Implied Learning Paths from Option Prices

Massimo Guidolin Allan Timmermann
University of California, San Diego

January 22, 2000

Abstract

This paper shows that the best known empirical biases of the Black and Scholes (1973) option pricing formula can be explained by investors learning the parameters of the underlying fundamental process. In the context of an equilibrium model where dividend news evolve on a binomial lattice we derive closed-form pricing formulas for European options under Bayesian learning. Learning effects are found to be able to generate asymmetric skews in the implied volatility surface and systematic patterns in the term structure of option prices. We also infer from S&P 500 index option prices the parameters characterizing the maintained recursive learning process. This allows us to estimate the dynamics of learning and to provide an empirical test for the model.

1 Introduction

Although Black and Scholes' (1973) formula (BS) remains the most commonly used option pricing model in financial markets[1], it is well established that it suffers from strong empirical biases. Most commonly, BS biases are identified with the appearance of systematic patterns (smiles or skews) in the implied volatility surface produced by matching observed market prices with theoretical BS prices and solving for the unknown volatility parameter (Rubinstein (1985, 1990) and Dumas et al. (1998)). Implied volatility also appears to be systematically related to the term structure of option contracts (Das and Sundaram (1999)).

[1] Quoting Rubinstein (1994): "This model is widely viewed as one of the most successful in the social sciences and has perhaps (...) the most widely used formula, with embedded probabilities, in human history.

In an attempt to improve on the empirical performance of BS, a plethora of pricing models have been proposed during the last two decades. These include stochastic volatility models (Hull and White (1987), Wiggins (1987), Melino and Turnbull (1990), Heston (1993)); models in which the conditional variance follows an ARCH process (Duan (1995)); models with jumps in the underlying price process (Merton (1976)); jump-diffusion models (Ball and Torous (1985), Amin (1993)); and models incorporating transaction costs (Leland (1985)). The empirical performance of these models is summarized in Bakshi, Cao and Chen (1997). Although most of these option pricing models fail to improve significantly on the empirical fit of the BS model, this literature has contributed significantly to our understanding of the requirements of an option pricing model that can fit observed option prices. Nevertheless, approaches that rely on modifying the stochastic process followed by the underlying asset price do not provide an *economic* explanation for the systematic shortcomings of BS.[2]

In this paper we relax the key assumption underlying the BS model that investors have complete knowledge of the stochastic process driving fundamentals. More specifically, we assume that fundamentals evolve on a binomial lattice with 'up' and 'down' probabilities unknown to investors who update their probability estimates through Bayes' rule. Importantly, we do not regard the underlying asset price process as exogenous since the learning mechanism is embedded in an equilibrium model in which asset prices reflect all possible perceived future distributions of the parameter estimates. In equilibrium there are no expected gains from implementing trading strategies based on the unfolding of the estimation uncertainty. While there are now many papers studying the implications of learning for asset pricing, this paper is to the best of our knowledge the first to explore the derivative pricing implications of Bayesian learning in the context of an equilibrium asset pricing model.

Our paper is related to a large literature that infers the market's probability beliefs from asset prices (Rubinstein (1994), Jackwerth and Rubinstein (1996), Aït-Sahalia and Lo (1998)). However, there are two major differences between this literature and our approach. First our results are based on an equilibrium model which accounts for investors' behavior. Secondly, our model delivers a rich set of testable restrictions on how the entire cross-section of option prices changes over time. This is implied by our assumption

[2] While transaction cost models potentially provide an economic explanation for BS biases, they have so far failed to improve systematically on the empirical fit of the BS model.

that the market updates its beliefs through Bayes rule. It provides an understanding of the dynamics of how implied volatility surface skews evolve over time.

We show that maintaining all other assumptions of the BS model but introducing learning effects, pricing biases similar to those displayed by observed market prices emerge. Consistent with recent empirical evidence on state-price densities implied by option prices (e.g. Aït-Sahalia and Lo (1998)), by adding to tail probabilities learning effects alter the shape of the state price density perceived by investors. Furthermore, learning effects can generate implied volatility smiles as well as a variety of non-constant term structures of implied volatility.[3] Finally, we infer the dynamics of the parameters characterizing our Bayesian learning scheme from a data set of S&P 500 index option prices. Independently of the time horizon over which theoretical option prices are matched with observed prices, we find that estimated parameters are remarkably stable over time and that our model provides quite a satisfactory in-sample fit.

The outline of the paper is as follows. Section 2 introduces our data set of option prices and confirms the presence of systematic biases in the BS option pricing model. Section 3 briefly presents the binomial lattice model without estimation uncertainty. Bayesian learning effects are introduced in Section 4 which also derives explicit formulae for European option prices. Section 5 presents analytical results that characterize the equilibrium effect of learning on option prices. Section 6 calibrates the option pricing model under learning and compares it to the option price data from Section 2. The parameters characterizing the maintained recursive learning process are then inferred from option prices. This allows us to estimate the dynamics of learning and to provide an empirical test for the model. Section 7 concludes.

2 Biases in the Black-Scholes Model

This section establishes a benchmark for the systematic pricing biases in the BS option pricing model. Our data set of option prices from the CBOE consists of S&P 500 index option prices covering the period Jan. 4, 1993 to Dec. 31, 1993.[4] Put prices were converted into call prices through the

[3] Das and Sundaram (1999) show that the most popular alternatives to BS — jump-diffusion and stochastic volatility models — fail to *simultaneously* generate implied volatility smiles and (ATM) term-structures that match the complex features of the data.

[4] As stressed by Rubinstein (1994) the market for S&P 500 index options on the CBOE provides a case study where the conditions required by BS seem to be well approximated in terms of volumes, continuity of the trading process, hedging opportunities, likelihood

put-call parity relation. The data set is identical to that used in Aït-Sahalia and Lo (1998) and has been filtered in exactly the same way.[5, 6]

2.1 Implied Volatility Surfaces

Initially we confirm the existence of a systematic skew in the relationship between BS implied volatility and moneyness. For this purpose we plot in Figure 1 the implied volatility surface against moneyness, keeping maturity constant. CBOE rules create option contracts with monthly maturities and Figure 1 in Appendix B include 12 plots covering the period February 1993 to January 1994.[7] For most of the days in the sample period, the curve relating BS implied volatilities to the strike price is skewed. This is or course inconsistent with the maintained assumption of constant diffusion in the BS model.

---

of jumps in prices, etc. Therefore it is natural that our empirical tests concentrate on this market.

[5] We thank Yacine Aït-Sahalia for making his data set kindly available to the public for research purposes.

[6] Since in-the-money options are thinly traded, their prices are notoriously unreliable and are discarded from the data set. Also, out-of-the-money and near-at the money (ATM) put prices are translated into call prices using the put-call parity for European options. All information contained in liquid put prices has thus been extracted and converted into corresponding, liquid call prices. The remaining put options are discarded from the data set without any loss of information. We explicitly take into account that the index pays out daily dividends and follow Aït-Sahalia and Lo (1998) in using the continous dividend yield implied by daily values of the index and prices of future contracts of given maturity. By the spot-futures parity

$$F_{t,\tau} = S_t e^{(r_{t,\tau} - \delta_{t,\tau})\tau} \Longrightarrow \delta_{t,\tau} = \frac{1}{\tau} \ln \left(\frac{S_t}{F_{t,\tau}} \right) + r_{t,\tau}$$

The implied value of $\delta_{t,\tau}$ is a measure of the continuously compounded dividend yield that, as of time t, is expected between t and $t+\tau$.

Finally, to guarantee the absence of arbitrage opportunities, the futures price for a given maturity $t+\tau$ is itself implied by the put-call parity relation for a dividend paying stock index:

$$C(S_t, K, \tau, r_{t,\tau}, \delta_{t,\tau}) + X e^{-r_{t,\tau}\tau} = P(S_t, K, \tau, r_{t,\tau}, \delta_{t,\tau}) + F_{t,\tau} e^{-r_{t,\tau}\tau}$$

To infer reliable futures prices, we need reliable option prices. Therefore ATM call and put closing prices are used.

[7] There are not enough data to build a representative plot for the maturity of January 1993, while there are still enough data to obtain January 1994, although the sample ends on Dec. 31, 1993.

2.2 Term Structure of Implied Volatility

The data also reveal a systematic term structure in the implied volatility. Using three alternative values of moneyness over the period Jan. 18, 1993 to Jan. 25, 1993, Figure 1 shows that the implied volatilities of at-the-money options and options with moneyness less than one clearly increase with time to expiration. For options with moneyness above one the pattern is somewhat weaker: some days implied volatility is an increasing and convex function of time to expiration; other days implied volatility is a concave function of moneyness and occasionally the pattern is constant or even declining.

Figure 1. Sample term structure of implied volatility. The three graphs plot implied volatility as a function of maturity for S&P 500 index options, period Jan. 18, 1993 - Jan. 25, 1993. The three graphs plot the term structure for three alternative moneyness levels: 0.96 (in the money), 1 (at-the-money), and 1.04 (out-of-the-money).

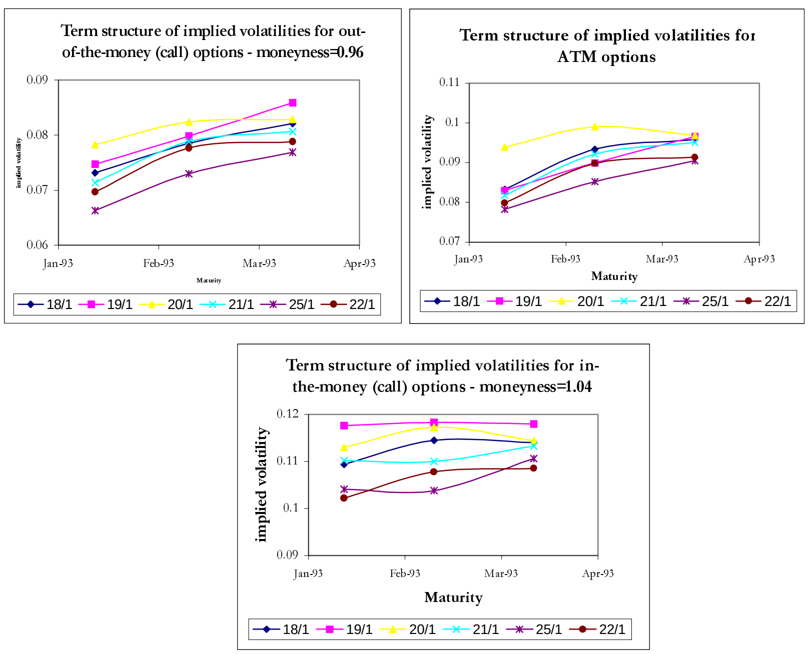

2.3 State Price Densities

Recently Rubinstein (1994) and Jackwerth and Rubinstein (1996) have proposed to extract state price densities (SPD) from implied binomial trees. This is another powerful procedure for demonstrating biases in the BS model whose assumption is that the state price density is log-normal. Using this approach, figure 2 shows the SPD inferred from option contracts with at least 50 calendar days to expiration and averaged across different maturities (a total of 334 estimated SPDs for different maturities).[8] To ensure that SPDs on different days are (roughly) comparable, all plots use standardized logarithmic returns. Particularly important in economic terms is the tail behavior of the SPD since this may provide information about the jump risk expected by markets (Bates (1991)). For this reason we plot in the bottom of Figure 2 the estimated tail behavior of the average SPD.

Figure 2. Average state-price density estimated from S&P 500 index options and the S&P 500 cash index over the period Jan. 4, 1993 - Dec. 31, 1993 compared to a lognormal SPD. The estimated SPD is the average of all SPDs estimated from options data with more than 50 calendar days to expiration (approximately 35 trading days), for a total of 334 estimated SPDs. SPDs are estimated using the nonparametric, implied binomial tree method of Rubinstein [1994] and Jackwerth and Rubinstein [1996]. The objective function is the maximum smoothness function:

$$\sum_{j=0}^{v} (P_{j-1} - 2P_j + P_{j+1})^2 \quad P_{-1} = P_{v+1} = 0$$

For every day in the sample and for each cross section of contracts (over strike) defined by maturity, we minimize the objective function subject to martingale restrictions imposing correct pricing of the options and of the index. The constraints are imposed by a penalty method that progressively raises the penalty parameter over various steps of the numerical optimization (see Judd [1998, 123-125]).

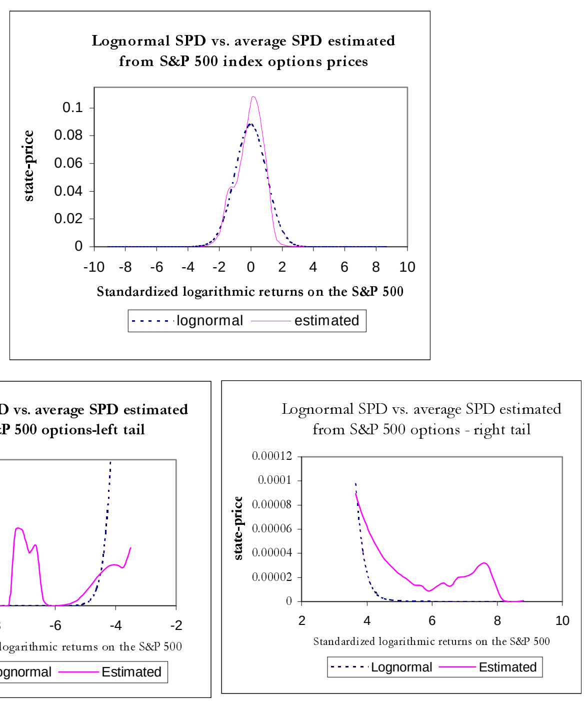

[8] Details on the relevant technical choices are given in the caption of the figure.

When compared to the lognormal benchmark, the SPD implied by the options data is clearly leptokurtic: there is more mass (higher state-prices) around the center of the SPD, less mass in the range [-4,-1]∪[+1,+4] standard deviations of logarithmic returns and fat tails occur in the form of lobes between [-8,-6] and [+6,+8] standard deviations of returns. Hence the markets appear to have very different perceptions of future S&P 500 returns than those implied by a lognormal SPD. We conclude that the markets seem to perceive state-price densities that are — at least on average and most of time — quite different from the lognormal density underlying the BS model.

In particular, market participants attach value to future 'extreme' outcomes — crashes as well as strong bull markets — that under a lognormal SPD would receive a much smaller state price. To demonstrate this point, we also compare no-arbitrage prices for a state-contingent asset paying off 1 dollar when the S&P 500 returns over a certain time span are below (above) a certain number of standard deviations of returns under the estimated SPDs and under a lognormal benchmark. Table 1 reports the results from this exercise. Estimated SPDs imply state prices for 'tail'-contingent securities that are larger by a factor of 10^8 when the asset pays out in case the S&P 500 return is below -7 standard deviations, and by a factor of 10^{10} when returns exceed 7 standard deviations. Higher prices for tail-contingent securities reflect the fact that under the estimated SPDs the market attaches a much higher risk-neutral price to a dollar paid out in either crash or strongly bullish states of the world.

3 Asset Prices under Full Information

The empirical findings reported in the last section confirm the presence of serious systematic biases in the BS model widely commented on in the literature, and suggest that a more general option pricing model is required. In this section we characterize option prices in a full information equilibrium model which has the BS model as a limiting case. This sets up a natural benchmark from which to evaluate option pricing biases and the effects of learning. Next section introduces learning and derives option prices in this setting.

Table 1. The table reports the price of a state-contingent claim that pays out 1 dollar when (demeaned) S&P 500 returns are below/above X times their standard deviation, calculated from the SPD estimated from option contracts with at least 50 calendar days to maturity. As a benchmark the table also reports the price of the same contingent claims based on a lognormal SPD. The SPDs are normalized by demeaning them and dividing them by $\sigma\sqrt{\tau}$ so that the standard deviation is set to 1. The price of the contingent-claims are calculated according to the formulae:

$$Q^{-}=\sum_{j=0}^{v}\lambda\left(\frac{1}{\sigma\sqrt{\tau}}\left[\ln\left(\frac{S^{j}_{t+v}}{S_t}\right)-\mu\right]\right)I_{\left\{\frac{1}{\sigma\sqrt{\tau}}\left[\ln\left(\frac{S^{j}_{t+v}}{S_t}\right)-\mu\right]<X\right\}},\quad Q^{+}=\sum_{j=0}^{v}\lambda\left(\frac{1}{\sigma\sqrt{\tau}}\left[\ln\left(\frac{S^{j}_{t+v}}{S_t}\right)-\mu\right]\right)I_{\left\{\frac{1}{\sigma\sqrt{\tau}}\left[\ln\left(\frac{S^{j}_{t+v}}{S_t}\right)-\mu\right]>X\right\}}$$

where the λ_js are the state-prices, $\mu=\sum_{j=0}^{v}\ln\left(\frac{S^{j}_{t+v}}{S_t}\right)\mathrm{Prob}\{S^{j}_{t+v}=u^{j}d^{v-j}S_t\}$, and

$\sigma^2=\sum_{j=0}^{v}\left[\ln\left(\frac{S^{j}_{t+v}}{S_t}\right)-\mu\right]^2\mathrm{Prob}\{S^{j}_{t+v}=u^{j}d^{v-j}S_t\}$ with $\mathrm{Prob}\{S^{j}_{t+v}=u^{j}d^{v-j}S_t\}$ a risk-neutral probability, either based on the normal density or estimated from options data. $I_{\{\bullet\}}$ is a standard indicator function. We report both the average the median state-price densities estimated from S&P 500 index options and the S&P 500 cash index over the period Jan. 4, 1993 - Dec. 31, 1993 using the nonparametric, implied binomial tree method of Rubinstein [1994] and Jackwerth and Rubinstein [1996]. The objective function is the maximum smoothness function:

	X	Average price		Median price	
		Lognormal SPD	Estimated SPD	Lognormal SPD	Estimated SPD
	-7<	$3\cdot10^{-13}$	0.0000126	$8\cdot10^{-14}$	$1\cdot10^{-10}$
	-6<	$1\cdot10^{-10}$	0.00002244	$1\cdot10^{-10}$	$4\cdot10^{-10}$
	-5<	$2\cdot10^{-7}$	0.00002402	$2\cdot10^{-7}$	$9\cdot10^{-10}$
Q^{-}	-4<	0.00001857	0.00003413	0.00001870	$2\cdot10^{-9}$
	-3<	0.001505	0.00026	0.001515	$7\cdot10^{-9}$
	-2<	0.01974	0.01264	0.01987	0.0000028
	-1<	0.1744	0.1742	0.1747	0.1203
	>1	0.1381	0.07197	0.1381	0.0679
	>2	0.02334	0.00376	0.02334	0.00159
	>3	0.000865	0.000849	0.000865	0.000173
Q^{+}	>4	0.000235	0.000428	0.0000235	0.0000391
	>5	$6\cdot10^{-8}$	0.000241	$6\cdot10^{-8}$	0.0000170
	>6	$1\cdot10^{-10}$	0.000186	$1\cdot10^{-10}$	0.0000117
	>7	$8\cdot10^{-14}$	0.000126	$8\cdot10^{-14}$	$7\cdot10^{-6}$

3.1 A Binomial Lattice Model

Before pricing options we introduce the fundamentals process determining the price of the underlying asset. Our starting point is a version of the infinite horizon, representative agent endowment economy studied by Lucas (1978). There are three assets: A one-period default-free, zero-coupon bond in zero net supply trading at a price of P_t and earning interest of $r_t = (1/P_t - 1)$; a stock traded at a price of S_t whose net supply is normalized at 1; and a European call option written on the stock with $\tau \equiv T - t$ periods to expiration, strike price K and current price C_t.

The stock pays out an infinite stream of real dividends $\{D_{t+k}\}_{k=1}^{\infty}$. Dividends are perishable and must therefore be consumed in the period in which they are received. They evolve on a binomial lattice so dividend growth rates $g_{t+k} = \frac{D_{t+k}}{D_{t+k-1}} - 1$ follow a Bernoulli process that is subject to change m times in each unit interval. Within the interval $[t, t+\tau]$ dividends thus follow a $v = \tau m$-step binomial process. Each period $[t+k, t+k+1]$ the dividend growth rate can be either g_h with probability π or g_l with probability $1-\pi$:

$$g_{t+k+1} = \begin{cases} g_h & \text{with prob. } \pi \\ g_l & \text{with prob. } 1-\pi \end{cases} \qquad \forall k \geq 0, \quad \pi \in (0,1) \tag{1}$$

Without loss of generality we assume that $g_h > g_l > -1$ so that dividends are non-negative provided $D_t > 0$. This gives a standard recombining binomial tree similar to the one adopted by Cox, Ross and Rubinstein (1979) for the underlying asset price process. We follow the literature in normalizing the parameters to the incremental time unit: $1+g_h = e^{\sigma\sqrt{\frac{dt}{v}}}$, $1+g_l = (1+g_h)^{-1}$, and $\pi = \frac{1}{2} + \frac{1}{2}\frac{\mu}{\sigma}\sqrt{\frac{dt}{v}}$. As $\frac{dt}{v} \to 0$, the distribution of dividends (weakly) converges to a Geometric Brownian motion with constant drift and diffusion (μ, σ).[9]

To price assets we assume a perfect capital market. There are unlimited short sales, perfect liquidity (no price impact from sales or purchases of securities), no taxes, no transaction costs or borrowing and lending constraints and markets are open at all points in time in which news on dividends are generated, i.e., on all the nodes of the binomial lattice.

The representative investor has power utility

$$u(C_t) = \begin{cases} \frac{C_t^{1-\gamma}-1}{1-\gamma} & \gamma < 1 \\ \ln C_t & \gamma = 1 \end{cases} \tag{2}$$

where C_t is real consumption at time t. We focus on the case where $\gamma \leq 1$; models where $\gamma > 1$ have the counter-intuitive property that stock prices decline when fundamentals are high.[10] The representative agent chooses bond, stock, and call option holdings to maximize the discounted value (at a rate of impatience ρ) of the infinite stream of expected future utilities derived from consumption:

[9] See for instance Neftci (1996).

[10] For further explanations of this property, see Abel (1988).

$$\max_{\{C_{t+k},w^s_{t+k},w^b_{t+k}\}_{k=0}^{\infty}} E_t\left[\sum_{k=0}^{\infty}\beta^k u(C_{t+k})\right]$$

$$s.t.\ C_{t+k}+w^s_{t+k}S_{t+k}+w^b_{t+k}P_{t+k}=w^s_{t+k-1}(S_{t+k-1}+D_{t+k-1})+w^b_{t+k-1} \quad (3)$$

where $\beta=\frac{1}{1+\rho}$ and w^s_{t+k} and w^b_{t+k} represent the number of stocks and bonds in the agent's portfolio as of period $t+k$.[11]

Standard dynamic programming methods yield the following Euler equations for stock and bond prices:

$$\begin{aligned} S_t &= E_t\left[Q_{t+1}(S_{t+1}+D_{t+1})\right] \\ P_t &= E_t\left[Q_{t+1}\right] \end{aligned} \quad (4)$$

where $Q_{t+1}=\beta\frac{u'(C_{t+1})}{u'(C_t)}=\beta\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}$ is the pricing kernel defined as the product of the discount factor and the intertemporal marginal rate of substitution in consumption.

Guidolin and Timmermann (1999) price the underlying stock and risk-free bond in this setting subject to transversality and no-arbitrage conditions.[12] For convenience, we state the result using the transformed parameters $g^*_l=(1+g_l)^{1-\gamma}-1$ and $g^*_h=(1+g_h)^{1-\gamma}-1$.

Proposition 1 (Guidolin and Timmermann (1999)) *There exists an equilibrium in which the full information rational expectations (FIRE) stock price is given by*

$$S^{FIRE}_t=\frac{1+g^*_l+\pi(g^*_h-g^*_l)}{\rho-g^*_l-\pi(g^*_h-g^*_l)}D_t,$$

[11] Since the call is a redundant asset which does not let agents expand the set of attainable consumption patterns, option holdings do not enter into the program and the equilibrium stock and bond prices can be determined independently of the option price.

[12] The transversality condition is,

$$\lim_{T\to\infty} E_t\left[\left(\prod_{k=1}^{T}Q_{t+k}\right)S_{t+T}\right]=0$$

and no-arbitrage conditions are

$$g^*_l+\pi(g^*_h-g^*_l)<\rho<\pi(1+g^*_h)+(1-\pi)\left[(1+g_h)(1+g_l)^{-\gamma}\right]-1$$

while the FIRE bond price is

$$P_t^{FIRE} = \frac{(1+g_l)^{-\gamma} + \pi\left[(1+g_h)^{-\gamma} - (1+g_l)^{-\gamma}\right]}{1+\rho}.$$

A property of the solution is that the stock price is homogeneous of degree one in dividends and that the *ex-dividend* stock price follows the same binomial lattice $\{g_h, g_l, \pi, m\}$ as dividends.[13] Guidolin and Timmermann (1999) also show that the stochastic process followed by dividends and the stock price can be alternatively characterized as a stationary Markov chain, a fact that will be useful later on.

3.2 Option Prices under Full Information

Pricing European calls is straightforward under full information. This follows from noting that (i) we have ruled out arbitrage opportunities; (ii) markets are complete; (iii) Ex-dividend stock prices inherit the binomial lattice structure $\{g_h, g_l, \pi, m\}$ from dividends; (iv) although the underlying asset pays out cash dividends, the option is European and early exercise is not possible. Therefore all possible contingent claims (consumption profiles) are marketable and we can apply the results in Cox et al. (1979) to price the option by no-arbitrage:

Proposition 2 *Suppose the dividend process is binomial. Under full information rational expectations, the price of a European call with τ periods to expiration and strike price K is given by*

$$C_t^{FIRE}(K,T,S_t) = (1+r)^{-v}\sum_{j=0}^{v}\max\left\{0, S_{t+v}^{j} - K\right\}\binom{v}{j}z^j(1-z)^{v-j}$$

where $v = \tau m$, $S_{t+v}^{j} = (1+g_h)^j(1+g_l)^{v-j}S_t$, and z, is a risk-neutral probability measure given by

$$z = \pi\frac{1+r}{1+\rho}(1+g_h)^{-\gamma}$$

[13] From this and the result in Pliska (1997, 133) that a multiperiod securities markets model is complete if and only if all possible sequences of single period models obtained by decomposing the binomial lattice are formed by complete models it follows immediately that the asset market in our model is complete. Therefore the (conjectured) redundancy of the call option does hold in the FIRE equilibrium and our approach is valid.

Proof. See Appendix A. ■

This proposition shows that the results of Cox et al. (1979) fully extend to our framework where dividends rather than stock prices are assumed to follow a binomial lattice. This is an implication of the fact that stock prices evolve on the same lattice as dividends. Notice, however, that while CRR take the process for the underlying price as exogenous, we derive the underlying stock price in an equilibrium model in which preferences matter. This result can also be related to Stapleton and Subrahmanyam (SS, 1985), who value options in an equilibrium model when markets are incomplete and the stock price follows an exogenous binomial lattice. In both cases preferences affect the equilibrium price of stocks and (indirectly) options.

Harrison and Kreps (1979) prove that in the absence of arbitrage opportunities and with complete markets there exists a unique risk-neutral probability measure $f(S_{t+T})$ such that the price of a European option with payoff function $h(S_{t+T})$ and τ periods to expiration can be represented as

$$\begin{aligned} C_t &= (1+r)^{-v} \sum_{j=0}^{v} h(S_{t+T}) P^{RND} \left\{ S_{t+T} = S_{t+T}^{j} \right\} \\ &= E_{\widetilde{P}} \left[h(S_{t+T}) \right] \end{aligned} \quad (5)$$

where $\widetilde{P}(S_{t+T})$ is the SPD. In our case this is given by

$$\widetilde{P} \left\{ S_{t+T} = S_{t+T}^{j} \right\} = (1+r)^{-v} \binom{v}{j} z^{j} (1-z)^{v-j} \quad (6)$$

and is simply a transformation of the binomial process $Bi(v, z)$. The risk-neutral measure thus characterized retains a Markov chain structure. Inspection of $\widetilde{P} \left\{ S_{t+T}^{j} \right\}$ reveals that the transition matrix is

$$\left[\widetilde{P} \left\{ X_{t+k+1} = j | X_{t+k} = i \right\} \right] = \widetilde{M} = \begin{bmatrix} 1-z & z & 0 & \dots & 0 \\ 0 & 1-z & z & \dots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & 0 & \dots & 1 \end{bmatrix} \quad (7)$$

where $\widetilde{P} \left\{ X_{t+k+1} = T | X_{t+k} = T \right\} = 1$ is set arbitrarily. Computation of the risk neutral probability of the final nodes associated with the binomial lattice for stock prices is hence very simple:

$$\widetilde{P}\left\{S_{t+v} = S^j_{t+T}|S_t\right\} = \widetilde{P}\left\{X_{t+v} = j|X_{t+k} = 0\right\}$$
$$= e_1'\left[\prod_{k=1}^{v}\widetilde{M}\right]e_{j+1} = e_1'\widetilde{M}^v e_{j+1} = \binom{v}{j} z^j(1-z)^{v-j} \tag{8}$$

which is the same expression as (6), apart from the constant of proportionality $(1+r)^{-v}$, the SPD for time $t+T$ stock prices.[14]

To establish the link between our option price under FIRE and the Black-Scholes price, we use the result of Cox et al. (1979, 246-251) that, provided the parameters are suitably adjusted as the number of increments to the lattice goes to infinity, the price of a European option converges to the BS value. Let r be the risk-free rate, δ the dividend yield, with μ and σ respectively the (annual) mean and volatility of the dividend growth rate. The following proposition restates this result stressing the mapping between the deep parameters of our model and the BS inputs.

Proposition 3 *Suppose that the parameters have been scaled as follows:*

$$g_h^{(m)} = e^{\sigma\sqrt{dt/v}} - 1, 1 + g_l^{(m)} = \left(1 + g_l^{(m)}\right)^{-1}, \pi^{(m)} = \frac{1}{2} + \frac{1}{2}\frac{\mu}{\sigma}\sqrt{\frac{dt}{v}}$$
$$\rho^{(m)} = (1+r)^{\frac{dt}{v}}\left\{(1+g_l^{(m)})^{-\gamma} + \pi^{(m)}\left[(1+g_h^{(m)})^{-\gamma} - (1+g_l^{(m)})^{-\gamma}\right]\right\} - 1$$
$$\delta = \frac{(1+r)\left\{(1+g_l^{(m)})^{-\gamma} + \pi^{(m)}\left[(1+g_h^{(m)})^{-\gamma} - (1+g_l^{(m)})^{-\gamma}\right]\right\}^{\frac{v}{dt}}}{\left\{(1+g_l)^{1-\gamma} + \pi^{(m)}\left[(1+g_h^{(m)})^{1-\gamma} - (1+g_l^{(m)})^{1-\gamma}\right]\right\}^{\frac{v}{dt}}} - 1.$$

Also assume that $\rho^{(m)} > 0$. *As* $m \to \infty$, *the price of the European call converges to its BS value:*

$$\begin{aligned} C_t^{BS} &= S_t\Phi(d_1)e^{-\delta dt} - e^{-rdt}K\Phi\left(d_2\right) \\ d_1 &= \frac{\ln\left(\frac{S_t}{K}\right) + \left(r-\delta\right)dt + \frac{1}{2}v\left[\ln(1+g_h^{(m)})\right]^2}{\sqrt{v}\ln(1+g_h^{(m)})} \\ d_2 &= d_1 - \sqrt{v}\ln\left(1+g_h\right) \end{aligned}$$

[14] From these definitions it is clear that

$$P^{RND}\left\{S^j_{t+T}\right\} = (1+r)^v\widetilde{P}\left\{S^j_{t+T}\right\}$$

where $v = \tau m$, and $\Phi(\cdot)$ is the c.d.f. of the standard normal distribution.

Proof. See Appendix A. ■

Since the BS option price obtains in the limit when no learning is present, our model is ideally suited to discuss the origin of BS pricing biases.[15] Proposition 3 is hardly surprising. As $m \to \infty$, $\frac{dt}{v} = \frac{1}{m} \searrow 0$, validating the conditions under which the binomial lattice weakly converges to a Geometric Brownian motion with parameters (μ, σ), the distributional assumption required by BS in continuous time. Alternatively, this point can be shown by realizing that as $m, v \to \infty$ the discrete SPD converges to a transformation of the lognormal $(r - \delta, v\left[\ln(1+g_h)\right]^2)$:

$$\begin{aligned} \widetilde{f}(S_{t+T}) &= e^{-r\tau} \frac{1}{\sqrt{2\pi}\sigma\sqrt{\tau}} \frac{1}{S_{t+T}} \\ &\quad \times \exp\left\{ -\frac{\left[\ln(S_t) - \left(r - \delta + \frac{1}{2}\left[\ln(1+g_h)\right]^2\right) v\right]^2}{2v\left[\ln(1+g_h)\right]^2} \right\} \end{aligned} \tag{9}$$

4 Option Prices on a Learning Path

It is common in the option pricing literature to assume a given process for the underlying asset price and then price the option as a redundant asset whose payoffs can be replicated in a dynamic hedging strategy invested in the risky asset and a riskfree bond. The standard setup assumes that the asset price process is stationary and hence that there are no learning effects. Once learning is introduced, in equilibrium the asset price process will also change, and a model for the underlying asset price is required. In this section we briefly describe the main properties of this model based on the more extensive analysis in Guidolin and Timmermann (1999).

Suppose that the proportion of times dividends move up on the binomial lattice, π, is unknown to agents who estimate this parameter using all

[15] Notice that the preference parameter ρ has to be made a function of the number of binomial steps 'per period' m. To see this, consider the effect of an incremental increase in m, holding ρ fixed. For a given τ, an increase in m implies that dividends are paid out more often. This does not in itself change the current value of the stock index as m does not enter the solution of $S_t^{FIRE} = \lim_{T\to\infty} \sum_{s=1}^{T} E_t\left[\left(\prod_{k=1}^{s} Q_{t+k}\right) D_{t+s}\right] + E_t\left[\left(\prod_{k=1}^{T} Q_{t+k}\right) S_{t+T}\right]$. However, the kernel Q_{t+k} $(k \geq 1)$ depends on the rate of time preference ρ. Increasing m makes the assumed ρ (that pertains to a single step of the lattice) apply to a *shorter* interval of time, and this will depress the current stock price by increasing β. One can show that as $m \to \infty$ the European call value converges towards zero. The adjustment of ρ as a function of m serves to ensure that the BS option price is independent of m.

available information. Agents are assumed to recursively gain knowledge on π through the simple maximum likelihood estimator:

$$\widehat{\pi}_{t+k,M} = \frac{\sum_{j=1}^{m(t+k)+M} I_{\left\{g_{m(t+k)+j}=g_h\right\}}}{m(t+k)+M} = \frac{n_{t+k,M}}{N_{t+k,M}} \quad M = 0, 1, ..., m-1 \tag{10}$$

where $I_{\left\{g_{m(t+k)+j}=g_h\right\}}$ is a standard indicator function taking the value 1 when at the step $m(t+k)+M$ of the lattice the high dividend growth rate attains, and zero otherwise. n_j denotes the number of high growth states recorded up to node j, while N_j is the total number of nodes since t. Investors are assumed to start out with prior beliefs $\{n_0, N_0\}$. The indices $m(t+k)+M$ $(M = 0, 1, ..., m-1)$ take into account that learning happens on the binomial tree and not over calendar time. This of course corresponds to a *Bayesian* learning process in which the agent updates a given prior on the distribution of (the random variable) π using Bayes' rule.

Despite the presence of learning effects, the same features that simplified the solution of asset prices under FIRE are still in place: (i) consumption and dividends must coincide in general equilibrium; (ii) if markets are complete, investors form portfolio choices based only on the stock index and the bond; (iii) being redundant assets, options can be priced by no-arbitrage, using the unique risk neutral probability measure.

Guidolin and Timmermann (1999) prove the following result:

Proposition 4 *Suppose the representative agent is on a Bayesian learning path and that the stock price is homogeneous of degree one in the level of real dividends,* $S_t = \Psi_t^{BL}(g_h, g_l, \pi, \gamma, \rho, n_t, N_t)D_t$. *Then the (ex-dividend) Bayesian learning (BL) stock price is*

$$S_t^{BL} = \Psi_t^{BL} D_t = D_t \lim_{v\to\infty} \left\{ \sum_{i=1}^{v} \beta^i \sum_{j=0}^{j} (1+g_h^*)^j (1+g_l^*)^{i-j} \Pr{}_t^{BL} \left(D_{t+i}^j | n_t, N_t \right) \right\}$$

where $\Pr{}_t^{BL} \left(D_{t+i}^j = (1+g_h)^j (1+g_l)^{i-j} D_t | n_t, N_t \right)$ *is given by*

$$\Pr{}_t^{BL} \left\{ D_{t+i}^j | n_t, N_t \right\} = \binom{i}{j} \frac{\prod_{k=0}^{j-1}(n_t+k) \prod_{k=0}^{i-j-1}(N_t - n_t + k)}{\prod_{k=0}^{i-1}(N_t+k)}$$

The BL bond price is

$$P_t^{BL}(\widehat{\pi}_t) = \widehat{E}_t\left[\beta(1+g_{t+1})^{-\gamma}\right] = \frac{(1+g_l)^{-\gamma} + \widehat{\pi}_t\left[(1+g_h)^{-\gamma} - (1+g_l)^{-\gamma}\right]}{1+\rho}.$$

Proposition 4 has many implications. First, the price-dividend ratio is no longer constant depending (through n_t and N_t) on the current estimate $\widehat{\pi}_t$. Dividend changes acquire a self-enforcing nature: Positive dividend shocks lead to an increase in the stock price not only through the standard proportional effect, but also through the revision of the dividend multiplier, Ψ_t^{BL}. Second, while under FIRE the risk-free rate was constant, on a learning path it also changes as a function of the state variables n_t and N_t.[16]

Under BL, the stock price process is described by a recombining, flexible binomial tree with time-varying estimated state probabilities. Since dividend yields, though time-varying themselves, are still a fixed proportion of the stock price, in principle binomial methods to obtain the no-arbitrage price of European options on flexible trees could find application (see for instance Chriss (1997)). However, our equilibrium model delivers risk-free rates that are not only time-varying, but also a function of the state variable $\widehat{\pi}_{t+k}$. Unfortunately the interest rate process cannot be characterized as a recombining, flexible lattice. The value today of one dollar in the future depends not only on the *number* of high and low growth states occurring between today and the future, but also on their *sequence*. In other words, the appropriate discount factors become *path-dependent*. Since risk-free rates show up in the general risk-neutral valuation formula proved by Harrison and Kreps (1979), it also means that also the induced lattice for the call option is non-recombining. Therefore no-arbitrage pricing under learning has complicated elements of path-dependency.

This path dependence means that no-arbitrage methods are more complex in an equilibrium model with BL. Nevertheless, European call options can still be priced by employing a risk neutral change of measure from the BL perceived probabilities. The following proposition derives the price of a European call and the SPD under Bayesian learning.

Proposition 5 *On a Bayesian learning path, the no-arbitrage price of a*

[16] More specifically, when $\gamma \leq 1$, $\left[(1+g_h)^{-\gamma} - (1+g_l)^{-\gamma}\right] < 0$, so high dividend growth raises the risk free rate by raising $\widehat{\pi}_{t+k+1}$ above $\widehat{\pi}_{t+k}$. Incidentally, the yield of zero-coupon bonds stops being risk-free.

European call with time-to-expiration τ and strike price K is

$$
\begin{aligned}
C_t^{BL}(K,T,S_t^{BL}) &= \sum_{j=0}^{v} \beta^v \left(\frac{D_{t+v}^j}{D_t}\right)^{-\gamma} \max\left\{0, S_{t+v}^{BL,j} - K\right\} Pr_t^{BL}\left\{S_{t+v}^j | n_t, N_t\right\} \\
&= \sum_{j=a}^{v} \max\left\{0, S_{t+v}^{BL,j} - K\right\} \widetilde{P}_t^{BL,SPD}\left\{S_{t+v}^j | n_t, N_t\right\}
\end{aligned}
$$

where $v = \tau m$, $S_{t+v}^{BL,j} = (1+g_h)^j(1+g_l)^{v-j} S_t^{BL} = (1+g_h)^j(1+g_l)^{v-j}\Psi_{t+v}^{BL}(n_t+j, N_t+v)D_t$ $(j=0,1,...,v)$, $D_{t+v}^j = (1+g_h)^j(1+g_l)^{v-j}D_t$ $(j=0,1,...,v)$, and

$$
\begin{aligned}
\widetilde{P}_t^{BL,SPD}\left\{S_{t+v}^j\right\} &= \widetilde{P}\left\{D_{t+v}^j | n_t, N_t\right\} = \beta^v \left(\frac{D_{t+v}^j}{D_t}\right)^{-\gamma} \times \\
&\times \binom{v}{j} \frac{\prod_{k=0}^{j-1}(n_t+k)\prod_{k=0}^{v-j-1}(N_t-n_t+k)}{\prod_{k=0}^{v-1}(N_t+k)}
\end{aligned}
$$

Proof. See Appendix A. ■

Under BL the Markov chain characterization of the stock price process and of the risk-neutralized stock price process are no longer stationary, since both the possible rates of change of the stock index and the (perceived) probabilities of these changes follow heterogenous Markov chains. Indeed[17]

$$
P^{BL}\left\{X_{t+k+1} = j | F_{t+k}\right\} = P^{BL}\left\{X_{t+k+1} = j | X_{t+k}\right\} \tag{11}
$$

where X measures the number of high dividend growth rates occurring between t and $t+k+1$. Therefore only information useful to predict realization of dividends up to time $t+k+1$ is the current level of dividends. The time-varying transition matrix determining how beliefs are updated is given by

$$
\begin{aligned}
&\left[P^{BL}\left\{X_{t+k+1} = j | X_{t+k} = i\right\}\right] = M_{t+k}(i+1, j+1) \\
&= \begin{bmatrix}
\frac{N_{t+k}-n_t}{N_{t+k}} & \frac{n_t}{N_{t+k}} & 0 & ... & 0 & 0 \\
0 & \frac{N_{t+k}-n_t+1}{N_{t+k}} & \frac{n_t+1}{N_{t+k}} & ... & 0 & 0 \\
... & ... & ... & & \vdots & \vdots \\
0 & 0 & 0 & ... & \frac{N_{t+k}-n_t-k}{N_{t+k}} & \frac{n_t+k}{N_{t+k}} \\
0 & 0 & 0 & ... & 0 & 1
\end{bmatrix}
\end{aligned} \tag{12}
$$

[17] More precisely, the information structure $F = \{F_{t+k}; k = 0, 1, ...\}$ is a filtration composed of an infinite, nested ($F_{t+k+1} \supseteq F_{t+k}$ $\forall k \geq 0$) sequence of $\sigma-$algebras, with F_{t+k} corresponding to the first $t+k$ movements of dividends, that is, associated to the partition $\mathcal{P}_{t+k}$ consisting of 2^k cells, one for each possible sequence of dividends growth rates up to time $t+k$. See Guidolin and Timmermann (1999) for further details.

where $P^{BL}\left\{X_{t+k+1}=T|X_{t+k}=T\right\}=1$ is simply a normalization. Obviously, the probabilities appearing in this matrix change over time.

Similarly, the risk neutral distribution becomes a heterogeneous Markov chain with transition matrix:

$$\left[\widetilde{P}^{BL,RND}\left\{X_{t+k+1}=n_t+j|X_{t+k}=n_t+i\right\}\right]=\widetilde{M}_{t+k}(i+1,j+1)$$

$$=\beta\begin{bmatrix} R_{t+k}^{i=0,d}(1+g_l)^{-\gamma}\frac{N_{t+k}-n_t}{N_{t+k}} & R_{t+k}^{i=0,u}(1+g_h)^{-\gamma}\frac{n_t}{N_{t+k}} & 0 & \ldots & 0 \\ 0 & R_{t+k}^{i=1,d}(1+g_l)^{-\gamma}\frac{N_{t+k}-n_t+1}{N_{t+k}} & R_{t+k}^{i=1,u}(1+g_h)^{-\gamma}\frac{n_t+1}{N_{t+k}} & \ldots & 0 \\ \ldots & \ldots & \ldots & & \vdots \\ 0 & 0 & 0 & \ldots & 1 \end{bmatrix} \tag{13}$$

where $\widetilde{P}\left\{X_{t+k+1}=T|X_{t+k}=T\right\}=1$ is set arbitrarily. $R_{t+k}^{i,d}=1+r_{t+k}^{i,d}$ is the gross interest factor applicable when at time $t+k$ $X_{t+k}=i$ and the dividend growth is low, etc. Although the computation of the risk neutral probability of the final nodes associated with the binomial lattice is a simple matter of matrix multiplication and matrix element extraction, this now requires keeping track of the risk-free rate as we move along the information tree, a path-dependent and therefore tedious task. Hence the Markov chain characterization of the risk neutral process no longer simplifies calculations a great deal.

5 Option prices on a learning path: Comparative statics

In view of the empirical biases of the BS option pricing model documented in section 2, it is important to establish conditions under which learning has systematic effects on option prices. Guidolin and Timmermann (1999, prop. 3) show that although estimation uncertainty reduces the underlying asset price when agents are risk averse, it also increases the asset price through the positive covariance between future asset payoffs and parameter estimates. When risk aversion is not 'too high' the second effect dominates. A similar result can be established comparing option prices under learning to BS option prices:

Proposition 6 *Suppose dividends follow a binomial process subject to the parameter restrictions* $\pi\geq\frac{1}{2}$, $\frac{1+g_h^*}{1+\rho}\geq 1$, $\frac{1+g_l^*}{1+\rho}<1$. *Conditional on a com-*

mon stock price, if agents have optimistic beliefs then[18]

$$C_t^{BL}(K) \geq C_t^{BS}(K).$$

Proof. See Appendix A. ■

More specifically one can show that $C_t^{BL}(K) - C_t^{BS}(K)$ is positive at $K = 0$, increases over some interval $(0, K^{MAX})$ and then decreases towards zero.

The technical conditions $\frac{1+g_h^*}{1+\rho} \geq 1$ and $\frac{1+g_l^*}{1+\rho} < 1$ are sufficient for the price-dividend ratio to be a monotonically increasing and convex function of the estimates $\widehat{\pi}_{t+k}$.[19] Second, optimistic beliefs $(\widehat{\pi}_{t+k} > \pi)$ are sufficient for the difference between option prices $C_t^{BL}(K) - C_t^{BS}(K)$ to be monotonically increasing for low strike prices $(K \in [0, \Psi_{t+k+v}^{BL}(n_{t+k}, N_{t+k} + v)(1 + g_l)^v D_t))$. For medium-range strike prices, the assumption $\pi \geq \frac{1}{2}$ guarantees that enough probability mass for the payoff of the underlying asset is redistributed to ensure a higher option price under learning.

In general there will be intervals for the strike price over which $C_t^{BL}(K) \geq C_t^{FIRE}(K)$ and others over which this inequality is reversed. Proposition 6 shows that subject to placing additional restrictions on the parameter space we can establish conditions under which learning induces systematic effects on option prices. These restrictions may provide valuable information about the markets' perception of fundamentals, once the model is confronted with option data (see Section 7). For instance, market prices for European calls systematically above BS predictions and an implied volatility shape similar to the one implied by proposition 6 would suggest that learning effects are important and that investors are somewhat optimistic.

5.1 Learning Effects and BS Anomalies

The significance of Proposition 6 can best be illustrated by relating it to the option pricing anomalies reported in the finance literature.

First consider implied volatility skews. BS implied volatility is a highly nonlinear function of option prices and a formal treatment is difficult. Instead we provide the following heuristic explanation. At-the-money BS option values are known to be sensitive — their *vega* is high — to the volatility input while deep out-of-the-money and in-the-money BS option

[18] The notation $C_t^{BS}(K)$ assumes we are focusing on a parameterization of our FIRE model which as Black-Scholes option pricing formula as its limit as $v \to \infty$.

[19] Convexity turns out to be important for the proof. However, to establish a local result we only need the restriction $\frac{1+g_h^*}{1+\rho} \geq 1$, as Ψ_{t+k}^{BL} is required only to be convex in a range of values of $\widehat{\pi}_{t+k}$.

values are not so sensitive to changes in σ. Now, when $K = 0$, C_t^{BL} starts out above the BS value and their difference initially increases as the strike price rises. Therefore, for $K << S_t$, the BS implied volatility satisfies $\hat{\sigma}^{BL} > \sqrt{m}\ln(1+g_h) = \sigma^{BS}$. For very low strikes, $\hat{\sigma}^{BL}$ can even be increasing, but as moneyness approaches one from above, $\hat{\sigma}^{BL}$ will turn out to be above σ^{BS} (a constant) and decreasing, as the vega increases and $C_t^{BL}(K) - C_t^{BS}(K)$ starts declining too. For $K >> S_t$, K exceeds K^{MAX} so that the difference $C_t^{BL}(K) - C_t^{BS}(K)$ declines towards zero and with it $\hat{\sigma}^{BL}$ also declines towards σ^{BS}. If traded strikes do not extend to levels for which the vega is very low, an implied volatility skew where $\hat{\sigma}^{BL}$ is decreasing everywhere results. However, for very high strikes, a vega approaching zero might also make a locally increasing profile for $\hat{\sigma}^{BL}$ possible, i.e. a volatility smile. In summary, under the assumption of proposition 6, plausible patterns of implied volatility can result.

To see how state price densities perceived on a learning path change option prices, notice that there are two distinct learning effects: (i) the SPD perceived by agents changes w.r.t. the FIRE case; (ii) BL widens the support of the set of time $t+T$ stock prices that are achievable in equilibrium. As for the first effect, by taking the log-ratio of the state price density under BL and under FIRE, we see that

$$\ln\left(\frac{\widetilde{P}_t^{BL}\left\{S_{t+T}^j|n_t, N_t\right\}}{\widetilde{P}_t^{FIRE}\left\{S_{t+T}^j\right\}}\right) = \left[\sum_{k=0}^{j-1}\ln\left(\frac{\hat{\pi}_{t+k}}{\pi}\right)\right] + \left[\sum_{k=0}^{v-j-1}\ln\left(\frac{1-\hat{\pi}_{t+k}}{1-\pi}\right)\right]$$

If $\widehat{\pi}_t = \pi$, a limiting case of the conditions imposed by proposition 6, it follows that

$$\ln\left(\frac{\widetilde{P}_t^{BL}\left\{S_{t+T}^v|\widehat{\pi}_t\right\}}{\widetilde{P}_t^{FIRE}\left\{S_{t+T}^v\right\}}\right) = \sum_{k=0}^{v-1}\ln\frac{\hat{\pi}_{t+k}}{\pi} > 0$$

implying that $\widetilde{P}_t^{BL}\left\{S_{t+T}^v|\hat{\pi}_t\right\} > \widetilde{P}_t^{FIRE}\left\{S_{t+T}^v\right\}$, as $\hat{\pi}_{t+k} \geq \pi$ for all $k \geq 0$ since the dividends grow in all periods at a high rate. Similarly,

$$\ln\left(\frac{\widetilde{P}_t^{BL}\left\{S_{t+T}^0|\widehat{\pi}_t\right\}}{\widetilde{P}_t^{FIRE}\left\{S_{t+T}^0\right\}}\right) = \sum_{k=0}^{v-1}\ln\left(\frac{1-\hat{\pi}_{t+k}}{1-\pi}\right) > 0$$

implies $\widetilde{P}_t^{BL}\left\{S_{t+T}^0|\hat{\pi}_t\right\} > \widetilde{P}_t^{FIRE}\left\{S_{t+T}^0\right\}$, as $(1-\hat{\pi}_{t+k}) > (1-\pi)$ for all $k \geq 0$ since dividends always grow at a low rate. This shows that if the initial beliefs on π are unbiased, under BL the SPD is above the FIRE SPD at both the extremes of the possibile realizations of dividend growth rates between t and $t+T$. Letting $\Upsilon(j) = \ln\left(\frac{\widetilde{P}_t^{BL}\{S_{t+T}^j|\widehat{\pi}_t\}}{\widetilde{P}_t^{FIRE}\{S_{t+T}^j\}}\right)$ be the log-price ratio, notice that

$$\frac{d\Upsilon(j)}{dj} = \ln\left(\frac{\hat{\pi}_{t+j-1}}{\pi}\right) - \ln\left(\frac{1-\hat{\pi}_{t+v-j-1}}{1-\pi}\right) \qquad j = 1, 2, ..., v-1$$

and $\frac{d\Upsilon(j)}{dj} > 0$ for $j \geq int(v\pi) + 1$[20] and $\frac{d\Upsilon(j)}{dj} < 0$ for $j < int(v\pi) - 1$, i.e. $\Upsilon(0), \Upsilon(v) > 0$, and $\Upsilon(j)$ at first decreases in j, it reaches a minimum and then increases. Hence there exists one lower and one upper point, $j_1 < int(v\pi)$ and $j_2 > int(v\pi)$, such that $\Upsilon(j)$ changes sign when j crosses these points. For $j < j_1$ and $j > j_2$ $\Upsilon(j) > 0$, that is, $\widetilde{P}_t^{BL}\left\{S_{t+T}^j|\hat{\pi}_t\right\} > \widetilde{P}_t^{FIRE}\left\{S_{t+T}^j\right\}$. Therefore for $\widehat{\pi}_t = \pi$ the BL-SPD has fatter tails than the FIRE-SPD and lacks mass at the center of the distribution of future stock prices. Extreme events drawn from either end of the tail are perceived as more likely on a Bayesian learning path than under a lognormal distribution. Adding the second effect of BL, this conclusion is strengthened, since a wider support for the SPD means that more extreme events now become possible.

When $\widehat{\pi}_t > \pi$ and $\widehat{\pi}_t$ decreases towards π ($\widehat{\pi}_t \searrow \pi$), BL continues to inflate the tails of the SPD perceived by the agents, relative to the lognormal benchmark. Hence higher call prices and the implied volatility skew still result. The fact that $\widehat{\pi}_t > \pi$ implies that the right tail receives more probability mass than the left one. However, the picture changes somewhat when $\widehat{\pi}_t >> \pi$. In this case, we may have that

$$\Upsilon(0) = \ln\left(\frac{\widetilde{P}_t^{BL}\left\{S_{t+T}^0|\widehat{\pi}_t\right\}}{\widetilde{P}_t^{FIRE}\left\{S_{t+T}^0\right\}}\right) = \sum_{k=0}^{v-1} \ln\left(\frac{1-\hat{\pi}_{t+k}}{1-\pi}\right) < 0$$

since the sum of the terms $\frac{1-\hat{\pi}_{t+k}}{1-\pi} < 1$ may dominate the sum of terms for which $\frac{1-\hat{\pi}_{t+k}}{1-\pi} > 1$. All other findings hold and $\widetilde{P}_t^{BL}\left\{S_{t+T}^j|\hat{\pi}_t\right\}$ has a fatter right-hand tail than the lognormal, though the left tail might now be thinner. When we take into account also the second effect described above, the net result is uncertain with respect to the probability mass in the left tail. Intriguingly, the net effect could be similar to the stylized facts for SPDs found for the S&P 500 option data: densities are located more to the

[20] $int(x)$ denotes the integer part of the real number x. These derivatives obviously ignores the fact that j is by construction an integer index.

right than a lognormal benchmark but also attach positive probability mass to some crash events.

6 Simulation exercises

This section investigates whether the option pricing model with learning is capable of explaining some of the empirical evidence on implied volatility surfaces and SPDs reported in Section 2 for S&P 500 index options. As a reference point we take the plot of the implied volatilities versus moneyness for the S&P500 option contracts expiring in April 1993. As a matter of fact, these results are representative of what we have found in other subperiods of our overall sample.

We calibrate the parameters of the dividend process and fix the risk-free rate at a plausible level. Dividends are assumed to be paid out daily ($m = 1$). For a wide market index such as the S&P 500 this assumption is a fairly good approximation. The annual dividend growth rate (μ) is set to 3%, while the volatility (σ) is set at 5%. The growth rate matches the average dividend growth rate during the sample period June 1992 - June 1995 around our option simulation period. While our choice of σ is lower than the standard deviation of the annual dividend growth rate (11.6%), this latter estimate is likely to be inflated by the presence of periods where no dividend changes show up, even though markets receive news in a much smoother manner. The annualized risk-free rate (r) is set to 4%, while the dividend yield is fixed at 3%, matching its estimate of 2.82% during the period 1992-1995. The (annualized) discount rate is 0.02.[21]

European call options with 50 days to maturity, $\tau = v = 50$ are considered. From proposition 2 it follows that the no-arbitrage option price in an equilibrium model with dividends follows a binomial lattice model with parameters

$$g_h = 0.00315,\ g_l = -0.00314,\ \pi = 0.5189,\ \rho = 7.96 \cdot 10^{-5},\ \gamma = 0.999 \tag{14}$$

We assume marginally biased initial beliefs and quite mild learning effects: $n_t = 42$ and $N_t = 80$, i.e. $\hat{\pi}_t = 0.525 > 0.519 = \pi$.[22] BS prices are calculated

[21] γ is determined according to proposition 3 to guarantee that $C_t^{FIRE} \to C_t^{BS}$ as $m \to \infty$.

[22] Under pessimistic beliefs the BL model is less able to produce option prices that fit the stylized facts of Section 2. For example the implied volatility surface tends to increase as a function of the strike price. Although the SPD still produces fatter tails than the log-normal model, most probability mass is shifted to the left of the return distribution when beliefs are pessimistic.

by using the FIRE approximation. Learning effects are studied by modifying the FIRE model as described in section 4. Figure 3 plots the difference between FIRE and BL call prices for a wide range of strike prices. Since all of the assumptions of proposition 3 are satisfied, $C_t^{BL}(K) \geq C_t^{FIRE}(K)$ for all strikes.

Figure 3. Difference between the price of a European call with 50 days to expiration (τ=50) calculated under full information rational expectations and with Bayesian learning. The parameters are: m=1, g_h=0.00315, g_l=0.00314, π=0.519, ρ=0.02 (annual), and γ=0.999. For this combination of parameters the FIRE call price approximates the Black-Scholes price, while the annual risk-free rate under FIRE is 4% and the dividend yield is 3%. For BL prices, we take n_t=42 and N_t=80, implying only a slightly biased initial belief $\hat{\pi}_t$ =0.525. The current price of the S&P 500 index is taken to be 436.38 dollars, the closing price on Feb. 22, 1993.

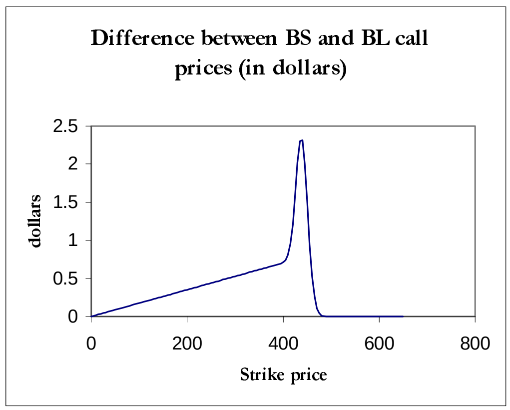

6.1 Implied Volatility Surfaces

Initially we study implied volatility as a function of moneyness. Figure 4 and 5 plot implied volatilities as a function of moneyness for the April 1993 maturity. The resemblance between the implied volatility of simulated prices under learning and market prices is striking. Noticeably, the lattice model under Bayesian learning prices options far more accurately than the BS model.

To further underline this point, Figure 6 compares option prices on February 22, 1993 and BL simulated prices when $n_t = 43$, $N_t = 80$ ($\hat{\pi}_t = 0.538$). The fit is even more striking than in Figure 5 and it is clearly indicative of the ability of the model to fit the stylized facts concerning the skews of implied volatility. Assuming that our model correctly captures the main dynamic features of the underlying process and that the markets really were on a Bayesian learning path on that day, it seems that the estimate

$\widehat{\pi}_t = 0.538$ with a precision of $N_t = 80$ accurately characterizes agents' beliefs during that period.

Figure 4. Implied Black-Scholes volatilities as a function of moneyness for a European call with 50 days to expiration (τ=50) calculated under full information rational expectations. The parameters are set as follows: m=1, g_h=0.00315, g_l=-0.00314, π=0.519, ρ=0.02 (annual), and γ=0.999. For this combination of parameters the BS call price approximates the Black-Scholes price, while the annual risk-free rate under FIRE is 4% and the dividend yield is 3%. The current price of the S&P 500 index is taken to be 436.38 dollars, the closing price on Feb. 22, 1993. The left panel reports for comparison implied volatilities as a function of moneyness for (put and call) options expiring in April 1993.

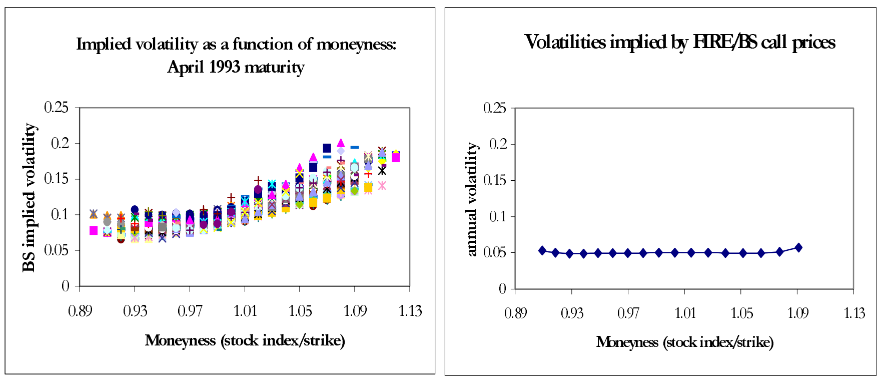

Figure 5. Implied Black-Scholes volatilities as a function of moneyness for a European call with 50 days to expiration (τ=50) calculated on a Bayesian learning path. The parameters are set as follows: m=1, g_h=0.00315, g_l=-0.00314, π=0.519, ρ=0.02 (annual), γ=0.999, n_t=42, and N_t=80, implying an initial belief $\hat{\pi}_t$ =0.525. The current price of the S&P 500 index is taken to be 436.38 dollars, the closing price on Feb. 22, 1993. The left panel reports for comparison implied volatilities as a function of moneyness for (put and call) options expiring in April 1993.

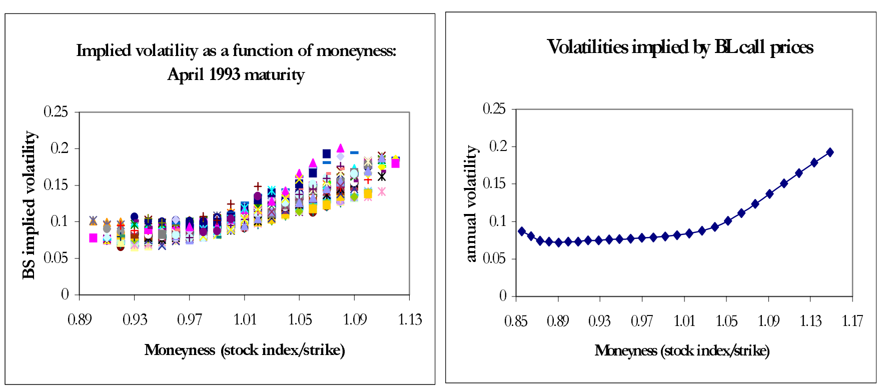

Figure 6. Implied Black-Scholes volatilities as a function of moneyness for a European call with 50 days to expiration (τ=50) calculated on a Bayesian learning path, and for market data on February 22, 1993. The parameters are set as follows: m=1, g_h=0.00315, g_l=0.00314, π=0.519, ρ=0.02 (annual), γ=0.999, n_t=43, and N_t=80, implying an initial belief $\hat{\pi}_t$ =0.538.

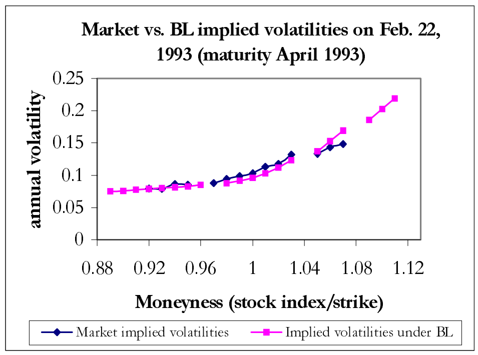

6.2 State Price Densities

Systematic differences between BS and BL European option prices must reflect differences in the underlying equivalent martingale measures employed by the market under the two alternative models or, equivalently, in the SPDs (cf. Harrison and Kreps (1979)). Sections 4 and 5 showed that the SPDs under BL can be very different from the Black-Scholes SPD. Figure 7 compares the SPDs of the two models from our previous exercise.

Bayesian learning produces SPDs that are skewed to the right and have fatter tails than a lognormal. Once again our model fits better the stylized facts for estimated (average) SPDs documented in Section 2. While estimated SPDs are leptokurtic with positive excess kurtosis essentially caused by side 'lobes' between six and eight standard deviations from the demeaned S&P500 returns. Furthermore, BL generates a SPD which displays sensibly lower state prices around the mean of the distribution of the index returns, with smooth fat tails.

Figure 7. State-price densities for the 50 days ahead (τ=50) values S&P 500 index derived from a FIRE/BS vs. a Bayesian learning model on Feb. 22, 1993. The parameters are set as follows: m=1, g_h=0.00315, g_l=0.00314, π=0.519, ρ=0.02 (annual), γ=0.999, n_t=42, and N_t=80, implying an unbiased initial belief $\hat{\pi}_t$ =0.525. The current price of the S&P 500 is 436.38 dollars, the closing price on Feb. 22, 1993. The third panel compares the FIRE SPD with the BL SPD adjusting for differences in their respective supports.

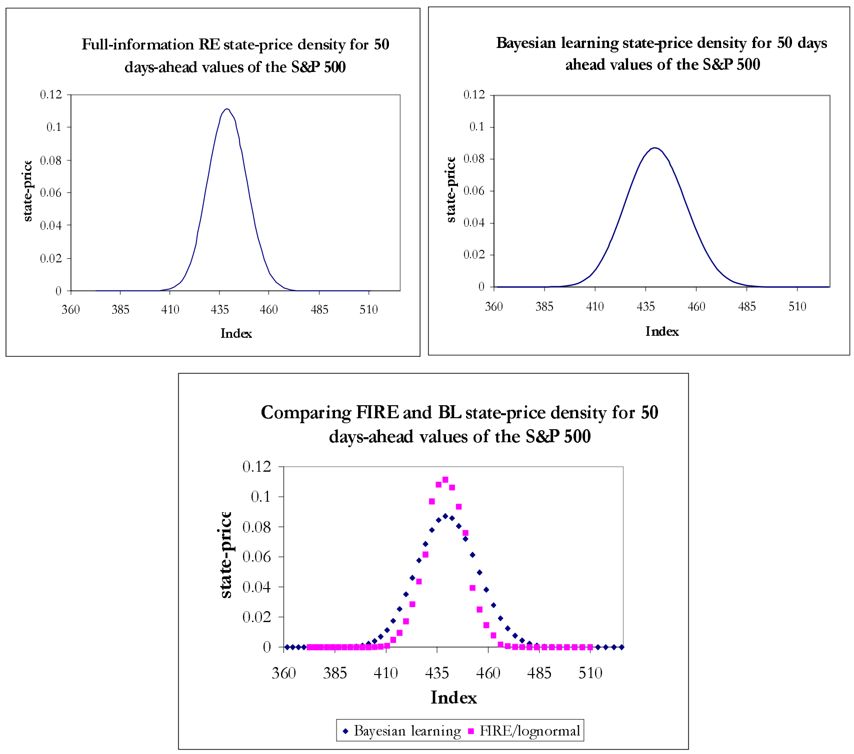

6.3 Term Structure of Implied Volatilities

Next we vary the time to maturity (τ) from 10 to 150 days in steps of 10 days to study which type of implied volatility term structure Bayesian learning implies. Figure 8 reports the outcome of this exercise for three sets of strike prices: $K_{ITM} = 420$ (moneyness 1.04), $K_{ATM} = 435$ (moneyness 1), and $K_{OTM} = 455$ (moneyness of 0.96). There is a strong resemblance with the stylized facts discussed in Section 2. Bayesian learning generates an upward sloping implied volatility term structure for ATM and out-of-the money call options, while the term structure at first decreases and then increases for in-the-money call options. These patterns are broadly consistent with what was found in the data, cf. figure 2. The increase of about 2 percentage points in implied volatility between close-to-expiration options and long-term options is also plausible.[23]

Figure 8. Volatility term structure plots for European calls calculated under full information rational expectations (squares) and with Bayesian learning (diamonds). In the simulations, m=1 while the other 'deep' parameters (g_h, g_l, π, ρ, and γ) are adjusted according to proposition 2 to guarantee that the FIRE option price is an approximation of the Black-Scholes value with an annual FIRE risk-free rate of 4% and a dividend yield of 3%. In the case of BL prices, we take n_t=42 and N_t=80, implying an slightly biased initial belief $\hat{\pi}_t$ =0.525. The current price of the S&P 500 is taken to be 436.38 dollars, the closing price on Feb. 22, 1993. The first panel sets K=455, the second panel K=435, and the third K=420.

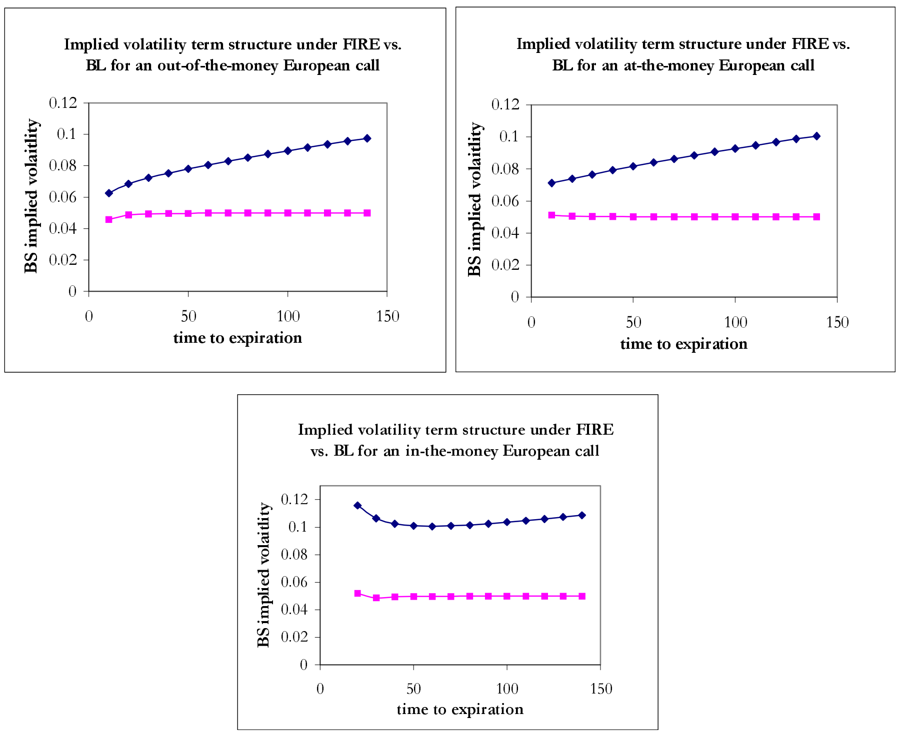

6.4 Weak Learning Effects

In our setup investors use a consistent estimator of π, and as the precision goes to infinity ($N_t \to \infty$), $\hat{\pi}_t \to \pi$ and learning effects diminish. To study the consequences of this, we set $n_t = 182$, $N_t = 350$, implying $\hat{\pi}_t = 0.52 \gtrsim 0.519 = \pi$. The representative agent now brings experience of over 16 months of trading and dividend realizations. Figure 9 shows the outcome of this new

[23] See Campa and Chang (1995).

set of simulations.

Figure 9. Weak learning effects — Plots of differences in dollars between call prices, implied volatilities as a function of the strike price, 50 days ahead state-price densities for the S&P 500 index on Feb. 22,1993 for the FIRE and BL asset pricing models. The first three graphs refer to a European call 50 days to expiration (τ=50). The fourth plot represents with squares the SPD calculated under full information rational expectations and with diamond the SPD calculated under Bayesian learning. The parameters are set as follows: m=1, g_h=0.00315, g_l=0.00314, π=0.519, ρ=0.02 (annual), γ=0.999, n_t=182, and N_t=350, implying an initial unbiased belief $\hat{\pi}_t$ =0.52 $\cong \pi$. The current price of the S&P 500 is taken to be 436.38 dollars, the closing price on Feb. 22, 1993.

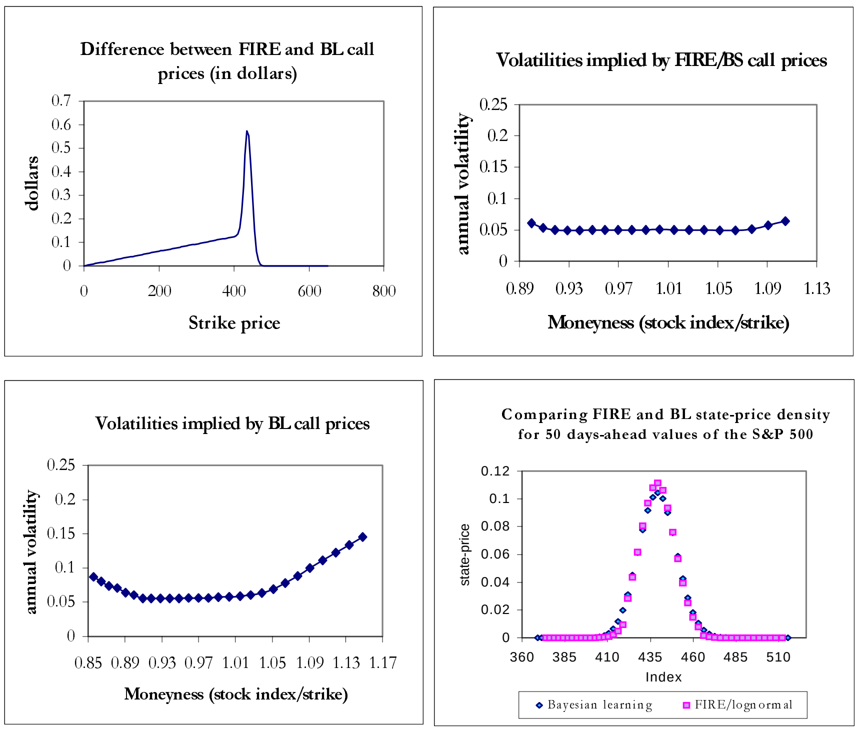

When learning effects are weak and agents have a more accurate estimate of π, BL option prices converge to BS prices. This is unsurprising and it follows since learning is the only source of non-stationarity in our model. The first panel of Figure 9 stresses that differences previously of the order of 1-2 dollars, now decline to a quarter of that range. Indeed panels two and three show that BL implied volatilities continue to display a precise pattern over moneyness. However, the implied volatility surface is flatter than the one exhibited in figure 5, with a left-hand tail that bends upwards. As $N_t \to \infty$ and $\hat{\pi}_t \to \pi$ (from above) a smile is obtained instead of the smirk in figure 5. This indicates that at times when option markets are characterized by smiles, learning effects are weak in the sense that agents attach high precision to their initial beliefs. On the other hand, smirks are indicative of markets with strong learning effects as agents are still highly uncertain about their beliefs. Finally, the fourth panel shows how BL and FIRE prices converge: The two underlying SPDs are now very close.[24]

7 Inferring Learning Effects from Option Prices

So far we have studied the ability of the BL model to generate option prices (BS implied volatilities) that match typical shapes and stylized facts on individual days. However, the Bayesian updating algorithm implies a set of dynamic restrictions on how implied volatility surfaces and term structures evolve over time as agents update $\widehat{\pi}_t$. Such testable restrictions do not have a counterpart in the BS model which does not consider the effect of changing probability beliefs. By tracking option prices on several consecutive days, not only do we get insights into how agents change their beliefs but we also get a more precise estimate of the initial beliefs.

Estimating the dynamics of beliefs from observed option prices is a very different inferential exercise that needs to be put in perspective. When asset markets are (dynamically) complete, equilibrium asset prices contain information about preferences and beliefs. Rubinstein (1985) observes that any of the following implies the third: (1) the preferences of a representative agent; (2) agents' beliefs; and (3) the state-price density (SPD). Therefore, a vast literature has attempted to use the observed prices of risky assets to infer preferences, the (objective) stochastic process of prices, or both. For instance, Bick (1990) and He and Leland (1993) impose parametric restrictions on the generating process of asset prices and infer the preferences of a representative agent in an equilibrium asset pricing model. Furthermore, Bates (1991), Jackwerth and Rubinstein (1996), and Aït-Sahalia and Lo (1998) (among others) back out the perceived (risk neutral) stochastic process of asset prices from observed market prices of options.

Bayesian learning provides an as far unexplored possibility to expand Rubinstein's list to a fourth and separate item: the dynamics of the learning process followed by a representative agent in an equilibrium model. Akin

[24] We are omitting the discussion of the impact of large N_t on the volatility term structure effects. What happens is easy to see: though the patterns remain as shown above, all the implied volatility schedules as a function of time-to-expiration get flatter and their level moves down towards the FIRE/BS value of $\sigma = 5\%$. Surprisingly, a N_t of the order of more than 16 months of observations implies effects that are still close to those observed in real option markets data.

to Rubinstein (1995) and Jackwerth and Rubinstein (1996), our objective is to fix preferences and infer a vector of unknown parameters from observed option prices on a standard binomial tree. However, our exercise also proceeds to impose restrictions on the stochastic process of the fundamentals (dividends) underlying stock prices. Since the dynamics of beliefs on a learning path determine the time-varying SPD, and the SPD ties down option prices to the underlying price in a precise manner, the very parameters characterizing the learning process can be inferred.

The model in section 4 imposes a precise structure on the temporal dynamics of beliefs which evolve according to

$$\widehat{\pi}_{t+k,M} = \frac{\sum_{j=1}^{m(t+k)+M} I_{\left\{g_{m(t+k)+j}=g_h\right\}}}{m(t+k)+M} = \frac{n_{t+k,M}}{N_{t+k,M}}$$

If only one piece of news arrives every period, this intertemporal structure rules out very volatile beliefs and jumps in the precision. This flexibility might be required since m is unknown and may vary over time. To account for such effects, we solve the program[25]

$$\min_{\{\pi_t\}_{t=1}^{T},N,m} \sum_{t=1}^{T} \sum_{\tau=\underline{\tau}_t}^{\bar{\tau}_t} \sum_{K_{\tau_t}=\underline{K}_{\tau_t}}^{\overline{K}_{\tau_t}} g\left(C^{BL}(\tau_t, K_{\tau_t}, S_t, \pi_t, N, m; \gamma, \beta), C(\tau_t, K_{\tau_t}, S_t)\right) \quad (15)$$

$$s.t. \quad \frac{\pi_t N}{N+m} \leq \pi_{t+1} \leq \frac{\pi_t N + m}{N+m}$$

$$0 \leq \pi_t \leq 1 \quad t = 1, ..., T \quad (16)$$

$$N > 0, \quad m > 0 \quad (17)$$

where $C^{BL}(\tau_t, K_{\tau_t}, S_t, \pi_t, N, m; \gamma, \beta)$ is the theoretical price of a call option on a learning path provided by proposition 5, $C(\tau_t, K_{\tau_t}, S_t)$ is the observed call price at time t for an option expiring in τ_t days, with strike K_{τ_t}, when the underlying stock index is S_t. The indexes appended to τ_t and K_{τ_t} reflect the fact that traded maturities and strike prices change over time, following the dynamics of the underlying price and the financial cycle. Finally, $g(\cdot)$ is a function measuring the distance between the cross-section of option prices observed in the market and the corresponding cross-section of theoretical BL

[25] Our approach is similar to Bates (1991), who imposes CRRA preferences to estimate by NLS the parameters of an asymmetric jump diffusion model in which jump risk is systematic (priced in equilibrium).

prices. For instance, we might choose to minimize the sum of the squared pricing errors across days/strikes/maturities.

Notice that the precision is a fixed constant N to be estimated and the program acquires an explicit intertemporal nature. Therefore, the estimation procedure provides an estimate $\hat{N}$ which represents how much precision the market assigns to its beliefs, an estimate $\hat{m}$ of the frequency with which these beliefs are updated over the sample period, and a $T \times 1$ vector $\hat{\pi}$ whose dynamics is constrained by Bayes rule. Notice that $\hat{N}$ can also be interpreted as an estimate of the length of the time window the agent is basing his knowledge about the stochastic process of the fundamentals on.

In the following we distinguish between two alternative empirical exercises that impose further restrictions on the fairly general program represented by (15). All of these estimation problems share a common set of inputs and extract some estimated outputs. Again, we assume that the annual volatility of the fundamentals is $\sigma = 5\%$. Given σ, g_h and g_l can be determined by applying the formulas in proposition 3. We assume $\gamma = 0.9$ and $\beta = \frac{1}{1.02} \simeq 0.98$ on an annual basis.[26] These choices are either based on the features of our data on the S&P 500 index and index options, or based on what seems plausible. Nevertheless we also explore the sensitivity of our results to these choices by experimenting with alternative values of σ, γ, and β.

7.0.1 Inferring learning effects from daily cross-sections of option prices

The objective of the first exercise is to infer from option prices the belief π_t, its precision level N_t, and the frequency with which beliefs are updated and new information arrives (m_t) for each day in the sample:

$$\min_{\pi_t, N_t, m_t} \sum_{\tau=\underline{\tau}_t}^{\overline{\overline{\tau}}_t} \sum_{K_{\tau_t}=\underline{K}_{\tau_t}}^{\overline{K}_{\tau_t}} \left[C^{BL}(\tau_t, K_{\tau_t}, S_t, \pi_t, m_t, N_t) - C(\tau_t, K_{\tau_t}, S_t)\right]^2 \quad (18)$$

[26] β also needs to be adjusted as a function of the frequency with which dividend news arrive (m) to avoid that the actual subjective discount rate varies as m varies. Therefore, with reference to each single step of our binomial lattice for dividends, we will be using a

$$\hat{\beta}_m = \beta^{\frac{1}{365m}}$$

Figure 10. Beliefs regarding the probability of good news (high growth of dividends, $\hat{\pi}_t$), precision of these beliefs (number of previously observed states, $\hat{N}_t$), and number of state realizations (dividends news, $\hat{m}_t$) implied by observed S&P 500 index option prices during the period Jan. 4, 1993 - June. 30, 1993 (125 trading days). $\{\hat{\pi}_t, \hat{N}_t, \hat{m}_t\}_{t=1}^{T}$ are estimated by solving the program:

$$\min_{\pi_t, N_t, m_t} \sum_{\tau=\underline{\tau}_t}^{\bar{\tau}_t} \sum_{K_{\tau_t}=\underline{K}_{\tau_t}}^{\bar{K}_{\tau_t}} \left[C_t^{BL}(\tau_t, K_t, S_t, \pi_t, N_t, m_t) - C_t^{BL}(\tau_t, K_t, S_t) \right]^2$$

i.e. by minimizing the daily (cross-sectional) sum of the squared pricing errors from the theoretical model with Bayesian learning. The exercise assumes β=.98 (annually), CRRA-power preferences with γ=0.9, and a volatility for dividends σ=5% (annually). Future payoffs are discounted employing the 3-months average T-bill rate, measured on a daily basis.

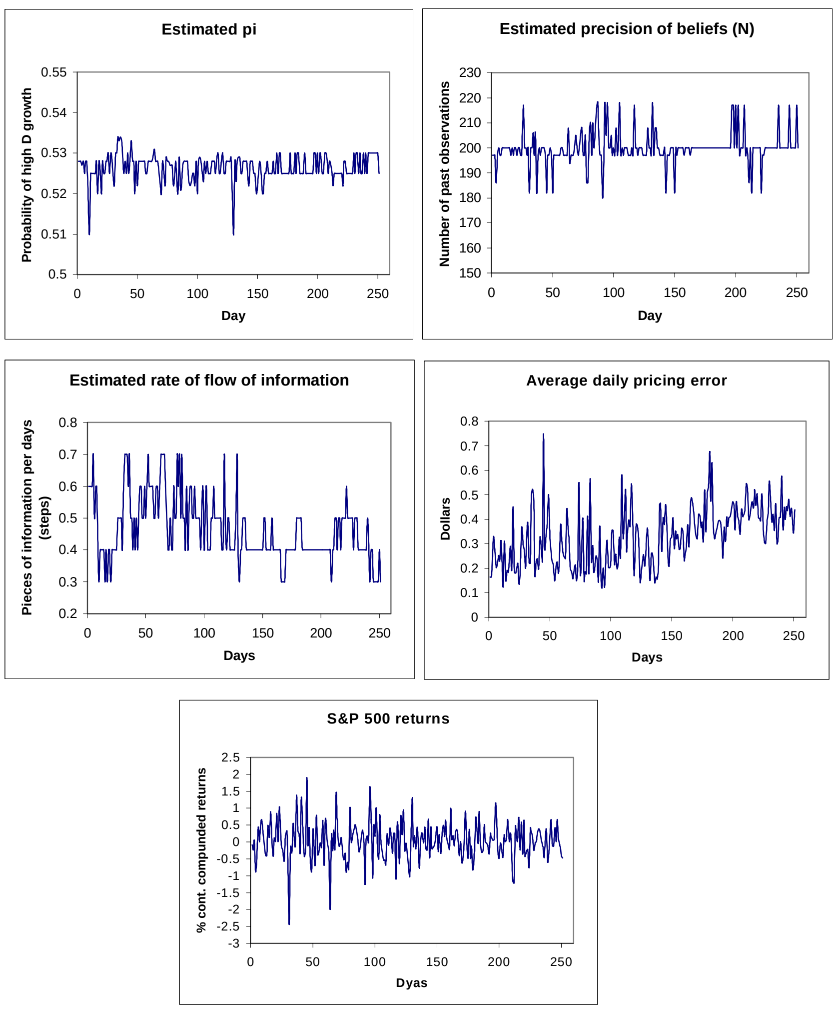

for $t = 1, 2, ..., T$. This amounts to minimizing the in-sample squared pricing errors produced by the model for each daily cross section of option prices. $\hat{\pi}_t, \hat{N}_t$, and $\hat{m}_t$ can be viewed as approximate NLS estimators. This exercise ignores the intertemporal restrictions imposed by our model on the updating of beliefs by agents and therefore does not provide the strongest possible test. We do not rule out high volatility of the estimated belief $\hat{\pi}_t$ or unbounded variation in the estimated precision level $\hat{N}_t$. On the other hand, the exercise is quite simple, requiring the estimation of only three parameters on a data set with a number of observations equal to the number of traded contracts in day t along the dimensions of maturity and strike price.

Figure 10 reports findings for our 1993 S&P500 index options data, using a total of 252 trading days. Program (18) was independently solved 252 times, i.e. for 252 different cross-sections of option prices yielding 252 vectors of daily estimates $\left[\hat{\pi}\ \hat{N}\ \hat{m}\right]'$. Figure 10 stresses that the estimated parameters are quite stable over time and that the model fits observed prices quite precisely. Following the same criteria as Dumas et al. (1998), this argues in favor of our model. In general option markets display weak optimism (i.e. $\hat{\pi}_t \geq \frac{1}{2}\ \forall t$). Over the period analyzed, news on the S&P 500 seem to flow at a quite slow and stable rate: option prices imply that on average the fundamentals change about every two days, $0.3 \leq \hat{m}_t \leq 0.7\ \forall t$. The flow of information seems to reach higher peaks between February and mid-April, when $\hat{m}_t$ is often above 0.5. On the contrary the flow of information perceived by option markets slows down between June and September. In general learning effects — and therefore also the slope of the implied volatility skew — are quite mild, since for most of the days $\hat{N}_t$ stays around 200. Such a level for $\hat{N}_t$ implies that market participants would tend to use a moving window of past observations of almost 20 months trading.[27] However in several periods (February, June, and October) the average precision drops on many occasions (to about 180-190) while in the same period $\hat{\pi}_t$ is often subject to drastic downward or upward revisions. Apparently infleuntial news induce the agents to feel much less confident about the precision of their knowledge than they normally are. However it is difficult to link the behavior of the estimated $\hat{\pi}_t$ to current and previous returns on the underlying stock index.

[27] Since $\hat{m}$ is on average around 0.5, $\hat{N} = 200$ implies that about 400 trading days have to elapse before receiving 200 pieces of information concerning the fundamentals.

Table 2. Summary statistics concerning the perceived probability ($\hat{\pi}_t$) of high dividend growth, its precision (number of previously observed states, $\hat{N}_t$), and the number of state realizations (dividends news, $\hat{m}_t$) implied by observed S&P 500 index option prices during the period Jan. 4, 1993 - June. 30, 1993 (125 trading days). $\{\hat{\pi}_t, \hat{N}_t, \hat{m}_t\}_{t=1}^{T}$ are estimated by solving the program:

$$\min_{\pi_t, N_t, m_t} \sum_{\tau=\underline{\tau}_t}^{\bar{\tau}_t} \sum_{K_{\tau_t}=\underline{K}_{\tau_t}}^{\bar{K}_{\tau_t}} \left[C_t^{BL}(\tau_t, K_t, S_t, \pi_t, N_t, m_t) - C_t^{BL}(\tau_t, K_t, S_t)\right]^2$$

i.e. by minimizing the daily cross-section of squared pricing errors from the theoretical model with Bayesian learning. The exercise assumes β=.98 (annually), CRRA-power preferences with γ=0.9, and a volatility of dividends σ=5% (annually). Future payoffs are discounted employing the 3-months T-bill rate, measured on a daily basis.
The innovation in $\hat{m}_t$ is measured as the residual from the autoregression $m_t = \alpha + \beta m_{t-1} + e_t$.

	Minimum	Maximum	Mean	Median	SD
$\hat{\pi}_t$	0.510	0.534	0.526	0.525	0.00313
$\hat{N}_t$	180	218	199.9	200	6.35
$\hat{m}_t$	0.3	0.7	0.463	0.4	0.098
Average daily pricing errors (dollars)	0.12	0.75	0.33	0.33	0.12
Correlations	Current return	Lagged Return	Absolute return	$\Delta\hat{N}_t$	$\hat{m}_t$ innovation
$\hat{\pi}_t$	-0.065	-0.037	0.000	0.054	0.352
$\Delta\hat{N}_t$	-0.041	0.051	-0.151	1.000	0.145
$\Delta\hat{m}_t$	-0.152	-0.281	-0.310	0.084	1.000
$\hat{m}_t$ innovation	-0.338	0.012	0.083	0.145	1.000
Average daily pricing errors	0.011	-0.034	0.100	-0.053	0.074

The fourth (in clockwise direction) panel of figure 10 highlights that our model achieves a very good in-sample fit. The fit deteriorates only around abrupt changes in the precision levels.[28] This is possibly a sign that our model is too simple to entirely capture the effects of fundamental shocks on beliefs and prices. Table 2 provides a few summary statistics: the mean (across maturities and strikes) of the average daily pricing error is of the order of 33 cents (median: 33 cents). This means that on average over the sequence of cross sections, each option was misspriced by only 33 cents.

[28] Furthermore, the in-sample fit worsen toeards the end of the sample period (Summer and Fall 1993). However this feature is likely to mainly depend on the appearance in the sample of longer maturity (around and above 200 days to expiration) options after July 1993.

There are also days in which the fit of the model is really impressive (for instance a mean of 12 cents is reached on May 14 over 32 different option contracts). In no day out of 252 the mean pricing error exceeds 1 dollar. The table stresses the extremely low standard deviation of the estimated parameters.[29]

7.0.2 Inferring learning effects from the dynamics of the implied volatility surface

Next, we infer the dynamics of learning from observed option market outcomes by imposing the intertemporal restrictions implied by the BL model. Fixing preferences, we solve the program:

$$\min_{\{\pi_t\}_{t=1}^{T}, N} \sum_{t=1}^{T} \sum_{\tau=\underline{\tau}_t}^{\bar{\tau}_t} \sum_{K_{\tau_t}=\underline{K}_{\tau_t}}^{\overline{K}_{\tau_t}} \left[C^{BL}(\tau_t, K_{\tau_t}, S_t, \pi_t, N, m; \gamma, \beta) - C(\tau_t, K_{\tau_t}, S_t)\right]^2$$

$$s.t. \quad \frac{\pi_t N}{N+1} \leq \pi_{t+1} \leq \frac{\pi_t N + 1}{N+1}$$

$$0 \leq \pi_t \leq 1 \quad t = 1, ..., T-1 \qquad N > 0 \tag{19}$$

on a bi-weekly basis, i.e. for the panel of option prices observed in the trading days of the first 15 days of January, then for the panel relative to the last 16 days of January, etc. The reason for limiting the problem to periods of 13-16 calendar days (hence 12 trading days at most) is to avoid the curse of dimensionality implicit in this NLS estimation program: as $T \to \infty$, the dimension of the parameter vector $\theta = [\{\pi_t\}_{t=1}^{T}, N]'$ tends to infinity making the optimization problem practically impossible. Limiting ourselves to a bi-weekly basis we thus approach the estimation of a vector with 10-13 parameters employing a number of observations usually well above 400. Notice that we are imposing much stronger requirements on our model than in the previous sub-section, narrowly restricting the temporal dynamics of

[29] To check the robustness of our results we repeated the same exercise for a few alternative choices of the deep parameters we have been calibrating, particularly the volatility of news on fundamentals σ and the CRRA coefficient γ. To save computing time, these robustness checks were applied to the shorter period Jan. - Feb. 1993 only, for a total of 39 trading days. The series of estimated parameters are still very stable (if any, they are even more stable) and the changes with respect to figure 12 are minor. In all three experiments, some degree of optimism prevails and the precision is around 200. Only the average level of the estimated m seems to strongly depend on the assumption on σ, although this is to be expected since the overall variability of the fundamentals over a fixed span of time depends both on σ and on the frequency with which news are assumed to flow in the market.

agents' beliefs with regard to the probability of good states. Therefore the ability of our model to adequately fit market prices should be considered a challenging test of its capability to effectively organize the available data.

Figure 11. Beliefs regarding the probability of good news (high growth of dividends, $\hat{\pi}_t$), precision of these beliefs (number of previously observed states, $\hat{N}_t$), and number of state realizations (dividends news, $\hat{m}_t$) implied by observed S&P 500 index option prices during bi-weekly periods, Jan. - Feb. 1993. $\{\hat{\pi}_t\}_{t=1}^{T}, \hat{N}_1$ are estimated by solving the program:

$$\min_{\{\hat{\pi}_t\}_{t=1}^{T},\hat{N}_1} \sum_{t=1}^{T} \sum_{\tau=\underline{\tau}_t}^{\bar{\tau}_t} \sum_{K_{\tau_t}=\underline{K}_{\tau_t}}^{\bar{K}_{\tau_t}} \left[C_t^{BL}(\tau_t,K_t,S_t,\pi_t,N_t,m_t)-C_t^{BL}(\tau_t,K_t,S_t)\right]^2$$

$$s.t \quad \frac{\pi_t N_1}{N_1+1} \le \pi_{t+1} \le \frac{\pi_t N_1+1}{N_1+1},\ 0\le \pi_t \le 1 \ \ t\ge 1$$

$$N_1>0, \quad N_1 \in \mathrm{N}$$

i.e. by minimizing the (cross-sectional) sum of the squared pricing errors implied by the model under BL and the implied volatilities of observed market prices. The exercise assumes β=.98 (annually) , CRRA-power preferences with γ=0.9, and a volatility for fundamental news σ=5% (annually). Future payoffs are discounted using the 3-months average T-bill rate, measured on a daily basis.

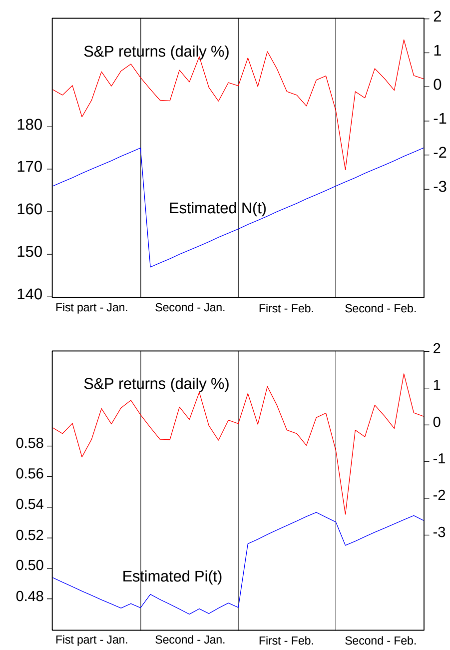

Figure 11 reports preliminary results for the two month period January - February 1993. Although the in-sample fit deteriorates[30], the estimated

[30] Due also to the fact that with $\sigma = 5\%$ we are *de facto* imposing $m = 1$ as implied

sequences $\{\hat{\pi}_t\}_{t=1}^T$ remain very stable and lie in the interval [0.51,0.54]. Also the estimation results concerning the data window used by the agents to infer the unknown process for fundamentals is consistent with the levels uncovered in the cross-sectional exercises and expands smoothly over time in a manner fully consistent with our learning-based explanation of BS mispricings. Although other and even more restrictive tests can be designed, we consider these findings *prima facie* evidence of an extraordinary ability of the model in section 4 to adequately fit observed option prices along many different contract dimensions.

8 Conclusion

This paper has proposed a new equilibrium model for asset prices under Bayesian learning and investigated its ability to match a variety of empirical biases in the Black-Scholes option pricing model.

The model with Bayesian learning proved to be able to match both the presence of skews in implied volatilities and the existence of a term-structure in implied volatility for given moneyness. In this respect, the model offers an approach to the correction of BS biases which is competing with the standard benchmarks in the literature, such as jump-diffusion and stochastic volatility models, which were found by Das and Sundaram (1999) to be unable to correctly fit both stylized facts for plausible parameters values.

When the model is fitted to daily prices for a panel of S&P 500 index options, it provides a satisfactory in-sample fit on a daily basis. The estimated parameters characterizing the state of learning each day are stable over time. Episodes of strong revision of the beliefs implied by options prices are rare and tend to involve the precision of these beliefs rather than their level, which we find plausible. The model is then fitted to short periods of time (two weeks) in order to test the intertemporal restrictions which are built-in our maintained learning scheme. Although the in-sample fit performance cannot but worsen, we still find that the estimated parameters are quite stable across periods and that they do not deviate much from the results obtained through the daily inferences.

by the model, although the estimates of the preceding subsection strongly stress that the information flow is likely to be much slower than that.

References

[1] 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.

[2] Aït-Sahalia, Y. and A. Lo, 1998, "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices", *Journal of Finance*, 53, 499-547.

[3] Amin K., 1993, "Jump Diffusion Option Valuation in Discrete Time", *Journal of Finance,* 48, 1833-1863.

[4] Ball C., and W. Torous, 1985, "On Jumps in Common Stock Prices and Their Impact on Call Option Pricing", *Journal of Finance,* 40, 155-173.

[5] Bakshi G., C. Cao and Z. Chen, 1997, "Empirical Performance of Alternative Option Pricing Models", *Journal of Finance*, 52, 2003-2049.

[6] Bates, D., 1991, "The Crash of '87: Was It Expected? The Evidence from Options Markets", *Journal of Finance*, 46, 1009-1044.

[7] Bick A., 1990, "On Viable Diffusion Price Processes of the Market Portfolio", *Journal of Finance,* 45, 673-689.

[8] Black, F. and M., Scholes, 1973, "The pricing of Options and Corporate Liabilities", *Journal of Political Economy,* 81, 637-659.

[9] Campa J., and K. Chang, 1995, "Testing the Expectations Hypothesis on the Term Structure of Volatilities", *Journal of Finance*, 50, 529-547.

[10] Chriss N., 1997, *Black-Scholes and Beyond. Option Pricing Models,* Irwin.

[11] Cox, J., S. Ross and M. Rubinstein, 1979, "Option Pricing : A Simplified Approach", *Journal of Financial Economics,* 7, 229-263.

[12] Das, S., and R. Sundaram, 1999, "Of Smiles and Smirks: A Term Structure Perspective", *Journal of Financial and Quantitative Analysis,* 34, 211-239.

[13] Duan, J.-C., 1995, "The GARCH Option Pricing Model", *Mathematical Finance,* 5, 13-32.

[14] Dumas, B., J. Fleming and R., Whaley, 1998, "Implied Volatility Functions: Empirical Tests", *Journal of Finance,* 53, 2059-2106.

[15] Guidolin M., and A. Timmermann, 1999, "Asset Prices on a Learning Path", *mimeo*, University of California, San Diego.

[16] Harrison, J. and D., Kreps, 1979, "Martingales and Arbitrage in Multiperiod Securities Markets", *Journal of Economic Theory,* 20, 381-408.

[17] He H., and Leland, H., 1993, "On Equilibrium Asset Price Processes", *Review of Financial Studies*, 6, 593-617.

[18] Heston S., 1993, "A Closed Form Solution for Options with Stochastic Volatility with Application to Bond and Currency Options", *Review of Financial Studies*, 6, 327-343.

[19] Hull, J. and A., White, 1987, "The pricing of Options on Assets with Stochastic Volatilities", *Journal of Finance,* 42, 281-300.

[20] Jackwerth, J. and M. Rubinstein, 1996, "Recovering probability Distributions from Option Prices", *Journal of Finance*, 51, 1611-1631.

[21] Leland, H., 1985, "Option Pricing and Replication with Transaction Costs", *Journal of Finance,* 40, 1283-1302.

[22] Lucas, R., 1978, "Asset prices in an Exchange Economy", *Econometrica,* 46, 1429-1445.

[23] Melino A., and S. Turnbull, 1990, "Pricing Foreign Currency Options with Stochastic Volatility", *Journal of Econometrics,* 45, 239-265.

[24] Merton, R., 1987, "Option Pricing When Underlying Stock returns are Discontinuous", *Journal of Financial Economics,* 3, 125-144.

[25] Neftci S., 1996, *Mathematics of Financial Derivatives*, Academic Press.

[26] Pliska S., 1997, *Introduction to Mathematical Finance*, Blackwell Publishers.

[27] Rubinstein, M., 1985, "Nonparametric Tests of Alternative Option Pricing Models Using All Reported Trades and Quotes on the 30 Most Active {CBOE} Option Classes from August 23, 1979 Through August 31, 1978", *Journal of Finance,* 40, 455-480.

[28] Rubinstein, M., 1994, "Implied Binomial Trees", *Journal of Finance,* 49, 771-818.

[29] Schachermayer, W., 1994, "Martingale Measures for Discrete-Time Processes with Infinite Horizon", *Mathematical Finance,* 4, 25-56.

[30] Stapleton, R., and M. Subrahmanyam, 1984, "The Valuation of Options When Asset Returns Are Generated by a Binomial Process", *Journal of Finance,* 39, 1525-1539.

[31] Wiggins, J., 1987, "Option Values Under Stochastic Volatility: Theory and Empirical Estimates", *Journal of Financial Economics,* 19, 351-372.

Appendix A.

Proof of proposition 2. The condition $g_l^* + \pi(g_h^* - g_l^*) < \rho < \pi(1+g_h^*) + (1-\pi)\left[(1+g_h)(1+g_l)^{-\gamma}\right]$ is invoked to guarantee the absence of arbitrage opportunities and hence[31] — with complete markets — the existence of a unique risk-neutral measure. Then a unique linear pricing measure exists and this implies that the law of one price holds, contingent claims with same payoff profile will receive the same price. Start noticing that

$$C_t^{FIRE}(K,T,S_t) = \sum_{j=0}^{v} \beta^v \left(\frac{D_{t+v}^j}{D_t}\right)^{-\gamma} \max\left\{0, S_{t+v}^j - K\right\} \pi^j (1-\pi)^{v-j} \tag{20}$$

is an equivalently good expression for the equilibrium price of a European call. This formula expresses the no-arbitrage price of a contingent claim paying off $\max\left\{0, S_{t+v}^j - K\right\}$ of the consumption good as of time $t+T$. Given that dividends follow a known binomial lattice $\{g_h, g_l, \pi, m\}$, (20) is equivalent to:

$$\begin{aligned}
C_t^{FIRE} &= \sum_{j=0}^{v} \left(\frac{1}{1+\rho}\right)^v \left[\prod_{i=1}^{j}(1+g_h)\right]^{-\gamma} \left[\prod_{i=1}^{v-j}(1+g_l)\right]^{-\gamma} \times \\
&\qquad \times \max\left\{0, S_{t+v}^j - K\right\} \binom{v}{j} \pi^j (1-\pi)^{v-j} \\
&= (1+\rho)^{-v} \sum_{j=0}^{v} \left[(1+g_h)^{-\gamma}\right]^j \left[(1+g_l)^{-\gamma}\right]^{v-j} \max\left\{0, S_{t+v}^j - K\right\} \times \\
&\qquad \times \binom{v}{j} \pi^j (1-\pi)^{v-j} \\
&= (1+r)^{-v} \sum_{j=0}^{v} \max\left\{0, S_{t+v}^j - K\right\} \binom{v}{j} \left[\pi \frac{1+r}{1+\rho}(1+g_h)^{-\gamma}\right]^j \times
\end{aligned}$$

[31] Under certain technical conditions (Schachermayer (1994)). Indeed for infinite horizon models the risk neutral measure may not exist even in the absence of arbitrage opportunities. Pliska (1997, 246-248) provides an example of such a model.

$$\times\left[(1-\pi)\frac{1+r}{1+\rho}(1+g_l)^{-\gamma}\right]^{v-j}$$
$$=(1+r)^{-v}\sum_{j=0}^{v}\max\left\{0,S_{t+v}^{j}-K\right\}\binom{v}{j}z^j(1-z)^{v-j}$$

where $z\equiv\pi\frac{1+r}{1+\rho}(1+g_h)^{-\gamma}$ is a risk-neutral probability measure. Indeed, calling $z^*\equiv(1-\pi)\frac{1+r}{1+\rho}(1+g_l)^{-\gamma}$,

$$z+z^*=\frac{1+r}{1+\rho}\left[\pi(1+g_h)^{-\gamma}+(1-\pi)(1+g_l)^{-\gamma}\right]$$
$$=\frac{1}{1+\rho}\frac{1+\rho}{\left[\pi(1+g_h)^{-\gamma}+(1\text{-}\pi)(1+g_l)^{-\gamma}\right]}\left[\pi(1+g_h)^{-\gamma}+(1\text{-}\pi)(1+g_l)^{-\gamma}\right]=1$$

from the expression for the equilibrium risk-free rate found in proposition 1. As such, $z^*=1-z$. Since z is a positive quantity, z defines a probability measure. We now check that z is a risk-neutral measure. In the affirmative, the discounted return process

$$\frac{r_{t+k}^s-r}{1+r}$$

(r_{t+k}^s is the stock return) is a martingale under the risk-neutral measure (see Pliska (1997, 97-98)):

$$E_t^Z\left[\frac{r_{t+1}^s-r}{1+r}\right]=0\Longleftrightarrow E_t^Z\left[r_{t+1}^s-r\right]=0$$

since r is a deterministic constant under FIRE. Then

$$E_t^Z\left[r_{t+1}^s-r\right]=E_t^Z\left[\frac{S_{t+1}+D_{t+1}}{S_t}-1-r\right]$$
$$=E_t^Z\left[\frac{(1+g_{t+1})\Psi D_t+(1+g_{t+1})D_t}{\Psi D_t}-1-r\right] \tag{21}$$
$$=\frac{1+\Psi}{\Psi}E_t^Z\left[(1+g_{t+1})\right]-(1+r) \tag{22}$$
$$=\frac{1+\Psi}{\Psi}\frac{1+r}{1+\rho}\left[\pi(1+g_h)^{1-\gamma}+(1\text{-}\pi)(1+g_l)^{1-\gamma}\right]-(1+r) \tag{23}$$
$$=\frac{(1+\rho)}{1+g_l^*+\pi(g_h^*\text{-}g_l^*)}\frac{1+r}{1+\rho}\left[1+g_l^*+\pi(g_h^*\text{-}g_l^*)\right]-(1+r) \tag{24}$$
$$=0 \tag{25}$$

where the line before the last derives from the definitions of g_l^* and g_h^*. Equivalently, the measure z induces an expected return from holding the stock index in each period which is exactly equal to the risk-free rate, which is what risk-neutralization implies. ■

Proof of proposition 3. The result sets the parameters $g_h^{(m)}$, $g_l^{(m)}$, $\pi^{(m)}$, $\rho^{(m)}$, and $\gamma^{(m)}$ as a function of the annual values for μ, σ, the risk-free rate r and the dividend yield δ. This is a system of five equations in five unknowns that has a recursive structure that makes its solution particularly easy. In fact $g_h^{(m)}$, $g_l^{(m)}$, and $\pi^{(m)}$ can be directly set μ and σ are fixed. Therefore the restrictions on the process for dividend growth involving $g_h^{(m)}$, $g_l^{(m)}$, and $\pi^{(m)}$ are the same as in Cox et al. (1979, 246-251). As for $\rho^{(m)}$ and $\gamma^{(m)}$ they jointly depend on r and δ. However this sub-system of equations in general admits a solution, i.e. a subjective discount factor and a coefficient of relative risk aversion that can be dynamically adjusted as $m \to \infty$ to hold the period $[t, t+\tau]$ risk-free and dividend rates constant and independent of m:

$$1 + r^{(m)} = \frac{1+\rho^{(m)}}{(1+g_l^{(m)})^{-\gamma}+\pi\left[(1+g_h^{(m)})^{-\gamma} - (1+g_l^{(m)})^{-\gamma}\right]} = (1+r)^{\frac{dt}{v}}$$

Therefore the gross interest rate over $[t, t+\tau]$, including v periods of length $\frac{1}{m}$, is $\left[(1+r)^{\frac{dt}{v}}\right]^v = (1+r)^{dt}$, or $(1+r)$ per year (dt=1). As $m \to \infty$, $\rho^{(m)}$ decreases as both $(1+r)^{\frac{dt}{v}}$ and $\left\{(1+g_l^{(m)})^{-\gamma} + \pi\left[(1+g_h^{(m)})^{-\gamma} - (1+g_l^{(m)})^{-\gamma}\right]\right\} \searrow 0$. Therefore the decrease in $r^{(m)}$ is exactly compensated by the increases of v. As for the dividend yield,

$$1 + \delta^{(m)} = 1 + \frac{1}{\Psi^{(m)}} = \frac{1+\rho^{(m)}}{1+(1+g_l)^{1-\gamma^{(m)}}+\pi^{(m)}\left[(1+g_h^{(m)})^{1-\gamma^{(m)}}-(1+g_l^{(m)})^{1-\gamma^{(m)}}\right]}$$

$$= \frac{(1+r)^{\frac{dt}{v}}\left\{(1+g_l^{(m)})^{-\gamma^{(m)}} + \pi\left[(1+g_h^{(m)})^{-\gamma^{(m)}} - (1+g_l^{(m)})^{-\gamma^{(m)}}\right]\right\}}{(1+g_l)^{1-\gamma^{(m)}} + \pi^{(m)}\left[(1+g_h^{(m)})^{1-\gamma^{(m)}} - (1+g_l^{(m)})^{1-\gamma^{(m)}}\right]} = (1+\delta)^{\frac{dt}{v}}$$

so that the gross dividend yield over $[t, t+\tau]$, is $\left[(1+\delta)^{\frac{dt}{v}}\right]^v = (1+\delta)^{dt}$, or $(1+\delta)$ per year (dt=1). Notice that $\delta^{(m)}$ is adjusted downwards, i.e. $\Psi^{(m)}$ is revised upwards as a function of m, allowing for a fixed current underlying price conditionally on which the option is priced. This completes the proof of the first statement.

As for the second part of the proposition, see Cox et al. (1979, 246-251). Notice that adjusting ρ and γ as a function of m brings us in their exact framework, with a periodal (say, yearly) risk-free rate (adjusted for the dividend payouts) which is independent of the number of nodes v modeled. Just replace $r' = \frac{1+r}{1+\delta}$ for the risk-free rate in their standard formulation and the result follows through the same steps. The result differs from theirs only for the fact that the BS formula is interpreted as an approximation in continuous time of our discrete time framework and it is therefore expressed entirely in terms of the parameters of the discrete time lattice model for dividends. In particular notice that

$$g_h^{(m)} = e^{\sigma\sqrt{dt/v}} - 1 \Longrightarrow \sigma = \sqrt{\frac{v}{dt}} \ln(1 + g_h^{(m)})$$

and $(1+\widehat{r})^{-dt} = (1+r^{(m)})^{-v} \Longrightarrow 1+\widehat{r} = (1+r^{(m)})^{\frac{v}{dt}} = \left[(1+r)^{\frac{dt}{v}}\right]^{\frac{v}{dt}} = 1+r$ so that the $\widehat{r}$ in the BS formula and r correspond. Analogously it can be shown that $(1+\widehat{\delta})^{-dt} = (1+\delta^{(m)})^{-v} \Longrightarrow \widehat{\delta} = \delta$ so that

$$\begin{aligned} d_1 &= \left[\sigma\sqrt{dt}\right]^{-1} \left[\ln\left(\frac{S_t}{K}\right) + \left(\widehat{r} - \widehat{\delta} + \frac{1}{2}\sigma^2\right) dt\right] \\ &= \left[\sqrt{v}\ln(1+g_h)\right]^{-1} \left[\ln\left(\frac{S_t}{K}\right) + (r-\delta)\, dt + \frac{1}{2}\left[\ln(1+g_h)\right]^2 v\right] \end{aligned}$$

as $v = \tau m$. Substituting this expression in the standard BS formula $C_t^{BS} = S_t \Phi(d_1) e^{-\widehat{\delta}\tau} - e^{-\widehat{r}\tau} K \Phi\left(d_1 - \sigma\sqrt{\tau}\right)$ gives the formula. ■

Proof of proposition 5. The equations easily follow from the expression for the no-arbitrage price of a contingent claim paying off $\max\left\{0, S_{t+v}^{BL,j} - K\right\}$ of the consumption good as of time $t + \tau = t + v$ once the relevant probabilities perceived on a learning path are employed.

We check that the probabilities appearing in the expression for the SPD are actually risk-neutralized in all the underlying single period models associated with the information structure. Pliska (1997, 98) shows that in a model with cash dividends, a risk-neutral probability measure must be such that the discounted price process augmented by the discounted sum of dividends paid out is a martingale. From the Euler equations under Bayesian learning, we know that (see Guidolin and Timmermann (1999))

$S_{t+k} = E\{Q_{t+k+1}(D_{t+k+1} + S_{t+k+1})|\hat{\pi}_{t+k}\}$. Dividing through by the price of the one-period zero coupon bond issued at time $t+k$ we have:

$$S_{t+k}\frac{(1+\rho)}{(1+g_l)^{-\gamma} + \widehat{\pi}_t\left[(1+g_h)^{-\gamma} - (1+g_l)^{-\gamma}\right]} =$$

$$= E\left\{Q_{t+k+1}(D_{t+k+1}+S_{t+k+1})\frac{(1+\rho)}{(1+g_l)^{-\gamma}+\widehat{\pi}_{t+k}\left[(1+g_h)^{-\gamma} - (1+g_l)^{-\gamma}\right]}|\hat{\pi}_{t+k}\right\}$$

where we have used the fact that by definition a zero coupon has unit price at expiration. Define now the discounted price and the cumulative dividend process, $S^*_{t+k} = S_{t+k}\frac{(1+\rho)}{(1+g_l)^{-\gamma}+\widehat{\pi}_t[(1+g_h)^{-\gamma}-(1+g_l)^{-\gamma}]}$ and $D^*_{t+k} = \sum_{s=0}^{k} D_{t+s}$ respectively. Adding D^*_{t+k} to both sides we obtain

$$\begin{aligned} S^*_{t+k}+D^*_{t+k} &= E\left[Q_{t+k+1}(S^*_{t+k+1}+D^*_{t+k+1})\times\right. \\ &\quad \left.\times\frac{(1+\rho)}{(1+g_l)^{-\gamma}+\widehat{\pi}_{t+k}\left[(1+g_h)^{-\gamma}\text{-}(1+g_l)^{-\gamma}\right]}|\widehat{\pi}_{t+k}\right\} \end{aligned}$$

showing that the process $S^*_{t+k}+D^*_{t+k}$ is a martingale under the (conditional) probability measure

$$\begin{aligned} &\widehat{P}\left\{S^j_{t+k+1}|\widehat{\pi}_{t+k}\right\} = Q_{t+k+1}\frac{(1+\rho)}{(1+g_l)^{-\gamma}+\widehat{\pi}_{t+k}\left[(1+g_h)^{-\gamma}\text{-}(1+g_l)^{-\gamma}\right]}\times \\ &\times P\left\{D^j_{t+k+1}|\widehat{\pi}_{t+k}\right\} = \beta\left(\frac{D^j_{t+k+1}}{D_{t+k}}\right)^{-\gamma}\left[1+r^{BL}_{t+k}(\widehat{\pi}_{t+k})\right]P\left\{D^j_{t+k+1}|\widehat{\pi}_{t+k}\right\} \end{aligned} \tag{26}$$

(26) represents the one-period risk neutral density. Since the one-period state-price density is simply

$$\begin{aligned} \widetilde{P}\left\{S^j_{t+k+1}|\widehat{\pi}_{t+k}\right\} &= \frac{1}{1+r^{BL}_{t+k}(\widehat{\pi}_{t+k})}\widehat{P}\left\{S^j_{t+k+1}|\widehat{\pi}_{t+k}\right\} \\ &= \beta\left(\frac{D^j_{t+k+1}}{D_{t+k}}\right)^{-\gamma}\left[I_{\{j=1\}}\widehat{\pi}_{t+k} + (1-I_{\{j=1\}})(1-\widehat{\pi}_{t+k})\right] \end{aligned} \tag{27}$$

it is immediate to derive the expression for the one-period ahead state-price measure, once it is recalled the form of (??). Notice that the advantage of working with the SPD is that it does not depend on the risk-free rate.

Well-known results — for instance in Pliska (1997, 96-98) — guarantee that the existence and uniqueness of the risk neutral measure in the underlying single-period models is sufficient for the existence and uniqueness of the risk neutral measure for the infinite horizon model. This risk-neutral measure can be found by 'pasting' together all the paths leading to a certain state of the world $t+v$ periods ahead and exploiting the fact that realizations of the dividend growth rate are i.i.d. notwithstanding the learning process. For instance, $\widetilde{\Pr}\left\{S^j_{t+v}|n_t, N_t\right\}$ corresponds to the product of the state-price densities of type (27) that a high dividend growth occurs j times out v, multiplied by the number of sample paths that can lead to this final outcome, $\binom{v}{j}$:

$$\begin{aligned}
\widetilde{\Pr}_t\left\{S^j_{t+v}\right\} &= \binom{v}{j}\prod_{k=1}^{v}\beta\left(\frac{D^j_{t+k+1}}{D_{t+k}}\right)^{-\gamma}\left[I_{\{j_k=1\}}\widehat{\pi}_{t+k}+(1\text{-}I_{\{j_k=1\}})(1\text{-}\widehat{\pi}_{t+k})\right] \\
&= \binom{v}{j}\beta^v\frac{n_t...(n_t+j-1)(N_t-n_t)...(N_t+v-j-n_t+1)}{N_t(N_t+1)...(N_t+j-1)...(N_t+v)}\times \\
&\qquad \times\prod_{k=1}^{v}\left(\frac{D^j_{t+k+1}}{D_{t+k}}\right)^{-\gamma} \\
&= \binom{v}{j}\beta^v\frac{n_t(n_t+1)...(n_t+j\text{-}1)(N_t\text{-}n_t)...(N_t+v\text{-}j\text{-}n_t+1)}{N_t(N_t+1)...(N_t+j-1)...(N_t+v)} \\
&\quad \times\left[\left(\frac{D^{j_1}_{t+1}}{D_t}\right)^{-\gamma}\left(\frac{D^{j_2}_{t+2}}{D^{j_2}_{t+1}}\right)^{-\gamma}...\left(\frac{D^{j_v}_{t+v}}{D^{j_{v-1}}_{t+v-1}}\right)^{-\gamma}\right] \\
&= \beta^v\left(\frac{D^j_{t+v}}{D_t}\right)^{-\gamma}\binom{v}{j}\frac{\prod_{k=0}^{j-1}(n_t+k)\prod_{k=0}^{v-j-1}(N_t-n_t+k)}{\prod_{k=0}^{v-1}(N_t+k)} \qquad (28)
\end{aligned}$$

as it is easily checked the correspondence between $\frac{\prod_{k=0}^{j-1}(n_t+k)\prod_{k=0}^{v-j-1}(N_t-n_t+k)}{\prod_{k=0}^{v-1}(N_t+k)}$ and $\frac{n_t...(n_t+j-1)(N_t-n_t)...(N_t+v-j-n_t+1)}{N_t(N_t+1)...(N_t+v)}$. Therefore (28) is the desired multi-period risk neutral measure. Notice that as expected from the general result in Guidolin and Timmermann (1999, prop. 1), the time t risk-neutral distribution of the time $t+v$ stock prices depends on the entire sequence of possible future probability beliefs. ■

Lemma 7 *Define* $P^{BL}_{t+k}(s) = P^{BL}_{t+k}\left\{D_{t+k+j} = (1+g_h)^s\ (1+g_l)^{j-s}|n_{t+k}, N_{t+k}\right\}$. *Suppose that* $\frac{1+g^*_h}{1+\rho} \geq 1$ *while* $\frac{1+g^*_l}{1+\rho} < 1$. *Under the same assumptions of* *proposition 4, the time $t+k$ BL pricing kernel*

$$\Psi^{BL}_{t+k}(n_{t+k}, N_{t+k}) = \underset{H\to\infty}{\ell im}\left\{\sum_{j=1}^{H}\beta^j\sum_{s=0}^{j}(1+g_h^*)^s(1+g_l^*)^{j-s}P^{BL}_{t+k}(s)\right\}$$

is an increasing and convex function of $\widehat{\pi}_{t+k} = \frac{n_{t+k}}{N_{t+k}}$ $\forall k$ *when* $\gamma < 1$.

Proof. Consider an alternative estimate $\widehat{\pi}'_{t+k} > \widehat{\pi}_{t+k}$. For $j = 1$ the probability of the two possible final dividend nodes under Bayesian learning are:

$$P^{BL}_{t+k}(s=1|\widehat{\pi}'_{t+k}) = \widehat{\pi}'_{t+k} > \widehat{\pi}_{t+k}$$

$$P^{BL}_{t+k}(s=0|\widehat{\pi}'_{t+k}) = (1-\widehat{\pi}'_{t+k}) < (1-\widehat{\pi}_{t+k})$$

and since for $\gamma < 1$ $(1+g_h^*) > (1+g_l^*)$, the probability mass is shifted from the bad to the good outcome,

$$\beta\sum_{s=0}^{1}(1+g_h^*)^s(1+g_l^*)^{j-s}P^{BL}_{t+k}(s|\widehat{\pi}'_{t+k}) > \beta\sum_{s=0}^{1}(1+g_h^*)^s(1+g_l^*)^{j-s}P^{BL}_{t+k}(s|\widehat{\pi}_{t+k})$$

Therefore the first summand in the definition of Ψ^{BL}_{t+k} is bigger under $\widehat{\pi}'_{t+k} > \widehat{\pi}_{t+k}$. For $j = 2$ and using the fact that the time $t+k$ estimator of π can be re-written in a recursive fashion as

$$\widehat{\pi}_{t+k+1} = \widehat{\pi}_{t+k} + \left[\frac{I_{\{(1+g_{t+k+1})=(1+g_h)\}}}{N_{t+k}+1} - \frac{1}{N_{t+k}+1}\widehat{\pi}_{t+k}\right]$$

we have:

$$P^{BL}_{t+k}(s{=}2|\widehat{\pi}'_{t+k}){=}\widehat{\pi}'_{t+k}\left\{\widehat{\pi}'_{t+k}{+}\frac{1}{N_{t+k}{+}1}(1\text{-}\widehat{\pi}'_{t+k})\right\}{>}\widehat{\pi}_{t+k}\left\{\widehat{\pi}_{t+k}{+}\frac{1}{N_{t+k}{+}1}(1\text{-}\widehat{\pi}_{t+k})\right\}$$

$$P^{BL}_{t+k}(s=1|\widehat{\pi}'_{t+k}){=}2\widehat{\pi}'_{t+k}\left\{1\text{-}\frac{N_{t+k}}{N_{t+k}{+}1}\widehat{\pi}'_{t+k}\right\}\underset{<}{\geq}2\widehat{\pi}_{t+k}\left\{1\text{-}\frac{N_{t+k}}{N_{t+k}{+}1}\widehat{\pi}_{t+k}\right\}$$

$$P^{BL}_{t+k}(s{=}0|\widehat{\pi}'_{t+k}) = (1-\widehat{\pi}'_{t+k})\left\{1\text{-}\frac{N_{t+k}}{N_{t+k}{+}1}\widehat{\pi}'_{t+k}\right\} < (1-\widehat{\pi}_{t+k})\left\{1\text{-}\frac{N_{t+k}}{N_{t+k}{+}1}\widehat{\pi}_{t+k}\right\}$$

Again, probability mass is shifted from the least favorable ($s = 0$) to the best outcome ($s = 2$), so that

$$\beta^2 \sum_{s=0}^{2}(1+g_h^*)^s(1+g_l^*)^{j-s}P^{BL}_{t+k}(s|\widehat{\pi}'_{t+k}) > \beta^2 \sum_{s=0}^{2}(1+g_h^*)^s(1+g_l^*)^{j-s}P^{BL}_{t+k}(s|\widehat{\pi}_{t+k})$$

Repeating this argument as $j \to \infty$ one can show that, provided $\gamma < 1$, under Bayesian learning the time $t+k$ pricing kernel increases with the current estimate $\widehat{\pi}_{t+k}$. As for convexity, observe that as $\widehat{\pi}_{t+k} \searrow 0$,

$$\Psi^{BL}_{t+k}(\widehat{\pi}_{t+k}) \searrow \lim_{H\to\infty}\left\{\sum_{j=1}^{H}\beta^j(1+g_l^*)^j\right\} = \frac{\beta(1+g_l^*)}{1-\beta(1+g_l^*)} > 0$$

under the hypothesis that $\frac{1+g_l^*}{1+\rho} = \beta(1+g_l^*) < 1$, while as $\widehat{\pi}_{t+k} \nearrow 1$

$$\Psi^{BL}_{t+k}(\widehat{\pi}_{t+k}) \nearrow \lim_{H\to\infty}\left\{\sum_{j=1}^{H}\beta^j(1+g_h^*)^j\right\} \to \infty$$

if $\frac{1+g_h^*}{1+\rho} = \beta(1+g_h^*) \geq 1$. For instance, for all $\rho \geq 0$ a $g_l < 0$ will be sufficient for the condition $\beta(1+g_l^*) < 1$ to hold. Now, $\Psi^{BL}_{t+k}(\widehat{\pi}_{t+k})$ is defined on the compact $[0,1]$, it is a monotonically increasing and continuous function, it is everywhere differentiable (it is just a sum of differentiable functions of $\widehat{\pi}$, the BL 'objective' probabilities), it goes to a finite positive quantity as $\widehat{\pi}_{t+k} \searrow 0$ and diverges as $\widehat{\pi}_{t+k} \nearrow 1$. Then $\Psi^{BL}_{t+k}(\widehat{\pi}_{t+k})$ must be convex. ■

Proof of proposition 6. To simplify the notation, define $\nabla^{BL}_{FIRE}(K) \equiv C^{BL}_{t+k}(K) - C^{FIRE}_{t+k}(K)$ and $P^{BL}_{t+k}(s) = P^{BL}_{t+k}\left\{D^s_{t+k+v}|n_{t+k}, N_{t+k}\right\}$. We start by observing that conditioning on a common, current stock price S_{t+k} has implications for the array of possible future stock prices in $t+k+v$. In fact $S^{BL}_{t+k+v}(0) = \Psi^{BL}_{t+k+v}(n_{t+k}, N_{t+k}+v)(1+g_l)^v D_t = \frac{\Psi^{BL}_{t+k+v}(n_{t+k},N_{t+k}+v)}{\Psi^{BL}_{t+k}(n_{t+k},N_{t+k})}(1+g_l)^v S_{t+k} << S^{FIRE}_{t+k+v}(0) = \Psi^{FIRE}(1+g_l)^v D_t = (1+g_l)^v S_{t+k}$ holds if and only if $\frac{\Psi^{BL}_{t+k+v}(n_{t+k},N_{t+k}+v)}{\Psi^{BL}_{t+k}(n_{t+k},N_{t+k})} < 1$ as $S^{FIRE}_{t+k} = S^{BL}_{t+k}$ by assumption. By the previous lemma, this is guaranteed by the fact that $\widehat{\pi}_{t+k+v}(0)=\frac{n_{t+k}}{N_{t+k}+v}<\widehat{\pi}_{t+k}=\frac{n_{t+k}}{N_{t+k}}$ so that $\Psi^{BL}_{t+k+v}(n_{t+k},N_{t+k}+v)<\Psi^{BL}_{t+k}(n_{t+k},N_{t+k})$. Similarly, $S^{BL}_{t+k+v}(v)=\Psi^{BL}_{t+k+v}(n_{t+k}+v,N_{t+k}+v)(1+g_h)^v D_t=\frac{\Psi^{BL}_{t+k+v}(n_{t+k}+v,N_{t+k}+v)}{\Psi^{BL}_{t+k}(n_{t+k},N_{t+k})}(1+g_h)^v S_{t+k}>S^{FIRE}_{t+k+v}(v)=\Psi^{FIRE}(1+g_h)^v D_t=(1+g_h)^v S_{t+k}$ holds if and only if $\frac{\Psi^{BL}_{t+k+v}(n_{t+k}+v,N_{t+k}+v)}{\Psi^{BL}_{t+k}(n_{t+k},N_{t+k})}>1$, which follows from the lemma. Therefore on a Bayesian learning path a few time $t+k+v$ stock prices become possible that could not be reached under FIRE, both below the minimum attainable FIRE stock price and above the maximum time $t+k+v$ FIRE price.

We study $\nabla^{BL}_{FIRE}(K)$ in five disjoint intervals for the strike price: $[0,S^{BL}_{t+k+v}(0))$, $[S^{BL}_{t+k+v}(0),S^{FIRE}_{t+k+v}(0))$, $[S^{FIRE}_{t+k+v}(0),S^{FIRE}_{t+k+v}(v))$, $[S^{FIRE}_{t+k+v}(v),S^{BL}_{t+k+v}(v))$, and $[S^{BL}_{t+k+v}(v),+\infty)$. As for the first interval, observe that when $K=0$

$$\begin{aligned}\nabla^{BL}_{FIRE}(0) &= \sum_{s=0}^{v}\beta^v\left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{-\gamma}S^{BL}_{t+k+v}(s)P^{BL}_{t+k}(s)+\\ &\qquad -\sum_{s=0}^{v}\beta^v\left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{-\gamma}S^{FIRE}_{t+k+v}(s)\binom{v}{s}\pi^s(1\text{-}\pi)^{v-s}\\ &= S_{t+k}\sum_{s=0}^{v}\beta^v\left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{1-\gamma}\Psi^{BL}_{t+k+v}(n_{t+k}+s,N_{t+k}+v)P^{BL}_{t+k}(s)+\\ &\qquad -S_{t+k}\sum_{s=0}^{v}\beta^v\left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{1-\gamma}\binom{v}{s}\pi^s(1\text{-}\pi)^{v-s}\end{aligned}$$

$$\begin{aligned}&= S_{t+k}\beta^v\widehat{E}_{t+k}\left\{\widehat{E}_{t+k+1}\left[...\widehat{E}_{t+k+v-1}\left(\Psi^{BL}_{t+k+v}\left(\frac{D_{t+k+v}}{D_{t+k}}\right)^{1-\gamma}\right)...\right]\right\}\\ &-S_{t+k}\beta^v E_{t+k}\left[\left(\frac{D_{t+k+v}}{D_{t+k}}\right)^{1-\gamma}\right]>0 \qquad (29)\end{aligned}$$

as $\Psi^{BL}_{t+k+v}\left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{1-\gamma}$ is an non-decreasing, convex function of $\widehat{\pi}_{t+k+v}$ (equivalently, of s, the number of high growth states between $t+k$ and $t+k+v$, once $\widehat{\pi}_{t+k}$ is given). Indeed both Ψ^{BL}_{t+k+v} and $\left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{1-\gamma}$ are positive, non-decreasing and convex functions of $\widehat{\pi}_{t+k+v}$, and as such it is straightforward to check that their product possesses the same properties. Then, for $\widehat{\pi}_{t+k}=\pi$, Jensen's inequality implies that

$$\Psi_{t+k}^{BL}\widehat{E}_{t+k}\left\{\frac{\Psi_{t+k+1}^{BL}}{\Psi_{t+k}^{BL}}\left(\frac{D_{t+k+1}}{D_t}\right)^{1-\gamma}\widehat{E}_{t+k+1}\left[\frac{\Psi_{t+k+2}^{BL}}{\Psi_{t+k+1}^{BL}}\left(\frac{D_{t+k+2}}{D_{t+k+1}}\right)^{1-\gamma}\dots\right.\right.$$

$$\left.\left.\dots\widehat{E}_{t+k+v-1}\left(\frac{\Psi_{t+k+v}^{BL}}{\Psi_{t+k+v-1}^{BL}}\left(\frac{D_{t+k+v}}{D_{t+k+v-1}}\right)^{1-\gamma}\mid\pi\right)\dots|\pi\right]|\pi\right\}$$

$$\geq\Psi_{t+k}^{BL}E_{t+k}\left[\frac{\Psi_{t+k+1}^{BL}}{\Psi_{t+k}^{BL}}\left(\frac{D_{t+k+1}}{D_{t+k}}\right)^{1-\gamma}\frac{\Psi_{t+k+2}^{BL}}{\Psi_{t+k+1}^{BL}}\left(\frac{D_{t+k+2}}{D_{t+k+1}}\right)^{1-\gamma}\dots\right.$$

$$\left.\dots\frac{\Psi_{t+k+v}^{BL}}{\Psi_{t+k+v-1}^{BL}}\left(\frac{D_{t+k+v}}{D_{t+k+v-1}}\right)^{1-\gamma}\right]$$

$$=\Psi_{t+k}^{BL}E_{t+k}\left[\frac{\Psi_{t+k+v}^{BL}}{\Psi_{t+k}^{BL}}\left(\frac{D_{t+k+v}}{D_{t+k}}\right)^{1-\gamma}\right]=\Psi_{t+k}^{BL}E_{t+k}\left[\left(\frac{D_{t+k+v}}{D_{t+k}}\right)^{1-\gamma}\right]\quad(30)$$

since under knowledge of π, $\Psi_{t+k+v}^{BL}=\Psi_{t+k}^{BL}$. For $\widehat{\pi}_{t+k}>\pi$, $C_{t+k}^{BL}(0)-C_{t+k}^{FIRE}(0)\geq 0$ holds as the actual

$$\widehat{E}_{t+k}\left\{\widehat{E}_{t+k+1}\left[\dots\widehat{E}_{t+k+v-1}\left(\Psi_{t+k+v}^{BL}\left(\frac{D_{t+k+v}}{D_{t+k}}\right)^{1-\gamma}|\widehat{\pi}_{t+k}\right)\dots|\widehat{\pi}_{t+k}\right]|\widehat{\pi}_{t+k}\right\}$$

is always bigger than the first line of (30) since for $\widehat{\pi}_{t+k}>\pi$ probability beliefs under Bayesian learning are no longer a pure mean-preserving spread of RE beliefs (as implied by Guidolin and Timmermann (1999, prop. 4)), but for $\gamma<1$ probability mass is transferred from bad to good states of the world, that is, states in which $\left(\frac{D_{t+k+v}^s}{D_{t+k}}\right)^{1-\gamma}\Psi_{t+k+v}^{BL}$ takes higher values.

Consider any $K\in(0,S_{t+k+v}^{BL}(0))$. Then:

$$\nabla_{FIRE}^{BL}(K)=\sum_{s=0}^{v}\beta^v\left(\frac{D_{t+k+v}^s}{D_{t+k}}\right)^{-\gamma}S_{t+k+v}^{BL}(n_{t+k}+s,N_{t+k}+v)P_{t+k}^{BL}(s)+$$

$$-\sum_{s=0}^{v}\beta^v\left(\frac{D_{t+k+v}^s}{D_{t+k}}\right)^{-\gamma}S_{t+k+v}^{FIRE}(s)\binom{v}{s}\pi^s(1\text{-}\pi)^{v-s}+$$

$$+ K\beta^v \sum_{s=0}^{v} \left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{-\gamma} \left[\binom{v}{s}\pi^s(1-\pi)^{v-s} - P^{BL}_{t+k}(s)\right]$$

$$> K\beta^v \sum_{s=0}^{v} \left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{-\gamma} \left[\binom{v}{s}\pi^s(1-\pi)^{v-s} - P^{BL}_{t+k}(s)\right] > 0$$

where we have recognized that the first two terms of $\nabla^{BL}_{FIRE}(K)$ correspond to $\nabla^{BL}_{FIRE}(0) > 0$. The last inequality obtains if and only if $\widehat{\pi}_{t+k} > \pi$, as assumed. Notice that for $\widehat{\pi}_{t+k} = \pi$ a Jensen's inequality argument would make us conclude that $K\beta^v \sum_{s=0}^{v} \left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{-\gamma} \left[\binom{v}{s}\pi^s(1-\pi)^{v-s} - P^{BL}_{t+k}(s)\right]$ is negative instead. If $\widehat{\pi}_{t+k} > \pi$, since $\left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{-\gamma}$ is now a decreasing function of the number of high dividend growth states, under BL probability mass is shifted from the states in which $\left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{-\gamma}$ is high to states in which it is lower, and as such the positive sign results. Formally,

$$K\beta^v \sum_{s=0}^{v} \left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{-\gamma} \left[\binom{v}{s}\pi^s(1-\pi)^{v-s} - P^{BL}_{t+k}\left\{D^s_{t+k+v} | \widehat{\pi}_{t+k} > \pi\right\}\right]$$

$$> K\beta^v \sum_{s=0}^{v} \left[\binom{v}{s}\pi^s(1-\pi)^{v-s} - P^{BL}_{t+k}\left\{D^s_{t+k+v} | \widehat{\pi}_{t+k} = \pi\right\}\right] = 0$$

as both sequences in square brackets (as a function of s) define a discrete probability measure. Therefore for $K \in (0, S^{BL}_{t+k+v}(0))$ $\nabla^{BL}_{FIRE}(K) > 0$ while our argument also implies that $\nabla^{BL}_{FIRE}(K) - \nabla^{BL}_{FIRE}(0) > 0$, and that $\nabla^{BL}_{FIRE}(K)$ increases in the interval $(0, S^{BL}_{t+k+v}(0))$.

It is convenient to proceed backwards for the rest of the proof. Suppose $K \geq S^{BL}_{t+k+v}(v)$. Since $\max\{0, S^{BL}_{t+k+v}(v) - K\} = 0 \Longrightarrow \max\{0, S^{FIRE}_{t+k+v}(v) - K\} = 0$ as $S^{BL}_{t+k+v}(v) > S^{FIRE}_{t+k+v}(v)$, the two call prices are both trivially nil and their difference is therefore nonnegative. Hence $\nabla^{BL}_{FIRE}(K) = 0$ in the $[S^{BL}_{t+k+v}(v), +\infty)$ interval.

Consider now the interval $K \in [S^{FIRE}_{t+k+v}(v), S^{BL}_{t+k+v}(v))$. Since $\max\{0, S^{FIRE}_{t+k+v}(v) - K\} = 0$ while $\max\{0, S^{BL}_{t+k+v}(v) - K\} \geq 0\, \forall K$, it follows that while $C^{FIRE}_{t+k}(K) = 0$ everywhere, $C^{BL}_{t+k}(K) > 0$. Hence $\nabla^{BL}_{FIRE}(K) > 0$. Also, $\nabla^{BL}_{FIRE}(K)$ is decreasing in $K \in [S^{FIRE}_{t+k+v}(v), S^{BL}_{t+k+v}(v))$ as $C^{BL}_{t+k}(K)$ is obviously decreasing in the strike price.

We can gain some insight by splitting up the third interval for the strike price, $[S^{FIRE}_{t+k+v}(0), S^{FIRE}_{t+k+v}(v))$, into two sub-intervals. Define $a^{FIRE}(K)$ as the smaller natural number s.t. $(1+g_h)^a(1+g_l)^{v-a}S_t > K$ and $a^{BL}(K)$ as the smaller natural number s.t. $\frac{\Psi^{BL}_{t+k+v}(n_{t+k}+a, N_{t+k}+v)}{\Psi^{BL}_{t+k}(n_{t+k}, N_{t+k})}(1+g_h)^a(1+g_l)^{v-a}S_t > K$. Observe that as for $s' = int(v\widehat{\pi}_{t+k}) + I_{\{v\widehat{\pi}_{t+k}-int(v\widehat{\pi}_{t+k})>0\}} \geq v\widehat{\pi}_{t+k}$ $\widehat{\pi}_{t+k+v} = \frac{N_{t+k}}{N_{t+k}+v}\widehat{\pi}_{t+k} + \frac{1}{N_{t+k}+v}s \geq \frac{N_{t+k}}{N_{t+k}+v}\widehat{\pi}_{t+k} + \frac{1}{N_{t+k}+v}v\widehat{\pi}_{t+k} = \widehat{\pi}_{t+k}$, the lemma implies that $\frac{\Psi^{BL}_{t+k+v}}{\Psi^{BL}_{t+k}} \geq 1\ \forall s \geq s'$. Therefore, $a^{BL}(K) \geq int(v\widehat{\pi}_{t+k}) + I_{\{v\widehat{\pi}_{t+k}-int(v\widehat{\pi}_{t+k})>0\}}$ implies that $a^{FIRE}(K) \geq a^{BL}(K)$, while for $a^{BL}(K) < int(v\widehat{\pi}_{t+k}) + I_{\{v\widehat{\pi}_{t+k}-int(v\widehat{\pi}_{t+k})>0\}}$ $a^{FIRE}(K) < a^{BL}(K)$ as $\frac{\Psi^{BL}_{t+k+v}}{\Psi^{BL}_{t+k}} < 1$. In other words, $\exists\overline{K} \in [S^{FIRE}_{t+k+v}(0), S^{FIRE}_{t+k+v}(v))$ s.t. $a^{BL}(\overline{K}) \geq int(v\widehat{\pi}_{t+k}) + I_{\{v\widehat{\pi}_{t+k}-int(v\widehat{\pi}_{t+k})>0\}}$ so that $\forall K \geq \overline{K}$ $a^{FIRE}(K) \geq a^{BL}(K)$ and $\forall s \geq a^{BL}(K)$ $S^{BL}_{t+k+v}(s) \geq S^{FIRE}_{t+k+v}(s)$, while $\forall K < \overline{K}$ $a^{FIRE}(K) < a^{BL}(K)$. This is consistent with the fact that when $K < S^{FIRE}_{t+k+v}(0)$, $a^{BL}(K) > a^{FIRE}(K) = 0$, while for $K > S^{FIRE}_{t+k+v}(v)$ $v = a^{FIRE}(K) \geq a^{BL}(K)$. Suppose first that $K \in [\overline{K}, S^{FIRE}_{t+k+v}(v))$. Then

$$
\begin{aligned}
\nabla^{BL}_{FIRE}(K) = & \sum_{s=a^{BL}(K)}^{v} \beta^v \left(\frac{D^s_{t+k+v}}{D_t}\right)^{-\gamma} [S^{BL}_{t+k+v} - K] P^{BL}_{t+k}(s) + \\
& - \sum_{s=a^{FIRE}(K)}^{v} \beta^v \left(\frac{D^s_{t+k+v}}{D_t}\right)^{-\gamma} [S^{FIRE}_{t+k+v} - K] \binom{v}{s} \pi^s (1-\pi)^{v-s} \\
= & \sum_{s=a^{BL}(K)}^{v} \beta^v \left(\frac{D^s_{t+k+v}}{D_t}\right)^{-\gamma} [S^{BL}_{t+k+v} - K] P^{BL}_{t+k}(s) + \\
& - \sum_{s=a^{BL}(K)}^{v} \beta^v \left(\frac{D^s_{t+k+v}}{D_t}\right)^{-\gamma} \max\left\{0, S^{FIRE}_{t+k+v} - K\right\} \binom{v}{s} \pi^s (1-\pi)^{v-s} > 0 \qquad \text{(31a)}
\end{aligned}
$$

as $[S^{BL}_{t+k+v} - K] > \max\{0, S^{FIRE}_{t+k+v} - K\}$ $\forall s \geq a^{BL}(K)$. As for the difference between probability beliefs on a learning path and under RE $f(s) = P^{BL}_{t+k}\left\{D^s_{t+k+v} | n_{t+k}, N_{t+k}\right\} - \binom{v}{s}\pi^s(1-\pi)^{v-s}$, Guidolin and Timmermann (1999, prop. 4) shows that for $\widehat{\pi}_{t+k} = \pi$ probability beliefs under learning are a mean-preserving spread of RE beliefs and that there exists a range of values for s such that

$$\pi^s(1-\pi)^{v-s} = \left(\frac{n_{t+k}}{N_{t+k}}\right)^s \left(\frac{N_{t+k}-n_{t+k}}{N_{t+k}}\right)^{v-s}$$
$$= \frac{\prod_{l=0}^{s-1}(n_{t+k}) \prod_{l=0}^{v-s-1}(N_{t+k}\text{-}n_{t+k})}{\prod_{l=0}^{v-1}(N_{t+k})} < \frac{\prod_{l=0}^{s-1}(n_{t+k}+l) \prod_{l=0}^{v-s-1}(N_{t+k}\text{-}n_{t+k}+l)}{\prod_{l=0}^{v-1}(N_{t+k}+l)}$$

Dividing through by the expression on the left-hand-side and taking logs, we obtain for the last inequality:

$$0 < \sum_{l=0}^{s-1} \ell n\left(1+\frac{l}{n_t}\right) + \sum_{l=0}^{v-s-1} \ell n\left(1+\frac{l}{N_t-n_t}\right) - \sum_{l=0}^{v-1} \ell n\left(1+\frac{l}{N_t}\right) \quad (32)$$

which can indeed be verified to hold when s is either very large ('close' to v) or very low. Differentiating (32) with respect to s it follows that the expression increases when $s > (v-1)\widehat{\pi}_{t+k}$ and decreases for $s < (v-1)\widehat{\pi}_{t+k}$. Therefore $\exists$ two values of s, $s_1 \geq 0$ and $s_2 \leq v$ such that $\forall s \in [0, s_1] \cup [s_2, v]$ $P^{BL}_{t+k}(s)$ is greater than $\binom{v}{s}\pi^s(1-\pi)^{v-s}$. Though the mean-preserving spread result implies that $\sum_{s=0}^{v} f(s) = 0$, for $\pi \geq \frac{1}{2}$ the function $f(s)$ contains a nonnegative weighting mass in the interval $[v\widehat{\pi}_{t+k}, v]$ in the sense that the positive values of $f(s)$ over $[s_2, v]$ at least compensate the negative values between $[v\widehat{\pi}_{t+k}, s_2)$ For $\widehat{\pi}_{t+k} > \pi$, the two values of s such that $f(s) > 0$ move 'right', to $s'_1 > s_1$ and $s'_2 > s_2$ as increasing probability mass is shifted from states with low final dividends to states with high final dividends. For the same reason, for $\pi \geq \frac{1}{2}$ $f(s)$ contains an even larger (positive) weighting mass in the interval $[v\widehat{\pi}_{t+k}, v]$. Since the summation in (31a) runs from $a^{BL}(K)$ to v, for K s.t. $a^{BL}(K) \geq s'_2 > v\widehat{\pi}_{t+k}$, $P^{BL}_{t+k}(s) > \binom{v}{s}\pi^s(1-\pi)^{v-s}$ $\forall s \geq a^{BL}(K)$, the result obtains. Furthermore, $\nabla^{BL}_{FIRE}(K)$ is decreasing in K. If K is s.t. $a^{BL}(K) \in [a^{BL}(\overline{K}), s'_2)$ the result follows from the fact that the wedge between $[S^{BL}_{t+k+v} - K]$ and $\max\{0, S^{FIRE}_{t+k+v} - K\}$ increases with s, while in $[v\widehat{\pi}_{t+k}, v]$ $f(s)$ contains a positive weighting mass. In this interval, $\nabla^{BL}_{FIRE}(K)$ can be either increasing or decreasing in K. When K increases, states of the world in which $P^{BL}_{t+k}(s) < \binom{v}{s}\pi^s(1-\pi)^{v-s}$ are 'discarded'. If for these states of the nature it was $[S^{BL}_{t+k+v} - X']P^{BL}_{t+k}(s) < \max\{0, S^{FIRE}_{t+k+v} - K\}\binom{v}{s}\pi^s(1-\pi)^{v-s}$ then the difference in option prices increases, otherwise it decreases.

Suppose instead that $K \in [S^{FIRE}_{t+k+v}(0), \overline{K})$. This implies $a^{FIRE}(K) < a^{BL}(K)$ while $\forall s < a^{BL}(\overline{K})$ $S^{FIRE}_{t+k+v}(s) > S^{BL}_{t+k+v}(s)$. By the same argument made above, $C^{BL}_{t+k}(K) - C^{FIRE}_{t+k}(K)$ can be either increasing or decreasing in K over this interval. First, note that $a^{BL}(S^{FIRE}_{t+k+v}(0))$ is s.t. $S^{BL}_{t+k+v}(a^{BL}(S^{FIRE}_{t+k+v}(0))) \geq S^{FIRE}_{t+k+v}(0)$, while $\forall s < a^{BL}(S^{FIRE}_{t+k+v}(0))\ S^{BL}_{t+k+v}(s) < S^{FIRE}_{t+k+v}(0)$. Given this, we check that $\nabla^{BL}_{FIRE}(S^{FIRE}_{t+k+v}(0))$ is nonnegative:

$$
\begin{aligned}
\nabla^{BL}_{FIRE}(S^{FRE}_{t+k+v}(0)) = {} & \nabla^{BL}_{FIRE}(S^{BL}_{t+k+v}(0)) + \\
& - \sum_{s=0}^{a^{BL}(S^{FIRE}_{t+k+v}(0))-1} \beta^v \left(\frac{D^s_{t+k+v}}{D_t}\right)^{-\gamma} S^{BL}_{t+k+v} P^{BL}_{t+k}(s) + \\
& + S^{BL}_{t+k+v}(0) \sum_{s=0}^{v} \beta^v \left(\frac{D^s_{t+k+v}}{D_t}\right)^{-\gamma} P^{BL}_{t+k}(s) +
\end{aligned}
$$

$$
\begin{aligned}
& - S^{FIRE}_{t+k+v}(0) \sum_{s=a^{BL}(S^{FIRE}_{t+k+v}(0))}^{v} \beta^v \left(\frac{D^s_{t+k+v}}{D_t}\right)^{-\gamma} P^{BL}_{t+k}(s) + \\
& - S^{BL}_{t+k+v}(0) \sum_{s=0}^{v} \beta^v \left(\frac{D^s_{t+k+v}}{D_t}\right)^{-\gamma} \binom{v}{s} \pi^s (1-\pi)^{v-s} + \\
& + S^{FIRE}_{t+k+v}(0) \sum_{s=0}^{v} \beta^v \left(\frac{D^s_{t+k+v}}{D_t}\right)^{-\gamma} \binom{v}{s} \pi^s (1-\pi)^{v-s}
\end{aligned}
$$

$$
> [S^{FIRE}_{t+k+v}(0) \text{-} S^{BL}_{t+k+v}(0)] \sum_{s=0}^{v} \beta^v \left(\frac{D^s_{t+k+v}}{D_t}\right)^{-\gamma} \left[\binom{v}{s} \pi^s (1\text{-}\pi)^{v-s} \text{-} P^{BL}_{t+k}(s)\right] + \tag{33}
$$

$$
+ \sum_{s=0}^{a^{BL}(S^{FIRE}_{t+k+v}(0))-1} \beta^v \left(\frac{D^s_{t+k+v}}{D_t}\right)^{-\gamma} [S^{FIRE}_{t+k+v}(0) - S^{BL}_{t+k+v}] P^{BL}_{t+k}(s) > 0 \tag{34}
$$

where the first inequality follows from the fact $\nabla^{BL}_{FIRE}(0) > 0$ and the second from the $[S^{FIRE}_{t+k+v}(0) - S^{BL}_{t+k+v}] > 0$ for every term in the summation, plus $\sum_{s=0}^{v} \left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{-\gamma} \left[\binom{v}{s}\pi^s(1-\pi)^{v-s} - P^{BL}_{t+k}\left\{D^s_{t+k+v}|n_{t+k}, N_{t+k}\right\}\right] > 0$ as already recognized. Consider now a strike $K = S^{FIRE}_{t+k+v}(0) + \delta$, $\delta > 0$ such that

$a^{BL}(S^{FIRE}_{t+k+v}(0)+\delta) = a^{BL}(S^{FIRE}_{t+k+v}(0))$. Then

$$\begin{aligned}
\nabla^{BL}_{FIRE}(K) = & \sum_{s=a^{BL}(K)}^{v} \beta^v \left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{-\gamma} [S^{BL}_{t+k+v} - K] P^{BL}_{t+k}(s) + \\
& - \sum_{s=1}^{v} \beta^v \left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{-\gamma} [S^{FIRE}_{t+k+v} - K] \binom{v}{s} \pi^s (1\text{-}\pi)^{v-s} \\
& = \nabla^{BL}_{FIRE}(S^{FIRE}_{t+k+v}(0)) - \delta \sum_{s=a^{BL}(K)}^{v} \beta^v \left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{-\gamma} P^{BL}_{t+k}(s) +
\end{aligned}$$

$$\begin{aligned}
& + \delta \sum_{s=1}^{v} \beta^v \left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{-\gamma} \binom{v}{s} \pi^s (1\text{-}\pi)^{v-s} \\
& = \nabla^{BL}_{FIRE}(S^{FIRE}_{t+k+v}(0)) + \delta \sum_{s=1}^{a^{BL}(K)-1} \beta^v \left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{-\gamma} \binom{v}{s} \pi^s (1\text{-}\pi)^{v-s} \\
& \quad + \delta \sum_{s=a^{BL}(K)}^{v} \beta^v \left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{-\gamma} \left\{ \binom{v}{s} \pi^s (1\text{-}\pi)^{v-s} - P^{BL}_{t+k}(s) \right\}
\end{aligned}$$

which is positive and increasing in K since

$$\sum_{s=a^{BL}(K)}^{v} \left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{-\gamma} \left\{ \binom{2}{s} \pi^s (1\text{-}\pi)^{v-s} - P^{BL}_{t}(s) \right\} > 0$$

from the fact that with $a^{BL}(K) < v\widehat{\pi}_{t+k} - \sum_{s=a^{BL}(K)}^{v} f(s)$ contains a non-negative probability mass. Repeating the same argument for all the subintervals of $[S^{FIRE}_{t+k+v}(0), \overline{K})$ formed by increasing the strike without letting $a^{BL}(K)$ to change and exploiting the continuity in K of $\nabla^{BL}_{FIRE}(K)$, we conclude that this function is everywhere positive and increasing in $[S^{FIRE}_{t+k+v}(0), \overline{K})$. The differential in call prices is at first increasing, though $\exists K^{MAX}$ s.t. for strike prices higher than K^{MAX} the wedge starts decreasing towards zero, level reached for $K = S^{BL}_{t+k+v}(v)$. Furthermore, by the argument just displayed we know that K^{MAX} must belong to $[\overline{K}, S^{FIRE}_{t+k+v}(v))$.

Finally, we are left with the interval $[S^{BL}_{t+k+v}(0), S^{FIRE}_{t+k+v}(0))$. However, $\nabla^{BL}_{FIRE}(K)$ is a continuous function and we have already proven that $\nabla^{BL}_{FIRE}(S^{BL}_{t+k+v}(0)) > 0$. Take now any $K' = S^{BL}_{t+k+v}(0) + \delta$, where $\delta > 0$ is a small number such that $a^{BL}(S^{BL}_{t+k+v}(0)) = a^{BL}(S^{BL}_{t+k+v}(0) + \delta) = 1$. When $v \to \infty$, δ must be picked very small. Notice that $\nabla^{BL}_{FIRE}(K')$ can be decomposed into:

$$\nabla^{BL}_{FIRE}(K') = \nabla^{BL}_{FIRE}(S^{BL}_{t+k+v}(0)) + \\ - \sum_{s=1}^{v} \beta^2 \left(\frac{D^s_{t+k+v}}{D_{t+k}}\right)^{-\gamma} \delta \left\{ P^{BL}_{t+k}(s) - \binom{v}{s} \pi^s (1-\pi)^{v-s} \right\} + \\ + \beta^2 \left(\frac{D^{(1)}_{t+k+v}}{D_{t+k}}\right)^{-\gamma} \delta (1-\pi)^v$$

implying that $\nabla^{BL}_{FIRE}(K') > \nabla^{BL}_{FIRE}(S^{BL}_{t+k+v}(0)) > 0$. Therefore $\nabla^{BL}_{FIRE}(K)$ is increasing and positive in $[S^{BL}_{t+k+v}(0), S^{BL}_{t+k+v}(0) + \delta]$. We can repeat the same argument for all the subintervals of $[S^{BL}_{t+k+v}(0), S^{FIRE}_{t+k+v}(0))$ in which although the strike price increases, $a^{BL}(K)$ does not change. Since $\nabla^{BL}_{FIRE}(K)$ is a continuous function this implies that it is everywhere increasing over $[S^{BL}_{t+k+v}(0), S^{FIRE}_{t+k+v}(0))$. Since $\nabla^{BL}_{FIRE}(S^{BL}_{t+k+v}(0)) > 0$ positivity everywhere follows. Summarizing, $\nabla^{BL}_{FIRE}(K)$ starts positive at $K = 0$, increases over $(0, K^{MAX})$ and then decreases between K^{MAX} and $S^{BL}_{t+k+v}(v)$, where it is zero, with $K^{MAX} \in [S^{FIRE}_{t+k+v}(0), S^{FIRE}_{t+k+v}(v))$. ■

Appendix B.

Figure 1. Implied volatility as a function of moneyness for S&P 500 index options with maturity over the period February 1993 - January 1994.

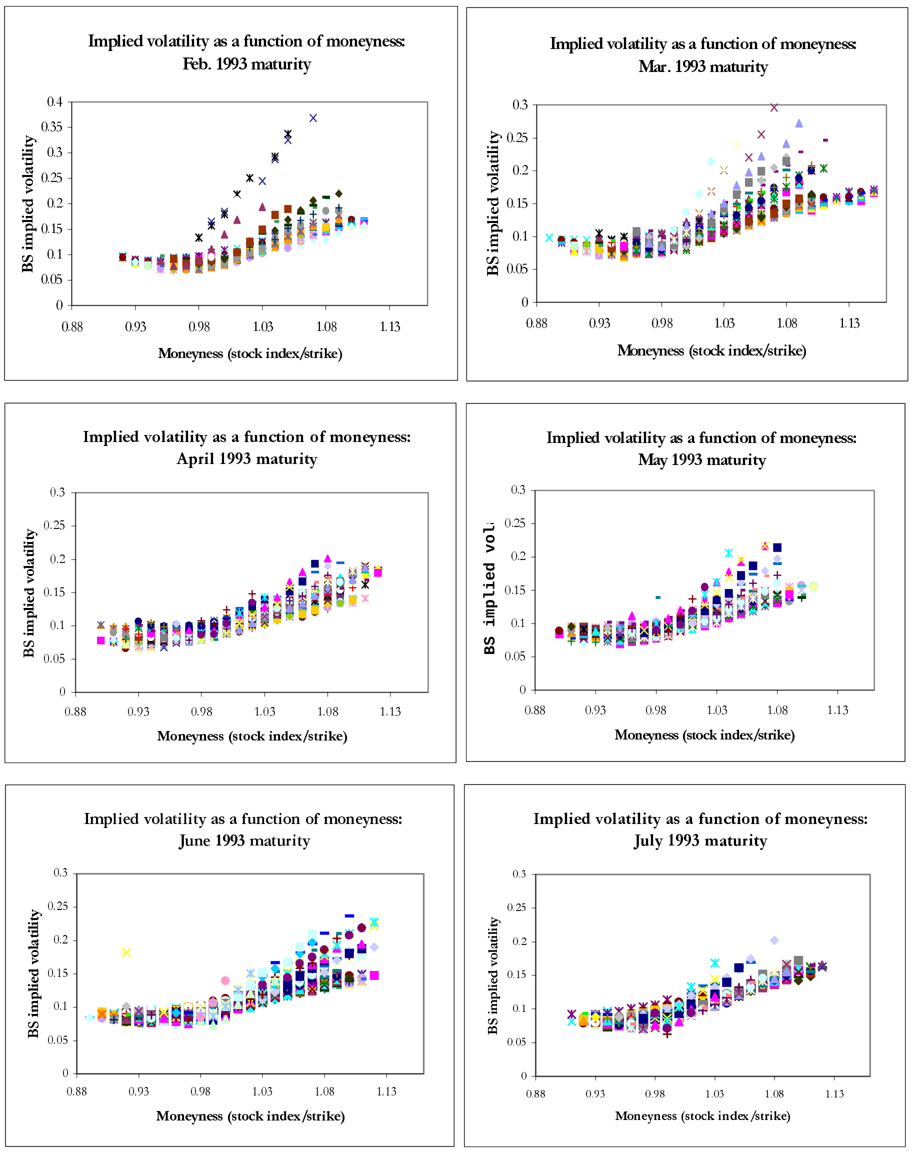

Figure 1 (cont'd). Implied volatility as a function of moneyness for S&P 500 index options.