Learning and Index Option Returns [1]

Alejandro Bernales   Gonzalo Cortazar

Luka Salamunic   George Skiadopoulos

This version: January 14th, 2017

Abstract

We offer for the first time an economic explanation to the patterns of index put options' expected returns, that of learning about economic fundamentals. Expected option returns are determined by the gap between $\mathbb{P}$ and $\mathbb{Q}$-probability measures. This gap is affected by parameter mis-estimation which evolves over time as the agent updates her beliefs. We formalize this intuition by employing a novel option pricing model with breaks in fundamentals where the representative agent updates her beliefs by Bayesian learning. We find that the patterns of the simulated under our economy put option returns are similar to the empirical S&P 500 put index options ones across different levels of moneyness and time-to-maturity. Results are robust even after controlling for leverage and market risk.

JEL classification: D83, G13, G14, G17.

Keywords: Asset pricing, Index put option returns, Equilibrium option pricing model, Bayesian learning, Breaks in fundamentals.

---

[1] Alejandro Bernales is at Universidad de Chile (Centro de Economía Aplicada and Centro de Finanzas), email: abernales@dii.uchile.cl. Gonzalo Cortazar is at Pontificia Universidad Católica de Chile, email: gcortaza@ing.puc.cl. Luka Salamunic is at Pontificia Universidad Católica de Chile, email: lsalamun@uc.cl. George Skiadopoulos is at University of Piraeus and Queen Mary University of London, email: gskiado@unipi.gr, g.skiadopoulos@qmul.ac.uk. We would like to thank Xiaohua Chen, Massimo Guidolin, Christian Hellwig, Eduardo Rodriguez, Marcela Valenzuela, Patricio Valenzuela, Thanos Verousis and seminar participants in the Pontificia Universidad Católica de Chile, Universidad de Chile and Banque de France for their comments on earlier versions of this paper. We are also thankful to Vincent Guegan for the IT support. Alejandro Bernales acknowledges financial support from Fondecyt project #11140628 and the Institute for Research in Market Imperfections and Public Policy (ICM IS130002). All remaining errors are ours.

1 Introduction

Index put option returns are found to be large. Bondarenko (2014) finds that, on average, at-the-money (out-of-the-money) naked index put returns have an impressive level of -40% (-90%) per month. Coval and Shumway (2001), Jones (2006), Santa-Clara and Saretto (2009) and Constantinides *et al*. (2013) also report large negative returns by examining various index put option strategies.

Broadie *et al*. (2009) and Chambers *et al*. (2014) show that the patterns of index put option returns may be the result of mis-estimating parameters in option pricing models. This is because the size of expected returns of a hold-to-maturity option strategy is a function of the gap between the $\mathbb{P}$ (real-world) and $\mathbb{Q}$ (risk-neutral) probability measures. The mis-estimation of parameters affects this gap. Broadie *et al*. (2009) state in their concluding section though "*Our results are silent on the actual economic sources 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 learning*". We fill this void by offering an economic explanation to the dynamics of the $\mathbb{P}$- $\mathbb{Q}$ gap, that of learning about fundamentals. Subsequently, we test whether our setting yields patterns of the S&P 500 index put option returns similar to these observed empirically.

We employ a representative agent's discrete-time endowment economy where stocks, bonds and options trade. The agent observes dividends whose mean dividend growth rate g_t is subject to breaks. Once a break occurs, the agent does not know the new true value of g_t . Yet, the agent starts learning about g_t recursively via a Bayesian updating scheme as new information arrives. This setting is a natural candidate to explain expected option returns because it allows modelling the gap between the $\mathbb{P}$ and $\mathbb{Q}$-probability measures as a function of the way the agent learns about estimation errors in fundamentals. In an economy under Bayesian learning, the wedge between $\mathbb{P}$ and $\mathbb{Q}$ probability measures differs from the one under full information. In the latter case, the stock price risk neutral probability measure $f^{\mathbb{Q}}\left(S_{t+1}\right)$ is related to the physical probability measure $f^{\mathbb{P}}\left(S_{t+1}\right)$ as $f^{Q}(S_{t+1}) = m_{t+1} f^{P}(S_{t+1}) / E_t(m_{t+1})$ where m_{t+1} is the stochastic discount factor. However, under Bayesian learning (BL), we show that this relation changes to $f^{\mathbb{Q},BL}\left(S_{t+1}|\Omega_t\right) = m_{t+1} f^{\mathbb{P}}(S_{t+1} \mid g_t) f(g_t \mid \Omega_t) / \mathrm{E}_t^{BL}(m_{t+1} \mid \Omega_t)$ where $f(g_t \mid \Omega_t)$ is the probability density function of g_t conditional on the information set Ω_t available at time *t* and $\mathrm{E}_t^{BL}(\bullet|\,\Omega_t)$ is the conditional expectation operator under Bayesian learning. As the agent updates her estimate for g_t, the physical probability is updated and as a result the risk-neutral probability is also updated.

We proceed as follows. First, we derive an option pricing model under breaks in fundamentals and full information where the agent knows the true value of g_t. Then, we use the model within the Bayesian learning setting to simulate option returns and we compare them to the empirically observed S&P 500 put index options' returns. Four are the main findings of our study. First, our setting generates large and statistically significant option returns and CAPM alphas for naked index put options as well as for option portfolios that control for leverage effects and market risk (delta-hedged portfolios). Second, the returns of the considered option strategies decline in magnitude as moneyness and time-to-maturity increase. These findings are in accordance with these reported by the previous literature (e.g., Coval and Shumway, 2001; Bondarenko, 2003; Broadie *et al*., 2009; Constantinides *et* *al*., 2013; and Chambers *et al*., 2014). Third, the option returns generated by our model have similar patterns to the ones computed from actual S&P 500 index option data. Fourth, we document that the volatility risk premium explains a significant portion of returns on option trading strategies. In the case of leveraged-adjusted and delta-hedged put option portfolios, CAPM alphas decrease significantly once the volatility risk premium is included as an additional explanatory factor. This is consistent with Broadie *et al*. (2009) and Chambers *et al*. (2014) who report that put option portfolios' returns can be explained by parameters' mis-estimation which induces a gap between the $\mathbb{P}$ and $\mathbb{Q}$ probability measures.

We conclude this section by discussing the differences between our study and related literature. Broadie *et al*. (2009) and Chambers *et al*. (2014) examine the impact of estimation errors by increasing/decreasing the $\mathbb{Q}$-measure parameters by one standard deviation from the $\mathbb{P}$-parameters. Their approach represents a reduced-form model. Our paper complements these two studies because we explain option returns by endogenizing the dynamics of parameters' mis-estimation as a function of learning over time. To the best of our knowledge, our study is the first that examines index options' expected returns by providing an economic explanation (learning about fundamentals) for their properties.

Our study is also associated with equilibrium models that employ learning to explain the existence of a non-flat implied volatility surface rather than option returns. David and Veronesi (2002) propose a model where the dividend drift follows a two-state stochastic process and the representative agent is uncertain about the current state of the economy. Guidolin and Timmermann (2003) present an equilibrium model where the dividend grows along a binomial path with an unknown state probability that is updated recursively.

Shaliastovich (2009) offers a model where investors learn about the consumption growth rate based on a recency-biased updating procedure and the consumption growth rate is uncertain and subject to breaks. Benzoni et al. (2011) extend a general equilibrium setting with an Epstein-Zin representative agent to include jumps and Bayesian updating. Finally, a remark is in order. Our paper examines *index* option returns and hence it is distinct from the literature that examines option returns for individual equity options (e.g., Goyal and Saretto, 2009, Ni, 2009, and Buraschi et al., 2014).

The remainder of the paper is structured as follows. Section 2 presents the learning model about fundamentals, how learning affects option returns and the model's properties. Section 3 describes the way to generate simulated option returns and discusses results for the naked index put options. Section 4 discusses results for the leverage-adjusted portfolios. Section 5 reports additional robustness checks. Section 6 concludes.

2 The learning model about fundamentals

We derive the option-learning model within a discrete-time representative agent's endowment economy. We assume that the mean dividend growth rate g_t is subject to breaks. First, we consider a full information case where the agent knows the true value of g_t once a break occurs. Then, we incorporate learning by relaxing the full information assumption. We assume that once the break occurs, the true value of g_t is unknown by the agent, yet she learns about it gradually by observing market signals (i.e. dividends).

2.1 An economy under full information

First, we assume an economy with full information, where a representative agent prices different types of assets. At time t, there is a one-period zero-coupon default-free bond $\boldsymbol{B_t}$ (in zero net supply), a stock S_t (with net supply normalized at one) and a set of European put option contracts $p_t\left(K,\tau\right)$ with the stock as underlying asset, where $\boldsymbol{K}$ is the strike price and $\boldsymbol{\tau}$ is the time-to-maturity. We assume a perfect capital market. The stock pays real dividends D_t which follow a geometric random-walk process with drift μ_t and volatility σ, i.e.

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

where the mean dividend growth rate g_t (and hence μ_t given that $\mu_t=\ln(1+g_t)-\sigma^2/2$) is subject to breaks. The time-period between any two consecutive breaks follows a geometric distribution defined by a parameter π. Therefore, g_t changes over time, yet its value is constant between breaks. As soon as a break occurs, a new value for g_t is drawn from a univariate distribution. In line with Pesaran *et al* (2006) and Koop and Potter (2007), we choose a geometric distribution because it is a memoryless stochastic process. This choice is consistent with the assumption that agents cannot predict the future and hence they cannot predict future breaks in fundamentals.

In this type of economy, the market is not complete due to the additional uncertainty generated by breaks in the mean dividend growth rate. We make the market dynamically complete by allowing the trading of change-of-state (COS) securities. The COS concept has been introduced by Guo (2001) as a vehicle to hedge uncertainty caused by regime shifts and hence it is a natural instrument to apply for the purposes of our analysis. We assume that at any point in time t, a COS security with price A_t trades which pays one unit in the period where a break in the mean dividend growth rate occurs and zero otherwise. The COS security becomes worthless after a break and a new COS security is issued to ensure the market is dynamically complete.

The representative investor's preferences are described by a power utility function

$$u\left(C_t\right)=\begin{cases}\dfrac{C_t^{1-\eta}-1}{1-\eta} & \eta\geq 0,\eta\neq 1\\ \ln C_t & \eta=1\end{cases} \tag{2}$$

where C_t is the real consumption at time t, and η is the coefficient of relative risk aversion. Dividends are the economy's single source of income and they are consumed as soon as they are received, i.e. $C_t = D_t$. The representative agent chooses holding of assets with prices S_t, B_t and A_t to maximize her lifetime expected utility:

$$\max_{\left\{w_{t+k}^S,w_{t+k}^B,w_{t+k}^A\right\}} E_t\left[\sum_{k=0}^{\infty}\beta^k u\left(D_{t+k}\right)\right] \tag{3}$$

where $\beta = 1/\left(1+\rho\right)$, 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. Notice that put option contracts are not considered in the agent's maximization problem described by equation (3). This is because markets are complete due to the existence of COS securities and hence options are redundant. The prices S_t, B_t and A_t are calculated by solving the following Euler equations

$$S_t = \mathrm{E}_t\left[m_{t+1}\left(S_{t+1} + D_{t+1}\right)\right] \tag{4}$$

$$B_t = \mathrm{E}_t\left[m_{t+1}\right] \tag{5}$$

$$A_t = \mathrm{E}_t\left[m_{t+1}\left(b_{t+1} - b_t\right)\right] \tag{6}$$

where $m_{t+1} = \beta\left(D_{t+1} / D_t\right)^{-\eta}$ is the stochastic discount factor and b_t is a break indicator which indicates the occurrence of breaks in the mean dividend growth rate at time t. $b_t = b_{t-1}$ in the case where there is no break at t and $b_t = b_{t-1} + 1$ if a break has taken place at t.

For the case of full information where breaks in the dividend price take place, the following Proposition I provides expressions for the equilibrium prices of S_t, B_t and A_t by solving equations (4), (5) and (6), respectively. In addition, we develop Proposition II, which provides no-arbitrage prices of a European put option written on the stock.

Proposition I: *The equilibrium prices of S_t, B_t and A_t, under full information and breaks in the dividend process described in equation* (1) *are given by:*

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

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

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

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

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

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

Proof: See Timmermann (*2001*).

Proposition II: *Under full information, the price* $p_t^{full}(K,\tau)$ *of European put option at time* t *with underlying asset the stock with price* S_t^{full}, *strike price* K *and time-to-expiration* τ *is*

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

with

$$\begin{aligned} f^{\mathbb{Q}}\left(S_{t+\tau}^{full}\right) = (1+r_{t+\tau})\beta^{\tau}\left(D_{t+\tau}/D_t\right)^{-\eta} f\left(\varepsilon_{t+\tau}|0,\sigma\right)\varphi\left(z|\tau,\pi\right)\cdot \\ \psi(h_1|\pi)\cdot\left[\psi(h_2|\pi)\varrho(g_2)\ldots\psi(h_z|\pi)\varrho(g_z)\right] \end{aligned} \tag{11}$$

where $f^{\mathbb{Q}}\left(S_{t+\tau}^{full}\right)$ *is the risk-neutral price density under full information,* $S_{t+\tau}^{full}$ is defined in (equation (7)), $D_{t+\tau} = D_t \exp\left(\sqrt{\tau}\sigma\varepsilon_{t+\tau} - \frac{\tau\sigma^2}{2}\right)(1+g_1)\prod_{i=2}^{\tau}\left(1+g_{i-1}+A_i\left(g^* - g_{i-1}\right)\right)^t$ is the risk-neutral process for dividends, where $A_i = 1\ (i = 1,\ldots,\tau)$ *with probability* π *and* g^* *is drawn from a uniform distribution defined on the support* $[\mathrm{g_d}, \mathrm{g_u}]$, *z is a random variable that counts the number of breaks between* t *and* $t+\tau$ *drawn from a Binomial distribution* $\phi(z\,|\,\tau,\pi)$ *with* *parameters τ and π, and* $\{h_i\}_{i=1}^{z}$ *are random variables which measure time periods between breaks drawn from geometric distributions* $\psi(h_i|\pi)$ *where* $\tau = \sum_{i=1}^{z} h_i$ *.*

$1 + r_{t,t+\tau}^{full} = \prod_{j=1}^{\tau}\left(1 + r_{j-1,j}^{full}\right)$ *with* $1 + r_{j-1,j}^{full} = 1 / B_{j-1}^{full}$ *where* B_{j-1}^{full} *is the price of the risk-free one-period bond at period* $j-1$ *under full information [*equation (8)*],* $\{g_i\}_{i=2}^{z}$ *are drawn from a univariate distribution* $G(\bullet)$ *with probability density function* $\varrho(\cdot)$ *defined on the support* $[g_d, g_u]$ *where* $g_1 = g_t$ (*the current mean dividend growth rate*) *and* $g_z = g_{t+\tau}$ (*the mean dividend growth rate at time-to-maturity*), $\varepsilon_{t+\tau}$ *is the innovation term of the dividends' geometric random walk characterised by a normal density* $\phi(\varepsilon_{t+\tau}|0,\sigma)$ *with mean zero and volatility σ.*

Proof: See Appendix B.

Three remarks are in order regarding Proposition II. First, equation (10) shows that to compute the put price under full information, one needs to integrate the option's payoff over the risk-neutral price density $f^{\mathbb{Q}}\left(S_{t+\tau}^{full}\right)$. This can be done by Monte Carlo simulation. To this end, we simulate dividends under the risk-neutral measure up to time *t+τ*. We run *M* simulated paths. For any given dividend simulated path, we calculate the respective stock price obtained from equation (7). This yields *M* respective simulated stock prices, i.e. *M* simulated respective option payoffs. Finally, we average these simulated payoffs over *M*. Note that the $1 + r_{t,t+\tau}^{full}$ term in equation (10) cancels out with the one in equation (11) and hence knowledge of $r_{t,t+\tau}^{full}$ is not required. Second, to simulate dividends, we need to know the size of the break (drawn from a uniform distribution) and the probability of the break (drawn from a geometric distribution). Third, the COS securities are implicit in the risk-neutral density function. This is because $g_t = (1 - A_t) g_{t-1} + A_t g^* = g_{t-1} + A_t (g^* - g_{t-1})$ where g^* is the post break value of g.

2.2 An economy under partial information and Bayesian learning

We relax the assumption of full information and we assume that once a break in g_t occurs at time *t*, its new true value is unknown by the representative agent. The agent knows that dividends evolve according to equation (1), yet she cannot estimate the new mean dividend growth rate accurately because there is no available historical information immediately after a break. The agent 'learns' about the new values of g_t once a break occurs via a Bayesian updating procedure by observing the n historical dividend returns $\{D_i / D_{i-1}\}_{i=t-n}^{t}$ (signals) paid by the stock, where $n+1$ is the number of periods since the last break. Similar to Timmermann (2001), we assume that the agent does not know *ex-ante* the future dates of breaks (memoryless stochastic process), yet she realizes that a break in g_t occurs as soon as this happens.

Under Bayesian learning, asset prices at time *t* are computed as the conditional expected value $\mathrm{E}_t^{BL}\left[\lambda_t(\mu_t)|\boldsymbol{\xi}_t\right]$ given by (Timmermann, 2001)

$$\mathrm{E}_t^{BL}\left[\lambda_t(\mu_t)|\boldsymbol{\xi}_t\right] = \frac{\int_{\mu_d}^{\mu_u} \lambda_t^{full}(\mu_t) f^{\mathbb{P}}(\boldsymbol{\xi}_t \mid \mu_t) f^{\mathbb{P}}(\mu_t) d\mu_t}{\int_{\mu_d}^{\mu_u} f^{\mathbb{P}}(\boldsymbol{\xi}_t \mid \mu_t) f^{\mathbb{P}}(\mu_t) d\mu_t} \tag{12}$$

where $\lambda_t^{full}(\mu_t)$ is the value of $\lambda_t(\mu_t)$ under full information, $\boldsymbol{\xi}_t = \left[\ln\left(D_t / D_{t-1}\right)\dots \ln\left(D_{t-n} / D_{t-n-1}\right)\right]$ is the vector of historical n signals used to learn the μ_t value from the most recent break. $f^{\mathbb{P}}\left(\mu_t\right)$ is the probability density function of μ_t from which a new value of μ is drawn once a break occurs, i.e. $f^{\mathbb{P}}\left(\mu_t\right)$ is transformation of the univariate distribution $G(\bullet)$ from which the mean dividend growth rate is drawn once a break occurs. Given that the uniform probability density function is $\rho(g_t) = 1/(g_u - g_d)$, the corresponding probability density function $y\left(\mu_t\right)$ of μ_t is $y\left(\mu_t\right) = \exp\left(\mu_t + \sigma^2/2\right)/\left(g_u - g_d\right)$ where $\mu_d = \ln\left(1+g_d\right) - \sigma^2/2$ and $\mu_u = ln(1+g_u) - \sigma^2/2$. The probability $f^{\mathbb{P}}(\boldsymbol{\xi}_t|\mu_t)$ is the sample likelihood function is assumed to be normal, i.e.

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

which is a normal probability density function with mean $\bar{\xi}_t = \left(1/n\right)\sum_{i=t-n+1}^{t} \xi_i$ and variance σ^2/n, because the agent knows that historical signals follow the geometric random walk described in equation (1). We use equation (12) to compute put option prices under Bayesian learning; we describe the implementation in Section 3.1.

2.3 Expected option returns under Bayesian learning

We explain why learning about the true mean dividend growth rate can be an economic explanation to the size of options' expected returns. The hold-to-maturity $t+\tau$ put options expected returns $R_{t+\tau}^{p}$ and $R_{t+\tau}^{p,BL}$ under full information and Bayesian learning, respectively, are defined as

$$R^{p}_{t+\tau}=\frac{E^{\mathbb{P}}_{t}\left[\max\left(K-S_{t+\tau},0\right)\right]}{p_{t}\left(K,\tau\right)}-1=\frac{E^{\mathbb{P}}_{t}\left[\max\left(K-S_{t+\tau},0\right)\right]}{E^{\mathbb{Q}}_{t}\left[e^{-r\tau}\max\left(K-S_{t+\tau},0\right)\right]}-1, \tag{14}$$

$$R^{p,BL}_{t+\tau}=\frac{E^{\mathbb{P},BL}_{t}\left[\max\left(K-S_{t+\tau},0\right)|\,\boldsymbol{\xi}_{t}\right]}{E^{\mathbb{Q},BL}_{t}\left[e^{-r\tau}\max\left(K-S_{t+\tau},0\right)|\,\boldsymbol{\xi}_{t}\right]}-1 \tag{15}$$

The numerator in equations (14) and (15) is obtained under the physical probability measure $\mathbb{P}$ whereas the denominator is under the risk-neutral probability measure $\mathbb{Q}$. Consequently, differences between $\mathbb{P}$ and $\mathbb{Q}$ probability measures affect the size of hold-to-maturity put option expected returns. Under full information, the risk neutral probability measure $f^{\mathbb{Q}}\left(S_{t+1}\right)$ is related to the physical probability measure $f^{\mathbb{P}}\left(S_{t+1}\right)$ as

$$f^{\mathbb{Q}}(S_{t+1})=\frac{m_{t+1}f^{\mathbb{P}}(S_{t+1})}{\mathrm{E}_{t}\left[m_{t+1}\right]} \tag{16}$$

where m_{t+1} is the stochastic discount factor. However, under partial information and Bayesian learning, the $\mathbb{P}$ probability measure $f^{\mathbb{P},BL}(S_{t+1}\,|\,\boldsymbol{\xi}_{t})$ is conditional on the information $\boldsymbol{\xi}_{t}$ received after a given break. Hence, the risk neutral probability measure $f^{\mathbb{Q},BL}\left(S_{t+1}\right)$ is also conditional on $\boldsymbol{\xi}_{t}$, i.e.

$$f^{\mathbb{Q},BL}(S_{t+1}\,|\,\boldsymbol{\xi}_{t})=\frac{m_{t+1}f^{\mathbb{P},BL}(S_{t+1}\,|\,\boldsymbol{\xi}_{t})}{\mathrm{E}^{BL}_{t}\left[m_{t+1}\,|\,\boldsymbol{\xi}_{t}\right]} \tag{17}$$

We prove (see Appendix B) that equation (17) can be rewritten as

$$f^{\mathbb{Q},BL}(S_{t+1}\,|\,\boldsymbol{\xi}_{t})=\frac{m_{t+1}f^{\mathbb{P}}(S_{t+1}\,|\,\mu_{t})f^{\mathbb{P}}(\mu_{t}\,|\,\boldsymbol{\xi}_{t})}{\mathrm{E}^{BL}_{t}\left[m_{t+1}\,|\,\boldsymbol{\xi}_{t}\right]} \tag{18}$$

where $f^{\mathbb{P}}(\mu_t|\boldsymbol{\xi}_t)$ is the probability density function of μ_t conditional on the available information.

Equation (18) shows that learning affects the gap between $\mathbb{P}$ and $\mathbb{Q}$ probability measures. The right hand side of (18) shows that the agent updates her beliefs on μ_t via $f^{\mathbb{P}}(\mu_t\,|\,\boldsymbol{\xi}_t)$ given the information set $\boldsymbol{\xi}_t$. This yields an updated $\mathbb{P}$-probability measure $f^{\mathbb{P}}(S_{t+1}\,|\,\mu_t)$ with respect to future events and hence an updated $f^{\mathbb{Q},BL}(S_{t+1}\,|\,\boldsymbol{\xi}_t)$.

2.4 Model's properties

We discuss the properties of our model with a view on the effects on learning of the number *n* of signals and some key parameters. Consider the posterior probability density function $f^{\mathbb{P}}(\mu_t|\boldsymbol{\xi}_t)$ with mean μ_t^* and variance $\sigma^2_{\mu,t}$. The prior distribution density function $f^{\mathbb{P}}(\mu_t)$ of the agent is a normal with mean μ_0 and variance σ_0^2. Then, $f^{\mathbb{P}}(\mu_t|\boldsymbol{\xi}_t)$ is also normally distributed with mean and variance μ_t^* and $\sigma^2_{\mu,t}$, respectively given by

$$\mu_t^* = \bar{\xi}_t \frac{\frac{n}{\sigma^2}}{\frac{n}{\sigma^2}+\frac{1}{\sigma_0^2}} + \mu_0 \frac{\frac{1}{\sigma_0^2}}{\frac{n}{\sigma^2}+\frac{1}{\sigma_0^2}} \tag{19}$$

and

$$\sigma^2_{\mu,t} = \frac{1}{\frac{n}{\sigma^2}+\frac{1}{\sigma_0^2}}. \tag{20}$$

where $\bar{\xi}_t$ is the sample arithmetic mean of the signals. μ_t^* is a weighted average of the prior mean μ_0 and $\bar{\xi}_t$ where weights depend on the amount of received information. We can see that the agent learns about the true value of the parameter as she receives more signals over time; $\sigma^2_{\mu,t}$ (the inaccuracy of the estimation) decreases as n increases. Moreover, $\sigma^2_{\mu,t}$ increases immediately after a break because the counter of number of signals n used to learn is reset to zero; $n = 0$ yields $\mu_t^* = \mu_0$ in equation (19). Interestingly, learning does not make the market incomplete. This is because $\sigma^2_{\mu,t}$ is not stochastic. Instead, it decreases in a deterministic way with a reduction factor of order $1/n$ as more signals are received [equation (20)].

The model contains three key parameters that affect the way that agent learns. These are the probability π of the occurrence of a break in μ, the magnitude σ of the noise of signal and the representative agent's risk aversion η. We examine their effects via equations (19) and (20). First, in the case where $\pi = 0$ i.e. there are no breaks in the mean dividend growth rate, the agent can use infinitely many observations (i.e. $n \to \infty$) as $t \to \infty$ because the value of n is never reset to zero. Then, $\mu_t^* = \bar{\xi}$ and hence, μ_t^* converges to the true value μ given that $\bar{\xi}$ is a consistent estimator of μ in large samples under the normality assumption. Therefore, once a break occurs, the agent learns the true value μ_t asymptotically as in Guidolin and Timmermann (2003). Thus, asset prices converge to the ones given by the full information expressions presented in Proposition I and Proposition II. In the other extreme case where $\pi = 1$, i.e. the mean dividend growth rate presents breaks at every point in time, n is always zero and hence, $\mu_t^* = \mu_0$. This implies that the agent cannot learn the true value of the unknown parameter because there are no available signals. In the intermediate case where $0 < \pi < 1$, the agent never learns the true value of μ_t even asymptotically because the learning process for the mean growth rate is reinitiated after each break. Immediately after a break, the counter n of the number of signals restarts with a value equal to zero.

Second, the variance σ^2 of the process followed by dividends affects the way that the agents learns about the true parameter once a break occurs. Consider the case $\sigma \to \infty$. Then, $\mu_t^* = \mu_0$. Thus, similar to the case where $\pi = 1$, the agent cannot learn from the information received (even asymptotically) because signals are extremely noisy. On the other hand, $\sigma = 0$ (i..e. signals are not noisy) yields $\mu_t^* = \overline{\xi}$ and $\sigma_{\mu,t}^2 = 0$. In this case, the agent learns the true value of μ just after the break since there is no uncertainty. Third, the coefficient of relative risk aversion η determines the extent to which learning about g_t affects asset prices. For instance, learning about g_t does not affect the stock price when $\eta = 1$. This is because the expression $(1+g_t)^{1-\eta}$ will equal one in the full information case [equation (7)]. Consequently, the option price under Bayesian learning will not be affected by learning about g_t via the stock price either since g_t (and equivalently μ_t) will not appear in equation (12). Learning though does affect the price of the other assets even when $\eta = 1$.

3. Index put option returns: Simulations and empirical evidence

In this section, we simulate put options returns under Bayesian learning about fundamentals. Next, we compare the simulated put option returns to the ones empirically observed in S&P 500 index option data to assess whether our model can explain the empirical patterns of index put option returns.

3.1 Simulation: The setting

We simulate option returns in an economy with breaks and Bayesian learning in the setting explained in Section 2. To this end, first we simulate dividends and the prices S_t, B_t and A_t under the physical probability measure $\mathbb{P}$. For a given set of parameters, we perform 10,000 simulation runs; all simulation runs start from the same initial (January 2, 1996) $\mu = \bar{\xi}$. For each simulation run, we generate 12 years (3,024 trading days) of daily dividends. We simulate daily dividends in two steps. First, we simulate a time series of 12 years of daily dividends via equation (1). Then, we generate breaks in g_t (and thus breaks in μ_t) that we impose on the simulated in the first step 12-years dividends time series. To generate breaks in g_t, time periods between breaks follow a geometric distribution with parameter π. Once a break occurs, a new value for g_t is drawn from a uniform distribution defined on the support $[g_d, g_u]$ and this corresponds to a value for μ_t. Across each one of the 10,000 simulated dividend paths, we obtain the simulated prices for S_t, B_t and A_t on each one of the 3,024 trading days from equation (12).

Next, we obtain simulated European put option prices $p_t^{BL}(K,\tau)$ under Bayesian learning using equation (12); the risk-neutral probability measure $\mathbb{Q}$ is now required. Given equations (12), (18) and Proposition II, the put option price under Bayesian learning is given by

$$p_t^{BL}(K,\tau)=\int_{\mu_d}^{\mu_u}\int_0^{\infty}\frac{1}{1+r_{t,t+\tau}^{BL}}\max\{K-S_{t+\tau},0\}f^{\mathbb{Q},BL}(S_{t+\tau}\,|\,\boldsymbol{\xi}_t)dS_{t+\tau}d\mu_t$$

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

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

with

$$\begin{aligned}f^{\circ}(S_{t+\tau}\,|\,\mu_t)&=\frac{m_{t+\tau}f^{\mathbb{P}}(S_{t+\tau}\,|\,\mu_t)}{\mathrm{E}_t^{BL}[m_{t+\tau}\,|\,\boldsymbol{\xi}_t]}\\&=(1+r_{t,t+\tau}^{BL})\beta^{\tau}\left(D_{t+\tau}/D_t\right)^{-\eta}f\left(\varepsilon_{t+\tau}|0,\sigma\right)\varphi\left(z|\tau,\pi\right)\cdot\\&\quad\psi\left(h_1|\pi\right)\cdot\left[\psi\left(h_2|\pi\right)f^{\mathbb{P}}\left(\mu_2\right)\ldots\psi\left(h_z|\pi\right)f^{\mathbb{P}}\left(\mu_z\right)\right],\end{aligned} \qquad (22)$$

We use equation (21) to compute option prices on a monthly basis to obtain non-overlapping one-month option returns. On each month of each one of the 10,000 12-year simulated paths, we solve equation (21) by a two-step procedure. First, we solve the inner integral in equation (21) that contains the option price under full information. This integral is solved by Monte Carlo simulations based on M=20,000 independent paths of the stock price to obtain the put price under full information as described in Proposition II. We generate each of the 20,000 simulated paths by using the probability density $f^{\circ}(S_{t+\tau}\,|\,\mu_t)$ in equation (22). For each one of the 20,000 paths, we use the same initial value for μ_t and we obtain an analytical expression for the option payoff which depends of μ_t. Second, we integrate the analytical expression for the option payoff with respect to μ_t by using numerical (adaptive Simpson quadrature) integration; thus we solve the exterior integral that depends on μ_t in each one of the Monte Carlo's paths. Next, we average the outcome from the 20,000 paths to obtain 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. Given the computational intensity of the numerical methods, the simulations were performed in parallel by the use of 20 cores in clusters provided by our institutions running continuously for around 120 days.

We assume the following parameter values quoted on a monthly frequency to conduct simulations in line with Timmermann (2001). We set the lower and upper bound of the uniform distribution $G(\bullet)$ from which the monthly post-break mean dividend growth rate is drawn to $g_l = -0.126\%$ and $g_u = 0.705\%$, respectively. These values are in accordance with the values observed in market data in *real* terms; the average of the S&P 500 real dividend growth over 1996-2007 equals 0.510%. We also set the rate of impatience at 0.713%. This choice is necessary to satisfy the condition $1+\rho > (1-\pi)(1+g_u)^{1-\alpha}$ in order to ensure a positive stock price in equation (7) for any combination of parameters. We set the volatility of the geometric random-walk process to 1.44%. We consider three different values for the coefficient of relative risk aversion $\eta = 0.2, 0.5, 5$. We set the probability of breaks $\pi = 0.056$. This follows from applying the Chu et al. (1996) dynamic test for structural breaks to daily real dividends from the S&P 500 index over the period 1996-2007. We detect eight breaks in the mean dividend growth rate over the 3,024 days of the 12-year period.

3.2 Index options returns: Evidence from simulated and actual option data

In line with Broadie *et al*. (2009) and Chambers *et al*. (2014), we create a time series of non-overlapping simulated hold-to-maturity index put options returns r_{t+T}^{p} for different strikes and maturities for each simulated dividend path as

$$r_{t+\tau}^{p} = \frac{\max\left(K - S_{t+\tau}, 0\right)}{p_t^{BL}\left(K, \tau\right)} - 1 \tag{23}$$

where $p_t^{BL}\left(K,\tau\right)$ is the price of a K-strike τ-maturity put option under Bayesian learning obtained from equation (21) where $S_{t+\tau}$ is the stock price at time $t+\tau$. Moreover, we compare the size and pattern of the simulated put option returns to the S&P 500 index put options empirical ones. To this end, we obtain daily S&P 500 European put index options data from the OptionMetrics Ivy DB database spanning 1996 – 2007. The database contains daily closing bid and ask option prices, Black-Scholes implied volatilities, time-to-maturity, strike price, closing underlying index price, and the risk-free interest rate. Option prices correspond to closing bid-ask midpoints. In line with Bernales and Guidolin (2014), we exclude options whose prices violate arbitrage bounds, their ask price is less than the bid price, their bid price equals zero, their price is less than $3/8 to avoid the effects of price discreteness since the smallest tick size is $1/16 for the S&P 500 index options and which have zero open interest.

Table 1 reports summary statistics for the simulated option returns for three different values of the agent's relative risk aversion ($\eta = 0.2$, 0.5 and 5.0, Panels A, B and C, respectively). Entries report the average moment (mean, volatility, skewness, kurtosis) over the 10,000 simulated dividend paths of hold-to-maturity put option returns for different moneyness levels (ranging from 0.96 to 1.02) and time-to-maturities (30 days and 90 days-to-maturity); we focus on this particular range of moneyness levels because most of the options trading activity occurs there (Broadie *et al*., 2009). *t*-statistics and *p*-values computed under the null hypothesis that returns equal zero are also reported. Table 1 also reports summary statistics for the empirical option returns (Panel D). We compute average option returns for the same moneyness and time-to-maturity levels used in the simulation setting. To compute the fixed moneyness level's option return, we interpolate linearly across the option returns of the two options that surround the targeted moneyness level. Then, to compute the fixed maturity option return, we interpolate linearly across the obtained fixed moneyness option returns of the two options whose times-to-maturity surround the targeted maturity level. From now onwards, we report results from the analysis on the 90-days returns by first converting the 90-days returns to 30-days returns for comparison purposes.

[Insert Table 1 here]

We can see that the simulated index put option returns are large and negative. They are statistically significant and their distribution is far from being normal. In addition, their magnitude decreases as the moneyness level increases. For example, a one-month-to-maturity put option with ratio $K/S = 0.96$ ($K/S = 1.00$) has an average monthly return of −95.71% (−54.82%), −97.08% (−38.29%) and −99.55% (− 92.69%) for coefficients of relative risk aversion at 0.2, 0.5 and 5.0, respectively. Option returns also decrease in magnitude as the time-to-maturity increases for any given moneyness level. For instance, an at-the-money put option has an average return of -54.82% and -12.93% when the time-to-maturity is 30- and 90-days, respectively, for the case where $\eta = 0.2$. These findings are in line with those reported in previous studies (see, e.g., Coval and Shumway, 2001, Bondarenko, 2003, Broadie *et al*., 2009; and Constantinides et *al*., 2013).

Furthermore, we can see that the patterns of the S&P 500 empirical put option returns are similar to the simulated ones. The magnitude of empirical option returns decreases as the moneyness level and time-to-maturity increases. For instance, a three-month-to-maturity S&P 500 put option contract with ratio $K/S = 0.96$ ($K/S = 1.00$) has an average return of -23.87% (-15.47%). In addition, an at-the-money option has an average return of -27.18% and -15.47% when the time-to-maturity is 30-days and 90-days. The empirical put returns are large, negative, and statistically different from zero, just as it was the case with the simulated put option returns under our economy with learning.

4. Returns on leverage-adjusted option portfolios

The returns of naked put options are affected by their leverage. We expect greater option returns for option contracts that have a greater leverage; this is the case especially for out-of-the-money option contracts with short-term time-to-maturities as shown in Table 1 for our simulated and empirical option average returns. In this section, we examine whether our results for naked puts carry over once we control for leverage. To this end, we construct leverage-adjusted portfolios.

4.1 Formation of option strategies

To construct leverage-adjusted portfolios, we assume that Black-Scholes (1973) and Merton's (1973) assumptions hold in line with previous studies which examine leverage-adjusted index option returns. Then, the instantaneous option returns are related to the underlying asset's returns as follows

$$\frac{df\left(S_t\right)}{f\left(S_t\right)}=rdt+\frac{S_t}{f\left(S_t\right)}\frac{\partial f\left(S_t\right)}{\partial S_t}\left(\frac{dS_t}{S_t}-\left(r-\delta\right)dt\right) \tag{24}$$

where S_t is the underlying asset price, r the risk-free rate, $f(S_t)$ the option price, and δ the continuous dividend yield paid by the underlying asset. Equation (24) is re-arranged to

$$\omega^{-1}\frac{df\left(S_t\right)}{f\left(S_t\right)}+\left(1-\omega^{-1}\right)rdt=\frac{dS_t}{S_t}+\delta dt \tag{25}$$

where $\omega=\left(df\left(S_t\right)/\partial S_t\right)\left(S_t\,/\,f\left(S_t\right)\right)$ is the option's elasticity with respect to the underlying asset price. Equation (25) shows that we can construct a leverage-adjusted portfolio that earns the return of the underlying asset plus its dividend yield by investing ω^{-1} in an option contract and $1-\omega^{-1}$ in the risk-free rate. This implies that 'leverage-adjusted' portfolios with the same underlying asset should earn the *same* return (i.e. that of the underlying asset) across different levels of moneyness and time-to-maturities. In addition, we construct delta-hedged portfolios. Delta-hedged portfolios with the same underlying asset are also free of leverage effects because they earn the same return (equal to the risk-free rate) across different levels of moneyness and time-to-maturities.

Inevitably, the construction of the leverage-adjusted portfolios under equation (25)) is not model-free and relies on the validity of the Black-Scholes-Merton assumptions which admittedly do not hold in our setting. However, this model error effect is not a concern for the purposes of our study. This is because the previous academic literature has also computed option returns of leverage-adjusted and delta-hedged portfolios by adopting the same assumptions (e.g., Jones, 2006, Gonçalves and Guidolin, 2006, and Constantinides *et al.*, 2013). This enables us to compare our subsequently obtained results to these of the previous literature on an equal footing. In addition, the construction of these strategies and hence the computation of option returns is in line with the industry practice. Interestingly, albeit equation (25) is not model-free, option's delta and option's elasticity formulas are model-free for a wide family of popular option pricing models that price plain vanilla options, including the Black and Scholes model, and option models with jump-diffusion and stochastic volatility (Alexander and Nogueira, 2007). This alleviates concerns on the effect of model error to the construction of the delta-hedged portfolios.

4.2 Evidence from simulated and S&P 500 options data

We calculate simulated and empirical S&P 500 non-overlapping hold-to-maturity returns of leverage-adjusted and delta-hedged option portfolios. Table 2 reports summary statistics for 30-days and 90-days simulated hold-to-maturity non-overlapping returns on leverage-adjusted and delta-hedged option portfolios (Panels A and B, respectively) as well as their empirical counterparts (Panels C and D, respectively). In the interests of brevity, we shall report results for the case where the relative risk aversion η equals 0.2; we obtain qualitatively similar results for the cases where $\eta = 0.5, 5$.

[Insert Table 2 here]

We can see that the magnitude of average simulated returns of leverage-adjusted and delta-hedged option portfolios decreases as the moneyness and time-to-maturity of put options employed in the portfolios increase. For instance, leverage-adjusted strategies for put options with ratio $K/S = 0.96$ ($K/S = 1.00$) have 3.09% (2.45%) and 1.37% (1.21%) one- and three-month average returns, respectively. This is in line with the results of previous literature on the patterns of leverage-adjusted (Constantinides *et al*., 2013) and delta-hedged (Bakshi and Kapadia, 2003, and Broadie *et al.,* 2009) option positions.

Moreover, we can see that the patterns of the empirically observed average returns in both option portfolios are in accordance with the ones generated by our model. For any given maturity, the magnitude of average returns decreases as the moneyness increases. Similarly, for a given moneyness level, average returns decrease as the time-to-maturity increases. Therefore, our setting generates put option returns similar to the empirically observed in terms of their size and patterns even when we control for leverage and market risk.

4.3 CAPM alphas and Sharpe ratios under learning

Previous literature documents that empirically observed index put option returns are too big to be explained by the CAPM model (e.g., Coval and Shumway, 2001; and Bondarenko, 2003). Typically, this is manifested by high Sharpe ratios and statistically significant CAPM alphas. We investigate whether our learning model also generates option returns that exhibit these properties. In case it does, then this would be additional evidence that option returns may be the result of a learning model with breaks in the mean dividend growth. Once we simulate option returns for a given option strategy (i.e. the naked put strategy, the leverage-adjusted put option portfolio or the delta-hedged put option portfolio), we estimate the CAPM alpha α by running the following regression:

$$(r_{F,t} - r_t) = \alpha_F + \beta_F (r_{m,t} - r_t) + \varepsilon_{F,t} \tag{26}$$

where $r_{F,t}$ is the simulated return of a given option portfolio at time *t*. In the case of simulations generated by our option learning model, r_m equals the excess return of the underlying asset (adjusted for dividends) obtained from equation (12) because this is the only stock in the economy and r is the risk free rate calculated as the inverse of the bond price under partial information obtained from equation (12). In the case of market option data, $r_{F,t}$ is the market return of the considered option portfolio, $r_{m,t}$ the S&P 500 index excess return (adjusted for dividends) and r_t is the one-month LIBOR rate.

Table 3 reports the CAPM alphas and Sharpe ratios estimated from simulated hold-to-maturity 30- and 90-days returns of index put option, leverage-adjusted put, delta-hedged put and straddle (Panels A, B and C, respectively). Panels D, E and F report the empirical counterparts for the three respective strategies. In the case of simulated data, averages are taken over the 10,000 simulated dividend paths. Simulations are carried out assuming a representative agent's risk aversion $\eta = 0.2$; results are qualitatively similar for $\eta = 0.5$, 5 and are available from the authors upon request. The alpha's *t*-statistic value is also reported. This is computed by Newey-West standard errors to correct for heteroscedasticity and serial correlation in the residual term of equation (26).

[Insert Table 3 here]

We can see that simulated CAPM alphas are statistically significant and Sharpe ratios are high in absolute terms ranging up to 3.78. They decrease as the moneyness and time-to-maturity increases. Moreover, the sign of empirical CAPM alphas and Sharpe ratios and their patterns are similar to these generated by our option learning model. These findings are also similar to the ones reported in previous empirical studies (e.g., Coval and Shumway, 2001, Bondarenko, 2003, Bollen and Whaley, 2004, Constantinides *et al.*, 2013, amongst others). In sum, our results on CAPM alphas and Sharpe ratios provide additional evidence that our learning model can provide an explanation for the behaviour of index put option strategies' empirical returns.

5 Further robustness tests

5.1 Is learning necessary to reproduce empirical option returns?

One may argue that the values of CAPM alphas and Sharpe ratios in Table 3 (especially in leverage-adjusted and delta-hedged portfolios) can be induced only by making g_t subject to breaks rather than by also introducing a learning mechanism. To assess this conjecture, we estimate CAPM alphas and Sharpe ratios for leverage-adjusted and delta-hedged portfolios under two cases: (i) full information with no breaks and (ii) full information with breaks. (i) and (ii) are special cases of the model presented in Section 2. The case of full information and no breaks is Lucas (1978) model where option prices of European contracts are obtained from the Black and Scholes (1973) model. In the case of full information with breaks, we calculate asset prices using Propositions I and II.

Table 4 reports CAPM alphas and Sharpe ratios obtained from simulated option returns, under (i) and (ii), for the leverage-adjusted and delta-hedged option portfolios (Panels A and B, respectively). Entries report average estimates over 10,000 simulations. We calculate 30-day and 90-day hold-to-maturity returns for option contracts with strike-to-price ratios ranging from 0.96 to 1.02.

[Insert Table 4 here]

We can see that a setting which does not incorporate learning about fundamentals and allows only for breaks in g_t cannot explain the empirically observed index put option returns. The *t*-statistics for CAPM alphas are marginally greater in the case of full information and breaks compared to the no-breaks case (signaling a small effect of breaks on returns of put option portfolios), yet there is still a small percentage of simulations with significant alphas. These findings hold for both leverage-adjusted and delta-hedged option portfolios and across the various moneyness and time-to-maturity levels. They are in sharp contrast to the ones reported in Table 3 for the CAPM alphas and Sharpe ratios obtained from simulated under the learning setting and empirical option data. This implies that a full information model with breaks in the mean dividend growth cannot explain the empirically observed index put strategy returns and learning should also be incorporated.

5.2 Option returns, multifactor models and learning

In Section 4.3, we documented that our model can explain the empirically estimated CAPM alphas and Sharpe ratios. In this section, we investigate whether our model generates option returns which can be explained by factors that have been documented to describe option returns. This will provide further support to our economic explanation for index option returns because it will highlight that learning can account for both the empirical regularities in index put option returns as well as for the relation between option returns and other market factors.

We examine whether option returns are related to the market risk premium, the volatility risk premium, the slope $Slope_{Mon}$ of implied volatilities as a function of moneyness, and the slope $Slope_{Mat}$ of implied volatilities as a function of time-to-maturity (implied volatility term-structure). In line with Bollerslev (2009), at the end of each month we estimate the volatility risk premium as the difference between the option implied volatility IV and the realized volatility RV prevailing over the previous month. We extract IV from an at-the-money European put option with one-month to maturity. We calculate RV as the annualized standard deviation of the daily stock log-returns over each month to avoid overlapping periods. Our model generates a 12%, 6% and 46% variance risk premium per month on average over 1996-2007 for respective coefficients $\eta = 0.2, 0.5, 5$ of relative risk aversion of the representative agent. We compute $Slope_{Mon}$ as the difference between the implied volatility of a 30-day put option contract with $K/S = 0.96$ and that of a 30-day option contract with $K/S = 1.04$. We compute $Slope_{Mat}$ as the difference between the implied volatility of an at-the-money contract with 30-days to maturity and that of an at-the-money contract with 90-days to maturity. In the case of the empirical analysis with S&P 500 option contracts, we compute the fixed moneyness and maturity implied volatilities by interpolating linearly across the implied volatilities extracted from options that surround the targeted levels.

Table 5 reports results from regressing simulated option returns on excess market returns, the variance risk premium VRP_t, $Slope_{Mon}$ and $Slope_{Mat}$ for the case of leverage-adjusted and delta-hedged portfolios (Panels A and B, respectively). Panels C and D report the factors' coefficients obtained from running regressions with empirical data for the case of leverage-adjusted and delta-hedged portfolios, respectively. We report results for 30-day time-to-maturity options for moneyness at 0.96 and 1.00. t-statistics for the empirical regressions are reported in square brackets.

[Insert Table 5 here]

We can see that the volatility risk premium is an important factor in explaining simulated option returns. In at least 80% of the simulations, VRP_t is statistically significant for both types of portfolios and for both moneyness levels. Moreover, the percentage of simulations where CAPM alpha is statistically significant as well as the size of alpha is reduced significantly in the case where VRP_t is incorporated as a regressor. For instance, in the case of a leverage-adjusted strategy with moneyness 1.00, alpha is reduced from 0.02 when VRP_t is not used (where 100% of simulations have significant alpha values) to 0.00 when we include VRP_t as a factor (where only 12% of simulations comprise significant alphas). This echoes Broadie *et al.* (2009) who argue that option returns can be explained by parameter mis-estimation that generates a difference between the risk-neutralized and physical probability measures; the variance risk premium also arises as a result of this difference. The $Slope_{Mon}$ and $Slope_{Mat}$ explain partially the option portfolio returns; they are statistically significant for a smaller percentage of simulations compared to the evidence for the variance risk premium's case. Regarding the evidence from empirical data, we can see that the volatility risk premium is statistically significant. This is in line with the results obtained from the simulated under our learning model option returns and in line with the results reported by Broadie *et al.* (2009). On the other hand, the coefficients of $Slope_{Mon}$ and $Slope_{Mat}$ are marginally significant in just a few cases.

6. Conclusions

The difference between $\mathbb{P}$ and $\mathbb{Q}$ probability measures affects expected option returns. One way to generate this gap and hence to explain expected option returns is by allowing for parameters' mis-estimation. Rather than taking a reduced form approach to model parameters' mis-estimaiton, we endogenize parameter uncertainty as a function of learning about economic fundamentals (dividends). This micro-founded approach induces gaps between $\mathbb{P}$ and $\mathbb{Q}$ probability measures and therefore it is a natural candidate to explain patterns in put index options returns. To the best of our knowledge, the study of expected option returns within a learning about fundamentals setting is novel.

We develop an equilibrium option pricing model where there are breaks in the mean dividend growth rate. We invoke the model within a Bayesian learning setting where the agent starts learning about the true mean of the dividend process once a break occurs. Our model delivers index put option returns and CAPM alphas that have the same patterns with the empirically observed S&P 500 put index options returns across different levels of moneyness and time-to-maturity. We also show that the volatility risk premium explains a significant portion of returns on option trading strategies. This is in accordance with Broadie *et al.*'s (2009) results who find that a difference between risk-neutral and physical probability measures induced by estimation risk can explain the returns of index put option portfolios. Results are robust across different option strategies that control for leverage and market risk and they are in line with previous empirical literature.

Appendix A: Proof of Proposition II

Proof of Proposition II: To obtain equations (10) and (11), we divide both sides of the Euler equation (4) used to price the stock at $t + k$, by the bond price (equation (8))

$$\frac{(1+\rho)S_{t+k}}{(1-\pi)(1+g_{t+k})^{-\eta} + \pi \int_{g_d}^{g_u}(1+g_{t+k})^{-\eta} dG(g_{t+k})} = \mathrm{E}_{t+k}\Bigg[\beta\left(\frac{D_{t+\mathrm{k}+1}}{D_{t+k}}\right)^{-\eta} \tag{A.1}$$
$$\cdot (S_{t+k+1} + D_{t+k+1}) \frac{(1+\rho)}{(1-\pi)(1+g_{t+k})^{-\eta} + \pi \int_{g_d}^{g_u}(1+g_{t+k})^{-\eta} dG(g_{t+k})}\Bigg]$$

Under full information, the future stock price and the future cumulative dividend values are given by:

$$S_{t+k}^{*} = \frac{(1+\rho)S_{t+k}}{(1-\pi)(1+g_{t+k})^{-\eta} + \pi \int_{g_d}^{g_u}(1+g_{t+k})^{-\eta} dG(g_{t+k})} \tag{A.2}$$

and

$$D_{t+k}^{*} = \sum_{\mathrm{s}=0}^{\mathrm{k}} D_{t+s} \frac{(1+\rho)}{(1-\pi)(1+g_{t+s})^{-\eta} + \pi \int_{g_d}^{g_u}(1+g_{t+s})^{-\eta} dG(g_{t+s})} \tag{A.3}$$

Under the power utility assumption, the definition of the pricing kernel yields:

$$\mathrm{E}_t\left[\beta\left(\frac{D_{t+\mathrm{k}+1}}{D_{t+k}}\right)^{-\eta} \frac{(1+\rho)}{(1-\pi)(1+g_{t+k})^{-\eta} + \pi \int_{g_d}^{g_u}(1+g_{t+k})^{-\eta} dG(g_{t+k})}\right] = 1 \tag{A.4}$$

Adding D_{t+k}^{*} to (A.1) and combining the resulting expression with (A.4) yields

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

$$= \mathrm{E}_{t+k}\left[\beta\left(\frac{D_{t+\mathrm{k}+1}}{D_{t+k}}\right)^{-\eta}\frac{(1+\rho)}{(1-\pi)(1+g_{t+k})^{-\eta}+\pi\int_{g_d}^{g_u}(1+g_{t+k})^{-\eta}dG(g_{t+k})}(S_{t+k+1}^{*} + D_{t+k+1}^{*})\right] \tag{A.5}$$

Equation (A.4) shows that $S_{t+k}^{full^*} + D_{t+k}^{*}$ follows a martingale under the probability measure. Thus, the risk-neutral density is:

$$f^{\mathbb{Q}}(S_{t+k})$$

$$= \beta\left(\frac{D_{t+\mathrm{k}+1}}{D_{t+k}}\right)^{-\eta}\frac{(1+\rho)}{(1-\pi)(1+g_{t+k})^{-\eta}+\pi\int_{g_d}^{g_u}(1+g_{t+k})^{-\eta}dG(g_{t+k})}f_t(D_{t+k+1}) \tag{A.6}$$

$$= (1+r_{t+k})\beta\left(\frac{D_{t+\mathrm{k}+1}}{D_{t+k}}\right)^{-\eta}f_t(D_{t+k+1})$$

where r_{t+k} is the one-period risk-free rate under full information.

Therefore, the risk-neutral measure for any single-period is unique and exists, which is a sufficient condition to have a unique risk-neutral measure on an infinite-period economy obtained by repetition of several single-periods (Pliska, 1997). Following this, we 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:

$$E_t[D_{t+\tau}] = D_t E_t\left[\frac{D_{t+1}}{D_t}E_{t+1}\left[\left(\frac{D_{t+2}}{D_{t+1}}\right)\dots E_{t+\tau-1}\left[\left(\frac{D_{t+\tau}}{D_{t+\tau-1}}\right)\right]\right]\right] \tag{A.7}$$

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

$$E_t[D_{t+\tau}] = D_t E_t \left[\exp\left(\sqrt{\tau}\sigma\varepsilon_{t+\tau} - \tau\sigma^2/2\right) \prod_{i=1}^{\tau} (1 + g_{t+i-1}) \right] \tag{A.8}$$

Let us consider z the number of breaks between t and $t+\tau$ drawn from a binomial distribution, $\varphi(z|\tau,\pi)$, with parameters τ and π; while $\{h_i\}_{i=1}^{z}$ are the time periods between breaks which are also random variables that follow a geometric distribution with parameter π, $\psi(h_i|\pi)$, where $\tau = \sum_{i=1}^{z} h_i$. Therefore, in each path we have:

$$D_{t+\tau}^{FI} = D_t \exp\left(\sqrt{\tau}\sigma\varepsilon_{t+\tau} - \tau\sigma^2/2\right) \prod_{i=1}^{z} (1 + g_i)^{h_i} \tag{A.9}$$

where g_i is constant between breaks, while post-break g_i is drawn from a continuous univariate density $G(\cdot)$ with pdf $\varrho(g_i)$ and g_l and g_u as the lower and upper bounds, in which g_1 (g_z) is the current (at time-to-maturity) level of the mean dividend growth rate. Therefore:

$$f_t(D_{t+\tau}) = \phi(\varepsilon_{t+\tau}|0,\sigma)\varphi(z|\tau,\pi)\psi(h_1|\pi)[\psi(h_2|\pi)\varrho(g_2) \ldots \psi(h_z|\pi)\varrho(g_z)] \tag{A.10}$$

Thus, from Equation (A.6) we have:

$$\begin{aligned} f^{\mathbb{Q}}(S_{t+\tau}) = (1 + r_{t+\tau})\beta^{\tau} \left(\frac{D_{t+\tau}}{D_t}\right)^{-\eta} \phi(\varepsilon_{t+\tau}|0,\sigma)\varphi(z|\tau,\pi) \\ \cdot \psi(h_1|\pi)[\psi(h_2|\pi)\varrho(g_2) \ldots \psi(h_z|\pi)\varrho(g_z)] \end{aligned} \tag{A.11}$$

where $1 + r_{t+\tau} = \prod_{j=1}^{\tau}(1 + r_{j-1,j})$ with $1 + r_{j-1,j} = 1/B_{j-1}$ where B_{j-1} is the price of the risk-free one-period bond in the period $j-1$ defined in Preposition I.

Appendix B: Proof of Equation (18)

$f^{\mathbb{P},BL}(S_{t+1}\,|\,\boldsymbol{\xi}_t)$ can be expressed as

$$f^{\mathbb{P},BL}(S_{t+1}\,|\,\boldsymbol{\xi}_t)=\frac{f^{\mathbb{P}}(\boldsymbol{\xi}_t\,|\,S_{t+1})f^{\mathbb{P}}(S_{t+1})}{f^{\mathbb{P}}(\boldsymbol{\xi}_t)} \tag{B.1}$$

Given that the $\mathbb{P}$ -probability measure is conditional on $\boldsymbol{\xi}_t$ through the unknown parameter μ_t, we use multivariate Bayes rule and re-write equation (B.1) as

$$f^{\mathbb{P},BL}(S_{t+1}\,|\,\boldsymbol{\xi}_t)=\frac{f^{\mathbb{P}}(\boldsymbol{\xi}_t\,|\,S_{t+1},\mu_t)f^{\mathbb{P}}(\mu_t\,|\,S_{t+1})f^{\mathbb{P}}(S_{t+1})}{f^{\mathbb{P}}(\boldsymbol{\xi}_t)} \tag{B.2}$$

since $f^{\mathbb{P}}(\boldsymbol{\xi}_t\,|\,S_{t+1})=f^{\mathbb{P}}(\boldsymbol{\xi}_t\,|\,S_{t+1},\mu_t)f^{\mathbb{P}}(\mu_t\,|\,S_{t+1})$. The Euler equation (4) shows that the stock price depends on μ_t. Hence, $f^{\mathbb{P}}\left(\xi_t|S_{t+1},\mu_t\right)=f^{\mathbb{P}}\left(\xi_t|\mu_t\right)$. Also, Bayes rule yields $f^{\mathbb{P}}\left(\mu_t|S_{t+1}\right)f^{\mathbb{P}}\left(S_{t+1}\right)=f^{\mathbb{P}}(S_{t+1}\,|\,\mu_t)f^{\mathbb{P}}\left(\mu_t\right)$. hen, equation (B.2) is rewritten as

$$f^{\mathbb{P},BL}(S_{t+1}\,|\,\boldsymbol{\xi}_t)=\frac{f^{\mathbb{P}}(S_{t+1}\,|\,\mu_t)f^{\mathbb{P}}(\boldsymbol{\xi}_t\,|\,\mu_t)f^{\mathbb{P}}(\mu_t)}{\int_{\mu_d}^{\mu_u}f^{\mathbb{P}}(\boldsymbol{\xi}_t\,|\,\mu_t)f^{\mathbb{P}}(\mu_t)d\mu_t} \tag{B.3}$$

since $f^{\mathbb{P}}\left(\xi_t\right)=\int_{\mu_d}^{\mu_u}f^{\mathbb{P}}(\xi_t\,|\,\mu_t)f^{\mathbb{P}}\left(\mu_t\right)d\mu_t$. $f^{\mathbb{P}}\left(\mu_t\right)$ is the prior distribution of μ_t and $f^{\mathbb{P}}(S_{t+1}\,|\,\mu_t)$ is also known given equation (7). Given equation (B.3), equation (17) becomes

$$\begin{aligned}f^{\mathbb{Q},BL}(S_{t+1}\,|\,\boldsymbol{\xi}_t)&=\frac{m_{t+1}}{\mathrm{E}_t^{BL}\left[m_{t+1}\,|\,\boldsymbol{\xi}_t\right]}f^{\mathbb{P}}(S_{t+1}\,|\,\mu_t)\frac{f^{\mathbb{P}}(\boldsymbol{\xi}_t\,|\,\mu_t)f^{\mathbb{P}}(\mu_t)}{\int_{\mu_d}^{\mu_u}f^{\mathbb{P}}(\boldsymbol{\xi}_t\,|\,\mu_t)f^{\mathbb{P}}(\mu_t)d\mu_t}\\&=\frac{m_{t+1}f^{\mathbb{P}}(S_{t+1}\,|\,\mu_t)f^{\mathbb{P}}(\mu_t\,|\,\boldsymbol{\xi}_t)}{\mathrm{E}_t^{BL}\left[m_{t+1}\,|\,\boldsymbol{\xi}_t\right]}\end{aligned} \tag{B.4}$$

We derive the last line in (B.4) by applying Bayes rule.

References

Abel, A., 1988, Stock prices under time-varying dividend risk: An exact solution in an infinite-horizon general equilibrium model, *Journal of Monetary Economics* 22, 375–393.

Alexander, C., and Nogueira, L.M. (2007), Model-free hedge ratios and scale-invariant models, Journal of Banking and Finance 31, 1839-1861.

Bakshi, G., and N. Kapadia, 2003, Delta-hedged gains and the negative market volatility risk premium, *Review of Financial Studies* 16, 527–566.

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., Guilolin, 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.

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.

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

Buraschi, A., F. Trojani, and A. Vedolin, 2014, When uncertainty blows in the orchard: Comovement and equilibrium volatility risk premia, *Journal of Finance*, 69, 101-137.

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.

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.

Gonçalves, S., and M. Guidolin, 2006, Predictable dynamics in the S&P 500 index options implied volatility surface, *Journal of Business* 79, 1591–1635.

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.

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.

Lettau, M., and S. Van Nieuwerburgh, 2008, Reconciling the return predictability evidence, *Review of Financial Studies* 21, 1607–1652.

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

Pesaran, M. H., D. Pettenuzzo, and A. Timmermann (2006). Forecasting time series subject to multiple structural breaks. *Review of Economic Studies* 73, 1057–1084.

Santa-Clara, P., and A. Saretto, 2009, Option strategies: good deals and margin calls, *Journal of Financial Markets* 12, 291–417.

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

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

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

Table 1: Summary statistics for put options simulated and empirical returns

This table displays summary statistics for the simulated and empirical returns of naked S&P 500 puts spanning 1996-2007. Simulations are based on a discrete-time economy with breaks in the stock's dividend process and partial information under a Bayesian recursive process. Entries report statistics for values of relative risk aversion $\eta = 0.2, 0.5, 5$. Average 30- and 90-day hold-to-maturity non-overlapping returns are computed for naked put options with strike-to-price ratios ranging from 0.96 to 1.02. p-values are computed under the null hypothesis that average returns are zero. Entries for simulated data report the average estimates over 10,000 simulations; for each of these simulations, we generate 12 years of daily dividends. The percentage of simulations with a significant mean return is reported in parentheses at a 5% level of significance. We report results from the analysis on the 90-days returns by first converting the 90-days returns to 30-days returns for comparison purposes.

	30 days to expiration				90 days to expiration			
K/S	0.96	0.98	1.00	1.02	0.96	0.98	1.00	1.02
	Panel A: Simulated R_{put}, η =0.2							
Mean	-95.71%	-83.80%	-54.82%	-22.78%	-27.82%	-22.04%	-12.93%	-6.69%
t-stat	-182.77	-18.73	-8.48	-3.85	-27.83	-3.72	-2.14	-1.21
p-value	0.00	0.00	0.00	0.01	0.13	0.12	0.19	0.32
	(100%)	(99%)	(99%)	(96%)	(70%)	(67%)	(46%)	(20%)
Volatility	102.84%	83.69%	76.07%	63.39%	90.93%	74.83%	65.56%	54.70%
Skewness	10.85	4.99	2.03	0.64	8.68	5.78	3.25	1.79
Kurtosis	107.26	31.11	7.88	2.88	82.77	35.75	13.14	4.28
	Panel B: Simulated R_{put}, η =0.5							
Mean	-97.08%	-73.18%	-38.29%	-10.33%	-9.58%	-7.95%	-5.02%	-2.04%
t-stat	-110.70	-9.95	-4.45	-1.51	-7.26	-1.58	-0.73	-0.28
p-value	0.02	0.01	0.02	0.02	0.24	0.36	0.43	0.53
	(96%)	(95%)	(94%)	(92%)	(57%)	(30%)	(14%)	(6%)
Volatility	93.13%	72.86%	102.77%	53.72%	80.85%	72.84%	60.11%	50.55%
Skewness	10.61	5.55	1.95	0.57	9.36	6.02	3.45	1.81
Kurtosis	116.46	37.38	6.91	2.74	90.19	41.08	13.63	4.16
	Panel C: Simulated R_{put}, η =5.0							
Mean	-99.55%	-95.50%	-92.69%	-64.65%	-57.55%	-57.25%	-53.22%	-20.42%
t-stat	-142.91	-146.08	-59.71	-32.16	-37.75	-26.43	-26.59	-10.23
p-value	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.02
	(100%)	(99%)	(98%)	(98%)	(98%)	(97%)	(96%)	(96%)
Volatility	80.56%	61.01%	30.06%	21.78%	71.21%	53.17%	23.06%	21.84%
Skewness	9.41	9.33	7.20	4.18	8.63	8.82	7.65	5.18
Kurtosis	97.92	93.45	64.64	33.09	79.17	76.70	65.70	37.05
	Panel D: Empirical S&P 500 *Rput*							
Mean	-68.72%	-36.02%	-27.18%	-22.17%	-23.87%	-18.56%	-15.47%	-11.37%
t-stat	-10.64	-4.61	-4.59	-4.27	-8.58	-6.88	-6.25	-3.83
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Volatility	139.41%	173.76%	136.07%	107.67%	35.19%	38.23%	37.30%	35.00%
Skewness	6.01	4.04	2.79	1.83	5.05	3.50	2.84	2.07
Kurtosis	45.80	22.59	13.71	7.69	33.22	17.52	11.97	7.64

Table 2: Summary statistics for simulated and empirical returns of leverage-adjusted and delta-hedged put options

This table displays summary statistics for simulated and empirical returns of put option portfolios that control for leverage and market risk spanning 1996-2007. Simulations are based on a discrete-time economy with breaks in the stock's dividend process and partial information under a Bayesian recursive process for a coefficient of relative risk aversion $\eta = 0.2$. Average 30- and 90-day hold-to-maturity non-overlapping returns R_{put} and R_{DHput} are computed for leverage-adjusted and delta-hedged put option portfolios, respectively, with strike-to-price ratios ranging from 0.96 to 1.02. p-values are computed under the null hypothesis that average returns are zero. Entries for simulated data report the average estimates over 10,000 simulations; for each of these simulations, we generate 12 years of daily dividends. The percentage of simulations with a significant mean return is reported in parentheses at a 5% level of significance. We report results from the analysis on the 90-days returns by first converting the 90-days returns to 30-days returns for comparison purposes.

	30 days to expiration				90 days to expiration			
K/S	0.96	0.98	1.00	1.02	0.96	0.98	1.00	1.02
	Panel A: Simulated R_{put} (leverage-adjusted)							
Mean	3.09%	2.98%	2.45%	1.66%	1.37%	1.30%	1.21%	1.09%
t-stat	47.47	25.50	13.16	7.70	18.69	9.26	6.57	5.49
p-value	0.00	0.00	0.00	0.00	0.02	0.00	0.00	0.00
	(100%)	(100%)	(100%)	(100%)	(98%)	(98%)	(98%)	(97%)
Volatility	2.84%	2.49%	2.18%	2.14%	2.80%	2.46%	2.22%	1.99%
Skewness	-1.92	-4.02	-1.81	-0.57	-6.69	-5.18	-3.20	-1.77
Kurtosis	25.90	25.01	6.16	2.93	55.89	32.66	12.10	3.98
	Panel B: Simulated R_{DHput}							
Mean	-2.23%	-2.16%	-1.66%	-0.86%	-0.67%	-0.35%	-0.18%	-0.05%
t-stat	-8.94	-8.92	-7.81	-0.75	-1.93	-1.64	-1.53	-1.16
p-value	0.00	0.00	0.00	0.02	0.22	0.14	0.23	0.32
	(99%)	(100%)	(99%)	(99%)	(94%)	(63%)	(39%)	(23%)
Volatility	1.12%	0.90%	0.49%	0.04%	1.41%	0.96%	0.50%	0.20%
Skewness	0.26	0.64	0.87	0.59	1.91	1.69	1.53	1.40
Kurtosis	4.34	4.07	4.17	3.75	12.91	6.78	4.83	4.50
	Panel C: Empirical S&P 500 R_{put} (leverage-adjusted)							
Mean	2.22%	1.36%	1.28%	1.29%	1.83%	1.58%	1.45%	1.18%
t-stat	10.89	5.49	6.02	5.51	9.39	8.10	7.93	5.17
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Volatility	4.41%	5.52%	4.89%	4.84%	2.46%	2.77%	2.75%	2.68%
Skewness	-6.11	-4.20	-2.98	-1.82	-4.32	-3.11	-2.72	-2.19
Kurtosis	52.98	27.05	16.84	8.43	27.51	15.09	11.86	8.89
	Panel D: Empirical S&P 500 R_{DHput}							
Mean	-1.05%	-0.42%	-0.17%	-0.05%	-0.75%	-0.57%	-0.42%	-0.22%
t-stat	-6.46	-2.70	-1.62	-0.57	-5.31	-4.75	-4.34	-2.14
p-value	0.00	0.01	0.11	0.57	0.00	0.00	0.00	0.03
Volatility	3.52%	3.42%	2.46%	1.74%	1.78%	1.72%	1.46%	1.23%
Skewness	2.02	2.54	1.74	1.05	1.33	1.13	1.06	0.96
Kurtosis	17.55	18.24	11.57	6.58	7.78	5.88	5.10	4.33

Table 3: Simulated and empirical CAPM alphas and Sharpe ratios

This table displays the simulated and empirical CAPM α and Sharpe ratio on 30-day and 90-day hold-to-maturity naked put option, leverage-adjusted put and delta-hedged put option returns ($R_{DHput}, R_{put}, R_{put}$, respectively) for option contracts with strike-to-price ratios ranging from 0.96 to 1.02. Simulations are based on a discrete-time economy with breaks in the stock's dividend process and partial information under a Bayesian recursive process for a coefficient of relative risk aversion $\eta = 0.2$. Associated t-statistics are computed under the null hypothesis that CAPM α is zero. Entries for simulated figures report the average estimates over 10,000 simulations; for each of these simulations, we generate 12 years of daily dividends. The percentage of simulations with a significant CAPM α is reported in parentheses at a 5% level of significance. We report results from the analysis on the 90-days returns by first converting the 90-days returns to 30-days returns for comparison purposes.

	Puts (30 days to expiration)				Puts (90 days to expiration)			
K/S	0.96	0.98	1.00	1.02	0.96	0.98	1.00	1.02
	Panel A: Simulated R_{put}							
CAPM α	-1.02	-0.86	-0.54	-0.23	-0.35	-0.25	-0.15	-0.07
t-stat	-110.21	-22.36	-15.35	-15.14	-10.13	-5.04	-3.76	-3.48
p-value	0.00	0.00	0.00	0.00	0.09	0.05	0.04	0.04
	(100%)	(100%)	(100%)	(100%)	(80%)	(85%)	(84%)	(83%)
Sharpe ratio	-3.78	-1.65	-0.74	-0.35	-1.18	-0.59	-0.35	-0.20
	Panel B: Simulated R_{put} (leverage-adjusted)							
CAPM α	0.02	0.02	0.02	0.01	0.01	0.01	0.00	0.00
t-stat	35.39	21.97	15.23	15.93	9.38	5.02	4.18	3.66
p-value	0.00	0.00	0.00	0.00	0.07	0.04	0.03	0.03
	(100%)	(100%)	(100%)	(100%)	(84%)	(87%)	(93%)	(86%)
Sharpe ratio	2.98	1.65	0.76	0.36	1.42	0.60	0.35	0.22
	Panel C: Simulated R_{DHput}							
CAPM α	-0.02	-0.02	-0.02	-0.01	-0.01	-0.01	0.00	0.00
t-stat	-37.17	-21.15	-15.62	-15.87	-9.63	-5.25	-3.89	-3.66
p-value	0.00	0.00	0.00	0.00	0.07	0.04	0.03	0.03
	(100%)	(100%)	(100%)	(100%)	(84%)	(87%)	(93%)	(86%)
Sharpe ratio	-1.16	-1.31	-1.34	-1.15	-0.60	-0.59	-0.56	-0.51
	Panel D: Empirical R_{put}							
CAPM α	-0.59	-0.28	-0.15	-0.09	-0.22	-0.17	-0.13	-0.08
t-stat	-10.78	-4.90	-4.28	-3.66	-10.07	-9.07	-8.86	-6.75
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Sharpe ratio	-0.50	-0.21	-0.20	-0.21	-0.71	-0.52	-0.45	-0.36
	Panel E: Empirical R_{put} (leverage-adjusted)							
CAPM α	0.02	0.01	0.00	0.00	0.00	0.00	0.00	0.00
t-stat	9.21	4.37	4.05	3.31	3.17	1.80	0.44	-1.61
p-value	0.00	0.00	0.00	0.00	0.00	0.07	0.66	0.11
Sharpe ratio	0.42	0.18	0.19	0.19	0.27	0.16	0.11	-0.01
	Panel F: Empirical R_{DHput}							
CAPM α	-0.02	-0.01	-0.01	0.00	-0.02	-0.02	-0.02	-0.01
t-stat	-9.78	-4.94	-4.86	-4.36	-15.57	-13.90	-13.15	-12.90
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Sharpe ratio	-0.40	-0.23	-0.22	-0.21	-1.09	-1.07	-1.04	-1.02

Table 4: Simulated CAPM alphas and Sharpe ratios under two scenarios: Full information with no breaks and full information with breaks

This table reports CAPM α and Sharpe ratios under two scenarios: full information with no breaks in the mean dividend growth rate g_t (Black-Scholes setting), and full information with breaks in g_t. Entries report results for leverage-adjusted put returns R_{put} and delta-hedged put option returns R_{DHput}. Average 30-day and 90-day hold-to-maturity returns are calculated for option contracts with strike-to-price ratios ranging from 0.96 to 1.02. Associated t-statistics are computed under the null hypothesis that CAPM α is zero. Entries for simulated figures report the average estimates over 10,000 simulations; for each of these simulations, we generate 12 years of daily dividends with a representative agent's risk aversion $\eta = 0.2$. The percentage of simulations with a significant CAPM α is reported in parentheses at a 5% level of significance. We report results from the analysis on the 90-days returns by first converting the 90-days returns to 30-days returns for comparison purposes.

		Puts (30 days to expiration)				Puts (90 days to expiration)			
K/S		0.96	0.98	1.00	1.02	0.96	0.98	1.00	1.02
		Panel A: Simulated R_{put} (leverage-adjusted)							
No Breaks Full Inf. (Black-Scholes)	CAPM α	0.00	0.00	0.00	0.00	-0.01	0.00	0.00	0.00
	t-stat	0.29	-0.26	0.42	0.52	-0.41	-0.02	0.13	0.80
	p-value	0.40	0.44	0.45	0.48	0.38	0.14	0.27	0.44
		(22%)	(17%)	(19%)	(27%)	(23%)	(16%)	(13%)	(20%)
	Sharpe ratio	0.04	-0.03	0.04	0.05	-0.01	-0.02	0.02	0.01
Breaks Full Inf.	CAPM α	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	t-stat	0.62	0.28	0.89	0.93	0.97	0.25	0.26	0.94
	p-value	0.52	0.47	0.53	0.67	0.44	0.29	0.30	0.45
		(27%)	(22%)	(21%)	(32%)	(24%)	(18%)	(14%)	(22%)
	Sharpe ratio	0.04	0.03	0.04	0.05	0.05	0.02	0.01	0.04
		Panel B: Simulated R_{DHput}							
No Breaks Full Inf. (Black-Scholes)	CAPM α	0.00	0.00	0.00	0.00	0.01	0.00	0.00	0.00
	t-stat	-0.35	0.24	-0.47	-0.52	0.48	0.20	-0.17	-0.81
	p-value	0.40	0.24	0.45	0.47	0.39	0.24	0.17	0.44
		(22%)	(16%)	(19%)	(24%)	(23%)	(14%)	(13%)	(26%)
	Sharpe ratio	-0.01	0.03	-0.03	-0.05	0.05	0.04	-0.01	-0.03
Breaks Full Inf.	CAPM α	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	t-stat	-0.64	-0.30	-0.85	-0.94	-0.98	-0.27	-0.30	-0.85
	p-value	0.49	0.28	0.42	0.44	0.44	0.25	0.30	0.44
		(27%)	(22%)	(22%)	(27%)	(23%)	(17%)	(14%)	(28%)
	Sharpe ratio	-0.03	-0.02	-0.07	-0.08	-0.02	-0.01	-0.01	-0.01

Table 5: Relation of model simulated index put option strategies returns to factors

Entries report the coefficients of factors employed to explain simulated and empirical 30-day leverage-adjusted and delta-hedged put option returns R_{put}, R_{DHput}, respectively. Simulations are based on a discrete-time economy with breaks in the stock's dividend process and partial information under a Bayesian recursive process for a coefficient of relative risk aversion of η=0.2. $R_{m,t}$ is the simulated (empirical) excess market's return equivalent to the excess stock's return in the model (S&P 500). VRP_t, $Slope_{Mon,t}$ and $Slope_{Mat,t}$ are the volatility risk premium, slope of the implied volatility skew and slope of the term structure of implied volatilities, respectively. Entries for simulated figures report average estimates over 10,000 simulated daily dividend paths. The percentage of simulations with significant statistics for the respective factors and t-statistics are reported in parentheses at a 5% level of significance.

	Moneyness 0.96 and expiration 30 days				Moneyness 1.00 and expiration 30 days			
	Panel A: Simulated R_{put} (leverage-adjusted)							
Constant	0.02	0.01	0.02	0.01	0.02	0.00	0.02	0.00
	(100%)	(75%)	(95%)	(71%)	(100%)	(12%)	(97%)	(18%)
$R_{m,t}$	0.08	0.07	0.04	0.07	0.94	0.92	0.92	0.92
	(39%)	(41%)	(45%)	(38%)	(99%)	(100%)	(100%)	(100%)
VRP_t		0.18				0.25		
		(85%)				(96%)		
$Slope_{Mon,t}$			0.09				0.02	
			(95%)				(23%)	
$Slope_{Mat,t}$				0.34				0.38
				(71%)				(66%)
	Panel B: Simulated R_{DHput}							
Constant	-0.02	-0.01	-0.02	-0.01	-0.02	0.00	-0.02	0.00
	(100%)	(74%)	(97%)	(73%)	(100%)	(13%)	(99%)	(18%)
$R_{m,t}$	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06
	(20%)	(19%)	(19%)	(21%)	(29%)	(29%)	(30%)	(27%)
VRP_t		-0.17				-0.25		
		(80%)				(93%)		
$Slope_{Mon,t}$			-0.08				-0.02	
			(96%)				(24%)	
$Slope_{Mat,t}$				-0.30				-0.37
				(69%)				(61%)
	Panel C: Empirical R_{put} (leverage-adjusted)							
Constant	0.02	0.01	0.01	0.01	0.00	0.00	0.01	0.00
	[9.21]	[7.00]	[1.89]	[7.11]	[4.05]	[1.33]	[2.96]	[3.19]
$R_{m,t}$	0.72	0.68	0.72	0.72	1.03	0.97	1.03	1.03
	[16.32]	[13.23]	[16.26]	[16.32]	[37.43]	[31.97]	[37.56]	[37.18]
VRP_t		0.08				0.13		
		[1.68]				[4.21]		
$Slope_{Mon,t}$			0.07				0.09	
			[0.96]				[1.81]	
$Slope_{Mat,t}$				0.09				0.02
				[0.79]				[0.33]
	Panel D: Empirical R_{DHput}							
Constant	-0.02	-0.01	-0.01	-0.01	-0.01	0.00	-0.01	0.00
	[9.78]	[7.50]	[2.02]	[7.57]	[4.86]	[2.06]	[3.12]	[3.87]
$R_{m,t}$	0.27	0.31	0.27	0.27	-0.03	0.02	-0.03	-0.03
	[6.37]	[6.34]	[6.41]	[6.27]	[1.20]	[0.80]	[1.28]	[1.23]
VRP_t		-0.08				-0.12		
		[1.66]				[4.16]		
$Slope_{Mon,t}$			-0.07				-0.08	
			[1.01]				[1.73]	
$Slope_{Mat,t}$				-0.08				-0.02
				[0.80]				[0.32]