Learning and Index Option Returns

Alejandro Bernales   Gonzalo Cortazar

Luka Salamunic   George Skiadopoulos [*]

Abstract

Little is known about the economic sources that may generate the abnormal returns observed in put index options. We show that the learning process followed by investors may be one such source. We develop an equilibrium model under partial information in which a rational Bayesian learner prices put option contracts. Our model generates put option returns similar to the empirical returns of S&P 500 put index options. This result is not obtained when we analyse alternative setups of the model in which no learning process exists.

JEL classification: D83, G13, G14, G17.
Keywords: Put option returns, equilibrium model, partial information, learning, structural breaks.

---

1   Introduction

An important goal of financial economics research is to explain anomalies observed in asset prices. One such anomaly is related to the fact that the returns of put index options have been too high relative to their risk (e.g., Coval and Shumway, 2001; Bondarenko, 2003; Jones, 2006; and Constantinides *et al.*, 2013). To date, the economic sources driving this abnormal behaviour have remained unexplored by the literature.

In this paper, we provide one potential explanation for the high abnormal returns observed in put index option contracts: the learning process followed by investors. The intuition behind this hypothesis is simple. The expected put option return of a hold-to-maturity naked strategy, $R_{t+\tau}^{Put}$, is defined as $R_{t+\tau}^{Put} = E_t^{\mathbb{P}}[\max(K - S_{t+\tau}, 0)]/P_t(K, \tau) - 1$, where $E_t^{\mathbb{P}}[\cdot]$ is the expectation operator under the $\mathbb{P}$ (real-world) probability measure at time t, $P_t(K, \tau)$ is the price of a put option contract, $S_{t+\tau}$ is the underlying asset at time $t + \tau$, K is the strike price, and τ is the option's time to maturity. Thus, the value of $R_{t+\tau}^{Put}$ can also be written as $R_{t+\tau}^{Put} = E_t^{\mathbb{P}}[\max(K - S_{t+\tau}, 0)]/E_t^{\mathbb{Q}}[\max(K - S_{t+\tau}, 0)] - 1$, with $E_t^{\mathbb{Q}}[\cdot]$ denoting the expectation operator under the $\mathbb{Q}$ (risk-neutral) probability measure, which implies that the potential differences between the $\mathbb{P}$ and $\mathbb{Q}$ probability measures determine $R_{t+\tau}^{Put}$. Moreover, there is evidence that the process followed by investors to learn about unknown market variables can induce differences between the $\mathbb{P}$ and $\mathbb{Q}$ probability measures (e.g., David and Veronesi, 2002; Guidolin and Timmermann, 2003). Thus, this learning process is a natural candidate to explain the abnormal returns observed in put index option contracts.

We develop a simple equilibrium model under partial information, in which a rational Bayesian learner prices a set of put option contracts. Our model extends Timmermann (2001) by including put options in the market. Similarly to Timmermann (2001), our model reflects a Lucas (1978) discrete-time endowment economy with a representative agent. In the model, there is partial information because the mean dividend growth rate, g_t, is unknown. However, the agent recursively learns about g_t through a Bayesian updating procedure as new signals arrive.

In the model, learning does not disappear asymptotically because the number of signals cannot increase infinitely. This is because the mean dividend growth rate g_t is subject to structural breaks. After each structural break, a new value for g_t is drawn from a uniform density and kept constant until the next break. Therefore, the new value of g_t after a break must be learned by using only post-break information.

We find that our model generates abnormal put option returns for a naked option strategy, as well as for a straddle option portfolio (which is not affected by changes in the underlying stock). The put option returns generated by our model exhibit the same behaviour as the corresponding returns from actual S&P 500 index option data. Conversely, we cannot obtain similar results when we analyse alternative scenarios that do not include a learning process.

We also show that the put option returns generated by our modelling setup under learning can be described by a linear factor model that includes the volatility risk premium (i.e., the difference between the volatilities under the $\mathbb{Q}$ and $\mathbb{P}$ probability measures) as a factor. This result is also observed in put option returns obtained from S&P 500 index

option contracts. This finding is consistent with our argument that learning generates differences between the $\mathbb{Q}$ and $\mathbb{P}$ probability measures, which may be one of the economic explanations for the abnormal returns observed in put index options.

The findings drawn from our theoretical model are related to the empirical literature that provides evidence of abnormal returns on put index option contracts. For example, Coval and Shumway (2001), who pioneer a focus on option returns rather than on implied volatilities or option prices, document strong negative average returns on zero-beta at-the-money straddles. Bondarenko (2003) shows that naked puts yield large returns for option sellers. In a multifactor analysis, Jones (2006) and Constantinides *et al.* (2013) find high abnormal returns associated with short-term out-of-the-money puts.

Our model is closely related to Broadie *et al.* (2009), who show that option returns can be explained by models that incorporate a jump risk premium and estimation risk (Chambers *et al.*, 2014, replicate the analysis of Broadie *et al.*, 2009, and find similar results). In particular, Broadie *et al.*'s (2009) estimation risk argument is closely linked to our model. Estimation risk appears when the true parameter values used in pricing models cannot be accurately estimated from short time series. Broadie *et al.* (2009) take a reduced-form approach to analyse the impact of estimation risk on put option returns by increasing/decreasing the parameters of the $\mathbb{Q}$ probability measure by one standard deviation compared to the $\mathbb{P}$-parameters. However, they state in their conclusion, "*Our results are silent on the actual <u>economic sources</u> of the gaps between the $\mathbb{P}$ and $\mathbb{Q}$ measures. It is important to test potential explanations that incorporate investor heterogeneity, discrete trading, model misspecification, or <u>learning</u>*". Therefore, our study complements Broadie *et al.* (2009), since our model micro-founds the possibility of

estimation risk through learning, which endogenously induces differences between the $\mathbb{P}$ and $\mathbb{Q}$ probability measures and yields abnormally high put option returns.

Our paper is also related to theoretical studies in which stock returns are explained by the learning process followed by a representative agent (e.g., Timmermann, 1993, 1996, 2001; Veronesi, 1999, 2000; Brandt *et al.*, 2004; Guidolin, 2006; and Guidolin and Timmermann, 2007). In particular, our paper is closely related to the model developed in Timmermann (2001). His model is based on a Lucas (1978) economy where the mean dividend growth rate has breaks and is not known by a representative agent. He shows that learning generates volatility clustering and serial correlation in stock returns, in combination with an increase in their skewness and kurtosis. Departing from Timmermann (2001), we analyse whether learning can explain put option returns rather than stock returns. Therefore, we present an extended model that contains put option contracts in addition to the bonds and stocks present in a Lucas (1978) economy.

Our study is also associated with equilibrium models in which learning is used to explain the implied volatility surface (e.g., David and Veronesi, 2002; Guidolin and Timmermann, 2003; Shaliastovich, 2009; and Benzoni *et al.*, 2011). The observable implied volatility surface (hereafter *IVS*) is the variation of the implied volatility of option contracts written on the same underlying asset as a function of the option's strike price and time to maturity. The *IVS* should not be observed under the assumption of Black and Scholes' (1973) model, which is that the volatility of the returns of the underlying asset is constant. For example, David and Veronesi (2002) propose a model where the dividend drift follows a two-state stochastic process, and a representative agent is uncertain about which state is currently governing the economy. Guidolin and Timmermann (2003) present an

equilibrium model based on a Lucas (1978) economy, which is closer to our modelling setup. In their model, dividends evolve based on a binomial path with unknown state probability, which is recursively updated. Notably, in Guidolin and Timmermann (2003), the impact of rational Bayesian learning on option prices asymptotically disappears as the number of signals increases, which does not occur in our model due to structural breaks in the unknown parameter. Shaliastovich (2009) also offers a model in which the consumption growth rate is uncertain and investors learn about its value. Benzoni *et al.* (2011) extend a general equilibrium setting with an Epstein-Zin representative agent by including jumps and a Bayesian updating process for learning the unknown probability of future jumps in the growth and volatility of consumption.

Our paper differs in two ways from the above studies that use learning to explain the *IVS*. First, rather than explaining the variation in implied volatilities of option contracts across strike prices and times to maturity, our objective is to provide an economic explanation for the anomalous behaviour of put option returns. Second, from a modelling perspective, our unknown parameter is simply the mean dividend growth rate, rather than the current state of the economy, the state probability, the consumption growth rate, or the probability of future jumps in consumption growth and volatility. This is important since our objective is to show that learning can induce abnormal put option returns even in a simple model, where the unknown parameter is the mean dividend growth rate in a Lucas (1978) economy (despite the fact that similar results may be obtained with more sophisticated models, such as the ones described above).

The remainder of this paper is structured as follows. Section 2 presents the model, Section 3 describes its implementation, Section 4 analyses the theoretical results, and Section 5 concludes.

2 The model

2.1 An economy under full information

We first consider an economy under full information about economic fundamentals, wherein a representative agent prices four asset types: a bond, a stock, put option contracts, and change-of-state (hereafter COS) securities. Thus, there is a one-period zero-coupon default-free bond, B_t, in zero net supply at any time t. There is also a stock, S_t, with net supply normalised at one. In each period, the stock pays a real dividend, D_t, which follows a geometric random-walk process with drift μ_t and volatility σ:

$$\ln\left(\frac{D_t}{D_{t-1}}\right) = \mu_t + \sigma\varepsilon_t, \quad \varepsilon_t \backsim IIN(0,1), \tag{1}$$

where the mean dividend growth rate, $\boldsymbol{g_t}$ (and thus, $\boldsymbol{\mu_t}$, given that $\boldsymbol{\mu_t = \ln(1 + g_t) - \sigma^2/2}$), is subject to breaks. Time periods between breaks follow a geometric distribution with parameter $\boldsymbol{\pi}$. Therefore, $\boldsymbol{g_t}$ changes over time, but its value remains constant between breaks. We assume that as soon as a break occurs, a new value for the mean dividend growth rate, $\boldsymbol{g_t}$, is drawn from a uniform distribution, $\boldsymbol{g_{t+m} \backsim G(\cdot)}$, with lower and upper bounds of $\boldsymbol{g_d}$ and $\boldsymbol{g_u}$, respectively.

In Internet Appendix A, we examine the time series of daily dividends for the S&P 500 index to assess the validity of the assumption of breaks in the mean dividend growth rate. The results reported in Internet Appendix A (based on the Chu *et al.*, 1996, and Bai

and Perron, 1998, tests) show evidence of breaks. In addition, we choose a geometric distribution to characterise the time between breaks because it is a memoryless stochastic process. We assume a memoryless process in order to be consistent with the assumption that agents cannot predict the future and are thus unable to predict future breaks in g_t. The previous literature has also used a memoryless process to characterise periods between breaks (e.g., Pesaran *et al.*, 2006; Koop and Potter, 2007).

At time t, there is a set of European put option contracts with price $P_t(K, \tau)$ written on the stock, where K is the strike price and τ is the time to maturity. There is also a COS security with price A_t at any time t, which pays one unit in any period in which there is a break in g_t. COS securities make the market dynamically complete, as they are used to hedge the uncertainty generated by the breaks in the mean dividend growth rate, because the uncertainty generated by the breaks in g_t and the term ε_t in equation (1) cannot be dynamically hedged with the stock and the bond alone (Harrison and Pliska, 1981; Duffie and Huang, 1985). The use of COS securities to complete markets has been considered by Guo (2001), Mamon and Rodrigo (2005), and Yuen and Yang (2010).

We define the break indicator b_t, which indicates the occurrence of a break in g_t at time t; b_t equals 1 if there is a break at t and zero otherwise. Since periods between breaks follow a geometric distribution with parameter π, b_t follows a Bernoulli distribution with parameter π, where $\mathbf{Pr}(b_t = 0) = (1 - \pi)$ and $\mathbf{Pr}(b_t = 1) = \pi$. Thus, we assume that the COS security A_t pays one unit at $t + m$ if $m = inf\{m \geq 0: b_{t+m} = 1\}$. The COS security is analogous to an insurance policy that pays out in the case of a particular event. The COS security becomes worthless after a break, and a new COS security is issued that pays one unit in the period following the next break.

We assume a perfect capital market in which there are no taxes or transaction costs, unlimited short-sale possibilities, perfect liquidity, and a lack of borrowing or lending constraints. We assume that the representative agent has preferences described by a power utility function:

$$u(C_t) = \begin{cases} \dfrac{C_t^{1-\eta} - 1}{1 - \eta} & \eta \geq 0, \eta \neq 1 \\ \ln C_t & \eta = 1 \end{cases},$$

(2)

where C_t is the real consumption at time t and η is the coefficient of relative risk aversion. We assume that dividends are the economy's single source of income and that they are consumed as soon as they are received (i.e., $C_t = D_t$). The representative agent chooses asset holdings in order to maximise her lifetime expected utility:

$$\max_{\{w_{t+k}^S, w_{t+k}^B, w_{t+k}^A\}} E_t \left[\sum_{k=0}^{\infty} \beta^k u(D_{t+k}) \right],$$

(3)

in which $\beta = 1/(1 + \rho)$, ρ is the rate of impatience and w_{t+k}^S, w_{t+k}^B, and w_{t+k}^A are the shares of assets S_t, B_t, and A_t in the agent's portfolio, respectively. Note that put option contracts are not considered in the agent's maximisation problem described in equation (3), because markets are complete due to the existence of COS securities; thus, options are redundant. Prices S_t, B_t and A_t can be calculated by solving the respective Euler equations:

$$S_t = E_t[m_{t+1}(S_{t+1} + D_{t+1})],$$

(4)

$$B_t = E_t[m_{t+1}],$$

(5)

$$A_t = E_t[m_{t+1}b_{t+1}],$$

(6)

where $m_{t+1} = \beta(D_{t+1}/D_t)^{-\eta}$ is the stochastic discount factor. In the case of full information, Proposition I provides expressions for the equilibrium prices S_t, B_t, and A_t, which are determined by solving equations (4)-(6).

Proposition I: *Under full information and breaks in the dividend process, the equilibrium prices S_t, B_t, and A_t are given by:*

$$S_t = \frac{D_t}{1+\rho-(1-\pi)(1+g_t)^{1-\eta}}\left\{(1-\pi)(1+g_t)^{1-\eta} + \pi\left(\frac{I_1+(1-\pi)I_2}{1-\pi I_3}\right)\right\} = D_t\psi(g_t), \tag{7}$$

$$B_t = \frac{1}{(1+\rho)}\left\{(1-\pi)(1+g_t)^{-\eta} + \pi\int_{g_d}^{g_u}(1+g_t)^{-\eta}dG(g_t)\right\}, \tag{8}$$

$$A_t = \frac{1}{(1+\rho)}\pi\int_{g_d}^{g_u}(1+g_t)^{-\eta}dG(g_t), \tag{9}$$

where $I_1 = \int_{g_d}^{g_u}(1+g_t)^{1-\eta}dG(g_t)$, $I_2 = \int_{g_d}^{g_u}\frac{(1+g_t)^{2-2\eta}}{1+\rho-(1-\pi)(1+g_t)^{1-\eta}}dG(g_t)$,

$$I_3 = \int_{g_d}^{g_u}\frac{(1+g_t)^{1-\eta}}{1+\rho-(1-\pi)(1+g_t)^{1-\eta}}dG(g_t),$$

with $1 + \rho > (1 + g_u)^{1-\eta}$ to obtain positive stock prices.

Proof: Timmermann (2001).

Proposition I shows that the price-dividend ratio and the bond price are time varying, since they depend on g_t. Note that the one-period interest rate is given by $1 + r_t = 1/B_t$; thus, the interest rate is also time varying. When there are no breaks in the economy (i.e., $\pi = 0$), stock and bond prices are $S_t = D_t(1 + g_t)^{1-\eta}/\{1 + \rho - (1 + g_t)^{1-\eta}\}$ and $B_t = (1 + g_t)^{-\eta}/(1 + \rho)$, whereas the COS security is priced at zero (i.e., $A_t = 0$). In relation to option contracts, Proposition II provides the put option price when there are breaks in the mean dividend growth rate and full information (i.e., the agent knows the true mean dividend growth rate).

Proposition II: *Under full information and breaks, the price $P_t^{full}(K,\tau)$ of a European put option issued on the stock at time t with strike price K and time to maturity τ is given by:*

$$P_t(K,\tau) = \int_0^\infty \frac{1}{1+r_{t,t+\tau}} max\{K - S_{t+\tau}, 0\} f^Q(S_{t+\tau}) dS_{t+\tau}, \tag{10}$$

with

$$f_t^Q(S_{t+\tau}) = \frac{m_{t+\tau} f^P(S_{t+\tau})}{E_t[m_{t+\tau}]} = \phi^Q(\varepsilon_{t+\tau}) \prod_{i=1}^{\tau} Pr^Q(b_{t+i}) \varrho^{\mathbb{Q}}(g_{t+i}|b_{t+i} = 1), \tag{11}$$

where $f^Q(S_{t+\tau})$ is the risk-neutral density under full information and $S_{t+\tau}$ is defined in equation (7), while $D_{t+\tau} = D_t exp(\sqrt{\tau}\sigma\varepsilon_{t+\tau} - \tau\sigma^2/2) \prod_{i=1}^{\tau}(1 + r_{t+i})$, with $1 + r_{t+i} = 1/B_{t+i-1}$, where B_{t+i-1} is the price of the risk-free one-period bond in period $t + i - 1$ (as defined in equation (8)). In addition, $1 + r_{t,t+\tau} = \prod_{j=1}^{\tau}(1 + r_{j-1,j})$. In equation (11), $\phi^{\mathbb{Q}}(\varepsilon_{t+\tau})$ is a normal density with mean zero and volatility $\sigma^Q = \sqrt{\tau}\sigma$; $Pr^{\mathbb{Q}}(b_{t+i})$ is the probability function of a break in the mean dividend growth rate, which is a Bernoulli distribution with parameter $\pi_{t+i}^Q = A_{t+i}/B_{t+i}$ (where A_{t+i} is defined in equation (9)); and $\varrho^{\mathbb{Q}}(g_{t+i}|b_{t+i} = 1)$ is a uniform density with lower and upper bounds g_l and g_u, respectively, which is used to obtain a new value of g_{t+i} in the case of a break.

Proof: Internet Appendix B.

Equation (10) shows that the option price is obtained by integrating the option's payoff over risk-neutral density $f^Q(S_{t+\tau})$. The prices of the stock, the bond, and the COS security are explicitly considered in equation (10), either in the option's payoff or in the risk-neutral density. Notably, the market is incomplete without COS securities in the economy, as mentioned earlier. In an incomplete market, the no-arbitrage conditions do not pin down a unique state price density (hereafter referred to as *SPD*), which can take

infinite forms. In this case, the *SPD* is not unique, which also implies that there are multiple risk-neutral ($\mathbb{Q}$) probability measures. For example, the exclusion of COS securities from our model would cause the probability of a break under the risk-neutral measure, π_t^Q, to take an infinite number of values.

The potential set of forms of the *SPD* can be reduced in an incomplete market by imposing bounds to rule out "good deals" (i.e., by reducing the profitability of investments) based on either the Sharpe ratio (see Cochrane and Saa-Requejo, 2000) or the gain-loss ratio (see Bernardo and Ledoit, 2000). However, the *SPD* is still defined in a continuum set under bounds to rule out "good deals". For instance, the price bounds of a European put option contract are given by $\left(\inf_{\mathbb{Q} \in Q^*} E[mX], \sup_{\mathbb{Q} \in Q^*} E[mX] \right)$, where X reflects the future option's payoff and Q^* is a subset of risk-neutral probability measures in a setup where "good deals" are ruled out.

In contrast, in a complete market, there exists a unique *SPD* (and thus, a unique $\mathbb{Q}$ probability measure), which means that the price bounds mentioned above coincide (i.e., $\inf_{\mathbb{Q} \in Q^*} E[mX] = \sup_{\mathbb{Q} \in Q^*} E[mX]$), since Q^* is a singleton. Thus, in our model, where markets are complete (thanks to COS securities), the probability of a break under the risk-neutral measure is defined by $\pi_t^Q = A_t/B_t$, where A_t and B_t are the prices of the COS security and the bond, respectively, defined in Proposition I, because the expected value of the COS security is the same under both the $\mathbb{P}$ and $\mathbb{Q}$ probability measures (i.e., $A_t^P = A_t^Q$). Under the $\mathbb{P}$ probability measure, A_t is defined by equation (6) in Proposition I, whereas under the $\mathbb{Q}$ probability measure, $A_t^Q = \pi_t^Q/(1 + r_t)$; thus, $\pi_t^Q = A_t/B_t$.

We cannot obtain a closed-form solution to equation (10) in Proposition II to calculate the prices of put options, since there are breaks in the economy. In the special case of no breaks and full information, option prices can be obtained through Rubinstein's (1976) discretised version of the Black-Scholes formula. However, when breaks are present, the underlying process of the dividends is non-stationary (i.e., the drift in equation (1) changes over time), and the interest rate is not constant. Thus, we use a numerical method, which we explain in Section 3. Moreover, in Section 3, we describe the additional complexity that the numerical method faces in an economy with partial information and Bayesian learning, which we develop in Section 2.2.

2.2   An economy under partial information and rational Bayesian learning

We now relax the full information assumption by assuming that g_t is unknown. Nevertheless, the agent observes the dividends paid by the stock and uses them to learn about the new value of g_t after a break. Note that historical information (before a break) does not help the agent to learn the new post-break value of g_t. In particular, the signals used to learn g_t are the n historical dividend returns, $\{D_i/D_{i-1}\}_{i=t-n}^{t}$, assuming that $n+1$ periods have passed since the most recent break. Moreover, the agent knows that dividends follow a geometric random walk, as in equation (1); thus, she knows that signals are noisy and only partially reveal the true value of g_t. Consequently, after each break, the agent has access to a short historical dataset that does not allow her to accurately estimate the unknown parameter. As more signals are received, the learning process increases the accuracy of the parameter estimation.

Similarly to Timmermann (2001), we assume that the agent does not know the future dates of breaks *ex ante* but recognises breaks as soon as they occur (the agent recognises the timing of the breaks), which allows us to isolate the impact of one source of learning (i.e., the process of learning g_t) on the put option return, rather than considering simultaneous learning about both the timing of the break and the value of g_t. Nevertheless, as a robustness check, in Internet Appendix C, we present an extension of the model where the representative agent knows neither the timing of the breaks nor the value of g_t. We find that the presence of both effects yields qualitatively similar results to the case where the agent learns only about g_t.

We assume that the representative agent uses the available information efficiently by updating her beliefs through a rational Bayesian updating procedure. We express the Bayesian updating process in terms of μ_t rather than g_t (given the relation between μ_t and g_t: $1 + g_t = \exp(\mu_t + \sigma^2/2)$). We do so because the signals used to learn about this parameter are the log returns of dividends, $\{\ln(D_i/D_{i-1})\}_{i=t-n}^{t}$, which are normally distributed (see equation (1)); thus, using μ_t makes the model mathematically simpler.

Under Bayesian learning, the representative agent views the unknown parameter μ_t as a random variable. The learning process starts with a prior belief about the density of μ_t, $f^{\mathbb{P}}(\mu_t)$, in the physical world $\mathbb{P}$. The prior density $f^{\mathbb{P}}(\mu_t)$ is the density function of μ_t, from which its new value is drawn after a break. Given that the new post-break value of g_t is drawn from a uniform density $\varrho^{\mathbb{P}}(g_t) = 1/(g_u - g_d)$, the new post-break value of μ_t has the following probability density: $f^{\mathbb{P}}(\mu_t) = \exp(\mu_t + \sigma^2/2)/(g_u - g_d)$, where $\mu_d = \ln(1 + g_d) - \sigma^2/2$ and $\mu_u = \ln(1 + g_u) - \sigma^2/2$.

The Bayesian agent recursively updates her prior belief as new information is received, thus obtaining a posterior belief regarding the density of μ_t, $f^{\mathbb{P}}(\mu_t|\boldsymbol{\xi}_t)$, through Bayes' rule:

$$f^{\mathbb{P}}(\mu_t|\boldsymbol{\xi}_t) = \frac{f^{\mathbb{P}}(\boldsymbol{\xi}_t|\mu_t)f^{\mathbb{P}}(\mu_t)}{f^{\mathbb{P}}(\boldsymbol{\xi}_t)}, \tag{12}$$

where $\boldsymbol{\xi}_t = [\ln(D_t/D_{t-1})\dots\ln(D_{t-n}/D_{t-n-1})]$ is the vector of n historical signals used to learn about μ_t from the most recent break, and $f^{\mathbb{P}}(\boldsymbol{\xi}_t|\mu_t)$ is the sample likelihood function:

$$f^{\mathbb{P}}(\boldsymbol{\xi}_t|\mu_t) = \frac{1}{\sqrt{2\pi(\sigma^2/n)}}\exp\left(\frac{-(\bar{\bar{\xi}}_t - \mu_t)^2}{2(\sigma^2/n)}\right), \tag{13}$$

which is a normal probability density function with mean $\bar{\bar{\xi}}_t = (1/n)\sum_{i=t-n+1}^{t}\xi_i$ and variance σ^2/n, since the agent knows that historical signals follow the geometric random walk as described in equation (1), where the most recent break happened $n+1$ periods ago. Thus, we can rewrite equation (12) as:

$$f^{\mathbb{P}}(\mu_t|\boldsymbol{\xi}_t) = \frac{f^{\mathbb{P}}(\boldsymbol{\xi}_t|\mu_t)f^{\mathbb{P}}(\mu_t)}{\int_{\mu_d}^{\mu_u} f^{\mathbb{P}}(\boldsymbol{\xi}_t|\mu_t)f^{\mathbb{P}}(\mu_t)d\mu_t}, \tag{14}$$

given that $f^{\mathbb{P}}(\boldsymbol{\xi}_t) = \int_{\mu_d}^{\mu_u} f^{\mathbb{P}}(\boldsymbol{\xi}_t|\mu_t)f^{\mathbb{P}}(\mu_t)d\mu_t$. Let $\lambda_t^{BL}(\mu_t)$ be the price of any asset under partial information and Bayesian learning (e.g., $\lambda_t^{BL}(\mu_t)$ can be the Bayesian value of the stock price, $S_t^{BL}(\mu_t)$, the bond price, $B_t^{BL}(\mu_t)$, or the price of put option contracts, $P_t^{BL}(K,\tau)$). Then, by using equation (14), the expected value of $\lambda_t^{BL}(\mu_t)$ can be obtained as follows:

$$\lambda_t^{BL} = \int_{\mu_d}^{\mu_u} \lambda_t^{full}(\mu_t)f^{\mathbb{P}}(\mu_t|\boldsymbol{\xi}_t)d\mu_t = \frac{\int_{\mu_d}^{\mu_u} \lambda_t^{full}(\mu_t)f^{\mathbb{P}}(\boldsymbol{\xi}_t|\mu_t)f^{\mathbb{P}}(\mu_t)d\mu_t}{\int_{\mu_d}^{\mu_u} f^{\mathbb{P}}(\boldsymbol{\xi}_t|\mu_t)f^{\mathbb{P}}(\mu_t)d\mu_t}, \tag{15}$$

where $\lambda_t^{full}(\mu_t)$ represents the value of the asset under full information (i.e., the asset prices under full information as defined in Proposition I and Proposition II), which depends on the unknown parameter μ_t.

2.3 Option returns under Bayesian learning

The expected hold-to-maturity return of a put option contract under full information is:

$$R_{t+\tau}^p = \frac{E_t^P[\max(K - S_{t+\tau}, 0)]}{P_t(K, \tau)} - 1 = \frac{E_t^P[\max(K - S_{t+\tau}, 0)]}{E_t^Q[e^{-r\tau}\max(K - S_{t+\tau}, 0)]} - 1. \tag{16}$$

Under Bayesian learning, the expected hold-to-maturity put option return is:

$$R_{t+\tau}^{p,BL} = \frac{E_t^{P,BL}[\max(K - S_{t+\tau}, 0)|\xi_t]}{E_t^{Q,BL}[e^{-r\tau}\max(K - S_{t+\tau}, 0)|\xi_t]} - 1. \tag{17}$$

The numerators of equation (16) and equation (17) are obtained under the physical probability measure $\mathbb{P}$, whereas their denominators are obtained under risk-neutral probability measure $\mathbb{Q}$. Consequently, differences between the $\mathbb{P}$ and $\mathbb{Q}$ probability measures affect the level of the expected hold-to-maturity put option returns. Under full information, the risk-neutral probability measure $f^{\mathbb{Q}}(S_{t+1})$ is obtained from the physical probability measure $f^{\mathbb{P}}(S_{t+1})$ as follows:

$$f^{\mathbb{Q}}(S_{t+1}) = \frac{m_{t+1}f^{\mathbb{P}}(S_{t+1})}{E_t[m_{t+1}]}, \tag{18}$$

where m_{t+1} is the stochastic discount factor. However, under partial information and Bayesian learning, the $\mathbb{P}$ probability measure is conditional on the information received after each break. Hence, $f^{\mathbb{P},BL}(S_{t+1}|\xi_t)$ can be written as:

$$f^{\mathbb{P},BL}(S_{t+1}|\xi_t) = \int_{\mu_d}^{\mu_u} f^{\mathbb{P}}(S_{t+1}|\xi_t,\mu_t) f^{\mathbb{P}}(\mu_t|\xi_t)\,d\mu_t. \tag{19}$$

Here, $f^{\mathbb{P}}(\mu_t|\xi_t)$ is the posterior belief of the Bayesian agent defined in equation (14). Equation (19) shows that in a learning environment, the conditional distribution of potential estimated values of the unknown parameter, given the information set ξ_t, affects the $\mathbb{P}$ probability measure of the future stock price.

The risk-neutral probability measure is also affected by the Bayesian learning process, as it is also conditional on the signals received. Accordingly, in the case of Bayesian learning, the risk-neutral probability measure $f^{\mathbb{Q},BL}(S_{t+1})$ is calculated as:

$$f^{\mathbb{Q},BL}(S_{t+1}|\xi_t) = \int_{\mu_d}^{\mu_u} f^{\mathbb{Q}}(S_{t+1}|\xi_t,\mu_t) f^{\mathbb{P}}(\mu_t|\xi_t)\,d\mu_t = \int_{\mu_d}^{\mu_u} \frac{m_{t+1}f^{\mathbb{P}}(S_{t+1}|\xi_t,\mu_t)}{E_t[m_{t+1}]} f^{\mathbb{P}}(\mu_t|\xi_t)d\mu_t. \tag{20}$$

Equation (20) shows that learning can generate a gap between the $\mathbb{P}$ and $\mathbb{Q}$ probability measures when changes in the parameter estimations (given the information available at t) are statistically associated with the agent's opinion regarding the dynamics of the stochastic discount factor. In fact, this is the case in our model because a change in the perception of the dividend drift induces a change in the agent's view about how the future stochastic discount factor evolves over time. Moreover, the gap between the $\mathbb{P}$ and $\mathbb{Q}$ probability measures is dynamic and depends on the number of signals received after each break. Consequently, learning should affect option returns (see equation (17)), given that the Bayesian learning process modifies the wedge between the $\mathbb{P}$ and $\mathbb{Q}$ probability measures, as described by equations (19) and (20).

2.4   Properties of the model that affect the learning process

In terms of the properties of the model, we analyse three key elements that affect the agent's learning process: (i) the presence of breaks in the mean dividend growth rate, (ii) the noise in the signals, and (iii) the representative agent's relative risk aversion. First, in terms of the presence of breaks, suppose that there are no breaks in the mean dividend growth rate (i.e., $\pi = 0$). In this case, as $t \to \infty$, the agent can use an infinite number of observations to learn about the unknown parameter (i.e., $n \to \infty$, since the value of n is never reset); thus, the estimated parameter value converges towards the true value. In this case, learning effects on option pricing vanish asymptotically, as in Guidolin and Timmermann (2003), and asset prices gradually converge towards those obtained under full information (as in Propositions I and II).

Conversely, when $0 < \pi < 1$, learning never ends, even asymptotically, because the learning process is reinitiated after a break in g_t, which can be seen in equation (13), where the number of signals n is immediately reset to zero after a break. Under the learning process in our model, the estimation error of the unknown parameter is proportional to $1/n$ (see equation (13)). Thus, the inaccuracy of the parameter estimation is reduced as more information is received (i.e., when the number of signals n increases). Note that the market is complete under partial information and learning, just as it is under full information, because the inaccuracy of the parameter estimation decreases following a deterministic path with a reduction factor of order $1/n$, as more signals are observed and processed via the learning mechanism.

Second, the signals used to learn the unknown parameter (i.e., log dividend returns) are noisy. The accuracy of the parameter estimation depends on the number of signals n as well as on the variance σ^2 of the dividend process (i.e., historical signals, $\{\ln(D_i/D_{i-1})\}_{i=t-n}^{t}$, are normally distributed with variance σ^2). Thus, the value of σ controls the level of noise contained in the historical information used by the agent to learn about the unknown parameter and, consequently, the speed at which the agent learns. For instance, suppose that $\sigma \to \infty$. In this case, the agent cannot learn from the information received (even asymptotically), since the signals are excessively noisy. Thus, the agent can use only the expected value of the parameter based on her prior belief. In contrast, when $\sigma = 0$ (i.e., the signals are not noisy at all), learning stops after the first signal is received, since signals are constant over time. Thus, μ_t equals any element of the vector of historical signals $\xi_t = [\ln(D_t/D_{t-1}) \ldots \ln(D_{t-n}/D_{t-n-1})]$. In this paper, we assume a value of σ based on market data (which is explained in the following section). As a robustness check, we use a different level of σ in the Internet Appendix C to analyse the effect of a change in this parameter on put option returns.

The coefficient of relative risk aversion, η, also plays an important role in the learning process. For example, learning does not affect the value of the stock price when $\eta = 1$, as documented in Guidolin and Timmermann (2003), because the expression $(1 + g_t)^{1-\eta}$ in the stock price equals one (see equation (7)). Thus, in this case, when the Bayesian updating process is used (i.e., through equation (15)), learning has no effect on the stock price because its value no longer depends on g_t (i.e., the unknown parameter). However, learning is still present in the bond price, thus affecting the price and the return of put option contracts. In the implementation of the model, we use different levels of

relative risk aversion to evaluate the effect of a change in the level of η on put option returns.

3   Implementation of the model

We use a simulation analysis based on the model presented in Section 2, with different parameter setups, to analyse the effects of learning on put option returns. We perform 10,000 simulations per combination of parameters. In each of these simulations, we generate 12 years (3,024 trading days) of daily dividends, which are the signals that the representative agent observes and uses to learn about g_t. Thus, our simulation set represents $10{,}000{\times}3{,}024 = 30{,}240{,}000$ simulated trading days for any given choice of parameter values.

We simulate daily dividends using two nested stochastic processes. First, we use the dividends' geometric random-walk process, $\ln(D_{t+1}/D_t) = \mu_t + \sigma\varepsilon_t$, to simulate the time series of 12 years of daily dividends. Second, in each 12-year time series, we generate random breaks in g_t (and thus breaks in μ_t), which affect the simulation of dividends by the random-walk process. In the generation of breaks, the times between breaks follow a geometric distribution with parameter π, and when each break occurs, a new value for g_t is drawn from a uniform distribution $g_t \sim G(\cdot)$.

Prices S_t and B_t on each of the 3,024 trading days (and in each of the 10,000 simulations) are obtained using equation (15), calculated by means of numerical integration through the adaptive Simpson quadrature. Between breaks, A_t is calculated using equation (9), since the COS security's price does not depend on g_t and thus is not

affected by learning. When there is a break in the economy, A_t is equal to one (as defined in the model). European put option prices are also calculated using equation (15). Thus, we can write the following equation for the price of a put option contract:

$$P_t^{BL}(K,\tau) = \frac{\int_{\mu_d}^{\mu_u}\left[\int_0^\infty \frac{1}{1+r_{t,t+\tau}}\max\{K-S_{t+\tau},0\}f^{\mathbb{Q}}(S_{t+\tau})\,dS_{t+\tau}\right]f^{\mathbb{P}}(\boldsymbol{\xi}_t|\mu_t)f^{\mathbb{P}}(\mu_t)d\mu_t}{\int_{\mu_d}^{\mu_u}f^{\mathbb{P}}(\boldsymbol{\xi}_t|\mu_t)\,f^{\mathbb{P}}(\mu_t)d\mu_t}. \qquad (21)$$

European put option prices are calculated using equation (21) on a monthly basis to obtain non-overlapping one-month option returns. In each month of the 12-year simulated period (and for each of the 10,000 simulations), equation (21) is solved through a two-step procedure. As a first step, we address the internal integral in equation (21), which contains the max($\cdot$) function. This integral is solved by Monte Carlo simulation, in which $j \in \{1,2,\dots,J\}$ is a simulated independent path of the stock price, where $J = 20,000$.

Each of the 20,000 paths in the Monte Carlo simulation is generated using the risk-neutral density, which is decomposed into a sequence of one-period risk-neutral probabilities. Thus, each path is generated through the repetition of several single-period steps. Suppose that we generate path j, in which the first break (after current time t) occurs in period $t + m$, where $m \leq \tau$. Thus, between periods t and $t + m - 1$, the dividend drift is equal to the current unknown value at time t, μ_t. In this path j, the Bayesian agent learns about μ_t from the set of signals $\boldsymbol{\xi}_t$. Since we also have to integrate with respect to μ_t in the exterior integral (due to the Bayesian learning process used to learn μ_t) in equation (21), we leave the expression generated in path j as a symbolic expression that depends on μ_t between periods t and $t + m - 1$. After period $t + m$, the new value of the dividend drift does not depend on the set of signals $\boldsymbol{\xi}_t$ (which is also the case if there are new breaks

between periods $t + m$ and τ). Thus, we obtain an analytical expression for the option payoff, which depends on μ_t in each path of the Monte Carlo simulation.

In the second step, we integrate the analytical expression for the option payoff with respect to μ_t (which is obtained in each path in the first step of the procedure) using numerical integration through the adaptive Simpson quadrature. Thus, in this second step, we solve the exterior integral that depends on μ_t in each of the Monte Carlo paths. We then average the outcomes from the 20,000 paths, thereby obtaining the put option price for a given simulated month. The denominator of equation (21) is also solved through numerical integration using the adaptive Simpson quadrature.

Notably, in a model under full information (i.e., when there is no learning, as in Section 2.1), the two-step procedure is not required to solve equation (10). Under full information, the integral in equation (10) is solved on a monthly basis using standard Monte Carlo simulation, again with 20,000 independent paths. Each of the 20,000 paths in this Monte Carlo simulation is generated by means of the risk-neutral density. Then, we average the outcomes from the 20,000 paths to obtain the simulated put option price for a given month.

In terms of the parameterisation of the model, we consider the following parameter setup. Since the new post-break mean dividend growth rate is taken from uniform distribution $G(\cdot)$, we assume that its lower and upper bounds are $g_l = -0.126\%$ and $g_u = 0.705\%$, respectively, on a monthly basis. The values of g_l and g_u are consistent with the real dividend growth rate of the S&P 500 index during the period of our empirical analysis. For example, the annual dividend log returns of the S&P 500 index in real terms between

1996 and 2007 (taken from Robert Shiller's database) have a mean and standard deviation of 3.39% and 5.27%, respectively. Thus, the interval with centre 3.39%, plus and minus 5.27%, is [-1.88%, 8.66%] on an annual basis, which is equivalent to the interval [-0.16%, 0.69%] on a monthly basis. This interval, based on market data, is close to the assumed lower and upper bounds of the uniform distribution $G(\cdot)$.

We set the rate of impatience, ρ, at 0.713% on a monthly basis or, equivalently, 8.900% on a yearly basis, because the rate of impatience is constrained in the model by $1 + \rho > (1 - \pi)(1 + g_u)^{1-\alpha}$ in order to obtain positive stock prices (see equation (7)), where the value of g_u was set in the previous paragraph. The rate of impatience is high in relation to the real interest rates observed during the period 1996–2007, but as mentioned above, we need this level of ρ to obtain positive stock prices. We could calibrate ρ using market data and then adjust g_u using $1 + \rho > (1 - \pi)(1 + g_u)^{1-\alpha}$. However, we prefer to select g_u and then adjust ρ (thus sacrificing the accuracy of its calibration to some degree), since g_t is the variable that must be learned by the Bayesian agent in our model. While our model does not match real interest rates closely, it is important to note that our focus is on providing a simple model of a learning environment, with the objective of offering a potential explanation for the abnormal put option returns rather than perfectly calibrating all variables in the economy.

The volatility of the geometric random walk is set at 1.44% on a monthly basis (i.e., 5.00% on an annual basis), which is also consistent with market data (the standard deviation between 1996 and 2007 of real dividend log returns on the S&P 500 index was 5.27%). For the coefficient of relative risk aversion, η, we use levels of 0.2, 0.5, and 5.0. In an attempt to reproduce reality, we obtain the probability of breaks, π, using a dynamic test

for structural breaks as proposed by Chu *et al.* (1996) with data on daily real dividends from the S&P 500 index during the period 1996–2007 (see Internet Appendix A). We detect eight breaks in the mean dividend growth rate over the 3,024 trading days of the 12 years analysed; thus, we set π at 0.056 (on a monthly basis).

4 Results

4.1 Index put option returns

We calculate returns obtained from a hold-to-maturity naked trading strategy. In each 12-year simulation, following Broadie *et al.* (2009) and Chambers *et al.* (2014), we generate time series of one-month non-overlapping put option returns, $r_{\tau+T}^p$, given by:

$$r_{t+\tau}^p = \frac{\max(K - S_{t+\tau}, 0)}{P_t(K, \tau)} - 1, \tag{22}$$

where $P_t(K, \tau)$ is the put option price. Figure 1 presents the results of one such 12-year simulation under three scenarios: full information with no breaks, full information with breaks, and partial information with breaks and learning. The first two scenarios are special cases of the model presented in Section 2. In the scenario under full information with no breaks, our model is equivalent to Rubinstein's (1976) discrete version of the Black-Scholes model. In the scenario under full information with breaks, we calculate put option prices using Proposition II. The third scenario is the full model presented in Section 2, which represents an economy with breaks, partial information and learning.

[Insert Figure 1 here]

Figure 1 reports option returns with a strike-to-price ratio of 1.00. In this simulation, the coefficient of relative risk aversion is set at 0.5. Figure 1 shows that in all three scenarios, a naked investment strategy generates put option returns that are far from normal, reflecting high levels of skewness and kurtosis. Interestingly, under partial information and learning, put option returns take negative values more frequently than they do under full information (with and without breaks).

Table 1 reports summary statistics of empirical and simulated put option returns. This table reports the average values of one-month hold-to-maturity put option returns for different moneyness levels (ranging from 0.96 to 1.02). We focus on this particular range of moneyness levels because most options' trading activity occurs within that range (see Broadie *et al.*, 2009).

[Insert Table 1 here]

Empirical put option returns, presented in Panel A of Table 1, are calculated using S&P 500 European put index options obtained from the OptionMetrics Ivy DB database spanning 1996–2007. We exclude option prices that violate arbitrage conditions, have an ask price that is lower than the bid price, have a bid price of zero, and/or have no option open interest (see Bernales and Guidolin, 2014). We obtain option returns for fixed moneyness and time-to-maturity levels by interpolating linearly across the returns of the four S&P 500 put option contracts that surround the required values of moneyness and time to maturity.

In Panels B-D of Table 1, as in Figure 1, simulations are based on an economy under three scenarios: full information with no breaks, full information with breaks, and partial

information with breaks and learning. We use three values for the coefficient of relative risk aversion ($\eta = 0.2, 0.5,$ and 5.0 in Panels B, C, and D, respectively).

Panel A of Table 1 shows that empirical put option returns are negative and significantly different from zero. In addition, the absolute values of empirical option returns decrease as the moneyness level increases, which is in line with the behaviour of put option returns derived in Coval and Shumway (2001). For instance, an S&P 500 put option contract with a ratio of $K/S = 0.96$ ($K/S = 1.00$) has an average return of -0.69 (-0.27).

Panels B-D of Table 1 report that simulated put option returns are also negative and that their absolute values decrease as the moneyness level increases, which is observed for all scenarios and coefficients of relative risk aversion. However, the simulated results obtained under the two scenarios without learning (i.e., full information with no breaks and full information with breaks) stand in sharp contrast to the empirical results presented in Panel A. The simulated put option returns with no learning are smaller (and in general not significant) than the empirical put option returns (which are significant with t-statistics of at least -4.27).

In contrast to the scenarios without learning, simulated put option returns are large in absolute value when there is learning in the economy (i.e., the last four columns on the right-hand side of Table 1). For example, a one-month-to-maturity put option with $K/S = 1.00$ has an average monthly return of $-0.55, -0.38,$ and -0.93 for coefficients of relative risk aversion 0.2, 0.5, and 5.0, respectively. In addition, the simulated put option returns

under learning are generally significant, similar to the observed put option returns of S&P 500 option contracts presented in Panel A.

Importantly, the difference in the levels and *t*-statistics of simulated put option returns marginally increases from the "full information, no breaks" scenario to the "full information with breaks" scenario (signalling a small effect of breaks on the returns on put option contracts). However, very few of the simulations for these first two scenarios have significant values in relation to the simulations under information with breaks and learning, which suggests that learning matters more than breaks in generating large returns on put option contracts.

4.2   Abnormal put option returns under the CAPM

There is empirical evidence that put option returns are too high to be explained by the capital asset pricing model (CAPM) – see, e.g., Bondarenko, 2003. Thus, we analyse whether our model generates put option returns that are also too high relative to the CAPM. In particular, we estimate the CAPM α, which reflects whether the expected return of an investment is abnormal relative to the expected return obtained by the CAPM. To this end, we run the following regression:

$$\left(r_{Opt} - r_F\right) = \alpha + \beta\left(r_m - r_F\right) + \varepsilon, \tag{23}$$

where r_{Opt} is the return on put option contracts, r_m is the return of the market portfolio, and r_F is the risk-free rate. Table 2 reports the average values of the CAPM α estimated from the empirical and simulated put option returns. Panel A reports the estimates of the CAPM α obtained from empirical option data. In this case, r_{Opt} is obtained from the S&P

500 put option contracts, r_F is the one-month LIBOR rate, and r_m is the S&P 500 index return.

[Insert Table 2 here]

Panels B-D of Table 2 report the estimated values of the CAPM α from simulated put option returns under the same three scenarios used in Figure 1 and Table 1 (i.e., full information with no breaks, full information with breaks, and partial information with breaks and learning). In the simulations generated by our model, r_{Opt} is calculated through option prices obtained from our model using equation (22). The risk-free rate, r_F, is obtained from the bond price (i.e., $r_F = 1/B - 1$). In the scenario without breaks, r_m is the market portfolio formed by the stock. In the scenarios with breaks, the market portfolio is formed by the stock and the COS security, in which the weight of the stock, ω_S, and the weight of the COS security, ω_A, are given by the tangency portfolio lying on the efficient frontier. Thus, $\omega_S = (\sigma_A^2 E(r_S - r_F) - \sigma_{S,A} E(r_A - r_F)/(\sigma_A^2 E(r_S - r_F) - \sigma_{S,A} E(r_A - r_F) + \sigma_S^2 E(r_A - r_F) - \sigma_{S,A} E(r_S - r_F))$ and $\omega_A = 1 - \omega_S$, where $E(r_S)$ and $E(r_A)$ are the expected returns of the stock and the COS security, respectively. The values of σ_S^2 and σ_A^2 are the variances of the returns on the stock and the COS security, respectively, while $\sigma_{S,A}$ is the covariance between these two assets. Expected returns, variances and the covariance are calculated in each simulation using daily simulated data. In Internet Appendix C, we also consider a scenario under partial information with breaks and learning in which the COS security is not included in the market portfolio when calculating the value of the CAPM α (despite the fact that COS security is part of the economy). Our results scarcely change when we omit the COS security because its weight is low in the market portfolio.

Panel A of Table 2 shows that the estimated CAPM α values from the empirical data are negative, with large and statistically significant absolute values. The absolute values of the CAPM α decrease as the moneyness level increases. In terms of simulated option returns, Panels B-D of Table 2 report that in an environment without learning (under full information with no breaks and full information with breaks), the simulated results of the CAPM α are negative but lower in absolute value than the empirical results presented in Panel A of Table 2. Most importantly, in the scenarios without learning (see Panels B-D), the simulated put option returns are generally not abnormal, since few of the simulations for either scenario have significant CAPM α values. Thus, these results suggest that the abnormal put option returns observed in S&P 500 option contracts cannot be explained by models in which no learning process exists.

Panels B-D of Table 2 show that in the scenario with learning, the simulated values of the CAPM α are negative and larger in absolute value than those in the other two simulated scenarios, where there is no learning in the economy. In addition, in the scenario with learning, the simulated values of the CAPM α are abnormal since their values are in most cases significant, similar to the empirical results of the CAPM α reported in Panel A.

We can observe in Table 2 that the simulated results of the CAPM α generated by our model in the scenario with learning are not exactly the same as those empirically observed in the S&P 500 option contracts. However, it is important to note that the main purpose of our study is to provide a potential explanation for the abnormal returns observed on put option contracts by using a simple model under a learning environment. We do not seek to fully describe all features of the option market, since the option pricing process could be affected by elements outside the scope of our model, such as jump risk

(e.g., Broadie *et al.*, 2009; Constantinides *et al.*, 2013) and individual investors' net option demand (e.g., Bollen and Whaley, 2004; Gârleanu *et al.*, 2009). Therefore, the objective of our study is not to propose a better model for option pricing but to suggest that the learning process followed by investors may be one of the reasons for the anomalous high returns of put option contracts.

4.3   A multifactor analysis of option returns

In this section, we analyse whether put option returns are related to other factors in addition to the market factor from the CAPM. In particular, we use the volatility risk premium and the price-to-dividend ratio (S/D). The volatility risk premium is the difference between the volatilities of the $\mathbb{Q}$ and $\mathbb{P}$ probability measures. As explained in equation (17), the volatility risk premium is a natural factor for describing option returns in an economy with learning, as it is a proxy for the differences between the $\mathbb{Q}$ and $\mathbb{P}$ probability measures. Regarding the use of the price-to-dividend ratio, there is theoretical evidence that this measure can explain stock returns when agents follow a learning process, which may also affect option returns (for an overview of the literature on the effect of learning on the relation between the price-to-dividend ratio and stock returns, see Pastor and Veronesi, 2009).

In line with Bollerslev *et al.* (2009), we calculate the volatility risk premium as the difference between the option's implied volatility (IV) and the realised volatility (RV). In the case of the empirical analysis using S&P 500 option contracts, we calculate the IV through the Black-Scholes model with at-the-money one-month-to-maturity put option

contracts. We use linear interpolation based on the *IV*s of the four contracts around the moneyness and the time to maturity required.

In the case of the simulated results, we obtain the *IV* (i.e., the volatility under the $\mathbb{Q}$ probability measure) numerically from the Monte Carlo simulation used to obtain the option prices each month (as explained in Section 3). Thus, the option's implied volatility is calculated as the annualised standard deviation of the simulated one-month stock returns under $\mathbb{Q}$ (namely, from the 20,000 Monte Carlo paths generated by the two-step procedure explained below equation (21)).

We calculate the *RV* as the annualised standard deviation of the daily stock log returns over each month to avoid overlapping periods. For the empirical results, we use daily S&P 500 index data; for the theoretical results, we calculate the *RV* using simulated daily stock log returns.

We explore whether the computed factors explain the empirical and simulated option returns by using a naked strategy (as in Tables 1 and 2) and a straddle portfolio. We also include the straddle portfolio because it is not affected by price changes in the underlying stock. The analysis of option returns using the straddle strategy is important as a robustness check, since the agent's learning process simultaneously affects the option contracts and the stock prices. Thus, one may conjecture that the effect of learning on option returns is mainly driven by changes in the underlying stock. However, the straddle portfolio allows us to isolate the impact of learning on option returns because this portfolio is free of risk derived from changes in the underlying stock price.

The straddle portfolio is formed by buying one European put option contract and one European call option contract with the same moneyness and time to maturity. The call option price is calculated using the same procedure that is used for the put option price (as explained in equation (21)); however, the call option payoff is used this time. We calculate only straddle portfolios with at-the-money contracts because no other straddle position represents a market-neutral strategy.

Table 3 presents results for the empirical and simulated option returns from the multifactor analysis. In our model, we use three coefficients of relative risk aversion ($\eta = 0.2$, $\eta = 0.5$, and $\eta = 5.0$ in Panels B, C, and D, respectively). Unlike Tables 1 and 2, Table 3 reports only the results for the scenario under partial information with breaks and learning, since the volatility risk premium ($IV - RV$) is equal to zero in the scenarios without learning. It is important to mention that the difference in volatility between the $\mathbb{Q}$ and $\mathbb{P}$ probability measures generated by our model is, on average, 6.9%, 3.4%, and 34.0% when the coefficient of relative risk aversion is set at 0.2, 0.5, and 5.0, respectively (see Internet Appendix D).

[Insert Table 3 here]

Table 3 reports that empirical and simulated straddle portfolios are abnormal under the CAPM (see the fifth column of results in the table), similar to the ones reported in Table 2 for a naked investment strategy under a learning environment. For example, the value of the CAPM α of the straddle portfolio using S&P 500 contracts (using our model with $\eta = 0.5$) is -0.08 (-0.33).

Most importantly, Table 3 shows that the volatility risk premium is a significant factor in explaining the empirical and simulated returns of both the naked investment strategy and the straddle portfolio. In addition, the values of the CAPM α estimated from empirical and simulated option returns, as well as their levels of significance, are strongly reduced in absolute value once a factor model that incorporates the volatility risk premium is used. This result further supports our study, as it highlights that learning generates differences between the $\mathbb{P}$ and $\mathbb{Q}$ probability measures, which affects put option returns (as explained in Section 2.3). Note that we show only the results of the volatility risk premium as a proxy for differences between the $\mathbb{P}$ and $\mathbb{Q}$ probability measures. However, learning generates differences in the entire distribution of the $\mathbb{P}$ and $\mathbb{Q}$ probability measures, including differences in volatility, skewness and kurtosis (see, e.g., David and Veronesi, 2002; Guidolin and Timmermann, 2003).

5   Conclusion

The economic sources that generate the empirically observed excessive put index option returns have not been studied extensively. We suggest that the learning process followed by investors may be one potential source. We present an equilibrium model under partial information about the mean dividend growth rate, in which a rational Bayesian learner prices put option contracts. We show that our model generates abnormal put option returns similar to those we compute from actual S&P 500 index option data. In addition, we document that this result is not obtained in an economy without learning. Our model is simple and intuitive and opens a number of potential avenues for future research. For instance, future research may address the impact of learning on the returns of other

derivative securities, the effect of learning in a setup with asymmetric information where informed agents reveal information to other market participants via their trading strategies, and the role of learning in an environment with liquidity shocks that also need to be learned.

References

Benzoni, L., P. Collin-Dufresne, and R. S. Goldstein, 2011, Explaining Asset Pricing Puzzles Associated with the 1987 Market Crash, *Journal of Financial Economics* 101, 552–573.

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

Bernardo, A. E., & Ledoit, O. (2000). Gain, loss, and asset pricing. *Journal of Political Economy* 108, 144–172.

Black, F., and M. Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–659.

Bollen, N., and R. Whaley, 2004, Does net buying pressure affect the shape of implied volatility functions? *Journal of Finance* 59, 711–53.

Bollerslev, T., Tauchen, G., Zhou, H., 2009. Expected stock returns and variance risk premia. *Review of Financial Studies* 22, 4463–4492.

Bondarenko, O. 2003. Why Are Put Options So Expensive? Working Paper, University of Illinois, Chicago.

Brandt, M., Q. Zeng, and L. Zhang, 2004, Equilibrium stock return dynamics under alternative rules of learning about hidden states. *Journal of Economics Dynamics and Control* 28, 1925–1954.

Broadie, M., M. Chernov, and M. Johannes, 2009, Understanding index option returns, *Review of Financial Studies* 22, 4493–4529.

Chambers, D. R., M. Foy, J. Liebner, and Q. Lu, 2014, Index option returns: Still puzzling, *Review of Financial Studies* 27, 1915–1928.

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

Cochrane, J. H., & Saa-Requejo, J., 2000, Beyond arbitrage: Good-deal asset price bounds in incomplete markets. *Journal of Political Economy*, 108(1), 79–119.

Constantinides, G.M., J.C. Jackwerth, and A. Savov, 2013, The puzzle of index option returns, *Review of Asset Pricing Studies*, 3, 229–257.

Coval, J., and T. Shumway, 2001, Expected option returns, *Journal of Finance* 56, 983–1009.

David, A., and P. Veronesi, 2002, Option prices with uncertain fundamentals, Working paper, University of Chicago.

Duffie, D., and C. Huang, 1985, Implementing Arrow–Debreu equilibria by continuous trading of few long lived securities, *Econometrica* 53, 1337–1356.

Garleanu, N., Pedersen, L. H., & Poteshman, A. M., 2009. Demand-based option pricing. *Review of Financial Studies* 22, 4259-4299.

Guidolin, M., 2006, High equity premia and crash fears. Rational foundations, *Economic Theory* 28, 693–708.

Guidolin, M., and A. Timmermann, 2003, Option prices under Bayesian learning: implied volatility dynamics and predictive densities, *Journal of Economic Dynamics and Control* 27, 717–769.

Guidolin, M., and A. Timmermann, 2007, Properties of equilibrium asset prices under alternative learning schemes, *Journal of Economic Dynamics and Control* 31, 161–217.

Guo, X., 2001, Information and option pricings. *Quantitative Finance* 1, 38–44.

Harrison, M. and Pliska, S., 1981, Martingales and stochastic integrals in the theory of continuous trading, *Stochastic Processes and Their Applications* 11, 215–260.

Jones, C., 2006, A nonlinear factor analysis of S&P 500 index option returns, *Journal of Finance* 41, 2325–2363.

Koop, G., and Potter, S. (2007), "Estimation and Forecasting in Models with Multiple Breaks," *Review of Economic Studies*, 74, 763–789.

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

Mamon, R. S. and Rodrigo, M. R. (2005). Explicit solutions to European options in a regime-switching economy. *Operations Research Letters*, 33(6):581–586.

Pastor, L., and P. Veronesi, 2009, "Learning in Financial Markets," *Annual Review of Financial Economics* 1, 361–381.

Pesaran, M. H., D. Pettenuzzo, and A. Timmermann (2006). Forecasting Time Series Subject to Multiple Structural Breaks. *Review of Economic Studies* 73, 1057–1084.

Rubinstein, M., 1976, The Valuation of Uncertain Income Streams and the Pricing of Options, *Bell Journal of Economics* 7, 407–425.

Shaliastovich, I., 2009, Learning, confidence and option prices, Working paper, University of Pennsylvania.

Timmermann, A., 1993, How learning in financial markets generates excess volatility and predictability of excess returns, *Quarterly Journal of Economics* 108, 1135–1145.

Timmermann, A., 1996, Excess volatility and predictability of stock prices in autoregressive dividend models with learning, *Review of Economic Studies* 63, 523–557.

Timmermann, A., 2001, Structural breaks, incomplete information, and stock prices, *Journal of Business & Economic Statistics* 19:3, 299–314.

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

Yuen, F.L., and H. Yang, 2010, Option pricing with regime switching by trinomial tree method, *Journal of Computational and Applied Mathematics* 233, 1821–1833.

Figure 1. Dynamics of put option returns under three scenarios: full information with no breaks, full information with breaks, and partial information with breaks and learning. This figure presents the time series of simulated put option returns (with a naked investment strategy) for one 12-year simulation under three scenarios: full information with no breaks, full information with breaks, and partial information with breaks and learning. The figure reports one-month hold-to-maturity returns with a strike-to-price ratio of 1.00. In this simulation, the coefficient of relative risk aversion is set at 0.5.

Table 1. Summary statistics of empirical and simulated put option returns. This table reports summary statistics of empirical and simulated returns on put option contracts using a naked investment strategy, R_{put}. The table reports average one-month hold-to-maturity option returns for non-overlapping intervals, with strike-to-price ratios ranging from 0.96 to 1.02 (with an increment of 0.02). Empirical option returns are obtained from S&P 500 option contracts between 1996 and 2007. Simulated option returns are obtained from simulat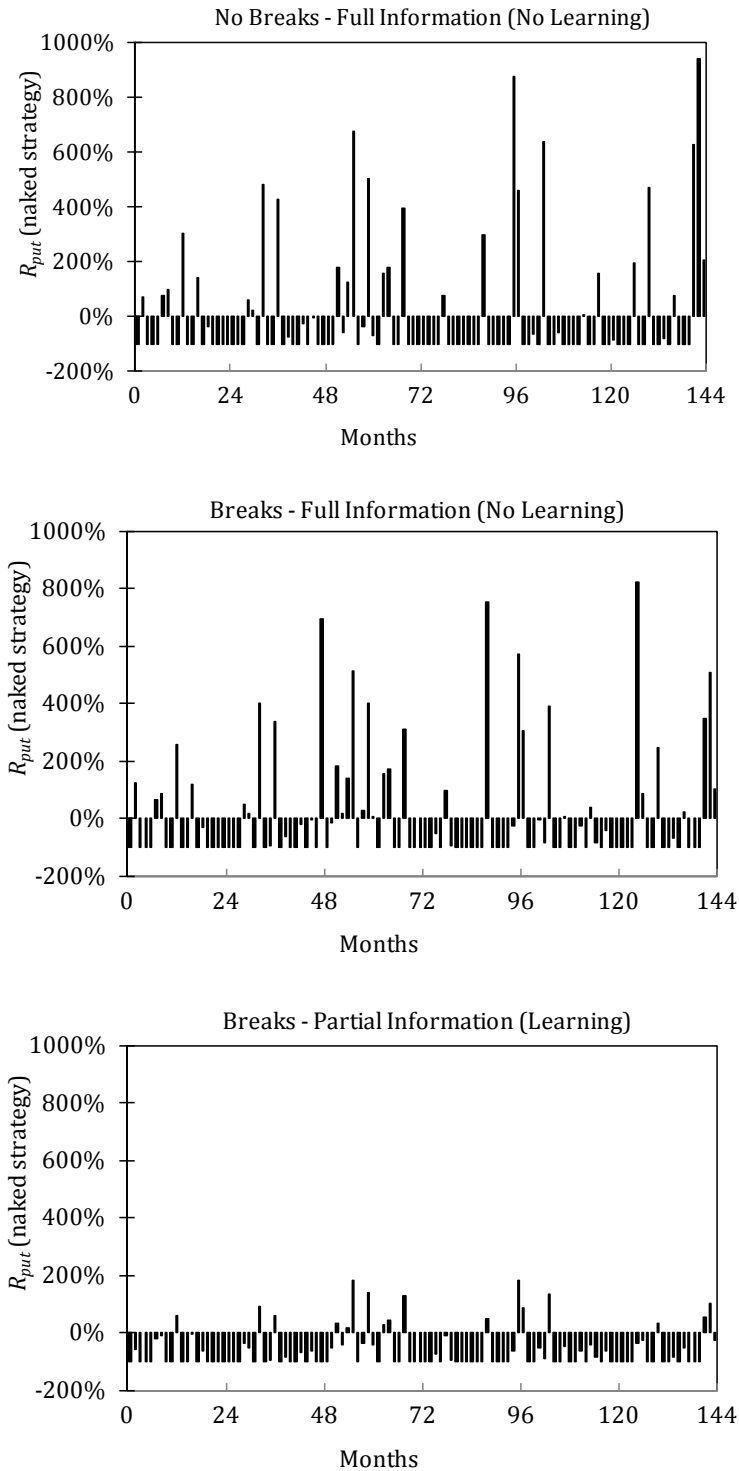ions of the model under three scenarios: full information with no breaks, full information with breaks, and partial information with breaks and learning. Simulations of the model are performed using three coefficients of relative risk aversion ($\eta = 0.2$, $\eta = 0.5$, and $\eta = 5.0$). The values for the simulated option returns are the averages over 10,000 simulations of 12 years each. The percentage of simulations with a significant mean return is reported in parentheses at 5% significance (since there is only one time series for empirical S&P 500 option returns, the value in parentheses can only be 0% or 100%).

K/S	0.96	0.98	1.00	1.02	0.96	0.98	1.00	1.02	0.96	0.98	1.00	1.02
Panel A. Empirical R_{put} (with S&P 500 Options)												
Mean					-0.69	-0.36	-0.27	-0.22				
t-stat					-10.64	-4.61	-4.59	-4.27				
p-value					0.00	0.00	0.00	0.00				
					(100%)	(100%)	(100%)	(100%)				
Volatility					1.39	1.74	1.36	1.08				
Skewnes					6.01	4.04	2.79	1.83				
Kurtosis					45.80	22.59	13.71	7.69				
Panel B. Simulated R_{put} with η =0.2												
	No Breaks - Full Inf. (No Learning)				Breaks - Full Inf. (No Learning)				Breaks - Partial Inf. (Learning)			
Mean	-0.23	-0.13	-0.09	-0.05	-0.29	-0.22	-0.14	-0.07	-0.96	-0.84	-0.55	-0.23
t-stat	-0.59	-0.62	-1.30	-1.27	-0.64	-0.81	-1.34	-1.41	-182.77	-18.73	-8.48	-3.85
p-value	0.59	0.57	0.26	0.21	0.51	0.39	0.18	0.17	0.00	0.00	0.00	0.01
	(11%)	(11%)	(34%)	(36%)	(11%)	(16%)	(35%)	(38%)	(100%)	(99%)	(99%)	(96%)
Volatility	8.90	4.95	1.62	0.79	11.56	6.56	2.30	1.10	1.03	0.84	0.76	0.63
Skewnes	10.36	7.14	2.11	0.49	14.62	9.18	3.29	0.67	10.85	4.99	2.03	0.64
Kurtosis	65.61	56.64	4.48	-0.32	86.06	76.55	6.90	-0.47	107.26	31.11	7.88	2.88
Panel C. Simulated R_{put} with η =0.5												
	No Breaks - Full Inf. (No Learning)				Breaks - Full Inf. (No Learning)				Breaks - Partial Inf. (Learning)			
Mean	-0.21	-0.11	-0.08	-0.04	-0.25	-0.15	-0.11	-0.04	-0.97	-0.73	-0.38	-0.10
t-stat	-0.57	-0.59	-1.15	-1.25	-0.64	-0.70	-1.18	-1.34	-110.70	-9.95	-4.45	-1.51
p-value	0.60	0.58	0.30	0.27	0.57	0.50	0.29	0.15	0.02	0.01	0.02	0.02
	(10%)	(11%)	(27%)	(31%)	(10%)	(15%)	(31%)	(36%)	(100%)	(99%)	(99%)	(96%)
Volatility	8.56	4.89	1.61	0.77	9.50	5.54	2.26	0.78	0.93	0.73	1.03	0.54
Skewnes	9.82	6.90	2.20	0.50	11.94	8.09	2.63	0.65	10.61	5.55	1.95	0.57
Kurtosis	70.28	59.76	4.61	-0.28	82.90	65.01	5.27	-0.37	116.46	37.38	6.91	2.74
Panel D. Simulated R_{put} with η =5.0												
	No Breaks - Full Inf. (No Learning)				Breaks - Full Inf. (No Learning)				Breaks - Partial Inf. (Learning)			
Mean	-0.25	-0.14	-0.10	-0.05	-0.42	-0.30	-0.16	-0.07	-1.00	-0.96	-0.93	-0.65
t-stat	-0.61	-0.72	-1.37	-1.43	-0.70	-0.78	-1.46	-1.49	-142.91	-146.08	-59.71	-32.16
p-value	0.49	0.54	0.19	0.16	0.47	0.55	0.16	0.13	0.00	0.00	0.01	0.00
	(12%)	(12%)	(37%)	(43%)	(10%)	(15%)	(39%)	(46%)	(100%)	(99%)	(99%)	(96%)
Volatility	9.13	4.98	1.62	0.73	13.05	8.43	2.65	1.14	0.81	0.61	0.30	0.22
Skewnes	11.30	7.25	2.17	0.52	17.05	12.29	3.43	0.93	9.41	9.33	7.20	4.18
Kurtosis	69.45	58.34	4.94	-0.24	104.17	89.00	8.30	-0.37	97.92	93.45	64.64	33.09

Table 2. Simulated and empirical values of the CAPM α of put option returns. This table displays the CAPM α values of empirical and simulated returns on put option contracts using a naked investment strategy, $\boldsymbol{R_{put}}$. The table reports average one-month hold-to-maturity option returns for non-overlapping intervals, with strike-to-price ratios ranging from 0.96 to 1.02 (with an increment of 0.02). Empirical option returns are obtained from S&P 500 option contracts between 1996 and 2007. Simulated option returns are obtained from simulations of the model under three scenarios: full information with no breaks, full information with breaks, and partial information with breaks and learning. Simulations of the model are performed using three coefficients of relative risk aversion ($\boldsymbol{\eta = 0.2}$, $\boldsymbol{\eta = 0.5}$, and $\boldsymbol{\eta = 5.0}$). The values for the simulated option returns are the averages over 10,000 simulations of 12 years each. This table presents averages of the alpha's t-statistics computed using Newey-West standard errors to correct for heteroscedasticity and serial correlation. The percentage of simulations with a significant CAPM α is reported in parentheses at 5% significance (since there is only one time series for the empirical S&P 500 option returns, the value in parentheses can only be 0% or 100%).

K/S	0.96	0.98	1.00	1.02	0.96	0.98	1.00	1.02	0.96	0.98	1.00	1.02
					Panel A. Dependent Variable: Empirical R_{put} (with S&P 500 Options)							
CAPM α					-0.59	-0.28	-0.15	-0.09				
t-stat					-10.78	-4.90	-4.28	-3.66				
p-value					0.00	0.00	0.00	0.00				
					(100%)	(100%)	(100%)	(100%)				
					Panel B. Dependent Variable: Simulated R_{put} with η =0.2							
	No Breaks - Full Inf. (No Learning)				Breaks - Full Inf. (No Learning)				Breaks - Partial Inf. (Learning)			
CAPM α	-0.13	-0.03	-0.02	-0.01	-0.20	-0.09	-0.03	0.00	-0.97	-0.95	-0.54	-0.24
t-stat	-0.89	-0.77	-0.52	-0.41	-0.96	-0.90	-0.65	-0.47	-116.31	-21.79	-14.37	-15.18
p-value	0.34	0.44	0.47	0.53	0.23	0.34	0.45	0.52	0.00	0.00	0.00	0.00
	(32%)	(23%)	(14%)	(11%)	(43%)	(26%)	(17%)	(11%)	(100%)	(100%)	(100%)	(100%)
					Panel C. Dependent Variable: Simulated R_{put} with η =0.5							
	No Breaks - Full Inf. (No Learning)				Breaks - Full Inf. (No Learning)				Breaks - Partial Inf. (Learning)			
CAPM α	-0.11	-0.02	-0.02	-0.01	-0.19	-0.07	-0.03	0.00	-0.95	-0.72	-0.42	-0.12
t-stat	-0.60	-0.59	-0.36	-0.32	-0.77	-0.68	-0.55	-0.34	-89.43	-12.23	-8.50	-8.00
p-value	0.27	0.48	0.57	0.64	0.32	0.38	0.45	0.70	0.01	0.00	0.00	0.00
	(29%)	(23%)	(10%)	(9%)	(37%)	(22%)	(14%)	(11%)	(99%)	(99%)	(99%)	(97%)
					Panel D. Dependent Variable: Simulated R_{put} and η =5.0							
	No Breaks - Full Inf. (No Learning)				Breaks - Full Inf. (No Learning)				Breaks - Partial Inf. (Learning)			
CAPM α	-0.21	-0.03	-0.04	-0.01	-0.29	-0.17	-0.06	0.00	-1.08	-0.99	-0.96	-0.63
t-stat	-1.34	-1.13	-0.70	-0.61	-1.62	-1.53	-1.14	-0.92	-109.38	-91.44	-79.42	-53.82
p-value	0.17	0.23	0.41	0.24	0.11	0.10	0.41	0.24	0.00	0.00	0.00	0.00
	(42%)	(32%)	(26%)	(15%)	(54%)	(29%)	(19%)	(22%)	(100%)	(100%)	(100%)	(100%)

Table 3. Relation between option returns and different factors (based on empirical and simulated option data). This table reports the coefficients of factors (in addition to the market factor from the CAPM) that may explain empirical and simulated returns on put option contracts using a naked investment strategy, R_{put}, and a straddle investment strategy, R_{Strdl}. In this table, R_m is the excess market return, and $IV - RV$ is the volatility risk premium where IV is the volatility under the $\mathbb{Q}$ probability measure and RV is the volatility under the $\mathbb{P}$ probability measure. S/D is the price-to-dividend ratio. Empirical option returns are obtained from S&P 500 option contracts between 1996 and 2007. Simulated option returns are obtained from simulations of the model under partial information with breaks and learning, since $IV - RV$ is only different from zero in this scenario. Simulations of the model are performed using three coefficients of relative risk aversion ($\eta = 0.2$, $\eta = 0.5$, and $\eta = 5.0$). The values for the simulated option returns are the averages over 10,000 simulations of 12 years each. This table presents averages of the alpha's *t*-statistics computed using Newey-West standard errors. The percentage of simulations with a significant factor is reported in parentheses at 5% significance. For S&P 500 options, we report *t*-statistics in square brackets.

	Dependent Variable R_{put}				Dependent Variable R_{Strdl}			
	Panel A. Empirical Results (with S&P 500 Options)							
Constant	-0.15	-0.08	-0.15	-0.08	-0.08	-0.02	-0.08	-0.02
	[4.28]	[1.83]	[4.26]	[1.81]	[2.64]	[0.61]	[2.64]	[0.62]
R_m	-26.90	-25.29	-26.89	-25.28	0.06	1.25	0.05	1.25
	[30.33]	[25.70]	[30.29]	[25.66]	[0.08]	[1.59]	0.07	[1.58]
$IV - RV$		-3.62		-3.63		-2.70		-2.70
		[3.61]		[3.61]		[3.36]		[3.35]
S/D			0.00	0.00			0.00	0.00
			[0.31]	[0.33]			[0.31]	[0.30]
Adj. R^2	0.64	0.64	0.63	0.64	0.00	0.02	0.00	0.02
	Panel B. Simulated Results with $\eta =0.2$							
Constant	-0.54	-0.07	-0.53	-0.16	-0.47	-0.06	-0.55	0.10
	(100%)	(19%)	(99%)	(6%)	(100%)	(18%)	(97%)	(8%)
R_m	-33.07	-34.34	-28.47	-31.68	4.46	4.61	4.68	4.72
	(100%)	(100%)	(100%)	(100%)	(60%)	(67%)	(65%)	(70%)
$IV - RV$		-6.76		-7.60		-7.07		-7.10
		(82%)		(82%)		(83%)		(87%)
S/D			-0.62	-3.41			0.06	-2.95
			(7%)	(7%)			(7%)	(8%)
Adj. R^2	0.65	0.66	0.65	0.71	0.06	0.13	0.07	0.14
	Panel C. Simulated Results with $\eta =0.5$							
Constant	-0.42	-0.02	-0.43	-0.39	-0.33	-0.01	-0.33	-0.26
	(99%)	(13%)	(97%)	(8%)	(100%)	(10%)	(86%)	(7%)
R_m	-51.16	-42.50	-49.17	-47.80	1.08	1.16	1.16	1.30
	(100%)	(100%)	(100%)	(100%)	(69%)	(72%)	(66%)	(71%)
$IV - RV$		-13.32		-11.59		-10.12		-10.13
		(71%)		(72%)		(76%)		(76%)
S/D			-3.78	-6.16			2.28	3.89
			(8%)	(7%)			(6%)	(7%)
Adj. R^2	0.62	0.63	0.60	0.66	0.07	0.12	0.07	0.12
	Panel D. Simulated Results with $\eta =5.0$							
Constant	-0.96	-0.41	-0.83	-0.33	-0.91	-0.47	-0.84	-0.31
	(100%)	(46%)	(82%)	(43%)	(100%)	(55%)	(83%)	(52%)
R_m	-11.34	-12.47	-12.66	-12.03	5.04	5.45	5.38	5.41
	(98%)	(98%)	(98%)	(98%)	(88%)	(92%)	(79%)	(83%)
$IV - RV$		-1.25		-1.04		-1.19		-1.29
		(87%)		(84%)		(89%)		(84%)
S/D			-2.42	-3.60			1.50	0.64
			(30%)	(32%)			(28%)	(23%)
Adj. R^2	0.33	0.45	0.40	0.48	0.30	0.39	0.32	0.40

Learning and Index Option Returns

(Internet Appendix)

Appendix A: Evidence of breaks in the mean dividend growth rate

Our model assumes a geometric random walk for dividends, whose drift is subject to breaks. In this appendix, we analyse whether this specification is a realistic representation of the dynamics of actual dividend data. Thus, as a first step, we apply the Chu *et al.* (1996) test to the daily dividend time series of the S&P 500 index between 1996 and 2007, which have been deseasonalised using the Hodrick-Prescott filter and adjusted by the consumer price index, as in Shiller (2000). It is important to analyse deseasonalised time series because there is evidence of seasonality in dividends throughout the year. This seasonality occurs because many firms issue their quarterly dividends at approximately the same time (e.g., Cornell and French, 1983; Lakonishok and Smidt, 1988; Golez, 2014).

Chu *et al.* (1996) propose a dynamic test to detect structural breaks in real time. We apply this test, and we identify eight breaks in the time series of log return dividends using a significance level of 5%, as shown in Figure A1. As a robustness check, we also implement the Bai and Perron (1998) methodology of consistent estimation of multiple breaks. They present tests for both the number and the timing of breaks, which we use to detect multiple structural breaks in the drift of the dividend random-walk process. The MATLAB codes for this test were obtained from Pierre Perron's webpage.

The number of breaks is identified through three separate criteria by setting the maximum potential number of breaks at 12. The first criterion represents a sequential break test, which uses a significance level of 5%. The other two criteria – the Bayes information criterion (BIC) and the Akaike information criterion (AIC) – are based on a penalised likelihood function. All three criteria provide multiple breaks in g_t. The

sequential approach chooses five breaks, the BIC leads to seven breaks, and the AIC chooses the maximum allowed number of breaks (i.e., 12 breaks). Therefore, the implementation of the Chu *et al.* (1996) and Bai and Perron (1998) tests suggests that the assumption of breaks in the mean dividend growth rate is reasonable in terms of describing market data.

Appendix B: Proof of Proposition II

Proof of Proposition II: To obtain equation (10), one needs the probabilities described by the state price density to be risk-neutralised. First, we divide both sides of the Euler equation (4) used to price the stock at $t + k$ by the bond price in equation (8):

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

$$\left. \cdot (S_{t+k+1} + D_{t+k+1})\frac{(1 + \rho)}{(1 - \pi)(1 + g_{t+k})^{-\eta} + \pi \int_{g_d}^{g_u}(1 + g_{t+k})^{-\eta}dG(g_{t+k})}\right]. \tag{B1}$$

Under full information, the forward stock price and the forward cumulative dividend process are given by:

$$S_{t+k}^* = \frac{(1 + \rho)S_{t+k}}{(1 - \pi)\left(1 + g_{t+k}\right)^{-\eta} + \pi \int_{g_d}^{g_u}\left(1 + g_{t+k}\right)^{-\eta}dG(g_{t+k})} \tag{B2}$$

and

$$D_{t+k}^* = \sum_{s=0}^{k} D_{t+s}\frac{(1 + \rho)}{(1 - \pi)\left(1 + g_{t+s}\right)^{-\eta} + \pi \int_{g_d}^{g_u}\left(1 + g_{t+s}\right)^{-\eta}dG(g_{t+s})}, \tag{B3}$$

while the definition of the pricing kernel yields:

$$E_t \left[\beta \left(\frac{D_{t+k+1}}{D_{t+k}} \right)^{-\eta} \frac{(1+\rho)}{(1-\pi)\left(1 + g_{t+k}\right)^{-\eta} + \pi \int_{g_d}^{g_u} \left(1 + g_{t+k}\right)^{-\eta} dG(g_{t+k})} \right] = 1. \tag{B4}$$

Adding D_{t+k}^* to equation (B1) and combining the resulting expression with equation (B4) yields:

$$S_{t+k}^* + D_{t+k}^*$$

$$= E_{t+k} \left[\beta \left(\frac{D_{t+k+1}}{D_{t+k}} \right)^{-\eta} \frac{(1+\rho)}{(1-\pi)(1 + g_{t+k})^{-\eta} + \pi \int_{g_d}^{g_u} (1 + g_{t+k})^{-\eta} dG(g_{t+k})} (S_{t+k+1}^* + D_{t+k+1}^*) \right]. \tag{B5}$$

Equation (B5) shows that $S_{t+k}^* + D_{t+k}^*$ follows a martingale under the risk-neutral ($\mathbb{Q}$) probability measure. We know that the $\mathbb{Q}$ probability measure can be expressed as:

$$f^{\mathbb{Q}}(S_{t+k+1}) = \frac{m_{t+k+1} f_t^{\mathbb{P}}(S_{t+k+1})}{E_t[m_{t+k+1}]}. \tag{B6}$$

Given that equation (7) in Proposition I shows that $S_{t+k} = D_{t+k}\psi(g_{t+k})$, we can write the one-period physical ($\mathbb{P}$) probability measure in equation (B6) as:

$$f_t^{\mathbb{P}}(S_{t+k+1}) = \phi^{\mathbb{P}}(\varepsilon_{t+k+1})\text{Pr}^{\mathbb{P}}(b_{t+k+1})\varrho^{\mathbb{P}}\left(g_{t+k+1} \middle| b_{t+k+1} = 1\right), \tag{B7}$$

In equation (B7), ε_{t+k+1} is the innovation term of the dividends' geometric random walk characterised by normal density $\phi^{\mathbb{P}}(\varepsilon_{t+k+1})$ with mean zero and variance σ^2, while $\text{Pr}^{\mathbb{P}}(b_{t+k+1})$ follows a Bernoulli distribution with parameter π. In addition, in the case of a break, g_{t+k+1} is drawn from uniform density $\varrho^{\mathbb{Q}}(g_{t+k+1}|b_{t+k+1} = 1)$ with g_l and g_u being the lower and upper bounds, respectively.

We can also rewrite equation (B6) as:

$$f_t^{\mathbb{Q}}(S_{t+k+1}) = \phi^{\mathbb{Q}}(\varepsilon_{t+k+1}|0, \sigma)\text{Pr}^{\mathbb{Q}}(b_{t+k+1}) \varrho^{\mathbb{Q}}(g_{t+k+1}) \tag{B8}$$

where

$$\phi^{\mathbb{Q}}(\varepsilon_{t+k+1}) = \int_{g_{t+k+1}} \int_{b_{t+k+1}} f_t^{\mathbb{Q}}(S_{t+k+1}) \, db_{t+k+1} dg_{t+k+1}$$

$$= \int_{g_{t+k+1}} \int_{b_{t+k+1}} \frac{m_{t+k+1} f_t^{\mathbb{P}}(S_{t+k+1})}{E_t[m_{t+k+1}]} \, db_{t+k+1} dg_{t+k+1}$$

$$= \int_{g_{t+k+1}} \int_{b_{t+k+1}} \frac{m_{t+k+1} \phi^{\mathbb{P}}(\varepsilon_{t+k+1}) \mathrm{Pr}^{\mathbb{P}}(b_{t+k+1}) \varrho^{\mathbb{P}}(g_{t+k+1}|b_{t+k+1}=1)}{E_t[m_{t+k+1}]} \, db_{t+k+1} dg_{t+k+1}$$

$$\mathrm{Pr}^{\mathbb{Q}}(b_{t+k+1})$$

$$= \int_{g_{t+k+1}} \int_{\varepsilon_{t+k+1}} \frac{m_{t+k+1} \phi^{\mathbb{P}}(\varepsilon_{t+k+1}) \mathrm{Pr}^{\mathbb{P}}(b_{t+k+1}) \varrho^{\mathbb{P}}(g_{t+k+1}|b_{t+k+1}=1)}{E_t[m_{t+k+1}]} \, d\varepsilon_{t+k+1} dg_{t+k+1}$$

$$\varrho^{\mathbb{Q}}(g_{t+k+1})$$

$$= \int_{\varepsilon_{t+k+1}} \int_{b_{t+k+1}} \frac{m_{t+k+1} \phi^{\mathbb{P}}(\varepsilon_{t+k+1}) \mathrm{Pr}^{\mathbb{P}}(b_{t+k+1}) \varrho^{\mathbb{P}}(g_{t+k+1}|b_{t+k+1}=1)}{E_t[m_{t+k+1}]} \, db_{t+k+1} d\varepsilon_{t+k+1}$$

Thus, after some algebra, we find that $\phi^{\mathbb{Q}}(\varepsilon_{t+k+1})$ is a normal density with mean zero and volatility $\sigma^{\mathbb{Q}} = \sigma$, $\mathrm{Pr}^{\mathbb{Q}}(b_{t+k+1})$ follows a Bernoulli distribution with parameter $\pi^{\mathbb{Q}}_{t+k+1} = A_{t+k+1}/B_{t+k+1}$, and $\varrho^{\mathbb{Q}}(g_{t+k+1})$ is a uniform density with g_l and g_u being the lower and upper bounds, respectively (alternatively, the value of parameter $\pi^{\mathbb{Q}}_{t+k+1}$ can be obtained from the fact that $A_{t+k+1}^{\mathbb{P}} = A_{t+k+1}^{\mathbb{Q}}$, $A_{t+k+1}^{\mathbb{Q}} = \pi^{\mathbb{Q}}_{t+k+1}/(1+r_{t+k+1})$ and $B_{t+k+1} = 1/(1+r_{t+k+1})$; thus, $\pi^{\mathbb{Q}}_{t+k+1} = A_{t+k+1}/B_{t+k+1}$).

Therefore, given that the market is complete, the risk-neutral measure for any single period is unique and exists, which is a sufficient condition for having a unique, risk-neutral measure in an infinite-period economy obtained by the repetition of several single periods

(for proof, see Pliska, 1997). We then define our infinite-period risk-neutral measure by considering all states that the mean dividend growth rate could achieve in $t + \tau$ and by using the independence property of breaks. Therefore, $f^{\mathbb{Q}}(S_{t+\tau})$ is the risk-neutral density of all paths that lead to a certain state in which the dividend is $D_{t+\tau}$, where its expected value is:

$$E_t^{\mathbb{Q}}[D_{t+\tau}] = D_t E_t^{\mathbb{Q}}\left[\frac{D_{t+1}}{D_t} E_{t+1}^{\mathbb{Q}}\left[\left(\frac{D_{t+2}}{D_{t+1}}\right) \dots E_{t+\tau-1}^{\mathbb{Q}}\left[\left(\frac{D_{t+\tau}}{D_{t+\tau-1}}\right)\right]\right]\right] \tag{B9}$$

We know that the innovation term of the random-walk process, ε_t, and the breaks in g_t are independent; thus, the expected value of $D_{t+\tau}$ can be written as:

$$E_t^{\mathbb{Q}}[D_{t+\tau}] = D_t E_t^{\mathbb{Q}}\left[\exp\left(\sqrt{\tau}\sigma\varepsilon_{t+\tau} - \tau\sigma^2/2\right)\prod_{i=1}^{\tau}(1 + r_{t+i})\right] \tag{B10}$$

Therefore, in each path we have:

$$D_{t+\tau} = D_t\exp\left(\sqrt{\tau}\sigma\varepsilon_{t+\tau} - \tau\sigma^2/2\right)\prod_{i=1}^{\tau}(1 + r_{t+i}) \tag{B11}$$

where $\varepsilon_{t+\tau}$ is the innovation term of the dividends' geometric random walk. Consequently, since we know that $S_{t+\tau} = D_{t+\tau}\psi(g_{t+\tau})$, similarly to equation (B8), we can write:

$$f_t^{\mathbb{Q}}(S_{t+\tau}) = \phi^{\mathbb{Q}}(\varepsilon_{t+\tau}|0, \sigma)\prod_{i=1}^{\tau}\text{Pr}^{\mathbb{Q}}(b_{t+i})\,\varrho^{\mathbb{Q}}(g_{t+i}) \tag{B12}$$

where $\phi^{\mathbb{Q}}(\varepsilon_{t+\tau}|0, \sigma)$ is a normal density with mean zero and volatility $\sigma^{\mathbb{Q}} = \sqrt{\tau}\sigma$, $\text{Pr}^{\mathbb{Q}}(b_{t+i})$ follows a Bernoulli distribution with parameter $\pi_{t+i}^{\mathbb{Q}} = A_{t+i}/B_{t+i}$, and in the case of a break,

g_{t+i} is drawn from a continuous uniform density $\varrho^{\mathbb{Q}}(g_{t+i})$ in which g_l and g_u are also the lower and upper bounds, respectively.

□

Appendix C: Robustness check using alternative setups of the model

In this appendix, we present the outcomes of a robustness check with respect to analysing three alternative setups of the model presented in Section 2. The first alternative setup is a model in an economy under partial information with breaks and learning. In this setup, we vary the level of σ to evaluate the effect of a change in this parameter on put option returns. It is important to analyse the impact of a change in σ in our model, since that level affects the noisiness of the signals received by the Bayesian representative agent. The level of noisiness of the signals affects the agent's learning process and hence put option returns (as explained in Section 2.4).

The second alternative setup is an extension of our model in an economy under partial information with breaks and learning. In this setup, the representative agent does not know the time at which a break in g_t has occurred. We extend the model presented in Section 2 by assuming that the agent does not know exactly when a break has occurred; however, she uses the dynamic test presented in Chu *et al.* (1996) to detect breaks in the mean dividend growth rate in real time.

The third alternative setup is the same model presented in Section 2 in an economy under partial information with breaks and learning. However, the COS securities are not

considered in the market portfolio when calculating the value of the CAPM α (despite the fact that this risky security is part of the economy).

Table C1 shows the average CAPM α for hold-to-maturity returns on naked strategies using the three alternative setups of the model as described above. In the first alternative setup, we use a level of σ that is six times higher than the level used in Table 2 (while keeping all other parameters the same). Table C1 shows that an increase in the level of noisiness of the signals increases the values of the CAPM α when those values are compared to the values presented in Table 2 (the last four columns on the right-hand side). Nevertheless, despite the increase in the values of the CAPM α, which are induced by a high level of σ, the results presented in Table C1 are qualitatively similar to the simulated and empirical results reported in Table 2. Thus, Table C1 shows that the features of the model (in relation to an economy with partial information and learning) used to explain put option returns are still valid after a change in the level of noisiness of the signals.

[Insert Table C1 here]

In the second alternative setup, the main difference between the extension of the model presented in Table C1 and the model described in Section 2 is the knowledge of the dates of breaks. On the one hand, the results in Table 2 (the last four columns on the right-hand side) are calculated using the model with breaks, partial information and learning in which the agent detects a break as soon as it happens. On the other hand, in Table C1 (under the second alternative setup), the agent uses the Chu *et al.* (1996) test to determine when a break has occurred in real time. Thus, the agent's beliefs regarding the break dates in Table C1 do not necessarily represent the "true" dates on which breaks happened.

Therefore, the agent does not know with certainty when to reset the counter of signals, n, in equation (13) to zero in order to start the new post-break learning process.

Table C1, under the second alternative setup, shows that simulated values of the CAPM α are significant for a naked put strategy, which is consistent with the results presented in Table 2 (the last four columns on the right-hand side). Notably, the absolute values of the CAPM α in Table C1 are slightly greater than those observed in the model in which the timing of breaks is known (see Table 2). The reported values of the CAPM α are larger in absolute terms in Table C1 than they are in Table 2, as there is a new source of learning associated with the detection of breaks. However, the results in Table C1 are in line with the theoretical and empirical results presented in Table 2.

In relation to the third alternative setup, Table C1 presents a similar analysis to that presented in Table 2, but the market portfolio includes only the stock this time. This analysis is important because we want to examine the impact of COS securities on the values of the CAPM α, given that COS securities may not be directly considered in the market portfolio (e.g., if we use the S&P 500 index as the market portfolio, this index includes only 500 assets despite the fact that we know there are more assets in the economy). Table C1 shows that the results hardly change when we omit the COS security from the market portfolio. The differences in absolute values between Table C1 (under the second alternative setup) and Table 2 (also in the last four columns on the right-hand side) are small because the weight of the COS securities in the market portfolio is low, as explained in Section 4.2. Thus, Table C1 suggests that not considering COS securities in the market portfolio does not significantly modify the results obtained from the model described in the body of our paper.

Appendix D: Simulated and empirical volatility risk premium

As explained in our study, a gap between the $\mathbb{Q}$ and $\mathbb{P}$ probability measures should affect the expected hold-to-maturity returns on put option contracts, because the expected hold-to-maturity return on a put option contract is defined as $R_{t+\tau}^{Put} = (E_t^{\mathbb{P}}[\max(K - S_{t+\tau}, 0)]/E_t^{\mathbb{Q}}[\max(K - S_{t+\tau}, 0)] - 1)$. Therefore, it is important to analyse whether our model under partial information and learning can generate differences between the $\mathbb{Q}$ and $\mathbb{P}$ probability measures. Thus, Table D1 presents summary statistics on the difference in volatilities obtained under the risk-neutral and physical probability measures (i.e., the volatility risk premium). We examine volatility risk premium obtained from the model presented in Section 2 and the volatility risk premium calculated from the S&P 500 option data between 1996 and 2007.

[Insert Table D1 here]

The calculation of volatility under the $\mathbb{Q}$ and $\mathbb{P}$ probability measures (using our model and S&P 500 option contracts) is explained in Section 4.3. Table D1 shows that the difference in volatility between the $\mathbb{Q}$ and $\mathbb{P}$ probability measures in our model is, on average, 6.9%, 3.4% and 34.0% when the coefficient of relative risk aversion is 0.2, 0.5 and 5.0, respectively. The difference in volatility between the $\mathbb{Q}$ and $\mathbb{P}$ probability measures estimated from S&P 500 option contracts is 2.5%.

Table D1 shows that our model under partial information and learning induces a gap in volatility between the $\mathbb{Q}$ and $\mathbb{P}$ probability measures. However, we can observe in Table D1 that the difference in volatilities under the $\mathbb{Q}$ and $\mathbb{P}$ probability measures generated by our model is not exactly the same as the difference obtained from the actual

S&P 500 option data. However, it is important to note (as mentioned in the paper) that our focus is on providing a simple model under partial information and learning, with the objective of explaining option returns rather than calibrating all variables in the economy perfectly. Therefore, our model can be considered a first step in terms of a research endeavour to relate put option returns to learning. We have left the analysis of a richer model that would include other elements that may also affect the option pricing process (e.g., jump risk, investors' net option demand, among others) for future research.

References

Bai, J., and Perron, P. (1998), "Estimating and Testing Linear Models with Multiple Structural Changes," *Econometrica*, 66, 47–78.

Black, F., and M. Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–659.

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

Cornell, B., and K. R. French, 1983, Taxes and the pricing of stock index futures, *Journal of Finance* 38, 675–694.

Golez, B. 2014. Expected returns and dividend growth rates implied by derivative markets. *Review of Financial Studies* 27:790–822.

Lakonishok, Josef, and Seymour Smidt, 1988, Are seasonal anomalies real: A ninety year perspective, *Review of Financial Studies* 1, 403–425.

Lettau, M., and S. Ludvigson, 2005. Expected returns and expected dividend growth. *Journal of Financial Economics*, 76(3):583–626.

Merton, R. C., 1976, Option Pricing When Underlying Stock Returns Are Continuous, *Journal of Financial Economics*, 3, 125–144.

Shiller, R., 2000, Irrational exuberance, Princeton University Press, Princeton, NJ.

Figure A1. **Structural breaks identified using the Chu *et al.* (1996) test**. This figure presents the outcomes of applying the Chu *et al.* (1996) test to dividend log returns for the S&P 500 index in the period 1996–2007. The solid line reflects the mean of the dividend log returns, calculated with a rolling window of 125 trading days. The dotted line shows the breaks identified through the test. Breaks were identified in Dec. 1996, Aug. 1999, Sept. 2000, Apr. 2001, Oct. 2001, Aug. 2002, Nov. 2003 and Oct. 2004.

Table C1. Robustness check using alternative setups of the model. This table presents a robustness check in terms of analysing three alternative setups of the model presented in Section 2 in an economy under partial information with breaks and learning. This table displays simulated CAPM α values of put option returns, using a naked option strategy, R_{put}, and three coefficients of relative risk aversion ($\eta = 0.2$, $\eta = 0.5$ and $\eta = 5.0$). The first alternative setup is a model with a high level of noisiness in the signals used in the learning process. In this first alternative setup, the volatility of the geometric random walk is equal to 8.6% on a monthly basis (which is equivalent to 30% on an annual basis) instead of a volatility of 1.44% (which is equivalent to 5% on an annual basis), as in the body of the paper. The second alternative setup is a model that includes the effect of not knowing the time at which a break occurred. The second alternative setup represents an extension of the model presented in Section 2 in which the representative agent does not know the dates of breaks. Thus, the agent uses the test presented in Chu *et al.* (1996) to detect the presence of a break in the dividend drift contemporaneously. The third alternative setup is the same model that is presented in Section 2; however, the COS securities are not considered in the market portfolio when calculating the value of the CAPM α. The numbers in the table for the simulated option returns are the average values over 10,000 simulations of 12 years. This table presents averages of the alpha's t-statistics computed using Newey-West standard errors to correct for heteroscedasticity and serial correlation. The percentages of simulations with a significant CAPM α are reported in parentheses at 5% significance.

	Model with the level of σ six times higher than the value used in our study				Extended model including uncertainty about dates of breaks				Model with market portfolio only includes the stock (without the COS security)			
K/S	0.96	0.98	1.00	1.02	0.96	0.98	1.00	1.02	0.96	0.98	1.00	1.02
	Panel A. Dependent Variable: Simulated R_{put} with η =0.2											
CAPM α	-1.92	-1.61	-0.88	-0.48	-1.18	-1.06	-0.67	-0.29	-1.02	-0.86	-0.54	-0.23
t-stat	-238.83	-45.72	-22.98	-25.66	-118.59	-24.51	-15.16	-18.22	-110.21	-22.36	-15.35	-15.14
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	(100%)	(100%)	(100%)	(100%)	(100%)	(100%)	(100%)	(100%)	(100%)	(100%)	(100%)	(100%)
	Panel B. Dependent Variable: Simulated R_{put} with η =0.5											
CAPM α	-1.63	-1.42	-0.74	-0.36	-1.07	-1.02	-0.77	-0.34	-0.93	-0.80	-0.43	-0.12
t-stat	-162.70	-34.01	-21.16	-21.08	-132.82	-28.95	-20.19	-19.89	-88.79	-11.70	-8.14	-8.03
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.02	0.00	0.00	0.00
	(100%)	(100%)	(100%)	(100%)	(100%)	(100%)	(100%)	(100%)	(97%)	(99%)	(99%)	(98%)
	Panel C. Dependent Variable: Simulated R_{put} with η =5.0											
CAPM α	-1.74	-1.69	-1.52	-1.31	-1.21	-1.16	-1.12	-0.71	-1.03	-1.00	-1.00	-0.69
t-stat	-182.86	-147.07	-166.25	-110.18	-123.91	-95.17	-80.73	-55.91	-104.72	-94.36	-74.66	-56.37
p-value	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	(100%)	(100%)	(100%)	(100%)	(100%)	(100%)	(100%)	(100%)	(100%)	(100%)	(100%)	(100%)

Table D1. Empirical and simulated volatility risk premium. This table reports the empirical and simulated volatility risk premium ($IV_t - RV_t$). IV_t and RV_t are explained in Section 4.3. In the case of simulated option returns, simulations of the model are based on an economy under Bayesian learning, with three coefficients of relative risk aversion (η=0.2, η=0.5 and η=5.0). The numbers in the table for the simulated option returns are the average values over 10,000 simulations of 12 years. Empirical option returns are obtained from S&P 500 option contracts between 1996 and 2007.

	Volatility Risk Premium			
	Empirical S&P 500 Options IV_t - RV_t	Simulated IV_t - RV_t (Breaks with Partial Inf. and Learning)		
		Relative Risk Aversion (η)		
		0.2	0.5	5.0
Mean	0.025	0.069	0.034	0.340
Std. Dev.	0.048	0.018	0.013	0.090
Skewness	-0.536	-0.274	-0.080	-0.569
Kurtosis	6.365	5.704	4.833	5.471