Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

Learning and forecasts about option returns through the volatility risk premium

Alejandro Bernales [a,*], Louisa Chen [b], Marcela Valenzuela [c]

[a] *Universidad de Chile (Centro de Finanzas y Centro de Economía Aplicada)*
[b] *School of Business, Management, and Economics, University of Sussex, United Kingdom*
[c] *Universidad de Chile (Centro de Finanzas and Centro de Economía Aplicada), and London School of Economics (the Financial Markets Groups)*

A R T I C L E   I N F O

A B S T R A C T

We use learning in an equilibrium model to explain the puzzling predictive power of the volatility risk premium (VRP) for option returns. In the model, a representative agent follows a rational Bayesian learning process in an economy under incomplete information with the objective of pricing options. We show that learning induces dynamic differences between probability measures $\mathbb{P}$ and $\mathbb{Q}$, which produces predictability patterns from the VRP for option returns. The forecasting features of the VRP for option returns, obtained through our model, exhibit the same behaviour as those observed in an empirical analysis with S&P 500 index options.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

Are option returns predictable? A recent empirical study by Goyal and Saretto (2009) sheds some light on the return predictability of option contracts. In particular, Goyal and Saretto (2009) show that the volatility risk premium (measured as the difference between implied and realized volatility) can predict option returns. They argue that large deviations of implied volatility (henceforth, IV) from realized volatility (henceforth, RV) are signals of option mispricing, which can be used to predict the returns on option investment strategies. We offer a theoretical explanation for this empirical relationship between the volatility risk premium (henceforth, VRP or (IV-RV)) and option returns. We use learning to explain both why IV deviates from RV and how this deviation generates predictive dynamics in the returns of option portfolios.

We extend the learning model proposed by Timmermann (2001), including option contracts to analyse the effects of learning on the relationship between the VRP and option returns. The model describes a discrete-time endowment economy

---

* Corresponding author.

E-mail addresses: abernale@gmail.com, abernales@dii.uchile.cl (A. Bernales), L.X.Chen@sussex.ac.uk (L. Chen), mvalenzuela@dii.uchile.cl (M. Valenzuela).

with a Bayesian representative agent. There are three assets in the economy: a bond, a stock and a group of European put and call options. In the model, the mean dividend growth rate, g_t, exhibits breaks. If a break takes place at time $t + m$, the new value of g_{t+m} is drawn from a univariate distribution $g_{t+m} \sim G(\cdot)$, and g_{t+m} does not change until the next break. Time periods between breaks are described by a memoryless distribution. Moreover, the representative agent faces an environment under incomplete information because g_t is unknown. Nevertheless, the agent can learn recursively about the new value of g_t after each break, by following a rational Bayesian updating process as new information arrives.

In the model, the Bayesian agent learns about g_t from daily dividends received over time. Hence, dividends are the signals used to obtain a parameter estimate of the mean dividend growth rate, $\hat{g}_t$. The Bayesian agent starts the learning process after a given break with a prior density $p^{\mathbb{P}}(g_t)$ regarding the probability distribution of g_t in the physical world $\mathbb{P}$ (which is equal to the density of $G(\cdot)$). Subsequently, prior beliefs are recursively updated to obtain a posterior density, $p^{\mathbb{P}}(\hat{g}_t|\Omega_t)$, using Bayes' rule given the information set Ω_t of signals received since the most recent break.

In early periods after a break, when no 'long' history of dividend realizations is available, there are large revisions in the value of $\hat{g}_t$, since the estimation accuracy is low due to insufficient information (i.e. the number of signals is not large enough for an accurate estimation of the unknown value of g_t). At time t, $\hat{g}_t$ is random and described by the posterior probability, $p^{\mathbb{P}}(\hat{g}_t|\Omega_t)$, with mean $\hat{g}_t^*$ and variance $\sigma_{\hat{g},t}^2$ (i.e. $\hat{g}_t^*$ and $\sigma_{\hat{g},t}^2$ are dynamically updated as new information is received). Thus, the value of $\hat{g}_t^*$ reflects the expected value of the parameter estimate, and $\sigma_{\hat{g},t}^2$ represents the level of inaccuracy of the parameter estimation.

In terms of option returns, the expected holding-period return of a European put option contract is defined by $R_{t+\tau}^{Put} = \frac{E_t^{\mathbb{P}}[\max(K-S_{t+\tau}, \, 0)]}{Put_t(K,\tau)} - 1 = \frac{E_t^{\mathbb{P}}[\max(K-S_{t+\tau},0)]}{E_t^{\mathbb{Q}}[\max(K-S_{t+\tau},0)]} - 1$, where S_t is the price of the underlying asset, K is the strike price and τ is the time-to-maturity (see Broadie et al., 2009). A similar expression can be written for the holding-period return on a European call option contract, but with an option payoff given by $\max(S_{t+\tau} - K, 0)$. Therefore, the magnitude of the expected holding-period return of an option is affected by differences between the physical, $\mathbb{P}$, and risk-neutral, $\mathbb{Q}$, probability measures. Under full information, the relationship between the physical probability measure, $p^{\mathbb{P}}(S_{t+1})$, and the risk-neutral probability measure, $p^{\mathbb{Q}}(S_{t+1})$, is given by $p^{\mathbb{Q}}(S_{t+1}) = m_{t+1} \cdot p^{\mathbb{P}}(S_{t+1})/E_t(m_{t+1})$, where m_{t+1} is the stochastic discount factor. However, under incomplete information and learning, the $\mathbb{P}$ and $\mathbb{Q}$ probability measures incorporate changes in the agent's perceptions about how the economy evolves. We show that, under Bayesian learning, the physical probability measure is conditional on the available information and is given by $p^{\mathbb{P},BL}(S_{t+1}|\Omega_t) = \int p^{\mathbb{P}}(S_{t+1}|\Omega_t, \hat{g}_t)p^{\mathbb{P}}(\hat{g}_t|\Omega_t)d\hat{g}_t$, while the risk-neutral probability measure is $p^{\mathbb{Q},BL}(S_{t+1}|\Omega_t) = \int m_{t+1}p^{\mathbb{P}}(S_{t+1}|\Omega_t, \hat{g}_t)p^{\mathbb{P}}(\hat{g}_t|\Omega_t)d\hat{g}_t/E_t^{BL}(m_{t+1})$. Therefore, the posterior density, $p(\hat{g}_t|\Omega_t)$, affects the agent's beliefs about the probabilities of future states of the economy. In the case of a change in the value of $\hat{g}_t$ (due to the agent's learning process) being statistically related to the agent's beliefs regarding future levels of the stochastic discount factor, learning can induce additional differences between densities $\mathbb{P}$ and $\mathbb{Q}$. In fact, this is what happens, since a change in the value of $\hat{g}_t$ modifies perceptions about the evolution of the stochastic discount factor and its expected value.

Moreover, Bayesian learning induces predictability patterns since future parameter estimates are progressively updated in a recursive updating procedure. For instance, when the representative agent observes a signal that the mean dividend growth rate is higher than $\hat{g}_t^*$, the agent updates her expectations upwards. In the scenario in which future signals are largely higher than the value of $\hat{g}_t^*$ (i.e. at time t the agent is wrongly very pessimistic in her beliefs regarding the mean dividend growth rate), the future values of $\hat{g}_t^*$ will gradually be revised upwards over time; thus, predictable dynamics emerge. Conversely, the same effect can be observed in the case of an optimistic scenario. The agent can be incorrectly pessimistic or optimistic immediately after breaks, as the number of signals (dividends) about the new unknown parameter is low, and the agent's parameter estimates are mainly based on her prior beliefs. Thus, learning can induce predictability patterns in all assets, including the returns on option contracts.[1]

Given that the agent's learning process simultaneously generates predictability patterns in option returns and dynamic gaps between probability measures $\mathbb{P}$ and $\mathbb{Q}$, the *VRP* (which reflects the difference between the volatilities under $\mathbb{P}$ and $\mathbb{Q}$) is a natural candidate to be a predictor of option returns.[2] However, the analysis of the impact of learning on the return predictability of options is not straightforward given the non-linearity of option contracts. It is not possible to obtain a closed-form solution in our model for the relationship between the *VRP* and option returns under a Bayesian learning process. Thus, we use a numerical approach based on a simulation analysis which is driven by our model.

We compute the returns of option portfolios that control for the effect of changes in the underlying stock price, such as delta-hedged and straddle option portfolios. The use of these option portfolios is important because they are not affected by the predictability patterns of the underlying asset, which can also be endogenously induced by the agent's learning process in the model. Thus, the use of these option portfolios allows us to isolate the impact of learning on option prices and their returns. In particular, we calculate delta-hedged portfolios and straddle portfolios which are at-the-money (henceforth, *ATM*) and have a time-to-maturity of one month, similarly to Goyal and Saretto (2009).[3]

---

[1] The same argument is used by Timmermann (1993, 1996) and Guidolin (2006b) to explain return predictability patterns in stock markets.

[2] Hence, learning may provide an explanation for the puzzling empirical regularity that implied volatilities are different from realized volatilities. Several recent studies have proposed explanations for the existence of the *VRP* within the context of equilibrium-based pricing models. See, e.g., Bollerslev *et al.* (2009), Carr and Wu (2009), Drechsler and Yaron (2011). However, these studies have not explored why the *VRP* can predict option returns, which is the main objective of our research.

[3] The put (call) delta-hedged portfolio is formed by buying one put (call) option contract, and buying (short-selling) delta shares in the underlying stock. The straddle portfolio is formed by buying one call option contract and one put option contract.

The simulation analysis effectively shows that the agent's learning process induces a dynamic relationship between the *VRP* and option returns. We show that the *VRP* consistently predicts a non-trivial fraction of the returns of option portfolios, as in Goyal and Saretto (2009). The results are robust to the inclusion of alternative predictor variables. We also compare the results generated by our learning model to the returns observed on S&P 500 index option contracts. We find that our model generates the same features of the *VRP*'s ability to forecast option returns as observed in actual S&P 500 option data.

To the best of our knowledge, there has been no theoretical study to date that explains, through a dynamic equilibrium model, the predictive power of the *VRP* for option returns. The study is organized as follows. Section 2 presents a literature review. Section 3 introduces the model. Section 4 describes the model implementation. Section 5 reports the main results of our model simulations, which are compared to S&P 500 option data. Section 6 concludes.

2. Literature review

Our study is connected to the literature on theoretical models, in which a learning process affects option pricing (e.g., David and Veronesi, 2002; Guidolin and Timmermann, 2003; and Shaliastovich, 2009). It should be noted that these studies use learning mainly to explain the 'existence' of the *IV* surface; as such, they do not explore the predictability patterns and features of option returns as we do in our study. For example, David and Veronesi (2002) introduce a model based on a regime-switching process, in which agents learn because they do not observe the drift of the dividend stochastic process. Guidolin and Timmermann (2003) develop a two-state model in which the dividend growth evolves on a binomial lattice with an unknown state probability which is learned.[4] Shaliastovich (2009) presents a model in which there is a learning process based on a recency-biased updating procedure regarding the unknown consumption growth rate.

Our paper also relates to theoretical studies in which learning is used to explain diverse empirical anomalies in stock returns (see, e.g., Veronesi, 1999, 2000; Timmermann, 2001; Brandt et al., 2004; Massa and Simonov, 2005; Guidolin, 2006a, b; Guidolin and Timmermann, 2007; Branch and Evans, 2010; and Branch, 2016). In particular, our paper is connected to studies in which structural breaks are modelled under a learning environment to explain the anomalous behaviour of stock prices (see Timmermann, 2001; and Guidolin, 2006a, b). For instance, Guidolin (2006a, b) presents an endowment economy based on Lucas (1978), in which the dividend process is characterized by a binomial lattice in which the unknown state probability is subject to structural breaks. He shows that rational learning induces high equity premia, low risk-free interest rates, volatility clustering and long-run predictability.

Our paper is particularly related to Timmermann (2001), who modifies the model presented in Lucas (1978) by assuming that g_t exhibits breaks and is unknown. He reports that learning induces skewness, kurtosis, volatility clustering, and serial correlation in stock returns. We depart from Timmermann (2001) in several ways. Firstly, we consider a more developed model in which the representative agent has to price not only a bond and a stock but also option contracts. Secondly, the focus of our research differs from that of Timmermann (2001), in that we examine the effects of learning on the predictive power of the *VRP* for option returns, rather than the impact of learning on features of stock returns. Moreover, we perform an empirical analysis to contrast our theoretical analysis with the empirical evidence, by analysing the *VRP* and option returns from S&P 500 index option contracts.[5]

As explained in the introduction, our study is associated in particular with the empirical literature on option return predictability. Goyal and Saretto (2009) show that the *VRP* can predict option returns, since it is a signal of option mispricing; this is the empirical study that inspired our theoretical learning model. Cao and Han (2013) also show that the *VRP* is related to delta-hedged option returns. In addition, Cao and Han (2013) report that the returns on delta-hedged option portfolios decrease monotonically as the idiosyncratic volatility of the underlying stock increases. Furthermore, our paper is connected to studies in which the predictive power of the *VRP* is used to forecast returns on stocks instead of option returns (e.g., Bollerslev et al., 2009; and Drechsler and Yaron, 2011).

3. The model

As mentioned in the introduction, there is empirical evidence that the *VRP* can be used to predict option returns. In order to examine this link between the *VRP* and option returns, in this section we use an equilibrium model in which a representative agent follows a learning process. In particular, we extend the learning model of Timmermann (2001), introducing option contracts to analyse the effects of learning on the predictive power of the *VRP* for the returns of option contracts.

Similarly to Timmermann (2001), our model is based on Lucas (1978), and we assume in it that the mean dividend growth rate, g_t, has breaks. As a first step, we assume asset prices to be characterized by full information (i.e. the agent knows the value of g_t over time). Subsequently, we consider an economy under incomplete information, where g_t is unknown. However, the agent follows a learning process regarding g_t as new information is received, which allows her to calculate asset prices, including the prices of European option contracts and their returns.

---

[4] In contrast to the model in Guidolin and Timmermann (2003), in our model the effects of learning do not asymptotically vanish over time. This is due to the 'existence' of breaks, which means that the agent is repeatedly learning a new value of g_t after each break in the economy.

[5] Our paper is also related to studies that examine and explain why the *VRP* increases with uncertainty, where this uncertainty is generated through the learning processes associated with regimes of economic and policy variables (e.g. David and Veronesi, 2014), or through investor disagreement (e.g. Buraschi et al., 2014).

3.1. Asset prices under full information

There is a discrete-time endowment economy with a representative agent; there is a stock, S_t, and a bond, B_t, with one period to maturity and without coupons. The stock has a net supply normalized at one, while the bond is default free and has a zero net supply. The owner of the stock receives real dividends, D_t, per period. Dividends follow a geometric random walk, $\ln(D_t/D_{t-1}) = \mu_t + \sigma\varepsilon_t$, with volatility σ and drift μ_t. However, the fundamental mean dividend growth rate g_t exhibits breaks. Thus, breaks are also observable in μ_t, since we know that $g_t = \exp(\mu_t + \sigma^2/2) - 1$. Periods between breaks are described through a geometric distribution, with parameter π reflecting the probability of a break on a given date.[6] We assume that, when a break happens at time $t + m$, the new value of g_{t+m} is drawn from a continuous univariate distribution $g_{t+m} \sim G(\cdot)$, defined by the support $[g_d, g_u]$, and g_{t+m} does not change until there is a new break in the economy.

There is also a group of European put and call option contracts, $Put_t(K, \tau)$ and $Call_t(K, \tau)$, respectively. The option contracts have the stock as underlying asset, strike price K, and time-to-maturity τ. We assume a complete and perfect capital market with no friction in terms of trading possibilities. There is a power utility with a level of relative risk aversion α, which describes the agent's utility in relation to her real consumption C_t:

$$u(C_t) = \begin{cases} \frac{C_t^{1-\alpha}-1}{1-\alpha} & \alpha \geq 0 \\ \ln C_t & \alpha = 1. \end{cases} \tag{1}$$

We assume that dividends are the only income source, are non-storable and are consumed when they arrive (i.e. $C_t = D_t$). Therefore, the representative agent chooses her asset holdings by maximizing the discounted value of her expected future utility, $\max\limits_{\{w_{t+k}^S, w_{t+k}^B\}} E_t[\sum_{k=0}^{\infty} \beta^k u(D_{t+k})]$, in which $\beta \equiv 1/(1+\rho)$, where ρ is the impatience rate, while w_{t+k}^S and w_{t+k}^B are respectively the quantities of assets S_t and B_t in her portfolio. This yields the following Euler equations for the stock and the bond: $S_t = E_t[m_{t+1}(S_{t+1} + D_{t+1})]$ and $B_t = E_t[m_{t+1}]$, where $m_{t+1} = \beta(D_{t+1}/D_t)^{-\alpha}$ is the stochastic discount factor. Proposition 1 presents price expressions for the bond and the stock, which are obtained by solving the Euler equations under full information.

Proposition 1. *(Stock and bond prices under full information): The rational expected prices of assets S_t and B_t under full information are given by:*

$$S_t^{FI} = \frac{D_t}{1 + \rho - (1-\pi)(1+g_t)^{1-\alpha}} \left\{ (1-\pi)(1+g_t)^{1-\alpha} + \pi\left(\frac{I_1 + (1-\pi)I_2}{1 - \pi I_3}\right) \right\} = D_t\Psi(g_t), \tag{2}$$

and

$$B_t^{FI} = \frac{1}{(1+\rho)} \left\{ (1-\pi)(1+g_t)^{-\alpha} + \pi \int_{g_d}^{g_u} (1+g_{t+1})^{-\alpha}dG(g_{t+1}) \right\}, \tag{3}$$

where

$$I_1 = \int_{g_d}^{g_u} (1+g_{t+1})^{1-\alpha}dG(g_{t+1}); \quad I_2 = \int_{g_d}^{g_u} \frac{(1+g_{t+1})^{2-2\alpha}}{1 + \rho - (1-\pi)(1+g_{t+1})^{1-\alpha}}dG(g_{t+1});$$

$$I_3 = \int_{g_d}^{g_u} \frac{(1+g_{t+1})^{1-\alpha}}{1 + \rho - (1-\pi)(1+g_{t+1})^{1-\alpha}}dG(g_{t+1})$$

with $1 + \rho > (1+g_u)^{1-\alpha}$.

Proof: See Timmermann (2001).

In the absence of breaks, option prices under full information can be calculated using the discretized version of the Black-Scholes formula, as in Rubinstein (1976).[7] However, the Black-Scholes formula cannot be used in our model, since breaks cause the processes followed by the underlying asset, the interest rate and the dividend yield to be non-stationary. Nevertheless, a no-arbitrage option price can be obtained by means of a change in probability measure, which is presented in Proposition 2.[8]

Proposition 2. *(Option prices under full information): The full-information rational expected prices of a European put option contract, $Put_t^{FI}(K, \tau)$, and a European call option contract, $Call_t^{FI}(K, \tau)$, are given by:*

$$Put_t^{FI}(K, \tau) = \int_0^{\infty} \frac{1}{(1+r_{t+\tau})} \max\left\{K - S_{t+\tau}^{FI}, 0\right\} p^{\mathbb{Q}}\left(S_{t+\tau}^{FI}\right)dS_{t+\tau}^{FI} \tag{4}$$

---

[6] Since the geometric distribution is memoryless, the agent cannot predict when the next break will occur.

[7] Since there are no breaks in Guidolin and Timmermann (2003), they also show that under full information their model converges asymptotically towards the Black-Scholes formula.

[8] It is important to note that the representative agent uses the aggregate endowment process to price option contracts. Under full information, the representative agent knows the current fundamental mean dividend growth rate in each period, which will affect the stock and bond prices. Therefore, the agent can replicate the option payoffs period-by-period, which makes options replicable claims. Hence, all potential pricing measures yield the same, unique, no-arbitrage price for option contracts.

and

$$Call_t^{FI}(K, \tau) = \int_0^\infty \frac{1}{(1 + r_{t+\tau})} \max\left\{S_{t+\tau}^{FI} - K, 0\right\} p^{\mathbb{Q}}\left(S_{t+\tau}^{FI}\right) dS_{t+\tau}^{FI} \tag{5}$$

with risk-neutral price density, $p^{\mathbb{Q}}(S_{t+\tau})$, given by:

$$p^{\mathbb{Q}}(S_{t+\tau}) = (1 + r_{t+\tau})\beta^\tau \left(\frac{D_{t+\tau}}{D_t}\right)^{-\alpha} \phi(\varepsilon_{t+\tau}|0, \sigma)\varphi(z|\tau, \pi)\eta(h_1|\pi)$$
$$\cdot \left[\eta(h_2|\pi)\varrho\left(g_{t+h_2}\right)\dots\eta(h_z|\pi)\varrho\left(g_{t+h_z}\right)\right] \quad .$$

Here, $S_{t+\tau}^{FI} = D_{t+\tau}\Psi(g_{t+\tau})$ is defined in Proposition 1, $D_{t+\tau} = D_t \exp(\sqrt{\tau}\sigma\varepsilon_{t+\tau} - \tau\sigma^2/2)\prod_{i=1}^z (1 + r_{t+h_i})^{h_i}$, where z represents the number of breaks until maturity, which is described by a binomial distribution $\varphi(z|\tau, \pi)$ with parameters τ and π, while $\{h_i\}_{i=1}^z$ is the time between breaks, which is described by a geometric distribution $\eta(h_i|\pi)$. Hence, the time-to-maturity of an option contract is $\tau = \sum_{i=1}^z h_i$. In addition, $\{g_{t+h_i}\}_{i=2}^z$ is the value of the mean dividend growth rate, which is described by a univariate distribution $g_{t+h_i} \sim G(\cdot)$ with pdf $\varrho(\cdot)$ defined on the support $[g_d, g_u]$, where $g_{t+h_1} = g_t$ and $g_{t+h_z} = g_{t+\tau}$; meanwhile, $\varepsilon_{t+\tau}$ is described by a normal distribution, having density $\phi(\varepsilon_{t+\tau}|0, \sigma)$ with mean zero and volatility σ. Moreover, $1 + r_{t+\tau} = \prod_{j=1}^\tau (1 + r_{j-1,j})$, with $1 + r_{j-1,j} = 1/B_{j-1}$ and B_{j-1} is the price of the risk-free one-period bond in period $j - 1$, which was defined in Proposition 1.

Proof. *See* Appendix A.

Proposition 2 presents expressions for option prices that are obtained from the expected payoffs of the options under the risk-neutral probability $p^{\mathbb{Q}}(S_{t+\tau})$. In addition, the stock price and the bond price from Proposition 1 are used in both the option payoffs and the risk-neutral price density in Proposition 2.

3.2. Asset prices under incomplete information and learning

Propositions 1 and 2 are obtained assuming full information. However, we modify this full-information environment to include learning in the model. Suppose that the value of g_t is not observable by the representative agent, but the agent receives signals concerning this unknown parameter, through the dividends paid on the stock. Thus, the n historical dividend returns since the most recent break, $\{D_i/D_{i-1}\}_{i=t-n}^t$, represent the information received and used to learn about g_t.

We assume that the representative agent can detect the time at which there is a break in the dividend growth rate, as in Timmermann (2001). This allows us to study how learning affects option prices and their returns in a clean setup, in which only one parameter is unknown. The assumption that the agent can identify the dates of breaks is supported by the fact that there are real-time tests that can be used to contemporaneously detect breaks with a reasonable degree of precision (see, e.g., Chu et al., 1996; and Leisch et al., 2000).

In what follows in this paper, for mathematical simplicity we will describe the learning process as a function of the dividend drift, μ_t, which is also unknown (given the relationship $1 + g_t = \exp(\mu_t + \sigma^2/2)$).

3.2.1. Learning schemes

There are various learning schemes that have been used in the economic literature (see Guidolin and Timmermann, 2007). In this section, we briefly review three well-known learning schemes: rational learning, adaptive learning and Bayesian learning. In particular, we describe how these learning setups can be implemented in our model, and we analyse their main features (i.e. properties and relationships).

Rational learning

Suppose that the representative agent follows a rational learning (henceforth *RL*) process to learn about μ_t. As in the other learning schemes, which will be discussed in this section, the agent updates her parameter estimate, $\hat{\mu}_t$, over time as new signals are received. However, *RL* has three characteristics, which together distinguish this form of learning from others.

Firstly, a rational learner does not consider the unknown parameter, μ_t, as a random variable. Secondly, *RL* is forward looking in the sense that the agent takes into account future updates in parameter estimates when she prices any asset. Thus, asset prices take into account the current and possible future values of parameter estimates (i.e. considering all potential future probability distributions of parameter estimates). Therefore, a future parameter estimate, $\hat{\mu}_{t+h}$, where h is a positive integer, is a random variable under *RL*. Thirdly, *RL* implies that the agent consistently updates her beliefs regarding the unknown parameter, using all available information. This means that Bayes' rule is recursively applied in current and future parameter estimates using the correctly specified likelihood function.[9] For instance, in our model under incomplete

---

[9] This represents a slight difference between *RL* and Bayesian learning (*BL*), which will be explained in the following paragraphs. Under *BL* (*RL*), μ_t is (is not) considered random. However, a rational learner uses Bayes' rule to update her beliefs, despite the fact that she does not view μ_t as a random variable (see Guidolin and Timmermann, 2007).

information and learning, the representative agent's forecast of the dividend to be observed in a future period τ is given by:

$$D_{t+\tau}^{RL} = \hat{E}_t\left[\ldots\hat{E}_{t+\tau-2}\left[\hat{E}_{t+\tau-1}\left[D_{t+\tau}|\hat{\mu}_{t+\tau-1},\boldsymbol{\xi}_{t+\tau-1}\right]|\hat{\mu}_{t+\tau-2},\boldsymbol{\xi}_{t+\tau-2}\right]\ldots|\hat{\mu}_t,\boldsymbol{\xi}_t\right] \tag{6}$$

where $\boldsymbol{\xi}_t = [\xi_{t-n+1},\ldots,\xi_t]$ is the vector of historical signals received between the time of the last break and time t, where $\xi_i = \ln(D_i/D_{i-1})$, while $\hat{E}_{t+h}[\cdot|\hat{\mu}_{t+h},\boldsymbol{\xi}_{t+h}]$ is the expectation operator at time $t+h$ conditional on the future estimate $\hat{\mu}_{t+h}$ given the information set $\boldsymbol{\xi}_{t+h}$, with $h \in \{0,1,\ldots,\tau-1\}$. Consequently, under *RL* the complete set of future parameter estimates $\{\hat{\mu}_t,\hat{\mu}_{t+1},\ldots,\hat{\mu}_{t+\tau-1}\}$ is considered by the agent when she is pricing assets (i.e. the agent considers the effect of learning on the sequence of equilibrium outcomes of state variables and asset prices).

Adaptive learning

Suppose that the representative agent follows an adaptive learning (henceforth, *AL*) process to learn about the unknown parameter μ_t. Also called least-squares learning, *AL* has similarities and differences in relation to *RL*. On the one hand, similarly to *RL*, the unknown parameter, μ_t, is not viewed as random by the agent under *AL*. In addition, an adaptive learner updates the parameter estimate, $\hat{\mu}_t$, over time as new information arrives; hence, $\hat{\mu}_t$ is a random variable.

On the other hand, unlike a rational learner, an adaptive learner ignores potential future changes in parameter estimates (i.e. $\mu_t = \mu_t^* = \mu_{t+h}^*$). Moreover, the agent does not use Bayes' rule to update μ_t^*. Instead, the agent is considered a frequentist and uses the maximum likelihood rule to update the parameter estimate recursively. Thus, the expected value of the parameter estimate is:

$$\mu_t^* = \bar{\xi}_t \tag{7}$$

with $\bar{\xi}_t = (1/n)\sum_{i=t-n+1}^{t}\xi_i$, where ξ_i is as defined in Eq. (6). Therefore, the mean dividend growth rate under *AL* is $g_t^{AL} = \exp(\bar{\xi}_t + \sigma^2/2) - 1$, while the stock and bond prices are:

$$S_t^{AL} = \frac{D_t}{1 + \rho - (1-\pi)\left(1+g_t^{AL}\right)^{1-\alpha}}\left\{(1-\pi)\left(1+g_t^{AL}\right)^{1-\alpha} + \pi\left(\frac{I_1 + (1-\pi)I_2}{1-\pi I_3}\right)\right\}, \tag{8}$$

$$B_t^{AL} = \frac{1}{(1+\rho)}\left\{(1-\pi)\left(1+g_t^{AL}\right)^{-\alpha} + \pi\int_{g_d}^{g_u}\left(1+g_t^{AL}\right)^{-\alpha}dG\left(g_t^{AL}\right)\right\}, \tag{9}$$

where I_1, I_2 and I_3 are the same integrals as defined in Proposition 1.[10] In the case of option prices, the expressions are:

$$Put_t^{AL}(K,\tau) = \int_0^\infty \frac{1}{(1+r_{t+\tau})}\max\left\{K - S_{t+\tau}^{AL},0\right\}p^{\mathbb{Q}}\left(S_{t+\tau}^{AL}\right)dS_{t+\tau}^{AL} \tag{10}$$

and

$$Call_t^{AL}(K,\tau) = \int_0^\infty \frac{1}{(1+r_{t+\tau})}\max\left\{S_{t+\tau}^{AL} - K,0\right\}p^{\mathbb{Q}}\left(S_{t+\tau}^{AL}\right)dS_{t+\tau}^{AL} \tag{11}$$

where $p^{\mathbb{Q}}(\cdot)$ is the risk-neutral price density described in Proposition 2, in which the mean dividend growth rate is equal to g_t^{AL}.

The main issue with a model based on an *AL* process is that asset prices are not rational. This is due to the fact that an adaptive learner does not take into account future updates in parameter estimates. Thus, our model under *AL* is misspecified since rational future optimal decisions have to be based on future parameter estimates.

Bayesian learning

Suppose that the representative agent follows a Bayesian learning (henceforth, *BL*) process to learn about μ_t. Under *BL*, similarly to *RL* and *AL*, the learner updates the parameter estimate, $\hat{\mu}_t$, as new signals arrive. However, unlike rational and adaptive learners, a Bayesian learner views the unknown parameter, μ_t, as a random variable.

Under *BL*, the agent starts the learning process with a prior density, $p^{\mathbb{P}}(\mu_t)$, regarding the probability distribution of μ_t in the physical world, i.e. under $\mathbb{P}$. Afterwards, prior beliefs are recursively updated to obtain posterior beliefs using Bayes' rule, $p^{\mathbb{P}}(\mu_t|\boldsymbol{\xi}_t)$, given the information set $\boldsymbol{\xi}_t$ of signals received since the last break. Thus, $p^{\mathbb{P}}(\mu_t|\boldsymbol{\xi}_t)$ is:

$$p^{\mathbb{P}}\left(\mu_t|\boldsymbol{\xi}_t\right) = \frac{p^{\mathbb{P}}\left(\boldsymbol{\xi}_t|\mu_t\right)p^{\mathbb{P}}(\mu_t)}{p^{\mathbb{P}}\left(\boldsymbol{\xi}_t\right)}. \tag{12}$$

In addition, we know that:

$$p^{\mathbb{P}}\left(\boldsymbol{\xi}_t\right) = \int_{\mu_d}^{\mu_u}p^{\mathbb{P}}(\boldsymbol{\xi}_t|\mu_t)p^{\mathbb{P}}(\mu_t)d\mu_t, \tag{13}$$

---

[10] Under *AL*, we have to assume that $1 + \rho > (1-\pi)(1+g_t^{AL})^{1-\alpha}$ to avoid negative stock prices.

where $p(\xi_t|\mu_t)$ is the sample likelihood function described by:

$$p^{\mathbb{P}}\big(\xi_t|\mu_t\big) = \frac{1}{\sqrt{2\pi\sigma^2/n}}\exp\left[\frac{-\big(\bar{\xi}_t-\mu_t\big)^2}{2\sigma^2/n}\right],$$ (14)

which is a normal distribution function with mean $\bar{\xi}_t = (1/n)\sum_{i=t-n+1}^{t}\xi_i$ and variance σ^2/n, since the agent knows that dividends follow a geometric random walk, $\ln(D_t/D_{t-1}) = \mu_t + \sigma\varepsilon_t$. Hence, we can rewrite the posterior density in Eq. (12) as:

$$p^{\mathbb{P}}\big(\mu_t|\xi_t\big) = \frac{p^{\mathbb{P}}\big(\xi_t|\mu_t\big)p^{\mathbb{P}}(\mu_t)}{\int_{\mu_d}^{\mu_u} p^{\mathbb{P}}(\xi_t|\mu_t)p^{\mathbb{P}}(\mu_t)d\mu_t}.$$ (15)

Therefore, the representative agent uses the posterior probability in Eq. (15) to calculate the expected value of assets under incomplete information and learning. Let $X_t^{BL}(\mu_t)$ be the price of any asset under *BL*, which depends on the unknown parameter μ_t (in other words $X_t^{BL}(\mu_t)$ can be the Bayesian price of the bond, the stock or any option contract in our model). Then, $X_t^{BL}(\mu_t)$ is:

$$\begin{aligned}X_t^{BL}(\mu_t) &= \int_{\mu_d}^{\mu_u} X_t^{FI}(\mu_t)p^{\mathbb{P}}\big(\mu_t|\xi_t\big)d\mu_t \\ &= \frac{\int_{\mu_d}^{\mu_u} X_t^{FI}(\mu_t)p^{\mathbb{P}}\big(\xi_t|\mu_t\big)p^{\mathbb{P}}(\mu_t)d\mu_t}{\int_{\mu_d}^{\mu_u} p^{\mathbb{P}}\big(\xi_t|\mu_t\big)p^{\mathbb{P}}(\mu_t)d\mu_t},\end{aligned}$$ (16)

in which $X_t^{FI}(\mu_t)$ is the asset price under full information.

3.2.2. Relationships between learning schemes

As explained by Guidolin and Timmermann (2007), there are relationships between the learning schemes. Firstly, parameter estimates and asset prices under *RL* and *AL* cannot be identical in the majority of cases. Under *RL,* the agent considers future updates in parameter estimates when she values any asset, while an adaptive learner assumes that $\mu_t^* = \mu_{t+h}^*$, rendering asset prices under *AL* non-rational. Secondly, parameter estimates and asset prices under *AL* and *BL* are equal only in the case of the Bayesian learner having a degenerate prior density which assigns unit mass to the value of μ_t that is equal to $\bar{\xi}_t$.[11]

Thirdly, parameter estimates and asset prices under *RL* and *BL* are equivalent only in the scenario in which the prior distribution of the Bayesian learner is consistent with the model structure (see Guidolin and Timmermann, 2007). In the context of our model, parameter estimates and asset prices under *RL* and *BL* are identical in the case of the Bayesian learner having as prior density, $p^{\mathbb{P}}(\mu_t)$, the univariate density from which the new values of μ_t are drawn after breaks. When the Bayesian learner has different prior beliefs, Bayesian asset prices are not rational. Thus, *BL* is not a sufficient condition for obtaining full rationality through the learning process.

Parameter estimates and asset prices under *RL*, *AL* and *BL* are equal only when the number of signals used to learn is infinite after a break, $n \to \infty$ (i.e. when the parameter estimate asymptotically converges to the unknown value of μ_t). However, it is unlikely that the parameter estimate will asymptotically converge to μ_t in our model, since there is a possibility of breaks in μ_t, and thus the number of signals starts again from zero after each break in the economy.

3.2.3. Predictability patterns under learning

In each of the learning schemes explained in Section 3.2.1 (*RL*, *AL* and *BL*), the agent updates her parameter estimate step by step as new signals become available. In each of them, the agent learns progressively using new signals which are combined with historical information. Thus, learning can induce predictability patterns since current and future parameter estimates are based on a recursive updating procedure. In this section, we explain the intuition behind these predictability patterns.

As a first step, we assume that we have three types of learners: an adaptive learner, a Bayesian learner and a rational learner, who learn about the unknown parameter μ_t. Thus, the expected value of the parameter estimate for the adaptive learner is given by $\mu_t^* = \bar{\xi}_t$, as explained in Eq. (7).

In the case of the Bayesian learner, suppose that she has as prior density, $p^{\mathbb{P}}(\mu_t)$, a normal distribution with mean μ_0 and variance σ_0^2, defined by the support $[-\infty, \infty]$.[12] In addition, suppose that, in our model, the univariate probability density from which the new values of μ_t are drawn after breaks is also a normal distribution with the same mean, variance and support as the agent's prior beliefs. Thus, as explained in Section 3.2.2, parameter estimates and asset prices under

---

[11] Parameter estimates and asset prices under *AL* and *BL* are identical when the Bayesian learner has a prior density equal to: $p^{\mathbb{P}}(\mu_t) = \{\begin{smallmatrix}1 & if \ \mu_t = \bar{\xi}_t \\ 0 & otherwise\end{smallmatrix}$

[12] We assume that the agent's prior distribution is a normal distribution defined by the support $[-\infty, \infty]$ only in this section; this is in order to explain, in the simplest way possible, the intuition behind predictability patterns under learning.

BL and *RL* are identical, because the prior density is consistent with the model structure. Following Eq. (15), the posterior density, after the receipt of n signals, is normally distributed with mean μ_t^* and variance $\sigma_{\mu,t}^2$, described by

$$
\mu_t^* = \frac{\int_{\mu_d}^{\mu_u} \mu_t\, p^{\mathbb{P}}(\xi_t|\mu_t) p^{\mathbb{P}}(\mu_t) d\mu_t}{\int_{\mu_d}^{\mu_u} p^{\mathbb{P}}(\xi_t|\mu_t) p^{\mathbb{P}}(\mu_t) d\mu_t}
$$

$$
= \bar{\xi}_t \frac{\frac{n}{\sigma^2}}{\frac{n}{\sigma^2} + \frac{1}{\sigma_0^2}} + \mu_0 \frac{\frac{1}{\sigma_0^2}}{\frac{n}{\sigma^2} + \frac{1}{\sigma_0^2}}
\tag{17}
$$

and

$$
\sigma_{\mu,t}^2 = \frac{\int_{\mu_d}^{\mu_u} (\mu_t - \mu_t^*)^2 p^{\mathbb{P}}(\xi_t|\mu_t) p^{\mathbb{P}}(\mu_t) d\mu_t}{\int_{\mu_d}^{\mu_u} p^{\mathbb{P}}(\xi_t|\mu_t) p^{\mathbb{P}}(\mu_t) d\mu_t}
$$

$$
= \frac{1}{\frac{n}{\sigma^2} + \frac{1}{\sigma_0^2}}.
\tag{18}
$$

Thus, the mean μ_t^* is the expected value of the parameter estimate under *BL* or *RL* (since these learning schemes are equivalent in this case), which is equal to the weighted average of the information gleaned since the last break, $\bar{\xi}_t$, and the mean of the prior density, μ_0.

We can rewrite the expected value of the parameter estimate under *AL*, *BL* and *RL* in a recursive expression as follows:

$$
\Delta \mu_t^* = \left(\ln\left(\frac{D_t}{D_{t-1}} \right) - \mu_{t-1}^* \right) V_t
\tag{19}
$$

$$
with\ \Delta \mu_t^* = \mu_t^* - \mu_{t-1}^*\ and\ V_t = \begin{cases} \frac{1}{n} & under\ AL \\ \frac{\sigma_0^2}{n\sigma_0^2 + \sigma^2} & under\ BL\ or\ RL \end{cases},
$$

where V_t is a variable associated with the inaccuracy of the parameter estimate.

Eq. (19) shows that learning can induce predictability patterns in μ_t^*, due to the recursive-learning updating process. For instance, if $\ln(D_t/D_{t-1}) > \mu_{t-1}^*$, the representative agent updates her expectations upwards, since she observes a 'new' signal about the dividend drift, $\ln(D_t/D_{t-1})$, that is higher than her previous expected value of the parameter μ_{t-1}^*.

Suppose that the current value of the expected dividend drift μ_{t-1}^* is a long way below the unknown value of μ_t (i.e. the agent is incorrect and very pessimistic about μ_{t-1}^* in relation to the unknown dividend drift). In this scenario, future values of μ_t^* will progressively be revised upwards as a result of the recursive learning process; therefore, predictable dynamics will emerge in any of the learning schemes (similar arguments are used in Guidolin, 2006b). An analogous example can be described in an optimistic scenario.

The agent can have very wrong expected values of the dividend drift μ_{t-1}^* (i.e. she can wrongly be either very pessimistic or optimistic), especially after breaks. In early periods after breaks, when no 'long' history of dividend realizations is available, the agent calculates μ_t^* based on just a few signals; thus, its value can be far away from the unknown value of μ_t.

Therefore, *AL*, *BL* and *RL* are able to generate predictability patterns in the returns of all assets, including those of option contracts. Moreover, the 'inaccuracy' of parameter estimates also plays an important role in the predictability patterns generated by learning. In Eq. (19), the expression $\ln(D_t/D_{t-1}) - \mu_{t-1}^*$ (which gives the signal to revise the expected value μ_t^* upwards or downwards) is multiplied by V_t. Thus, upward or downward revisions are larger when the value of V_t is high, which reflects how inaccurate the agent's estimation is. A high value of V_t is likely after a break, when there are few available signals with which to obtain an accurate estimation of the unknown parameter.

3.3. Stock and bond prices under learning

In our study we assume that the agent follows a *BL* process, in which the agent's prior density, $p^{\mathbb{P}}(\mu_t)$, is described by a univariate probability density, which is equal to the density from which the new values of μ_t are drawn after breaks. We assume *BL* for several reasons. Firstly, the model under *AL* is misspecified, since in this learning scheme asset prices are not rational by definition. Secondly, models under *RL* are difficult to handle, because it is necessary to take into account all path-dependent sequences of potential probability distributions in relation to future parameter estimates.[13] Thirdly, asset prices under *BL* are rational when the prior density is consistent with the model structure, which is the case for our model (as assumed at the beginning of this paragraph). Fourthly, asset prices under *BL* can be obtained using Eq. (16), which provides an expression with which to obtain the price of any asset, including option contracts. Corollary I provides expressions for bond and stock prices under *BL*.

---

[13] Only in special cases is it possible to obtain a closed-form expression for the sequences of future probability distributions of the parameter estimates. For instance, Guidolin and Timmermann (2007) introduce a model based on a binomial lattice setup in which they present a closed-form solution for asset prices under *RL*.

Corollary I. *(Stock and bond prices under Bayesian learning): The prices of the assets, S_t and B_t, under incomplete information and Bayesian learning are given by:*

$$S_t^{BL} = \frac{\int_{\mu_d}^{\mu_u} S_t^{FI} L\left(\mu_t | \boldsymbol{\xi}_t\right)_n f(\mu_t) d\mu_t}{\int_{\mu_d}^{\mu_u} L\left(\mu_t | \boldsymbol{\xi}_t\right)_n f(\mu_t) d\mu_t},$$
(20)

and

$$B_t^{BL} = \frac{\int_{\mu_d}^{\mu_u} B_t^{FI} L\left(\mu_t | \boldsymbol{\xi}_t\right)_n f(\mu_t) d\mu_t}{\int_{\mu_d}^{\mu_u} L\left(\mu_t | \boldsymbol{\xi}_t\right)_n f(\mu_t) d\mu_t}$$
(21)

It is worth noting that the agent knows that new signals regarding parameter estimates incorporate some level of noise. This noise stems from the geometric random walk followed by dividends, specifically from the term $\sigma \varepsilon_t$. Thus, signals from new dividends provide partial information regarding the unknown value of μ_t after a break. Therefore, asset prices under *BL* differ somewhat from asset prices under full information. However, these differences narrow as more signals are received and learned.

3.4. The effect of bayesian learning on option prices and option returns

Under full information the risk-neutral probability measure, $p^{\mathbb{Q}}(S_{t+1})$, can be written as:

$$p^{\mathbb{Q}}(S_{t+1}) = \frac{m_{t+1} p^{\mathbb{P}}(S_{t+1})}{\mathrm{E}_t[m_{t+1}]}$$
(22)

where $p^{\mathbb{P}}(S_{t+1})$ is the physical probability measure, and m_{t+1} is the stochastic discount factor. However, under incomplete information both $p^{\mathbb{Q}}(S_{t+1})$ and $p^{\mathbb{P}}(S_{t+1})$ are affected by the *BL* process. The physical probability measure under *BL*, $p^{\mathbb{P},BL}(S_{t+1}|\boldsymbol{\xi}_t)$, is conditional on the information received after a break, which can be written as:

$$p^{\mathbb{P},BL}(S_{t+1}|\boldsymbol{\xi}_t) = \int_{\mu_d}^{\mu_u} p^{\mathbb{P}}(S_{t+1}|\boldsymbol{\xi}_t, \mu_t) p^{\mathbb{P}}(\mu_t|\boldsymbol{\xi}_t) d\mu_t$$
(23)

where $p^{\mathbb{P}}(\mu_t|\boldsymbol{\xi}_t)$ is the posterior probability distribution of the dividend drift estimate (explained in Eq. (15)). Eq. (23) shows that, under incomplete information, the *BL* process affects the agent's beliefs about the probabilities of current and future states of the economy. Similarly, the risk-neutral probability measure under *BL*, $p^{\mathbb{Q},BL}(S_{t+1}|\boldsymbol{\xi}_t)$, can be written as

$$
\begin{aligned}
p^{\mathbb{Q},BL}\left(S_{t+1}|\boldsymbol{\xi}_t\right) &= \int_{\mu_d}^{\mu_u} p^{\mathbb{Q}}(S_{t+1}|\boldsymbol{\xi}_t, \mu_t) p^{\mathbb{P}}(\mu_t|\boldsymbol{\xi}_t) d\mu_t \\
&= \int_{\mu_d}^{\mu_u} \frac{m_{t+1} p^{\mathbb{P}}(S_{t+1}|\boldsymbol{\xi}_t, \mu_t)}{\mathrm{E}_t[m_{t+1}]} p^{\mathbb{P}}(\mu_t|\boldsymbol{\xi}_t) d\mu_t.
\end{aligned}
$$
(24)

Therefore, Eqs. (23) and (24) show that *BL* can generate additional differences between the probability measures $\mathbb{P}$ and $\mathbb{Q}$ when changes in the estimated dividend drift are statistically associated with the agent's beliefs regarding the dynamics of the stochastic discount factor $m_{t+1} = \beta(D_{t+1}/D_t)^{-\alpha}$. This statistical association is present under a learning environment, since the agent's expectation regarding μ_t modifies her expectations about all variables which depend on μ_t, including how both the stochastic discount factor and its expected value evolve over time.

Corollary II. *(Option prices under Bayesian learning): The prices of a European put option contract, $Put_t^{FI}(K, \tau)$, and a European call option contract, $Call_t^{FI}(K, \tau)$, under incomplete information and Bayesian learning are given by:*

$$
\begin{aligned}
Put_t^{BL}(K, \tau) &= \int_0^{\infty} \frac{1}{\left(1 + r_{t+\tau}^{BL}\right)} \max\{K - S_{t+\tau}, 0\} p^{\mathbb{Q},BL}\left(S_{t+\tau}|\boldsymbol{\xi}_t\right) dS_{t+\tau} \\
&= \frac{\int_{\mu_d}^{\mu_u} \left[\int_0^{\infty} \frac{1}{(1+r_{t+\tau}^{BL})} \max\{K - S_{t+\tau}, 0\} p^{\diamond}\left(S_{t+\tau}|\boldsymbol{\xi}_t, \mu_t\right) dS_{t+\tau} \right] p^{\mathbb{P}}(\boldsymbol{\xi}_t|\mu_t) p^{\mathbb{P}}(\mu_t) d\mu_t}{\int_{\mu_d}^{\mu_u} p^{\mathbb{P}}(\boldsymbol{\xi}_t|\mu_t) p^{\mathbb{P}}(\mu_t) d\mu_t}
\end{aligned}
$$
(25)

and

$$
\begin{aligned}
Call_t^{BL}(K, \tau) &= \int_0^{\infty} \frac{1}{\left(1 + r_{t+\tau}^{BL}\right)} \max\{S_{t+\tau} - K, 0\} p^{\mathbb{Q},BL}\left(S_{t+\tau}|\boldsymbol{\xi}_t\right) dS_{t+\tau} \\
&= \frac{\int_{\mu_d}^{\mu_u} \left[\int_0^{\infty} \frac{1}{(1+r_{t+\tau}^{BL})} \max\{S_{t+\tau} - K, 0\} p^{\diamond}\left(S_{t+\tau}|\boldsymbol{\xi}_t, \mu_t\right) dS_{t+\tau} \right] p^{\mathbb{P}}(\boldsymbol{\xi}_t|\mu_t) p^{\mathbb{P}}(\mu_t) d\mu_t}{\int_{\mu_d}^{\mu_u} p^{\mathbb{P}}(\boldsymbol{\xi}_t|\mu_t) p^{\mathbb{P}}(\mu_t) d\mu_t}
\end{aligned}
$$
(26)

where

$$p^\diamond\left(S_{t+\tau}|\boldsymbol{\xi}_t,\mu_t\right) = \frac{m_{t+\tau}\,p^{\mathbb{P}}\left(S_{t+\tau}|\boldsymbol{\xi}_t,\mu_t\right)}{E_t^{BL}[m_{t+\tau}|\boldsymbol{\xi}_t]}$$

$$= \left(1+r_{t+\tau}^{BL}\right)\beta^\tau\left(\frac{D_{t+\tau}}{D_t}\right)^{-\alpha}\phi(\varepsilon_{t+\tau}|0,\sigma)\varphi(z|\tau,\pi)\eta(h_1|\pi)$$

$$\cdot\left[\eta(h_2|\pi)p^{\mathbb{P}}\left(\mu_{t+h_2}\right)\dots\eta(h_z|\pi)p^{\mathbb{P}}\left(\mu_{t+h_z}\right)\right],$$

where $S_{t+\tau} = D_{t+\tau}\Psi(g_{t+\tau})$, which depends on the mean dividend growth rate at time $t+\tau$, where $g_{t+\tau}$ is calculated from the dividend drift $\mu_{t+\tau}$ since $g_{t+\tau} = \exp(\mu_{t+\tau}+\sigma^2/2)-1$). In addition, $D_{t+\tau} = D_t\exp(\sqrt{\tau}\sigma\varepsilon_{t+\tau}-\tau\sigma^2/2)\prod_{i=1}^{z}(1+r_{t+h_i})^{h_i}$, where z represents the number of breaks until maturity, which is described by a binomial distribution $\varphi(z|\tau,\pi)$ with parameters τ and π; meanwhile $\{h_i\}_{i=1}^{z}$ is the time period between breaks, which is described by a geometric distribution, $\eta(h_i|\pi)$. Hence, the time-to-maturity of an option contract is $\tau = \sum_{i=1}^{z}h_i$. Moreover, $\{\mu_{t+h_i}\}_{i=2}^{z}$ are the new dividend drift values, which are described by a univariate density $\mu_{t+h_i} \sim p^{\mathbb{P}}(\cdot)$ with support $[\mu_d,\mu_u]$, where $\mu_{t+h_1} = \mu_t$ and $\mu_{t+h_z} = \mu_{t+\tau}$, and where $p^{\mathbb{P}}(\cdot)$ is obtained from a transformation of $G(\cdot)$ using the expression $1+g_t = \exp(\mu_t+\sigma^2/2)$. The dividend drift at time t, μ_t, is used as the integration variable in the integral that depends on μ_t. In addition, $1+r_{t+\tau}^{BL} = \prod_{j=1}^{\tau}(1+r_{j-1,j}^{BL})$ with $1+r_{j-1,j}^{BL} = 1/B_{j-1}^{BL}$, where B_{j-1}^{BL} is the price of the risk-free one-period bond in period $j-1$, which was defined in Eq. (21), *while $\varepsilon_{t+\tau}$ is described by a normal distribution with density $\phi(\varepsilon_{t+\tau}|0,\sigma)$, with mean zero and volatility σ.*

Corollary II shows that the Black-Scholes formula cannot be used to calculate option prices in an economy with breaks under *BL*. However, it is relevant to analyse our model's links to the option price formula in Guidolin and Timmermann (2003). They present an equilibrium model with incomplete information and *BL*, although without breaks. In the absence of breaks, the option prices in Eqs. (25) and (26) are equivalent to those in Guidolin and Timmermann (2003) for two reasons. First, both models are based on Lucas (1978) and use similar assumptions. For instance, our model is constructed in discrete time and dividends follow a geometric random walk; Guidolin and Timmermann's (2003) model is also constructed in discrete time, and the dividend process is characterized by a binomial lattice, which is equivalent to the geometric random walk used in our study. Thus, in an economy under full information and without breaks (and as mentioned in Section 3.1), both models provide option prices equivalent to the discretized version of the Black-Scholes formula, as in Rubinstein (1976). Second, under incomplete information and learning, both models are based on a rational *BL* procedure. Guidolin and Timmermann (2003) use a prior density that is consistent with their model structure (as described in Guidolin and Timmermann, 2007). This is also the case in our model (as explained at the beginning of Section 3.3); therefore, option prices are rational learning in Eqs. (25) and (26) and in Guidolin and Timmermann (2003).

In relation to the effect of learning on option returns, we know that the expected holding-period option returns of a European put option contract, $R_{t+\tau}^{Put}$, and a European call option contract, $R_{t+\tau}^{Call}$, are given by (see Broadie et al., 2009)

$$R_{t+\tau}^{Put} = \frac{E_t^{\mathbb{P}}[\max(K-S_{t+\tau},0)]}{Put_t(K,\tau)}-1 = \frac{E_t^{\mathbb{P}}[\max(K-S_{t+\tau},0)]}{E_t^{\mathbb{Q}}[\max(K-S_{t+\tau},0)]}-1 \tag{27}$$

and

$$R_{t+\tau}^{Call} = \frac{E_t^{\mathbb{P}}[\max(S_{t+\tau}-K,0)]}{Call_t(K,\tau)}-1 = \frac{E_t^{\mathbb{P}}[\max(S_{t+\tau}-K,0)]}{E_t^{\mathbb{Q}}[\max(S_{t+\tau}-K,0)]}-1. \tag{28}$$

Eqs. (27) and (28) show that the expected holding-period returns of options are described by the ratio of the expected option payoff calculated under the physical probability measure ($\mathbb{P}$) to the option price obtained under the risk-neutral probability measure ($\mathbb{Q}$). Thus, Eqs. (27) and (28) show that a change in the gap between probability measures $\mathbb{P}$ and $\mathbb{Q}$ can modify the magnitude of the expected returns of a holding-period option strategy. Consequently, the difference between the volatilities under measures $\mathbb{Q}$ and $\mathbb{P}$ (i.e. the *VRP*) can serve as a useful predictor of option returns. This is due to the fact that the *BL* process can simultaneously produce predictability patterns (see Eq. (19)) and differences between the two probability measures (see Eqs. (23) and (24)), which generate dynamic adjustments in option returns.

However, an analysis of the impact of learning on the predictability patterns of option returns is not straightforward. A closed-form expression for the predictable relationship between the *VRP* and option returns under a learning environment cannot be obtained, given the non-linearity of option contracts. Therefore, a numerical exercise is required, in which we use simulations based on our model to observe the effect of learning on the *VRP* and the returns of option strategies. The simulation analysis and results will be explained, and compared with actual market data on option returns, in the following sections.

4. Model implementation

We calculate the returns of option portfolios that control for changes in the underlying stock price, namely, delta-hedged and straddle portfolios. The use of these types of portfolios is important for our analysis, since they are not affected by

predictability patterns in the underlying asset, which can also be endogenously generated by learning effects. Hence, delta-hedged and straddle portfolios allow us to isolate the impact of learning on option prices and their returns.

Similarly to Goyal and Saretto (2009), we obtain delta-hedged portfolios and straddle portfolios from option contracts which are *ATM* and have a time-to-maturity of one month. The put-delta-hedged portfolio (call-delta-hedged portfolio) is formed by buying one put (call) option contract, and buying (short-selling) delta shares in the underlying stock. The straddle portfolio is calculated formed by buying one call option contract and one put option contract. The delta-hedged portfolios (for call and put options) and the straddle portfolio are obtained using a holding-period trading strategy on a monthly basis.

We compute the monthly returns of option portfolios following a procedure akin to that in Ni (2009) and Broadie et al., (2009), where returns have non-overlapping intervals. Put and call holding-period returns are calculated, respectively, using the following equations:

$$r_{t+\tau}^{Put} = \frac{\max(K - S_{t+\tau}, 0)}{Put_t^{BL}(K, \tau)} - 1 \tag{29}$$

and

$$r_{t+\tau}^{Call} = \frac{\max(S_{t+\tau} - K, 0)}{Call_t^{BL}(K, \tau)} - 1, \tag{30}$$

where $Put_t^{BL}(K, \tau)$ and $Call_t^{BL}(K, \tau)$ are the prices of put and call options, respectively, written at time t.

We produce 10,000 simulations based on our model. In each of the 10,000 simulations, we generate daily dividends over 12 years (3024 trading days). Thus, we simulate $10{,}000 * 3024 = 30{,}240{,}000$ trading days, with their dividends, which is the information received and learned by the representative agent.

In each one of the 10,000 simulations, we simulate daily dividends over the 12 years (3024 trading days) through the geometric random walk $\ln(D_t/D_{t-1}) = \mu_t + \sigma\varepsilon_t$. However, inside each 12-year simulation, we generate breaks in g_t, with times between breaks described by a geometric distribution with parameter π. The new mean dividend growth rate after a break is obtained from a uniform distribution $G(\cdot)$ with support $[g_d, g_u]$. Thus, μ_t also exhibits breaks, given that $g_t = \exp(\mu_t + \sigma^2/2) - 1$, which affects the generation of daily dividends in each simulation through the dividend geometric random walk.[14]

On each of the 3024 trading days in the 10,000 simulations, we use Eqs. (20) and (21) to obtain the prices of assets S_t and B_t, respectively. We use numerical integration (the adaptive Simpson quadrature) to perform the calculations.

We use Eqs. (25) and (26) to obtain monthly prices for call and put option contracts, respectively, in order to calculate non-overlapping one-month option returns. The numerators of Eqs. (25) and (26) are calculated through numerical methods, in two steps which are associated with the two integrals in these expressions. First, we solve the integral which depends on $S_{t+\tau}$ in Eqs. (25) and (26) by using Monte Carlo simulations with 20,000 paths. Each of the 20,000 paths is generated through the risk-neutral probability measure described below Eqs. (25) and (26). For each path, we obtain an analytical expression for the option payoff based on the dividend drift at time t, μ_t. Second, we integrate this analytical expression that depends on μ_t using numerical integration based on the adaptive Simpson quadrature. Thus, we solve the exterior integral (which depends on μ_t) for each of the 20,000 paths. Then, in order to obtain a price for a given option contract, we average the values obtained from the 20,000 paths. Therefore, in order to calculate the option prices, we generate 20,000 paths in the Monte Carlo simulation in each month of the 10,000 simulations (where each simulation reflects 12 years of market data).[15] The denominators of Eqs. (25) and (26) are also obtained by using the adaptive Simpson quadrature to solve the integrals.

Following Goyal and Saretto (2009), Bollerslev et al., (2009) and Cao and Han (2013), we calculate the *VRP* as the difference between the *IV* and the *RV*. The *IV* is the average of the implied volatilities of the European put and call option contracts (annualized), both of which are *ATM* and have time-to-maturity of one month. We obtain implied volatilities and option deltas by inverting the Black and Scholes (1973) model.[16] The *RV* is the standard deviation (annual basis) of the daily log stock returns in each month (avoiding overlapping periods).

In the model implementation we use the following model parameters. We use two values for the relative risk aversion coefficient: $\alpha = 0.20$ and $\alpha = 5.0$. The new mean dividend growth rate after a break is obtained from a uniform distribution $G(\cdot)$ with support $[g_d, g_u]$ with $g_u = 0.705\%$ and $g_d = -0.126\%$ (monthly basis). The uniform distribution support $[g_d, g_u]$ is

---

[14] In the model implementation, μ_t has probability density $f(\mu_t) = \exp(\mu_t + \sigma^2/2)/(g_u - g_d)$, where $\mu_d = \ln(1 + g_d) - \sigma^2/2$ and $\mu_u = \ln(1 + g_u) - \sigma^2/2$, which is obtained using the expression $g_t = \exp(\mu_t + \sigma^2/2) - 1$, and the fact that new values of g_t are described by a uniform distribution.

[15] In each month of the 10,000 simulations, we also compute the expected dividend yield and the expected rate of a zero-coupon bond with maturity equal to the expiration of the options; we use them to calculate the options' IVs.

[16] Despite the fact that the assumptions of Black and Scholes (1973) are not fully respected by our learning model, we obtain implied volatilities and option deltas by inverting the Black and Scholes (1973) model. This is consistent with other researchers (e.g. Guidolin and Timmermann, 2003; and Gonçalves and Guidolin, 2006) and practitioners, who also compute implied volatilities and option deltas using the Black and Scholes (1973) model, despite the very well-known fact that market data do not reflect the assumptions of this model. Actually, relationships and predictability patterns between assets and the VRP have been observed empirically by means of Black and Scholes' (1973) implied volatilities and option deltas. Consequently, one can look at our use of Black and Scholes' (1973) model as a way of controlling and making our results comparable with the previous empirical evidence and the evidence presented here based on S&P 500 option contracts. Consequently, the spirit of our analysis is to replicate with our learning model what academics and practitioners do, and thereby analyse whether we can obtain the same predictability patterns and behaviours in option returns.

Table 1

Summary statistics of model simulations. This table contains summary statistics for the returns on option portfolios and the *VRP* from simulations performed using our learning model. $R_{DHput,t}$ is the excess return on a put-delta-hedged portfolio that is formed by buying one *ATM* put option contract with time-to-maturity one month, and delta shares in the stock. $R_{STRD,t}$ is the excess return on a straddle portfolio that is calculated as a combination of one call option contract and one put option contract, where both contracts are *ATM* and have time-to-maturity one month. The implied volatility, IV_t, is calculated as the average of the implied volatilities of European put and call option contracts, where both contracts are *ATM* and have time-to-maturity one month, while RV_t is the standard deviation (annual basis) of the daily log stock returns in each month, avoiding overlapping periods. The values in the table reflect the averages of the respective variables over the 10,000 simulations performed using our model. The proportions of simulations with significant statistics (at 5% significance) are shown in round brackets.

Scenario		$R_{DHput,t}$	$R_{STRD,t}$ $\alpha = 0.2$	$IV_t - RV_t$	$R_{DHput,t}$	$R_{STRD,t}$ $\alpha = 5.0$	$IV_t - RV_t$
	Mean	−0.017	−0.479	0.092	−0.084	−0.927	0.364
	Std. dev.	0.005	0.382	0.017	0.002	0.292	0.082
	Skewness	0.866	1.134	−0.266	2.282	6.380	−0.589
	Kurtosis	4.172	4.778	5.606	13.567	60.303	5.496
Breaks - Inc. Inf. (Learning)				Correlations			
	$R_{DHput,t}$	1.00	0.94	−0.33	1.00	0.77	−0.80
		(100.00)	(100.00)	(94.70)	(100.00)	(100.00)	(96.72)
	$R_{STRD,t}$		1.00	−0.27		1.00	−0.38
			(100.00)	(82.34)		(100.00)	(88.40)
	$IV_t - RV_t$			1.00			1.00
				(100.00)			(100.00)

consistent with the real dividend growth rate in the period of our empirical analysis, based on the S&P 500 option returns (which will be described in the following section), and is also congruent with the values used in Timmermann (2001). For instance, if we take the annual dividend returns of the S&P 500 index in real terms (from Robert Shiller's database), they have a mean and standard deviation of 3.39% and 5.27%, respectively, in the period between 1996 and 2007. Hence, the interval of the mean of the annual returns on dividends at one standard deviation is [−1.88%, 8.66%] on an annual basis, or the equivalent of [−0.16%, 0.69%] on a monthly basis. These values are very close to the support $[g_d, g_u]$ used in our model. The dividend process volatility, σ, is set at 1.44% on a monthly basis (5.00% annual basis), which is in line with the empirical data (i.e., as mentioned above, the standard deviation of the real dividend returns of the S&P 500 index in the period 1996–2007 is 5.27%) and with Timmermann (2001).

The rate of impatience, ρ, is constrained by the model structure. In particular, ρ is constrained by the value g_u (which was set in the previous paragraph), since we need to have positive stock prices. To maintain positive stock prices, we need to impose that $1 + \rho > (1 + g_u)^{1-\alpha}$. Thus, we set ρ to 0.713% on a monthly basis. One may argue that this rate of impatience is high in relation to the real interest rates observed between 1996 and 2007. However, as mentioned above, we need this value to obtain positive stock prices. Moreover, it is worth pointing out that our focus is mainly to provide a learning model to ' explain' the predictive power of the *VRP* for option returns, rather than perfectly calibrating all variables. Thus, we prefer to closely calibrate the support of the uniform distribution from which we extract the new value of g_t after a break (as explained in the previous paragraph), and then to adjust ρ based on the constraint $1 + \rho > (1 + g_u)^{1-\alpha}$. This decision was made on the grounds that g_t is the ' unknown' parameter that the representative agent has to learn in our model.

In relation to the probability of breaks, π, we use the test introduced by Chu et al., (1996) to detect breaks using market data. We use real daily dividends from the S&P 500 index between 1996 and 2007 (after applying the Hodrick-Prescott filter to the data). We identify eight breaks in the market data; hence we assume that π is equal to 0.056 (on a monthly basis).

Table 1 presents summary statistics of the data simulated using our model under incomplete information and learning. In Table 1, as in the remainder of the paper, we limit ourselves to reporting the analysis for the returns on the put-delta-hedged portfolio and on the straddle portfolio, due to space limitations. The outcomes for the returns on the call-delta-hedged portfolio (available upon request) are quantitatively and qualitatively similar to the results presented here, and also congruent with the empirical literature. In Table 1, the mean excess returns on the put-delta-hedged (straddle) portfolios are equal to −1.7% and −8.4% (−47.9% and −92.7%) for relative risk aversion coefficient $\alpha = 0.20$ and $\alpha = 5.0$, respectively, which are consistent with the levels reported in previous empirical studies (see, e.g., Gonçalves and Guidolin, 2006; and Bernales and Guidolin, 2014).

In a Black-Scholes economy, the *VRP* should be equal to zero; however, we can see in Table 1 that our learning model generates a divergence between *IV* and *RV*. Table 1 also shows that the *VRP* is negatively correlated to the option returns, which is consistent with Goyal and Saretto (2009) and Cao and Han (2013).[17] In 94.70% and 96.72% (82.34% and 88.40%) of

---

[17] Whereas Goyal and Saretto (2009) and Cao and Han (2013) calculate the *VRP* as $RV - IV_t$, we use $IV_t - RV_t$ for our calculation. Nevertheless, our results are congruent with both papers. Goyal and Saretto (2009) find that portfolios with a large positive difference between *RV* and *IV* produce an economically and statistically significant monthly option return. Cao and Han (2013) show that the difference between *RV* and *IV* has a significant positive relationship with returns on call- and put-delta-hedged portfolios.

Table 2

Volatility risk premium regressions based on model-simulated option returns and S&P option returns. This table presents single-variable regressions of the lagged *VRP* on holding-period put-delta-hedged portfolio returns (R_{DHput}), and holding-period straddle portfolio returns (R_{STRD}) over 1-, 3-, 6-, 9- and 12-month forecasting horizons. The variables ($IV_t − RV_t$), $R_{DHput,t}$ and $R_{STRD,t}$ are defined as in Table 1. The results of the first two panels are based on model-simulated option returns, in an economy under Bayesian learning, while the results for the last panel are obtained from S&P 500 option returns between 1996 and 2007. The values in the table, in the case of simulated option returns, reflect the averages of the respective variables over the 10,000 simulations performed using our model. To adjust for heteroscedasticity and serial correlation, robust Newey-West (1987) *t*-statistics are used in the *t*-tests. In the case of the simulated option returns, the proportions of simulations with significant statistics (at 5% significance) based on one-sided *t*-tests are shown in round brackets. In the case of the S&P 500 option returns, the *t*-statistics are reported in square brackets.

Monthly forecasting horizons	1	3	6	9	12
	Breaks - Inc. Inf. (Learning) and $\alpha = 0.2$				
	Dependent Variable $R_{DHput,t}$				
Constant	−0.01	−0.01	−0.01	−0.01	−0.01
	(53.25)	(65.99)	(77.16)	(85.21)	(87.49)
IV - RV	−0.11	−0.09	−0.06	−0.04	−0.02
	(51.18)	(35.24)	(28.21)	(16.72)	(14.40)
Adj. R^2 (%)	3.05	2.03	1.08	0.58	0.36
	Dependent Variable $R_{STRD,t}$				
Constant	−0.35	−0.40	−0.42	−0.44	−0.49
	(70.77)	(77.18)	(83.44)	(87.30)	(95.62)
IV - RV	−1.74	−1.49	−1.04	−0.71	−0.40
	(26.57)	(23.41)	(15.49)	(10.62)	(11.72)
Adj. R^2 (%)	0.91	0.66	0.32	0.17	0.07
	Breaks - Inc. Inf. (Learning) and $\alpha = 5.0$				
	Dependent Variable $R_{DHput,t}$				
Constant	−0.02	−0.03	−0.05	−0.06	−0.07
	(45.22)	(83.83)	(95.30)	(92.28)	(96.51)
IV - RV	−0.18	−0.14	−0.09	−0.06	−0.03
	(97.61)	(98.56)	(71.67)	(32.40)	(17.51)
Adj. R^2 (%)	45.91	25.88	13.42	7.84	4.28
	Dependent Variable $R_{STRD,t}$				
Constant	−0.66	−0.76	−0.72	−0.75	−0.78
	(91.27)	(95.07)	(92.39)	(97.13)	(97.47)
IV - RV	−0.71	−0.58	−0.47	−0.38	−0.29
	(44.98)	(40.23)	(33.38)	(28.10)	(26.36)
Adj. R^2 (%)	4.10	2.94	2.03	1.29	0.95
	S&P 500 option returns				
	Dependent Variable $R_{DHput,t}$				
Constant	0.00	0.00	0.00	−0.01	0.00
	[0.65]	[1.06]	[1.73]	[2.39]	[1.49]
IV - RV	−0.15	−0.11	−0.03	0.04	−0.06
	[2.86]	[2.10]	[0.46]	[0.67]	[0.98]
Adj. R^2 (%)	0.05	0.02	−0.01	0.00	0.00
	Dependent Variable $R_{STRD,t}$				
Constant	−0.01	−0.03	−0.07	−0.09	−0.05
	[0.18]	[0.44]	[1.03]	[1.38]	[0.67]
IV - RV	−2.70	−2.25	−0.33	0.23	−1.73
	[2.00]	[1.68]	[0.24]	[0.17]	[1.22]
Adj. R^2 (%)	0.02	0.01	−0.01	−0.01	0.00

the simulations, the returns on the put-delta-hedged (straddle) portfolio have significant negative correlations with the *VRP* under relative risk aversion coefficients of $\alpha = 0.20$ and $\alpha = 5.0$, respectively.

5. Predictability patterns of option returns under a learning scheme

In order to test our hypothesis that learning can explain why the *VRP* predicts option returns, we firstly employ a single-variable OLS regression analysis to account for several forecasting horizons using our simulated data. Subsequently, we use multivariate OLS regressions to analyse the predictive power of the *VRP* in relation to other factors that can affect option returns.

Table 2 reports the results for single-variable regressions using option returns as the dependent variable, with one- to twelve-month-lagged values of the *VRP* as the explanatory variable. In the case of simulated option returns, note that in this table the percentage of simulations with significant statistics is reported in round brackets and is based on one-sided *t*-tests

at the 5% significance level (since we expect the forecasting power to be reflected in negative coefficients as in Goyal and Saretto, 2009; and Cao and Han, 2013).

In Table 2, we also show the results of an empirical analysis comparing the results of the model simulations to the actual behaviour observed in the market. We use data from the S&P 500 index options over the period 1996–2007, obtained from the OptionMetrics Ivy DB database. That database contains the daily closing bid and ask option prices (where option prices are calculated as closing bid-ask midpoints), the Black-Scholes implied volatilities, the time-to-maturity, the strike price, the closing underlying index price, and the risk-free interest rate.[18] Thus, Table 2 also presents an empirical analysis of the forecasting features of the *VRP* in respect to option returns, with the returns obtained from the S&P 500 index option contracts. In the case of the regressions using S&P option returns, Table 2 reports the *t*-statistics of the estimated coefficients in square brackets, instead of the percentages of simulations with significant coefficients as in the regressions based on simulated option returns.

In the case of the regressions using model-simulated data, Table 2 shows that the *VRP* has predictive power for the returns on option portfolios. The estimated coefficients of (*IV-RV*) are negative for R_{DHput} and R_{STRD}, implying that a higher level of the *VRP* forecasts more negative returns for option holders. For instance, Table 2 shows that 97.61% and 44.98% of the simulations with $\alpha = 5.0$ produce significant estimated coefficients of the one-month-lagged *VRP* with respect to the excess returns on put-delta-hedged portfolios (R_{DHput}) and straddle portfolios (R_{STRD}), respectively.

Table 2 shows that the predictive power of the *VRP* for option returns is reduced as the forecasting horizon increases. The percentage of simulations with significant estimated coefficients declines as the forecasting period becomes longer. For example, for simulations with $\alpha = 5.0$, when the twelve-month-lagged level of $IV - RV$ is used to predict R_{DHput} and R_{STRD}, only 17.51% and 26.36% of the simulations have a significant estimated coefficient, respectively. In addition, in model simulations with $\alpha = 5.0$, the average estimated coefficients of the one-month-lagged *VRP* start at -0.18, and -0.71 for R_{DHput} and R_{STRD}, respectively. However, the average absolute value of the coefficients gradually decreases as the forecasting horizon increases. For the same level of risk aversion, the average estimated coefficient of the twelve-month-lagged *VRP* for R_{DHput} is -0.03 and for R_{STRD} is -0.29.

In the case of the regressions based on S&P 500 option returns shown in Table 2, we calculate the returns on portfolios of option contracts close to *ATM* and with time-to-maturity one month. Then, we use a linear interpolation of these returns around the specified moneyness and maturity to obtain a proxy for the returns *ATM* with time-to-maturity of one month, as in the simulated data. Similarly, we use linear interpolation of the implied volatilities close to the moneyness and maturity targets in order to obtain a proxy for the implied volatility *ATM* and with time-to-maturity of one month. The implied volatility for a given moneyness and maturity is the average of the implied volatilities of a call contract and a put contract. The realized volatility is the standard deviation of the daily log-returns of the S&P 500 index in each month.

Table 2 shows that the *VRP* also has predictive power for the returns of the S&P 500 index options. The estimated coefficients of the one-month-lagged *VRP* with respect to the returns on the put-delta-hedged and straddle portfolios are significant, with *t*-statistics of 2.86 and 2.00, respectively. In addition, the forecasting power of the *VRP* for option returns is reduced as the forecasting horizon increases, following a similar pattern to the coefficients reported based on our model simulations. However, the predictive power of the *VRP* seems to increase again with a 12-month forecasting horizon (probably due to seasonality effects that can appear in the same month of each year), although the coefficients reported are not significant, with *t*-statistics of 0.98 and 1.22 for R_{DHput} and R_{STRD}, respectively. Thus, the results presented in Table 2 show that learning is able to explain the predictive relationship between the *VRP* and the option returns, observed in real option market data.

In Tables 3 and 4, we extend the regression analysis from Table 2 by including additional potential predictor variables lagged by one and three months, respectively. We include the realized volatility, *RV*, the slope in the moneyness dimension, $Slope_{Mon}$, and the market excess return, R_m. The *RV* was explained in regards to Table 1. The value of $Slope_{Mon}$ is obtained as the difference between the *IV* from contracts with $K/S = 0.96$ and time-to-maturity one month (i.e. the average of the call and put contracts) and the *IV* from contracts with $K/S = 1.04$ and time-to-maturity one month (again the average of the call and put contracts). The value of R_m is equal to the excess return per month on the underlying stock, obtained from Eq. (20), since this is the single stock in the economy under our option pricing model. Unlike in Table 2, where the percentages of simulations with significant statistics are based on one-sided *t*-tests, in Tables 3 and 4 they are based on two-sided *t*-tests since we do not want to impose any directional relationship on the analysis.

In the case of the regressions with simulated option returns, Table 3 reports that the *VRP* has more significant forecasting power for option returns than other factors, when they are used in 'single-variable' regressions. This result is observed for delta-hedged and straddle option portfolios, with both $\alpha = 0.2$ and $\alpha = 5.0$, which supports the evidence presented in Table 2. The highest forecasting power for one-month-lagged factors, without considering $IV - RV$, comes from $Slope_{Mon}$ in the case of the model simulations with $\alpha = 0.2$, and from *RV* in the case where $\alpha = 5.0$, but the numbers of simulations with significant coefficients are lower in both cases than in the case of the *VRP*.

Regarding the multivariate regressions reported in Table 3, we select predictor variables in such a way that multicollinearity problems are reduced. Table 3 shows that, even after including $R_{m,t}$ in the regressions involving the *VRP*, the *VRP* still

---

Table 3

Monthly return regressions based on model-simulated option returns and S&P option returns. This table reports single-variable and multi-variable regressions used to forecast holding-period put-delta-hedged portfolio returns, R_{DHput}, and holding-period straddle portfolio returns, R_{STRD}, using one-month-lagged predictors. The results in the panels on the left side are based on model-simulated option returns under Bayesian learning, while the results in the panels on the right side are obtained from S&P 500 option returns between 1996 and 2007. The variables $(IV_t - RV_t)$, $R_{DHput,t}$ and $R_{STRD,t}$ are defined as for Table 1. The slope on the moneyness dimension of the implied volatility surface, $Slope_{Mon}$, is calculated as the difference between the implied volatility of contracts with $K/S = 0.96$ and time-to-maturity one month and the implied volatility of contracts with $K/S = 1.04$ and time-to-maturity one month. $R_{m,t}$ is the excess return on the stock. The values in the table, in the case of the simulated option returns, reflect the averages of the respective variables over the 10,000 simulations performed using our model. To adjust for heteroscedasticity and serial correlation, robust Newey-West (1987) t-statistics are used in the t-tests. In the case of the simulated option returns, the proportions of simulations with significant statistics (at 5% significance) based on two-sided t-tests are shown in round brackets. In the case of the S&P 500 option returns, the t-statistics are reported in square brackets.

Monthly return regression

	Breaks - Inc. Inf. (Learning) and α =0.2									Breaks - Inc. Inf. (Learning) and α =5.0					S&P 500 option returns						
Dependent Variable $R_{DHput,t}$																					
Constant	−0.01 (44.69)	−0.01 (65.22)	−0.01 (96.17)	−0.02 (100.00)	−0.01 (11.54)	−0.02 (18.77)	−0.02 (16.85)	−0.02 (30.30)	−0.08 (93.65)	−0.09 (100.00)	−0.07 (97.51)	0.00 (6.33)	−0.14 (70.15)	−0.03 (49.33)	0.00 [0.65]	−0.01 [1.32]	−0.01 [1.43]	−0.01 [2.32]	0.00 [0.03]	0.00 [0.25]	−0.01 [0.61]
$IV_{t-1} - RV_{t-1}$	−0.11 (44.64)				−0.11 (40.10)			−0.18 (99.69)		−0.19 (97.26)					−0.15 [2.86]				−0.13 [2.51]		
RV_{t-1}		−0.03 (11.45)				−0.02 (10.69)			−0.12 (92.61)			−0.13 (87.42)				0.01 [0.31]				0.00 [0.13]	
$Slope_{Mon,t-1}$			−0.02 (25.64)				−0.02 (24.10)			−0.10 (75.02)			−0.13 (72.49)				−0.07 [0.73]				−0.06 [0.66]
$R_{m,t-1}$				0.00 (6.73)	−0.01 (7.17)	−0.01 (9.02)	0.00 (7.04)			−0.13 (35.49)	−0.05 (33.54)	−0.12 (23.07)		0.08 (43.15)				−0.10 [1.98]	−0.07 [1.42]	−0.09 [1.82]	−0.09 [1.86]
Adj. R^2 (%)	3.05	0.30	1.95	0.11	3.56	0.70	2.14	45.91	4.58	25.00	2.17	47.08	17.04	29.19	0.05	−0.01	0.00	0.02	0.05	0.01	0.01
Dependent Variable $R_{STRD,t}$																					
Constant	−0.35 (61.62)	−0.46 (68.67)	−0.48 (100.00)	−0.50 (100.00)	−0.60 (12.81)	−0.70 (16.44)	−0.67 (15.97)	−0.66 (79.99)	−0.95 (98.85)	−0.96 (100.00)	−0.97 (96.63)	−0.73 (48.41)	−1.35 (73.51)	−1.13 (62.70)	−0.01 [0.18]	−0.14 [1.05]	−0.26 [1.41]	−0.08 [1.40]	0.09 [0.42]	0.02 [0.07]	−0.12 [0.42]
$IV_{t-1} - RV_{t-1}$	−1.74 (18.92)				−1.72 (16.40)			−0.71 (24.36)		−0.71 (17.50)					−2.70 [2.00]				−2.24 [1.65]		
RV_{t-1}		−0.53 (9.18)				−0.44 (8.11)			−0.07 (17.10)			0.00 (7.49)				0.34 [0.43]				0.18 [0.22]	
$Slope_{Mon,t-1}$			−0.34 (14.54)				−0.30 (14.22)			−0.27 (8.44)			−0.19 (3.35)				−2.56 [1.04]				−2.18 [0.90]
$R_{m,t-1}$				−0.07 (7.06)	−0.11 (6.67)	−0.18 (7.56)	−0.08 (6.96)			−0.48 (7.82)	−0.37 (4.40)	−0.63 (3.86)		−0.27 (5.00)				−2.29 [1.85]	−1.75 [1.36]	−2.10 [1.64]	−2.15 [1.69]
Adj. R^2 (%)	0.91	0.10	0.45	0.03	1.01	0.24	0.56	4.10	−0.65	1.05	−0.42	3.83	0.42	1.50	0.02	−0.01	0.00	0.02	0.02	0.01	0.01

Table 4
Three-month-return regressions based on model-simulated option returns and S&P option returns. This table reports single-variable and multi-variable regressions used to forecast holding-period put-delta-hedged portfolio returns, R_{DHput}, and holding-period straddle portfolio returns, R_{STRD}, using three-month-lagged predictors. The results of the panels on the left side are based on model-simulated option returns under Bayesian learning, while the results of the panels on the right side are obtained from S&P 500 option returns between 1996 and 2007. The variables ($IV_t - RV_t$), $R_{DHput,t}$ and $R_{STRD,t}$ are defined as for Table 1. The slope on the moneyness dimension of the implied volatility surface, $Slope_{Mon}$, is calculated as the difference between the implied volatility of contracts with $K/S = 0.96$ and time-to-maturity one month and the implied volatility of contracts with $K/S = 1.04$ and time-to-maturity one month. $R_{m,t}$ is the excess return on the stock. The values in the table, in the case of the simulated option returns, reflect the averages of the respective variables over the 10,000 simulations performed using our model. To adjust for heteroscedasticity and serial correlation, robust Newey-West (1987) t-statistics are used in the t-tests. In the case of the simulated option returns, the proportions of simulations with significant statistics (at 5% significance) based on two-sided t-tests are shown in round brackets. In the case of the S&P 500 option returns, the t-statistics are reported in square brackets.

Three-month return regression

	Breaks - Inc. Inf. (Learning) and α = 0.2								Breaks - Inc. Inf. (Learning) and α = 5.0						S&P 500 option returns						
								Dependent Variable $R_{DHput,t}$													
Constant	−0.01 (52.96)	−0.02 (73.98)	−0.02 (97.28)	−0.02 (98.91)	−0.01 (8.21)	−0.01 (15.16)	−0.01 (12.79)	−0.03 (73.95)	−0.08 (100.00)	−0.08 (100.00)	−0.08 (100.00)	−0.02 (11.77)	−0.12 (64.90)	−0.04 (40.05)	0.00 [1.06]	0.00 [0.28]	0.00 [0.34]	0.00 [2.17]	−0.01 [0.69]	0.00 [0.21]	0.00 [0.41]
$IV_{t-1} - RV_{t-1}$	−0.09 (31.11)				−0.09 (28.44)			−0.14 (84.74)				−0.14 (80.85)			−0.11 [2.10]				−0.10 [1.76]		
RV_{t-1}		−0.01 (11.59)				−0.01 (9.43)			−0.10 (80.31)				−0.10 (77.27)			−0.04 [1.39]				−0.05 [1.50]	
$Slope_{Mon,t-1}$			−0.01 (18.85)				−0.01 (16.64)			−0.08 (75.23)				−0.11 (71.85)			0.04 [0.42]				0.03 [0.34]
$R_{m,t-1}$				0.00 (6.66)	0.00 (5.86)	0.00 (6.05)	0.00 (5.43)				−0.09 (37.55)	−0.02 (36.79)	−0.06 (19.34)	0.08 (37.41)				−0.10 [2.04]	−0.09 [1.72]	−0.11 [2.08]	−0.11 [2.03]
Adj. R^2 (%)	2.03	0.19	0.97	0.00	1.98	0.42	1.14	25.88	3.01	16.19	1.22	32.10	12.35	22.33	0.02	0.01	−0.01	0.02	0.03	0.02	0.01
								Dependent Variable $R_{STRD,t}$													
Constant	−0.40 (74.24)	−0.52 (73.81)	−0.46 (93.23)	−0.54 (100.00)	−0.34 (9.67)	−0.51 (13.04)	−0.45 (12.94)	−0.76 (82.52)	−0.92 (98.89)	−0.96 (100.00)	−0.90 (98.08)	−0.69 (56.60)	−1.23 (76.28)	−1.08 (58.22)	−0.03 [0.44]	0.09 [0.69]	−0.05 [0.27]	−0.07 [1.26]	−0.18 [0.83]	0.02 [0.07]	−0.19 [0.68]
$IV_{t-1} - RV_{t-1}$	−1.49 (16.56)				−1.43 (15.95)			−0.58 (18.77)				−0.58 (17.47)			−2.25 [1.68]				−2.00 [1.46]		
RV_{t-1}		−0.16 (8.78)				−0.10 (7.97)			−0.08 (17.04)				0.01 (7.77)			−1.12 [1.41]				−1.13 [1.41]	
$Slope_{Mon,t-1}$			−0.23 (12.97)				−0.24 (10.80)			−0.22 (8.36)				−0.18 (3.17)			0.47 [0.19]				0.18 [0.07]
$R_{m,t-1}$				−0.09 (6.67)	−0.03 (5.56)	−0.06 (5.30)	0.00 (4.58)				−0.37 (5.03)	−0.29 (6.27)	−0.40 (2.79)	−0.14 (5.36)				−1.86 [1.48]	−1.73 [1.32]	−2.10 [1.63]	−2.06 [1.59]
Adj. R^2 (%)	0.66	0.07	0.33	−0.06	0.66	0.13	0.37	2.94	−0.61	0.72	−0.43	2.80	0.18	0.94	0.01	0.01	−0.01	0.01	0.01	0.01	0.00

has important predictive power for option returns. For instance, in the case of delta-hedged portfolios, 40.10% and 97.26% of the simulations have significant coefficients for the *VRP,* with $\alpha = 0.2$ and $\alpha = 5.0$, respectively, when a lagged value of $R_{m,t}$ is included.

The three-month forecasting regressions presented in Table 4 confirm the findings reported in Table 3. On the one hand, in the single-variable regressions in Table 4, there are high percentages of simulations with significant estimated coefficients for the *VRP*. On the other hand, the percentage of simulations with significant estimated coefficients is still high after including the lagged excess market return. Nevertheless, the predictive power of the three-month-lagged *VRP* in Table 4 is lower than the predictive power reported for the one-month-lagged regressions in Table 3. This reduction in predictive power is reflected in the percentages of simulations with significant statistics, and the adjusted R^2. This reinforces our findings in Table 2, where we showed that the predictive power of the *VRP* declines with longer forecasting horizons.

In the case of the regressions using the S&P 500 option returns, we include in Table 3 (Table 4) the same forecasting variables as in the regression analysis using simulated data, lagged by one month (three months). In this empirical analysis, the value of $Slope_{Mon}$ is calculated using *IV*s for different option contracts around the required moneyness and time-to-maturity levels, and we then use linear interpolation to obtain the *IV*s for one-month-to-maturity contracts with $K/S = 0.96$ and $K/S = 1.04$. The value of R_m is equal to the excess return on the S&P 500 index in each month.

In Tables 3 and 4, we can see in the empirical analysis of the S&P 500 option returns that the *VRP* is the factor with the highest predictive power for option returns, in comparison to *RV*, $Slope_{Mon}$ and R_m. These results are in line with the results provided by Goyal and Saretto (2009). Most importantly, the results from the market data echo the results generated by our learning model. Therefore, the results in Tables 3 and 4 provide additional evidence that our learning model can offer an explanation for the puzzling forecasting features of the *VRP* with respect to option returns.

6. Conclusions

We use learning in a dynamic equilibrium model to explain the forecasting features of the *VRP* with regards to returns on option contracts. We extend the *BL* model proposed by Timmermann (2001), including options to analyse the effects of the learning process on option return predictability. We argue that the learning process can generate gaps between the probability measures $\mathbb{P}$ and $\mathbb{Q}$, which evolve over time and dynamically affect option returns. Moreover, the Bayesian updating procedure generates predictability patterns, given its recursive characteristics. Thus, the *VRP* (which reflects the difference between the volatility under the risk-neutral measure $\mathbb{Q}$ and the volatility under the physical measure $\mathbb{P}$) is a natural candidate to be a predictor of option returns in a learning environment.

To evaluate the effects of learning on the relationship between the *VRP* and option returns, we use a simulation analysis driven by our model. We show that, under a *BL* environment, the *VRP* consistently predicts a non-trivial fraction of returns on option strategies. Subsequently, we show that the forecasting features of the *VRP* for the option returns generated by our learning model are similar to those observed in an empirical analysis using S&P 500 index options. Thus, we provide evidence that learning offers an explanation for this option return puzzle, in which the *VRP* has predictive power regarding the returns on option contracts.

Finally, our paper opens the door for future studies into the effects of learning on option return predictability, since other interesting issues remain to be addressed. For instance, interesting avenues for future research include studying the effects of learning on the returns on options with different levels of moneyness and times-to-maturity, studying the relationship between learning and option returns when there is asymmetric information among agents, analysing a model in which there are noisy and irrational traders, and investigating cognitive mechanisms involving market makers in a microstructure setup using different types of derivative contracts.

Acknowledgments

Appendix A

Proof of Proposition 2: Option prices can be obtained using no-arbitrage conditions, but it is necessary to provide a proof that we have risk-neutral probabilities that characterize the state-price density. As a first step, we take the Euler equation of the stock price:

$$S_{t+k} = E_{t+k}\left[\beta\left(\frac{D_{t+k+1}}{D_{t+k}}\right)^{-\alpha}(S_{t+k+1} + D_{t+k+1})\right] \tag{A1}$$

Then, we divide both sides of Eq. (A1) by the bond price obtained from Eq. (3):

$$\frac{(1+\rho)S_{t+k}}{(1-\pi)(1+g_{t+k})^{-\alpha}+\pi\int_{g_d}^{g_u}(1+g_{t+k})^{-\alpha}dG(g_{t+k})}=E_{t+k}\left[\beta\left(\frac{D_{t+k+1}}{D_{t+k}}\right)^{-\alpha}\right.$$
$$\left.\cdot(S_{t+k+1}+D_{t+k+1})\frac{(1+\rho)}{(1-\pi)(1+g_{t+k})^{-\alpha}+\pi\int_{g_d}^{g_u}(1+g_{t+k})^{-\alpha}dG(g_{t+k})}\right]. \tag{A2}$$

In addition, we know that:

$$S_{t+k}^{*}=\frac{(1+\rho)S_{t+k}}{(1-\pi)(1+g_{t+k})^{-\alpha}+\pi\int_{g_d}^{g_u}(1+g_{t+k})^{-\alpha}dG(g_{t+k})} \tag{A3}$$

and

$$D_{t+k}^{*}=\sum_{s=0}^{k}D_{t+s}\frac{(1+\rho)}{(1-\pi)(1+g_{t+s})^{-\alpha}+\pi\int_{g_d}^{g_u}(1+g_{t+s})^{-\alpha}dG(g_{t+s})}, \tag{A4}$$

where S_{t+k}^{*} is the forward stock price, while D_{t+k}^{*} reflects the forward cumulative dividend process. We also know from the pricing kernel that:

$$E_t\left[\beta\left(\frac{D_{t+1}}{D_t}\right)^{-\alpha}\frac{(1+\rho)}{(1-\pi)(1+g_{t+k})^{-\alpha}+\pi\int_{g_d}^{g_u}(1+g_{t+k})^{-\alpha}dG(g_{t+k})}\right]=1. \tag{A5}$$

If we add D_{t+k}^{*} to both sides of Eq. (A2) and we use Eq. (A5), then we obtain:

$$S_{t+k}^{*}+D_{t+k}^{*}=E_{t+k}\left[\beta\left(\frac{D_{t+k+1}}{D_{t+k}}\right)^{-\alpha}\frac{(1+\rho)}{(1-\pi)(1+g_{t+k})^{-\alpha}+\pi\int_{g_d}^{g_u}(1+g_{t+k})^{-\alpha}dG(g_{t+k})}\cdot\left(S_{t+k+1}^{*}+D_{t+k+1}^{*}\right)\right]. \tag{A6}$$

Through Eq. (A6) we are demonstrating that $S_{t+k}^{*}+D_{t+k}^{*}$ is described by a martingale under the conditional probability measure. Hence, the risk-neutral density can be written as:

$$p^{\mathbb{Q}}(S_{t+k})=\beta\left(\frac{D_{t+1}}{D_t}\right)^{-\alpha}\frac{(1+\rho)}{(1-\pi)(1+g_{t+k})^{-\alpha}+\pi\int_{g_d}^{g_u}(1+g_{t+k})^{-\alpha}dG(g_{t+k})}p_t(D_{t+k})$$
$$=\beta\left(\frac{D_{t+1}}{D_t}\right)^{-\alpha}\left(1+r_{t+k}^{FI}\right)p_t(D_{t+k}) \tag{A7}$$

where r_{t+k}^{FI} is the interest rate of the bond. Thus, the multi-period risk-neutral measure can be expressed as a sequence of multiple single-period risk-neutral measures [see Pliska (1997) for proof]. Therefore, $p^{\mathbb{Q}}(S_{t+\tau})$ is the risk-neutral density which reflects all paths that reach the state in which the dividend is $D_{t+\tau}$.

Suppose that z represents the number of breaks until maturity, which is described by a binomial distribution $\varphi(z|\tau,\pi)$ with parameters τ and π, while $\{h_i\}_{i=1}^{z}$ are the times between breaks, described by a geometric distribution $\eta(h_i|\pi)$. Hence, the time-to-maturity is $\tau=\sum_{i=1}^{z}h_i$. Therefore, for each path we have:

$$D_{t+\tau}^{FI}=D_t\exp\left(\sqrt{\tau}\sigma\varepsilon_{t+\tau}-\tau\sigma^2/2\right)\prod_{i=1}^{z}\left(1+g_{t+h_i}\right)^{h_i}. \tag{A8}$$

Here, the value of $\{g_{t+h_i}\}_{i=2}^{z}$ is the value of the mean dividend growth rate, which is described by a univariate distribution $g_{t+h_i}\sim G(\cdot)$ and pdf $\varrho(g_{t+h_i})$ defined on the support $[g_d,g_u]$, while $g_{t+h_1}=g_t$ and $g_{t+\tau}=g_{t+h_z}$. Consequently,

$$p_t(D_{t+\tau})=\phi(\varepsilon_{t+\tau}|0,\sigma)\varphi(z|\tau,\pi)\eta(h_1|\pi)\left[\eta(h_2|\pi)\varrho\left(g_{t+h_2}\right)\ldots\eta(h_z|\pi)\varrho\left(g_{t+h_z}\right)\right]. \tag{A9}$$

Hence, from Eq. (A8) we can write:

$$p^{\mathbb{Q}}(S_{t+\tau})=\beta^{\tau}\left(\frac{D_{t+\tau}}{D_t}\right)^{-\alpha}\phi(\varepsilon_{t+\tau}|0,\sigma)\varphi(z|\tau,\pi)\eta(h_1|\pi)\left[\eta(h_2|\pi)\varrho\left(g_{t+h_2}\right)\ldots\eta(h_z|\pi)\varrho\left(g_{t+h_z}\right)\right]. \tag{A10}$$

□

References

Bernales, A., Guilolin, M., 2014. Can we forecast the implied volatility surface dynamics of equity options? Predictability and economic value tests. J. Banking Finance 46, 326–342.

Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. J. Political Econ. 81, 637–659.

Bollerslev, T., Tauchen, G., Zhou, H., 2009. Expected stock returns and variance risk premia. Rev. Financ. Stud. 22, 4463–4492.

Branch, William, Evans, G.W., 2010. Asset return dynamics and learning. Rev. Financ. Stud. 23, 1651–1680.

Branch, William A., 2016. Imperfect knowledge, liquidity and bubbles. J. Econ. Dyn. Control 62, 17–42.

Brandt, M., Zeng, Q., Zhang, L., 2004. Equilibrium stock return dynamics under alternative rules of learning about hidden states. J. Econ. Dyn. Control 28, 1925–1954.

Broadie, M., Chernov, M., Johannes, M., 2009. Understanding index option returns. Rev. Financ. Stud. 22, 4493–4529.

Buraschi, A., Trojani, F., Vedolin, A., 2014. When uncertainty blows in the orchard: Comovement and equilibrium volatility risk premia. J. Financ. 69, 101–137.

Cao, J., Han, B., 2013. Cross section of option returns and idiosyncratic stock volatility. J. Financ. Econ. 108, 231–249.

Carr, P., Wu, L., 2009. Variance risk premiums. Rev. Financ. Stud. 22, 1311–1341.

Chu, C.S.J., Stinchcombe, M., White, H., 1996. Monitoring structural change. Econometrica 64, 1045–1065.

David, A., Veronesi, P., 2002. Option Prices with Uncertain Fundamentals. Working paper. University of Chicago.

David, A., Veronesi, P., 2014. Investor's and central bank's uncertainty embedded in index options. Rev. Financ. Stud. forthcoming.

Drechsler, I., Yaron, A., 2011. What's vol got to do with it. Rev. Financ. Stud. 24, 1–45.

Gonçalves, S., Guidolin, M., 2006. Predictable dynamics in the S&P 500 index options implied volatility surface. J. Bus. 79, 1591–1635.

Goyal, A., Saretto, A., 2009. Cross-section of option returns and volatility. J. Financ. Econ. 94, 310–326.

Guidolin, M., 2006a. High equity premia and crash fears. Rational foundations. Econ. Theor. 28, 693–708.

Guidolin, M., 2006b. Pessimistic beliefs under rational learning: quantitative implications for the equity premium puzzle. J. Econ. Bus. 58, 85–118.

Guidolin, M., Timmermann, A., 2003. Option prices under Bayesian learning: Implied volatility dynamics and predictive densities. J. Econ. Dyn. Control 27, 717–769.

Guidolin, M., Timmermann, A., 2007. Properties of equilibrium asset prices under alternative learning schemes. J. Econ. Dyn. Control 31, 161–217.

Leisch, F., Hornik, K., Kuan, C.M., 2000. Monitoring structural changes with the generalized fluctuation test. Econometric Theor. 16, 835–854.

Lucas, R., 1978. Asset prices in an exchange economy. Econometrica 46, 1429–1445.

Massa, Massimo, Simonov, A., 2005. Is learning a dimension of risk? J. Banking Finance 29, 2605–2632.

Ni, S.X., 2009. Stock Option Returns: A puzzle. Hong Kong University of Science and Technology working paper.

Pliska, S.R., 1997. Introduction to Mathematical Finance. Blackwell, Oxford.

Rubinstein, M., 1976. The valuation of uncertain income streams and the pricing of options. Bell J. Econ. 7, 407–425.

Shaliastovich, I., 2009. Learning, Confidence and Option Prices. University of Pennsylvania Working paper.

Timmermann, A., 1993. How learning in financial markets generates excess volatility and predictability of excess returns. Q. J. Econ. 108, 1135–1145.

Timmermann, A., 1996. Excess volatility and predictability of stock prices in autoregressive dividend models with learning. Rev. Econ. Stud. 63, 523–557.

Timmermann, A., 2001. Structural breaks, incomplete information, and stock prices. J. Bus. Econ. Stat. 19, 299–314.

Veronesi, P., 1999. Stock market overreaction to bad news in good times: a rational expectations equilibrium model. Rev. Financ. Stud. 12, 975–1007.

Veronesi, P., 2000. How does information quality affect stock returns? J. Finance 55, 807–837.