_t \tag{25}$$

For finite-maturity assets:

$$Q_t^{\{(i,I),RE\}} = \left(\sum_{j=1}^{i}(\delta\mu_c\gamma_\phi)^j + (\delta\mu_c\gamma_\phi)^i \bar{Q}_0\right)\phi_t \tag{26}$$

The rational expectations growth rate is given by:

$$\frac{Q_t^{\{(i,I),RE\}}}{Q_{t-1}^{\{(i,I),RE\}}} = \frac{\phi_t}{\phi_{t-1}} = \gamma_\phi \varepsilon_t, \ \forall_{i\in[1,\mathcal{T}]} \tag{27}$$

Notice that under rational expectations the price of nominal bonds remains constant, thus, the covariance between bond nominal prices and the actual inflation rate will be zero, as the expected inflation is the same in every period.

[18] We assume parameter values guarantee that the price of the consol bond remains positive and finite.

Case II: Stochastic inflation and learning about bond price real growth rates
The euler equation for consol bonds (Eq. 3) becomes:

$$Q_t^{\{(\infty,II),RE\}} = \delta\mu_c\gamma_\pi\gamma_\phi + E_t\left[\bar{\phi} + Q_{t+1}^{\{(\infty,II),RE\}}\right] \tag{28}$$

solving forward we obtain:[19]

$$Q_t^{\{(\infty,II),RE\}} = \frac{\delta\mu_c\gamma_\pi\gamma_\phi}{1-\delta\mu_c\gamma_\pi\gamma_\phi}\bar{\phi} \tag{29}$$

The rational expectations real growth rate of bonds will be given by:

$$\frac{P_{t-1}}{P_t}\frac{Q_t^{\{(\infty,II),RE\}}}{Q_{t-1}^{\{(\infty,II),RE\}}} = \frac{P_{t-1}}{P_t} = \gamma_\pi\varepsilon_t^\pi \tag{30}$$

For finite-maturity assets:

$$Q_t^{\{(i,II),RE\}} = \left(\sum_{j=1}^{i}(\delta\mu_c\gamma_\pi\gamma_\phi)^j + (\delta\mu_c\gamma_\pi\gamma_\phi)^i\bar{Q}_0\right)\bar{\phi} \tag{31}$$

The rational expectations real growth rate is given by:

$$\frac{P_{t-1}}{P_t}\frac{Q_t^{\{(i,II),RE\}}}{Q_{t-1}^{\{(i,II),RE\}}} = \frac{P_{t-1}}{P_t} = \gamma_\pi\varepsilon_t^\pi,\ \forall_{i\in[1,\mathcal{T}]} \tag{32}$$

B. Properties of proposed learning mechanisms

Here we study the properties of the proposed learning mechanisms.
Case I: Stochastic coupon growth and learning about nominal bond price growth rates.

We derive the observed behaviour of finite-maturity bonds. Starting from the one-period to maturity asset in Eq. 16:

$$\frac{Q_t^{1,I}}{Q_{t-1}^{1,I}} = \gamma_\phi\varepsilon_t \tag{33}$$

now, for the two-period to maturity we have from Eq. 4 and our result for the one-period maturity in Eq. 15:

$$\frac{Q_t^{2,I}}{Q_{t-1}^{2,I}} = \frac{\gamma_\phi\phi_t + \hat{\beta}_t^1 Q_t^{1,I}}{\gamma_\phi\phi_{t-1} + \hat{\beta}_{t-1}^1 Q_{t-1}^{1,I}}$$

[19]As in the previous case, we assume parameter values are such that the price of the consol bond remains positive and finite.

Substituting the process for the coupon in Eq. 4 and our result for the one-period to maturity asset in Eq. 33:

$$\frac{Q_t^{2,I}}{Q_{t-1}^{2,I}} = \frac{\gamma_\phi(\gamma_\phi \phi_{t-1}\varepsilon_t) + \hat{\beta}_t^1(\gamma_\phi Q_{t-1}^{1,I}\varepsilon_t)}{\gamma_\phi \phi_{t-1} + \hat{\beta}_{t-1}^1 Q_{t-1}^{1,I}}$$

$$= \left(\gamma_\phi + \frac{\gamma_\phi \Delta\hat{\beta}_t^1}{\gamma_\phi \frac{\phi_{t-1}}{Q_{t-1}^{1,I}} + \hat{\beta}_{t-1}^1}\right)\varepsilon_t \tag{34}$$

which yields the actual law of motion for the growth rate of two-period to maturity asset. Now for the next maturity we have from Eq. 15:

$$\frac{Q_t^{3,I}}{Q_{t-1}^{3,I}} = \frac{\gamma_\phi \phi_t + \hat{\beta}_t^2 Q_t^{2,I}}{\gamma_\phi \phi_{t-1} + \hat{\beta}_{t-1}^2 Q_{t-1}^{2,I}}$$

Once again, from the assumed process for the coupon in Eq. 4 and our result for the two-period to maturity asset in Eq. 34:

$$\frac{Q_t^{3,I}}{Q_{t-1}^{3,I}} = \frac{\gamma_\phi(\gamma_\phi \phi_{t-1}\varepsilon_t) + \hat{\beta}_t^2\left[\left(\gamma_\phi + \frac{\gamma_\phi \Delta\hat{\beta}_t^1}{\gamma_\phi \frac{\phi_{t-1}}{Q_{t-1}^{1,I}} + \hat{\beta}_{t-1}^1}\right)\right] Q_{t-1}^{2,I}\varepsilon_t}{\gamma_\phi \phi_{t-1} + \hat{\beta}_{t-1}^2 Q_{t-1}^{2,I}}$$

$$= \left(\gamma_\phi + \frac{\gamma_\phi \Delta\hat{\beta}_t^2 + \gamma_\phi \hat{\beta}_t^2\left[\left(\frac{\Delta\hat{\beta}_t^1}{\gamma_\phi \frac{\phi_{t-1}}{Q_{t-1}^{1,I}} + \hat{\beta}_{t-1}^1}\right)\right]}{\gamma_\phi \frac{\phi_{t-1}}{Q_{t-1}^{2,I}} + \hat{\beta}_{t-1}^2}\right)\varepsilon_t \tag{35}$$

which is the actual law of motion for the growth rate of the three-period to maturity asset. Following this process we arrive to:

$$\frac{Q_t^{(i,I)}}{Q_{t-1}^{(i,I)}} = \left(\gamma_\phi + \frac{\gamma_\phi \Delta\hat{\beta}_t^{(i-1)} + \gamma_\phi \hat{\beta}_t^{(i-1)}\Omega_t^{(i-1)}}{\gamma_\phi \frac{\phi_{t-1}}{Q_{t-1}^{(i-1,I)}} + \hat{\beta}_{t-1}^{(i-1)}}\right)\varepsilon_t, \quad \forall_{i\in[2,\mathcal{T}]}, \tag{36}$$

where:

$$\Omega_t^{(i)} = \frac{\Delta\hat{\beta}_t^{(i-1)} + \hat{\beta}_t^{(i-1)}\Omega_t^{(i-1)}}{\gamma_\phi \frac{\phi_{t-1}}{Q_{t-1}^{(i-1,I)}} + \hat{\beta}_{t-1}^{(i-1)}}, \quad \forall_{i\in[2,\mathcal{T}]}. \tag{37}$$

which are the equations 17 and 18 in the main text. Additionally:

$$\Omega_t^{(1)} = 0, \tag{38}$$

As in the case of Adam et al. (2011) an equilibrium of the model corresponds to the case of rational expectations.

Case II: Stochastic inflation and learning about bond price real growth rates

We start describing the consol bond equations. Substituting assumptions in 7, 8 and 9 into Eq. 3 we obtain:

$$Q_t^{(\infty,II)} = \delta\gamma_\pi\mu_c\gamma_\phi\bar{\phi} + \delta\mu_c\hat{\beta}_t Q_t^{(\infty,II)} \tag{39}$$

without loss of generality, we assume a flat structure for coupons and normalize their value to 1. Thus, substituting $\gamma_\phi = 1$ and $\bar{\phi} = 1$ into Eq. 39 yields:

$$Q_t^{(\infty,II)} = \delta\gamma_\pi\mu_c + \delta\mu_c\hat{\beta}_t Q_t^{(\infty,II)} \tag{40}$$

solving for $Q_t^{(\infty,II)}$:

$$Q_t^{(\infty,II)} = \frac{\delta\gamma_\pi\mu_c}{1-\delta\mu_c\hat{\beta}_t} \tag{41}$$

thus the which is the *actual law of motion* for the consol bond price. We calculate the observed real growth rate:

$$\frac{P_{t-1}}{P_t}\frac{Q_t^{(\infty,II)}}{Q_{t-1}^{(\infty,II)}} = \left(1 + \frac{\delta\mu_c\Delta\beta_t}{1-\delta\mu_c\hat{\beta}_t}\right)\gamma_\pi\varepsilon_t^\pi = T(\beta,\Delta\beta)\varepsilon_t^\pi \tag{42}$$

where we have used the process for prices given by Eq. 8. The T-map shows properties similar to the ones in the previous case. A change in agents' beliefs about the real growth rate of prices will generate self-exciting dynamics.

Now for the case of finite maturity bonds, we use 2.

$$Q_t^{(i,II)} = \delta\gamma_\pi\mu_c + \delta\mu_c\hat{\beta}_t^{(i-1)} Q_t^{(i-1,II)}, \ \forall_{i\in[2,\mathcal{T}]} \tag{43}$$

which is the *actual law of motion* for the i-period to maturity bond. The behaviour of the actual real price growth is given by:

$$\frac{P_{t-1}}{P_t}\frac{Q_t^{(i,II)}}{Q_{t-1}^{(i,II)}} = \frac{\delta\gamma_\pi\mu_c + \delta\mu_c\hat{\beta}_t^{(i-1)}Q_t^{(i-1,II)}}{\delta\gamma_\pi\mu_c + \delta\mu_c\hat{\beta}_{t-1}^{(i-1)}Q_{t-1}^{(i-1,II)}}\gamma_\pi\varepsilon_t^\pi \tag{44}$$

Now we start from the first maturity:

$$\frac{P_{t-1}}{P_t}\frac{Q_t^{(1,II)}}{Q_{t-1}^{(1,II)}} = \frac{\delta\gamma_\pi\mu_c + \delta\mu_c\bar{Q}^{(0,II)}}{\delta\gamma_\pi\mu_c + \delta\mu_c\bar{Q}^{(0,II)}}\gamma_\pi\varepsilon_t^\pi = \gamma_\pi\varepsilon_t^\pi \tag{45}$$

The growth rate for the price of the second maturity will be given by:

$$\frac{P_{t-1}}{P_t}\frac{Q_t^{(2,II)}}{Q_{t-1}^{(2,II)}} = \frac{\delta\gamma_\pi\mu_c + \delta\mu_c\hat{\beta}_t^{(1)}Q_t^{(1,II)}}{\delta\gamma_\pi\mu_c + \delta\mu_c\hat{\beta}_{t-1}^{(1)}Q_{t-1}^{(1,II)}}\gamma_\pi\varepsilon_t^\pi \tag{46}$$

Replacing the result in 45:

$$\frac{P_{t-1}}{P_t}\frac{Q_t^{(2,II)}}{Q_{t-1}^{(2,II)}} = \frac{\delta\gamma_\pi\mu_c + \delta\mu_c\hat{\beta}_t^{(1)}\left(\frac{P_t}{P_{t-1}}Q_{t-1}^{(1,II)}\gamma_\pi\varepsilon_t^\pi\right)}{\delta\gamma_\pi\mu_c + \delta\mu_c\hat{\beta}_{t-1}^{(1)}Q_{t-1}^{(1,II)}}\gamma_\pi\varepsilon_t^\pi \tag{47}$$

$$= \frac{\delta\gamma_\pi\mu_c + \delta\mu_c\hat{\beta}_t^{(1)}Q_{t-1}^{(1,II)}}{\delta\gamma_\pi\mu_c + \delta\mu_c\hat{\beta}_{t-1}^{(1)}Q_{t-1}^{(1,II)}}\gamma_\pi\varepsilon_t^\pi \tag{48}$$

$$= \left(1 + \frac{\Delta\hat{\beta}_t^{(1)}}{\frac{\gamma_\pi}{Q_{t-1}^{(1,II)}} + \hat{\beta}_{t-1}^{(1)}}\right)\gamma_\pi\varepsilon_t^\pi \tag{49}$$

which is similar to the one found for the previous case. Further substitutions leads to:

$$\frac{P_t}{P_{t-1}}\frac{Q_t^{(i,II)}}{Q_{t-1}^{(i,II)}} = \left(1 + \frac{\Delta\hat{\beta}_t^{(i-1)} + \hat{\beta}_t^{(i-1)}\Omega_t^{(i-1,II)}}{\frac{\gamma_\pi}{Q_{t-1}^{(i-1,II)}} + \hat{\beta}_{t-1}^{(i-1)}}\right)\gamma_\pi\varepsilon_t^\pi, \quad \forall_{i\in[2,\mathcal{T}]}, \tag{50}$$

where:

$$\Omega_t^{(i,II)} = \frac{\Delta\hat{\beta}_t^{(i-1)} + \hat{\beta}_t^{(i-1)}\Omega_t^{(i-1,II)}}{\frac{\gamma_\pi}{Q_{t-1}^{(i-1,II)}} + \hat{\beta}_{t-1}^{(i-1)}}, \quad \forall_{i\in[2,\mathcal{T}]}. \tag{51}$$

Agents will carry over forecast errors from shorter maturities.