

Option Prices with Uncertain Fundamentals*

Alexander David

John M. Olin School of Business
Washington University in St. Louis

Pietro Veronesi

Graduate School of Business
University of Chicago

December 2002

*We thank Gurdip Bakshi, Michael Brennan, John Cochrane, George Constantinides, Gabrielle Demange, Darrell Duffie, Bernard Dumas, Rob Engle, Rick Green, Hayne Leland, Massimo Massa, Joe Ostroy, Rene Garcia, Ken Singleton, Allan Timmermann, Robert Whitelaw, Hong Yan, an anonymous referee, and the seminar participants at Carnegie Mellon, University of Chicago, CIRANO, Federal Reserve Board, INSEAD, London Business School, London School of Economics, McGill, Princeton, Stanford, Stockholm School of Economics, UCLA Economics, UCLA Finance, UCSD Economics, Washington University, 9th Annual Derivatives Securities Conference held at Boston University, 1999 European Summer Symposium in Financial Markets at Gerzensee, the 1999 Rodos Economic Theory Conference, the 2000 Western Finance Association meetings, and the 1st World Congress of the Bachelier Finance Society for their comments and discussion. We also thank Gary Anderson for invaluable advice on Mathematica programming, Yubo Wang for his help with options data, Shannon Hart and Robert Paul for research assistance, and Jim Hamilton for sharing his Gauss code on regime-switching models. Address (David): One Brookings Drive, Campus Box 1133, St. Louis, MO 63130-4899. Phone (314) 935-6431. Fax: (314) 935-5296. E. Mail: david@olin.wustl.edu. Address (Veronesi): Graduate School of Business, University of Chicago, Chicago, IL 60637. E. Mail: fpverone@gsbux1.uchicago.edu. Phone: (773) 702-6348. Fax: (773) 702-0458. Veronesi gratefully acknowledges support from the Dimensional Fund Advisors Research Fund at the Graduate School of Business, University of Chicago.

Option Prices with Uncertain Fundamentals

Abstract

In an incomplete information model, we show that investors' uncertainty about the drift of a firm's fundamentals affects option prices through its effect on stock volatility and the covariance between returns and volatility. We provide an option pricing formula using Fourier Transforms. Filtered investor beliefs from real earnings growth are able to explain time-variation in implied volatility, the skewness premium, and the kurtosis premium embedded in option prices. Option-implied investor beliefs vacillated rapidly before the crash of 1987, remained highly uncertain for a year afterwards, and except for the 1990-91 recession, strongly favored rapid growth until 1996. We fail to reject the hypothesis that the sequence of option's implied beliefs is generated by the transition density implied by the regime-switching model fitted to fundamentals. Our model calibrated to earnings data provides fits to cross-sectional option prices that are statistically insignificantly different from Heston's (1993) stochastic volatility model, which is recalibrated every month, while its variance forecasting performance is superior.

JEL Classification: G12, G13

Key Words: Earnings Uncertainty, Changing Return-Volatility Correlation, Belief Risk, Conditional Skewness Premium, Conditional Kurtosis Premium, Option-Implied Beliefs, Variance Forecasting

Introduction

A number of stylized facts about the intertemporal behavior of option prices have been uncovered in the empirical financial literature. One of the most interesting aspects of option prices is that on average the prices of out-of-the-money (OTM) put options carry a substantial premium relative to OTM call options, a situation that generates a negatively sloped implied volatility curve.¹ In addition, and perhaps most interestingly, this relative premium of OTM puts compared to OTM calls, called the *skewness premium*, shows a good deal of time-variation. For example, the top panel of Figure 1 shows that the skewness premium for options 3% out of the money has been as low as 0.8 in 1986-87 and as high as 2.0 in 1994-1995, while the Black-Scholes model, for instance, would have predicted a constant 0.97 (see also e.g. Bates (1991)). Indeed, the analysis of the time-series and cross-section of option prices suggests that the entire “risk-neutral” return distribution, which determines the prices of contingent claims on stocks, changes over time. We contribute to the options literature by providing a theory and supporting evidence that the dynamics of the risk-neutral return distribution are intimately related to the dynamics of investor’s perceived distribution of the growth of economic fundamentals, such as corporate earnings.

Several models for stock prices have been proposed in the literature that have been useful in fitting option price dynamics. However, since most option pricing models start with exogenously specified functions for stock volatility and jumps, they have been unable to provide a link between innovations to fundamentals and options prices. Such models cannot provide answers to questions like: What news on the economy or the basic earnings process can lead to wild fluctuations in options implied volatility curves? What mechanism is at play? If for example market participants receive bad news on future earnings, how does the implied volatility curve of options react? Do investors expectations of the growth of fundamentals have any effect on the risk-neutral distribution of stock prices? Taking a step back from the standard approach, we model explicitly the process for fundamentals and obtain endogenously the effects that investors’ uncertainty about the fundamental growth has on stock returns. We find that the more uncertain investors are, the more volatile are stock prices. In addition, when

¹The implied volatility of an option refers to the volatility level that has to be inserted into the Black and Scholes formula in order for it to match the price of a traded option using as inputs the current stock price, dividend yield, interest rate, and the strike price and maturity of that option contract. The implied volatility curve refers to the levels of implied volatilities *across* strike prices. Black and Scholes assumptions predict a constant implied volatility curve.

investors think earnings growth is more likely to slow down rather than pick up, they assess that stock prices are more likely to fall rather than rise and are willing to pay a premium for puts over calls. However, as the beliefs about earnings growth change over time, the size of this premium fluctuates, leading to a rationale to the time-varying skewness premium.

The need for an economic model able to explain the dynamics of implied volatility curves has been realized in the literature. For example, in their classic and influential paper, Hull and White (1987) show that if stock return volatility is stochastic, then a negative correlation between returns and volatility tends to generate a negatively sloped implied volatility curve, or, equivalently, a high skewness premium. However, when they turn to discuss the empirical results in Rubinstein (1985) they also write:

The observed implied volatility patterns in relation to S/X are consistent with a situation in which during the 1976-77 period the volatility was positively correlated with the stock price, while, in the 1977-78 period, the correlation was negative. [...] In order for [these results] to support this model it is necessary to posit that from one year to the next, the correlation between stock prices and the associated volatility reversed sign. It is difficult to think of a convincing story why this event should occur. (Page 297-298, Hull and White 1987)

Indeed, this difficulty is apparent also in the bottom panel of Figure 1, which plots the correlation between returns and volatility necessary at each date t to fit the Heston (1993) stochastic volatility model to the cross-section of options on that date. Clearly, the changes in correlation are very large and hence inconsistent with standard reduced-form models for returns that assume a constant correlation between returns and volatility.

Our model, by contrast, has a built-in economic mechanism that generates time-variation in the correlation between returns and volatility. In fact, we show that during good times, disappointing news on fundamentals lead to both negative returns and an increase in uncertainty about fundamental growth, which in turn increases return volatility. Hence, during good times investors rationally expect a negative covariance between returns and volatility, leading to a negatively skewed (risk-neutral) return distribution. The skewness premium is then high and the implied-volatility curve is downward sloping. During bad times, by contrast, good news on fundamentals lead to positive returns, but it also increases the uncertainty on whether a transition to an expansionary phase of fundamental growth has occurred. This higher uncertainty

increases the reaction of beliefs to news and hence return volatility. Thus, *ex ante* the return distribution is positively skewed and the implied volatility curve is upward sloping. Our model is then able to predict changes in the slope of the implied-volatility curves as well as a high time-variation in the skewness premium.²

The specifics of the model, developed in Section 1, are as follows: we assume that dividends of a firm (or of the economy as a whole) are generated by a standard lognormal process, whose drift jumps between two unobservable states. These can be interpreted as “booms” and “recessions,” although later we will see that shifts in earnings growth states are more frequent than the NBER-dated recession-boom cycle. Investors use past realizations of dividends and estimate the current drift rate of the economy. Under these assumptions, we obtain closed form formulas for the stock price and show that equilibrium return volatility is indeed stochastic and it has a stochastic correlation with returns themselves. The intuition is provided above.

In Section 2, we provide a fast Fourier series approximation for option prices. This method has been useful for providing analytical results to several option pricing problems in recent years, such as those with stochastic jumps, and stochastic volatility (some examples of this methodology are in Heston 1993, Bates 1996, Bakshi, Cao, and Chen 1997, Scott 1997). As in these papers, we reduce the option pricing problem to the numerical Fourier Inversion of the solution of a boundary value problem (BVP). The BVP itself has no explicit analytical solution, but its solution can be well approximated by a (small) finite series of polynomials. A numerical Fourier inversion yields the Fourier Transform of the risk-neutral return density function. Option prices are then calculated as the sum of the value of four primitive securities, each of which is obtained by evaluating the Fourier Transform at chosen points.

A calibration of the model is provided in Section 3. Real earnings growth on the S&P 500 for the period 1960:1 - 1998:9 is fitted with a 2-state regime-switching filter. The estimation procedure allows us to obtain a time series of “beliefs” over high and low growth states and hence to compute proxies for investors’ expected drift rate of earnings, their uncertainty, and the implied theoretical return volatility and covariance between returns and volatility. The latter is shown to be mainly negative (70% of the time), with occasional switches to a positive

²Indeed, as we shall see, earnings data suggests that a change in slope – from positive to negative – should have occurred at the end of 1977, when unexpected good news increased the probability of a high growth state. The finding in the quote by Hull and White reported above is therefore consistent with our model.

sign. This helps explain the fact that the implied volatility curve is mainly negatively sloped (positive skewness premium) with occasional switches to positive.

The first set of empirical tests, also in Section 3, uses the parameter estimates for the structural model and the belief process obtained from the filter described in the previous paragraph. From these, we can construct the time-series of “model-implied” options. We then show empirically (*i*) that the second, third, and, fourth moment of the model-implied risk-neutral return density are related to their empirical counterparts implicit in the prices of traded options, and (*ii*) that on average, the “model-implied” prices of options are of the same magnitude of those of traded options. The latter result is important in view of the fact that the model parameters were estimated without any information from actual traded options data.

It is worth emphasizing that each of the moments of the risk-neutral return distribution has an appealing interpretation within our model: The second moment provides a positive association between investors’ uncertainty about fundamentals and the implied volatility in traded options. The third moment test, constructed as the aforementioned *skewness premium* embedded in the ratio of prices of out-of-the-money puts to calls, shows that when investors are confident that fundamentals growth is strong, and therefore price-earnings ratios are high, they value highly the possibility that bad news will be disappointing and cause a sharp decline in the stock price. Our model then provides an alternative explanation for the skewness premium to the one put forward in the current literature, which essentially agrees that it is due to a “crash-o-phobia” (see e.g. Bates (1999)). Our explanation, instead, is based on investors’ views that the rapid growth of earnings in the 1990s simply could not persist indefinitely. Finally, the fourth moment test, constructed as the *kurtosis premium* embedded in the prices of well-practiced straddle strategies in the option market, shows that when investors are highly uncertain about fundamentals growth, they are willing to bet that upcoming stock price moves will not be ‘very’ large in either direction, because given their high uncertainty, they are unlikely to be surprised by a large shock to fundamentals.

In Section 4, we provide a second set of statistical tests of the uncertainty model: Following the rich tradition of “reverse financial engineering,” we use traded options at any date t in the sample to back-out *option’s implied beliefs*, that is, the level of beliefs that best fit the cross-section of option prices at t . The belief process obtained by the econometrician from this

exercise is consistent with investors' trading on signals other than the history of past earnings growth.

Qualitatively, the options' implied belief process has several interesting features, including a dramatic increase in economic uncertainty in the aftermath of the crash of October 1987 and during the 1990-91 recession. Besides its qualitative features, the options' implied belief process has also the statistical properties of a rational belief process generated by a regime-switching model. In other words, we cannot reject the joint hypothesis that the sequence of options' implied beliefs are generated by the transition density of beliefs satisfying the hypothesized stochastic differential equation under the parameters obtained by fitting the earnings process.

Finally, in order to gauge the size of the pricing errors in our model, we also compare them with those obtained from more standard models, such as the Heston (1993) stochastic volatility (SV) model or Black-Scholes model. We find that our model produces similar pricing errors as the SV model, and lower than the BS model. However, as the bottom panel of Figure 1 shows, the low pricing errors obtained in the SV model are obtained at the expense of a very high instability of the implied correlation parameter (and other parameters as well). By contrast, we can obtain similar pricing errors by keeping all the parameters of the model constant. We also conduct a hedging comparison of the various models and find that our model produces lower hedging errors than the BS models, and similar ones to the SV model. Finally, we use the implied parameters of the models and compare the forecasts of realized stock variance to the maturity of the option. From historical experience, stock market volatility has been high following periods of sporadic earnings growth leading in our model to high uncertainty. Under this metric, we find that our model performs best.

Review of the Literature

Learning models have been successful in providing intuition for the dynamics of key financial variables in standard finance problems such as asset-allocation and the term-structure of interest rates. We do not attempt an exhaustive survey of this literature. For financial markets with learning in continuous-time, prominently Detemple (1986), Dothan and Feldman (1986), Gennotte (1986), Brennan (1998), and Brennan and Xia (2001) use the results of the Kalman filter to characterize the learning of agents when the unobservable drift of returns followed a

diffusion process. To circumvent the deterministic variance property of the estimation error process in Gaussian models, David (1997), Veronesi (1999) and Veronesi (2000) developed pricing models with unobservable regime switches. These models are able to generate stochastic volatility of financial variables purely from the fluctuating uncertainty in the learning process.³ The added contribution of this article is to show theoretically and test empirically, that the time-varying *correlation* between returns and volatility induced by uncertainty about corporate earnings, is related not only to (implied) volatilities, but also to the conditional skewness and conditional kurtosis of investors' forward looking return distribution.

Since we first wrote this paper, we have become aware of four other working papers that explore the impact of investors' learning processes on option prices. Yan (1999) studies option pricing in a Kalman filtering model where despite the convergence of estimation error, return volatility is stochastic due to variations in the riskless rate. Campbell and Li (1999) study option prices in a model where investors learn about the underlying volatility regime in returns. While option prices in their model also depend on beliefs of agents, due to the lack of correlation between returns and volatility, their model implies symmetric implied volatility smiles. Garcia, Luger, and Renault (1999) provide a model where investors' learn about the underlying drift and volatility regime of exogenously specified stock returns and the stochastic discount factor process and generate asymmetric smiles through an exogenous and fixed correlation between returns and volatility. Finally, as in our paper, Guidolin and Timmermann (2000) study similar learning effects in an economy where the stock process is endogenously determined each period in an equilibrium framework. However, in their model, the underlying drift follows an i.i.d. process leading to a lack of persistence in investors' uncertainty, and in the continuous time limit, option prices converge to the Black-Scholes formula.

³We make a slight but often used distinction between risk and uncertainty: we reserve the term uncertainty for situations where the risky process generating fundamentals is not fully known by agents. Whether this is the key distinction between risk and uncertainty made by Frank Knight is a subject of interpretation and debate; see Jones and Ostroy (1984) for a further discussion and references to the literature.

1 The Model

We consider an endowment economy and begin by describing the processes for dividends and investors' endowment. Let $\mathbf{W}_t = (W_{1,t}, W_{2,t})'$ be a two-dimensional Brownian motion defined on a complete probability space $(\Omega, \mathcal{P}^0, \mathcal{F}^0)$. We make the following assumptions:

Assumption 1 (*real dividends*). The real dividend rate D_t evolves according to the stochastic process

$$\frac{dD_t}{D_t} = \theta_t dt + \sigma_D d\mathbf{W}_t, \quad (1)$$

where σ_D is a 1×2 constant vector. The dividend drift θ_t changes over time according to a process described in next assumption:

Assumption 2 (*regime shifts*). θ_t follows a continuous time, two-state regime-shift model with infinitesimal generator

$$\mathbf{\Lambda} = \begin{pmatrix} -\lambda_{12} & \lambda_{12} \\ \lambda_{21} & -\lambda_{21} \end{pmatrix}.$$

The two states are denoted by θ_1 and θ_2 with $\theta_1 < \theta_2$.

We assume that besides dividends, investors also receive an endowment over time. Following a number of recent articles (see e.g. Berk, Green, and Naik 1999, Brennan and Xia 2001, Cecchetti, Lam, and Mark 1993, Campbell and Cochrane 1999, Barberis, Huang, and Santos 2001), we also find it convenient to model directly the total consumption process C_t , which equals endowment plus dividends, and then define the endowment rate as a residual. Assuming an additional source of income in addition to dividends is also motivated by the empirical facts that total market dividends make up only a mere 5% of investors' total income and that consistently with this, the correlation between consumption growth and dividend growth is extremely low.⁴ Hence, we make the following assumption:

⁴ Following Veronesi (2000), it is possible to solve the model in Assumption 1 to 5 within the framework of a standard Lucas (1978) exchange economy, where in equilibrium consumption equals dividends. However, as in any exchange economy where investors have power utility and expected consumption growth is positively autocorrelated, such a model would imply that the price dividend ratio decreases following good news on earnings (and vice versa), which is counterfactual in the data (see e.g. Campbell 1999). Similarly to the articles cited in the text, we also find that the introduction of an endowment process greatly helps to obtain a price dividend ratio that is positively correlated with news on future cash flows. Other approaches that would yield this result include habit formation (see e.g. Campbell and Cochrane 1999) and recursive utility with a high elasticity of intertemporal substitution.

Assumption 3 (*total consumption*). Investors' consumption C_t evolves according to the stochastic process

$$\frac{dC_t}{C_t} = gdt + \sigma_C d\mathbf{W}_t, \quad (2)$$

where σ_C is also a 1×2 constant vector and g is constant.⁵

We finally describe investors' information set.

Assumption 4 (*incomplete information*). Investors do not observe the drift rate θ_t , but they know the parameters of the model.

Some comments about Assumptions 1 to 4 are in order. First, the use of Markov regime-switching models to capture the cyclical features of real macro-economic variables has been advocated by a number of recent articles, starting with the classic Hamilton (1989) paper about real GDP growth (see e.g., Kim and Nelson (1999) for a review). Indeed, in the empirical section we will see that real earnings growth shows a strong cyclical pattern, as we will be able to strongly reject a single-state model in favor of a two-state model. Although others, such as Brennan and Xia (2001), have favored the use of a linear Ornstein-Uhlenbeck (OU) process for θ_t , we find that the regime-switching model compares favorably with respect to the OU model, as it produces slightly lower forecasting errors. More recently, Cagetti, Hansen, Sargent, and Williams (2002) also use a regime-switching model for the drift rate of real earnings. One key additional advantage of our modeling strategy is that the conditional variance process of investors' estimates is stochastic under the regime-switching assumption, but it follows a deterministic process in Gaussian models (c.f. David 1997). This feature, as will be seen in Sections 3 and 4, helps provide an economic explanation for time-variation in investors' forward-looking state return density function.

Second, Assumption 4 formalizes Hamilton (1989) discussion that the econometrician is not aware of the current regime at any point in time and hence must use some filtering method to assign a probability to the two states. Basically, with assumption 4 we are assuming that investors' have the same "filtering" problem as does the econometrician in Hamilton (1989), so that they must use past data to obtain the likelihood to be in each regime.

Finally, we need an assumption on investors' preferences:

⁵ In a previous version of the model we assumed that the drift rate g of the consumption process (2) also evolved according to a two-state regime shift as in Assumption 2 above. However, when estimated with monthly consumption data it turned out that $g_1 \approx g_2$. For simplicity, we then assume that the growth rate of consumption is constant. This also implies that the interest rate is constant, allowing us to better compare our results to other models, such as Black and Scholes (1973) and Heston (1993).

Assumption 5 (*preferences*). Investors are endowed with a iso-elastic instantaneous utility function of consumption

$$U(C, t) = e^{-\phi t} \frac{C^{1-\gamma}}{1-\gamma}.$$

1.1 The Dynamics of Investors' Beliefs

Let the probability that investors assign to the high state (θ_2) be

$$\pi_t = \Pr(\theta_t = \theta_2 | \mathcal{F}_t),$$

where, by definition, $\mathcal{F}_t = \sigma(D_\tau, C_\tau : \tau \leq t)$ is the sigma-algebra generated by the observable processes D_τ and C_τ . We then have the following:

Lemma 1 *Given an initial condition $\pi_0 = \hat{\pi} \in [0, 1]$, the probability, π_t , satisfies the stochastic differential equation:*

$$d\pi_t = \kappa(\pi^* - \pi_t) dt + \sigma_\pi(\pi_t) d\widetilde{\mathbf{W}}_t, \quad (3)$$

where $\kappa = \lambda_{12} + \lambda_{21}$, $\pi^* = \frac{\lambda_{12}}{\lambda_{12} + \lambda_{21}}$, [later we will refer to the drift as $\mu_\pi(\pi_t) = \kappa(\pi^* - \pi_t)$], $\sigma_\pi(\pi_t) = \pi_t(1 - \pi_t) \mathbf{b}$, $\mathbf{b} = (\theta_2 - \theta_1, 0) (\sigma'_D, \sigma'_C)^{-1}$, and

$$d\widetilde{\mathbf{W}}_t = \begin{pmatrix} \sigma_D \\ \sigma_C \end{pmatrix}^{-1} \begin{pmatrix} \frac{dD_t}{D_t} - E_t \left(\frac{dD_t}{D_t} \right) \\ \frac{dC_t}{C_t} - E_t \left(\frac{dC_t}{C_t} \right) \end{pmatrix} \quad (4)$$

is a standard two-dimensional Brownian motion defined on a probability space $(\Omega, \mathcal{P}, \mathcal{F})$, where \mathcal{F} is the filtration generated by $\left(\frac{dD_\tau}{D_\tau}, \frac{dC_\tau}{C_\tau} \right)$. Moreover, for every $t > 0$, $0 < \pi_t < 1$.

Proof: See Liptser and Shirayev (1977).

The model implies that investors' posterior probability on the current state follows a diffusion process with mean-reverting drift and quadratic volatility. As we shall see, it is this last component that feeds back into the pricing equation, generating (endogenously) a stochastic volatility process for stock returns, which in turn affects option prices. In (4), $d\widetilde{\mathbf{W}}_t$ is simply the orthogonalized innovation process. Its impact on belief π_t is described by the diffusion part $\pi_t(1 - \pi_t)$ which is high when there is high uncertainty about the true current state (i.e. $\pi_t \approx 1/2$). Finally, the 1×2 vector \mathbf{b} regulates the “precision of signals”:

When either of the two components is very large, the posterior probability π_t converges quickly to the “true” underlying value.

The following corollary will be used below:

Corollary 1 *Conditional on the information filtration $\{\mathcal{F}_t\}$ obtained in Lemma 1, innovations in consumption growth and beliefs are orthogonal to each other:*

$$\text{Cov}_t \left(\frac{dC_t}{C_t}, d\pi_t \right) = \sigma(\pi_t) \cdot \sigma'_C = 0$$

Proof. In Appendix A.

The intuition is as follows: Agents can observe neither the drift θ_t nor the vector Wiener processes $\mathbf{W}_t = (W_{1,t}, W_{2,t})$ in the earnings process (1). Since the drift of the consumption process in (2) does not depend on the regime, its innovations contain no information about θ_t but only about the vector \mathbf{W}_t , which is then incorporated into the posterior π_t . The new vector Wiener process $\widetilde{\mathbf{W}}_t = (\widetilde{W}_{1,t}, \widetilde{W}_{2,t})$ obtained in Lemma 1 contains only the residual orthogonal part of dC_t , as it is an orthogonalized expectation error. As a consequence, conditional on investors’ information, $\text{Cov}_t(dC_t/C_t, d\pi_t) = 0$.

We finally report a result about the stationary density distribution of the belief process π_t — which is the long-run distribution over beliefs. By a simple adaptation of Property 2 in David (1997), the stationary distribution over beliefs is given by

$$\Psi(\pi) = C_1 \cdot \exp \left(\frac{-2 \cdot (\lambda_{12}\pi + \lambda_{21}(1 - \pi))}{\mathbf{b} \cdot \mathbf{b}' \cdot \pi(1 - \pi)} \right) \cdot \left(\frac{1 - \pi}{\pi} \right)^{\frac{2 \cdot (\lambda_{12} - \lambda_{21})}{\mathbf{b} \cdot \mathbf{b}'}} \cdot \frac{1}{\mathbf{b} \cdot \mathbf{b}' \cdot (1 - \pi)^2 \cdot \pi^2}, \quad (5)$$

where we recall that $\mathbf{b} = (\theta_2 - \theta_1, 0) (\sigma'_D, \sigma'_C)^{-1}$, and C_1 is a constant that ensures that the density integrates to 1. The shape of the stationary distribution depends on the characteristics of the learning process. We will comment on its properties in the empirical section.

1.2 Stock Returns and Volatility

Given assumption 3 and 5, we can define the pricing kernel $m_t = U_c(C_t, t) = e^{-\phi t} C_t^{-\gamma}$ to price any real claim yielding a payoff flow h_τ at time τ by using the standard formula

$$m_t \times P_t = E_t \left[\int_t^\infty m_\tau h_\tau d\tau \right]. \quad (6)$$

By applying this formula to bonds and stocks, where $h_\tau = 1_{\tau=T}$ and $h_\tau = D_\tau$, respectively, we obtain the following proposition:

Proposition 1 (a) *The real interest rate, r , is given by $r = \phi + \gamma g - \frac{1}{2}\gamma(\gamma + 1)\sigma_C\sigma'_C$.*

(b) *Let $\mu_S = r + \gamma\sigma_D\sigma'_C$, and let $\mu_S > \theta_i$ for $i = 1, 2$. Then, for every t , the stock price is given by:*

$$S_t = D_t \cdot (A\pi_t + B), \quad (7)$$

where A and B are the two constants

$$\begin{aligned} A &= \frac{\theta_2 - \theta_1}{(\mu_S - \theta_1 + \lambda_{12})(\mu_S - \theta_2 + \lambda_{21}) - \lambda_{12}\lambda_{21}} > 0, \quad \text{and,} \\ B &= \frac{\mu_S - \theta_2 + \lambda_{21} + \lambda_{12}}{(\mu_S - \theta_1 + \lambda_{12})(\mu_S - \theta_2 + \lambda_{21}) - \lambda_{12}\lambda_{21}} > 0. \end{aligned}$$

(c) *The stock return process is given by:*

$$\frac{dS_t}{S_t} = (\mu_S - \delta_t) dt + \sigma_{S,t} d\widetilde{\mathbf{W}}_t,$$

where expected return μ_S is given above in (b) and the dividend yield δ_t and the volatility $\sigma_{S,t}$ are given by:

$$\delta_t = \frac{D_t}{S_t} = \frac{1}{A\pi_t + B}, \quad \text{and,} \quad (8)$$

$$\sigma_{S,t} = \sigma_D + \frac{A\pi_t(1 - \pi_t)}{A\pi_t + B} \mathbf{b}. \quad (9)$$

Proof. In Appendix A.

First of all, notice that notwithstanding the ‘‘jumps’’ in the drift rate of dividends, equation (7) entails that the stock price follows a continuous-path process. Second, equation (9) shows that the volatility of stock returns depends on the probability π_t . Since σ_S is a 1×2 vector, the volatility of stock returns can be characterized more effectively by its *variance* as

$$V(\pi_t) \equiv \sigma_{S,t}\sigma'_{S,t} = \sigma_D\sigma'_D + 2\frac{A\pi_t(1 - \pi_t)(\theta_2 - \theta_1)}{A\pi_t + B} + \frac{A^2\pi_t^2(1 - \pi_t)^2}{(A\pi_t + B)^2} \mathbf{b}\mathbf{b}', \quad (10)$$

where we used the fact that $\mathbf{b}\sigma'_D = (\theta_2 - \theta_1)$. We have the following corollary:

Corollary 2 *The function $V(\pi)$ is convex in π , with a unique maximum $\pi^m = \arg \max_{\pi} (V(\pi))$. Hence, $V_{\pi} < 0$ if and only if $\pi > \pi^m$.*

Proof. In Appendix A.

By Ito's lemma, the stock return variance then evolves stochastically according to the following process:

$$dV_t = \mu_{V,t} dt + \sigma_{V,t} d\widetilde{\mathbf{W}}_t, \quad (11)$$

where $\mu_{V,t} = \lambda V_{\pi,t} (\bar{\pi} - \pi_t) + \frac{1}{2} V_{\pi\pi,t} \pi_t^2 (1 - \pi_t)^2 \mathbf{b}\mathbf{b}'$ and $\sigma_{V,t} = V_{\pi,t} \pi_t (1 - \pi_t) \mathbf{b}$. In addition, the covariance between volatility and returns changes over time as well, because

$$\text{Cov}_t \left(dV_t, \frac{dS_t}{S_t} \right) = \sigma_{V,t} \sigma'_{S,t} = V_{\pi,t} \pi_t (1 - \pi_t) \left((\theta_2 - \theta_1) + \frac{A\pi_t (1 - \pi_t)}{A\pi_t + B} \mathbf{b}\mathbf{b}' \right). \quad (12)$$

Since all the terms are positive, the sign of the covariance depends on the sign of $V_{\pi,t}$. From corollary 2 this implies a negative covariance for high π ($> \pi^m$) and positive for low π . The intuition for this result is given in the Introduction: Disappointing news in good times ($\pi_t \approx 1$) lead both to negative returns and to an increase in uncertainty, and hence of return volatility. The covariance between returns and volatility is then negative. During bad times ($\pi_t \approx 0$) the opposite argument leads to a positive covariance between returns and volatility.

1.3 The Model Conditional on Investors' Information

Before turning to the valuation of contingent claims, some readers may find it useful to see the formulation of the fundamental's processes in equations (1) and (2) *conditional* on investors' information filtration $\{\mathcal{F}_t\}$, which is described in Lemma 1. Indeed, most models in finance start out with a probability space $(\Omega, \mathcal{P}, \mathcal{F})$ and an *exogenous* filtration $\{\mathcal{F}_t\}$ on which all stochastic processes are defined. Lemma 1 simply endogenizes the filtration $\{\mathcal{F}_t\}$ and obtains the dynamics of a belief state-variable as implied by rational learning (see David (1997)). By inverting (4) we can rewrite

$$\frac{dD_t}{D_t} = \bar{\theta}_t dt + \sigma_D d\widetilde{\mathbf{W}}_t \quad \text{and} \quad \frac{dC_t}{C_t} = g dt + \sigma_C d\widetilde{\mathbf{W}}_t, \quad (13)$$

where $\bar{\theta}_t = E_t[\theta] = \pi_t \theta_2 + (1 - \pi_t) \theta_1$, follows $d\bar{\theta}_t = \kappa (\theta^* - \bar{\theta}_t) dt + (\theta_2 - \bar{\theta}_t) (\bar{\theta}_t - \theta_1) \mathbf{b} d\widetilde{\mathbf{W}}_t$ with $\theta^* = \pi^* \theta_2 + (1 - \pi^*) \theta_1$ and π^* given in Lemma 1.

In the empirical Section 3.2 we will find support for a regime-switching model for earnings and we will show that the resulting uncertainty is able to partly explain the dynamics of option prices and the risk-neutral return density. As a result, the original model described in Assumptions 1 to 4 not only provides an economic intuition for the dynamics of option prices, but it also proves useful to guide our estimation strategy and to interpret the empirical results.

2 The Value of Contingent Claims

We can use (6) to determine the value of any contingent claim written on the stock price S_t . Since the evolution of S_t is non-Markovian, we need to add the state variable π_t in order to describe the economy. The risk-neutral processes for the stock $\{S_t\}$ and the belief $\{\pi_t\}$ are

$$\frac{dS_t}{S_t} = (r - \delta_t) dt + \sigma_{S,t} d\widetilde{\mathbf{W}}_t^*, \quad \text{and}, \quad (14)$$

$$d\pi_t = \kappa(\bar{\pi} - \pi_t) dt + \pi_t(1 - \pi_t) \mathbf{b} d\widetilde{\mathbf{W}}_t^*, \quad (15)$$

where $\widetilde{\mathbf{W}}_t^*$ is a two-dimensional Wiener process. We remark that Corollary 1 implies $\text{Cov}(dm_t/m_t, d\pi_t) = 0$, and hence that variation in π_t does not require a market price of risk.⁶ As a consequence, the risk-neutralized process for the belief, $\{\pi_t\}$, equation (15), is the same as the one under the objective measure (3).

By standard no-arbitrage arguments (Duffie 1996, Chapter 5.G) every contingent claim $f(t, S, \pi)$ must satisfy the following partial differential equation (PDE):

Proposition 2 *Each security $f(t, S, \pi)$ satisfies the following PDE:*

$$\begin{aligned} rf = & \frac{\partial f}{\partial t} + \frac{\partial f}{\partial S} S(r - \delta_t) + \frac{\partial f}{\partial \pi} (\kappa(\bar{\pi} - \pi)) + \\ & + \frac{1}{2} \frac{\partial^2 f}{\partial S^2} S^2 \left(\sigma_D \sigma_D' + 2 \frac{A\pi(1-\pi)(\theta_2 - \theta_1)}{A\pi + B} + \frac{A^2 \pi^2 (1-\pi)^2}{(A\pi + B)^2} \mathbf{b}\mathbf{b}' \right) + \\ & + \frac{\partial^2 f}{\partial S \partial \pi} S \left(\pi(1-\pi)(\theta_2 - \theta_1) + \frac{A\pi^2(1-\pi)^2}{A\pi + B} \mathbf{b}\mathbf{b}' \right) + \frac{\partial f^2}{\partial \pi^2} \pi^2 (1-\pi)^2 \mathbf{b}\mathbf{b}', \quad (16) \end{aligned}$$

⁶Theorem 3 in Cox, Ingersoll, and Ross (1985) shows that for pricing contingent claims, the drift of an arbitrary state variable must be reduced by its equilibrium market price of risk. In their representative agent production economy with time-separable preferences, the market price of risk is proportional to the covariance between the factor and agent's wealth. See also eq. (12.10) in Hull (1993).

and a terminal condition determined by its final payoff.

Proof. In Appendix B.

An immediate implication of this proposition is the following:

Corollary 3 *Under perfect information and no regime shifts, the Black and Scholes option pricing formula obtains.*

Corollary 3 is immediate from the fact that under its assumptions the transition matrix $\Lambda = 0$ and $\pi_t = 1$ or 0 . Hence, in (16) all the partial derivatives with respect to π vanish and we obtain the standard PDE of Black and Scholes formula.

2.1 A Series Solution for Option Prices

In this section we solve for option prices using Fourier Transform methods (some examples of this methodology are in Stein and Stein 1991, Heston 1993, Bates 1996, Bakshi, Cao, and Chen 1997, Scott 1997). The price of a call option can be obtained as follows:

Proposition 3 *Let $z_t = \log(S_t)$. Then, the call option price can be unbundled into a portfolio of Arrow-Debreu securities as*

$$C^U(z_t, K, \pi_t, T-t) = G(z_t, \pi_t, T-t) \cdot \Pi_1(z_t, K, \pi_t, T-t) - e^{-r(T-t)} \cdot K \cdot \Pi_2(z_t, K, \pi_t, T-t), \quad (17)$$

each of which can be calculated from $\bar{f}(z_t, \omega_1, \pi_t, T-t)$, the Fourier transform of the risk-neutral return density computed at point ω_1 and conditional on z_t , π_t , and time-to-maturity, $T-t$. $G(z_t, \pi_t, T-t)$ is the forward price for a unit of stock, and $\Pi_1(z_t, K, \pi_t, T-t)$ and $\Pi_2(z_t, K, \pi_t, T-t)$ are, respectively, the time- t Arrow-Debreu securities that pay a dollar if the option finishes in the money under two different measures. $\bar{f}(z_t, \omega_1, \pi_t, T-t)$ can be obtained from the Fourier Inversion in the time variable of the solution to the boundary value problem (BVP)

$$-e^{i\omega_1 z_t} = h_1(\pi) \hat{f}_{\pi\pi} + h_2(\pi) \hat{f}_\pi + h_3(\pi) \hat{f}, \quad (18)$$

with natural boundaries at $\pi = 0$ and $\pi = 1$, where $h_i(\pi)$, $i = 1, 2, 3$ are simple polynomial functions in π provided in the proof.

Proof. See Appendix B.

To perform the inversion as well as the computation of option prices we define the following functions as in Scott (1997):

$$g_1(z_t, \omega_1, \pi, T - t) = \bar{f}(z_t, 1/i + \omega_1, \pi_t, T - t) / \bar{f}(z_t, 1/i, \pi_t, T - t), \quad \text{and}, \quad (19)$$

$$g_2(z_t, \omega_1, \pi_t, T - t) = \bar{f}(z_t, \omega_1, \pi_t, T - t) / \bar{f}(z_t, 0, \pi_t, T - t). \quad (20)$$

Then, by the Fourier Inversion formula (c.f., for example, Kendall and Stuart 1977) the probability distribution functions can be written as

$$\Pi_j(z_t, K, \pi_t, T - t) = \frac{1}{2} + \frac{1}{\Pi} \int_0^\infty \text{Re} \left[\frac{e^{-i\omega_1 \log K} g_j(z_t, \omega_1, \pi_t, T - t)}{i\omega_1} \right] d\omega_1. \quad (21)$$

The BVP (18) has no known explicit analytical solution, however we provide its M term polynomial approximate solution, $\bar{f}^M(z_t, \omega_1, \pi, T - t)$, in (49) of Appendix B. The M -term approximate option price is obtained using the Fourier Inversion formula in (21). In our numerical exercises, we have found that convergence of option prices occurs rapidly, and $M = 5$ is generally sufficient for accuracy up to 4 decimal places for option prices on an underlying stock price of \$100.⁷

The call option pricing formula in (17) and the analogous formula for puts are convenient enough to provide various partial-derivatives of call option prices with little additional numerical computation. The results of the following corollary will be used in Section 4.2, where we study the hedging properties of U-options:

Corollary 4 *The partial delta of the call option price in Proposition 3 with respect to the stock price, C_S , and beliefs, C_π , are given by:*

$$\begin{aligned} C_S^U(z_t, K, \pi_t, T - t) &= \frac{G(z_t, \pi_t, T - t)}{S} \Pi_1(z_t, K, \pi_t, T - t), \quad \text{and}, \\ C_\pi^U(z_t, K, \pi_t, T - t) &= G_\pi(z_t, \pi_t, T - t) \Pi_1(z_t, K, \pi_t, T - t) + G(z_t, \pi_t, T - t) \Pi_{1\pi}(z_t, K, \pi_t, T - t) \\ &\quad - K \cdot \Pi_{2\pi}(z_t, K, \pi_t, T - t), \end{aligned}$$

where the partial derivatives $\Pi_{i\pi}$, $i = 1, 2$, and G_π are provided in the proof in Appendix B.

⁷All prices reported in the paper have been verified by Monte-Carlo simulations.

Proof. See Appendix B.

The series solution in Proposition 3 turns out to be also very tractable in order to obtain variance forecasts over the life of the option, conditional on the belief π_t . The following result will be used in Section 4.2 to gauge the ability of this model to forecast future realized volatility:

Proposition 4 *An M -term series approximation of the variance forecast for the U model, conditional on the set of parameters Θ , and the belief at t , π_t , is given by*

$$\bar{V}^U(\omega, T; \Theta, \pi_t) = E_t \left[\int_t^T V(\pi_s) ds \right] = \sum_{m=1}^M a_m (T-t) (\pi_t - .5)^m, \quad (22)$$

the solution of the PDE

$$0 = -\bar{V}_t^U(t, T; \Theta, \pi_t) + V(\pi_t) + \bar{V}_\pi^U(t, T; \Theta, \pi_t) \mu_\pi(\pi_t) + \frac{1}{2} \bar{V}_{\pi\pi}^U(t, T; \Theta, \pi_t) \sigma_\pi(\pi_t)^2, \quad (23)$$

where, $V(\pi_t)$ is the variance function in eq. (10), and $\mu_\pi(\pi_t)$ and $\sigma_\pi(\pi_t)$ are as in Lemma 1.

Proof in Appendix B.

3 Empirical Analysis

In this section we use real earnings data to estimate the model in Section 1. We use the results of the estimation to obtain implications for the dynamics of options' prices and then test whether the conditional moments of the risk-neutral distribution implied by the model for fundamentals are consistent with the ones implied by options' data. We start from a description of the data used.

3.1 Data

Earnings per share data from 1960:01 and 1998:09 are obtained from Standard and Poor's, and the consumption deflator from the Bureau of Economic Analysis.

For options, we use S&P 500 European call and put options between 1986:04 and 1996:05 traded on the Chicago Board Options Exchange. Because news on fundamentals (earnings and inflation) arrives roughly by the middle of the month, we test relationships on the 15th (the first trading day following) of the month. For out-of-sample and hedging exercises, we consider

similarly filtered prices on the 16th (the following trading day) of the month. The options data are filtered for errors exactly as in Ait-Sahalia and Lo (1998). Specifically, observations with BS implied volatility larger than 70 percent, those with price less than 1/8, and those permitting obvious arbitrage are dropped. To avoid redundancy caused by the prices of calls and puts of the same strike, we use prices of out-of-the-money (OTM) options, which are known to be more liquid than in-the-money (ITM) options.

To minimize the effects of the term-structure of implied volatilities, we construct approximately constant-maturity time series of option prices of three maturities: one month, three months (the maturity closest to 90 days between 65 and 115 days) and six months (the maturity closest to 180 days between 150 and 210 days). We drop one-month observations that are more than 10 percent OTM although this last step provides only second-order improvements in fits of all models. For the time-series tests in Section 3.4 and the implicit fitting exercise in Section 4.2, we also used a sub-sample of the data that excluded the six months following the crash of October 1987. Overall, we used 7,500 option prices over the roughly 10-year sample period. The number of options of each maturity and the periods covered are summarized in Table 1.

For fitting option prices to data we, in addition, use Treasury yields to match the maturity of the options, provided by the Federal Reserve Board, and dividend yields from Standard and Poor's.

3.2 Estimation of the regime-switching model

Table 2 contains maximum likelihood estimates for the parameters of a two-state regime-switching model for real earnings (see Hamilton (1994)).⁸ In the remainder of the paper, we will refer to these set of parameters as $\hat{\Theta}^U$, and the filtered belief process generated by these parameters as $\{\hat{\pi}_t\}$.

⁸David (1993) proves that the belief process in a two state Hamilton discrete-time model converges almost surely to the continuous-time belief process in Liptser and Shirayev (1977), which has been used in Lemma 1. We realize that the estimates of a continuous-time diffusion process for earnings growth based on discrete approximations of the likelihood function may lead to inconsistent estimates for parameters (c.f. for example Lo 1988, Ait-Sahalia 1998). Nonetheless, we argue that for monthly estimation of parameters of the order of magnitude in this paper, the inconsistency will be 'small'. By Example 3 in Lo (1988), the MLE of a discretized log-normal diffusion with constant drift θ and sampled at intervals of length h converges (as the number of observations increases) to $\hat{\theta}^h = \frac{1}{h}(e^{\theta h} - 1)$. For $h = 1/12$, and for $\theta \in [-.2, .2]$, the inconsistency is of the order $.04 \cdot 1/12 + O(\theta^2 h^2) \simeq .003$. Therefore, our results are still informative for the goals of the present paper.

These parameter estimates imply a relatively fast learning process of earnings growth with a double-humped stationary distribution of estimated states, shown in the bottom left panel of Figure 2. The states are asymmetric with a larger mass in the higher state of growth. This finding is consistent with the well known stylized fact that booms are of longer duration than recessions. Figure 4 plots the monthly time-series for real earnings growth (Panel A) and beliefs, $\hat{\pi}_t$ (Panel B). The dark bars indicate NBER dated recessions. As seen in the figure, $\hat{\pi}_t$, the filtered probability of being in the high state, dips during and around each recession of the US economy in the sample, although there have been episodes when earnings growth slowed even when the overall economy was in an expansionary phase. Panels C of Figure 4 plots the time series of the index $\hat{\pi}_t \cdot (1 - \hat{\pi}_t)$, which can be interpreted as “earnings growth uncertainty” as it is maximized for $\hat{\pi}_t = .5$. As discussed earlier, this quantity should be correlated with option’s implied volatility and we will test this statistical relationship below. Finally, Panel D of Figure 4 plots the implied instantaneous covariance between returns and volatility, given by equation (12). The time variation in $\hat{\pi}_t$ induces a substantial variation in this (implied) covariance, which in turn affects the slope the implied volatility curves (see discussion in the Introduction). This time-variation gives a rationale to Hull-White’s puzzling statement quoted in the Introduction about the sudden change in the slope of the implied volatility curve in the 77-78 period. Indeed, a close examination of the pattern of this covariance from Panel D of Figure 4 reveals a positive implied covariance in the 1976-77 period and a negative one in 1977-78, as actually observed in the options data. This drastic change occurred because of good news in earnings realizations that increased the probability of being in a high earnings growth state, as shown in Panel B.

Finally, Table 2 also shows that a Likelihood Ratio test strongly rejects a single state univariate model for earnings growth in favor of a two-state model. Since standard asymptotic results for the LR statistics do not apply to determine the number of regimes, as the nuisance parameters are not identified under the null hypothesis (see e.g. Hamilton (1994), page 698), we apply the critical values in Table 1A of Garica (1998), which are specifically calculated for the choice of one versus two states in the means of uncorrelated and homoskedastic noise process.

3.3 Calibration of Implied Volatilities

We now use the parameter estimates in Table 2 to obtain calibrated option prices from the model in the previous section. This enables us to show that the size of the effects implied by the regime switching model and uncertainty are economically significant. We report the results in terms of Black-Scholes “implied volatilities:” That is, for given interest rate, r , and maturity T , we can use the pricing formula in Proposition 3 to obtain the price of a call option for every strike price K , stock value S and belief π . Given this price, we can invert the Black and Scholes option pricing formula to obtain an estimate of the implied volatility for the same strike and stock value. By doing this exercise for various strike-to-price ratios K/S we obtain a map between implied volatility curves and beliefs π_t . In the right-hand side panels of Figure 3, we plot three of such curves, for initial beliefs $\pi = 0.9$ (top panel), $\pi = 0.5$ (center panel) and $\pi = 0.1$ (bottom panel). To better understand the intuition of the results, the left-hand side panels also plot the risk-neutral return density functions for each of the three initial beliefs for a two-month horizon. For comparison, we also show a Gaussian density with the same first and second moments as of the return process of the uncertainty model. It is worth noticing that the density and moments are easily obtained in this framework by inverting the Fourier Transform of the risk-neutral return density function given by $\bar{f}_r^M = \sum_{n=0}^M b_n(\omega_1, .167) \cdot (\pi - .5)^n$, and the b_n are in (48) in Appendix 2.

The return densities for the three cases are markedly different: for a high π , the return density exhibits a fat left tail, for a low π , the density has a fat right tail, and for the medium beliefs, the density has two humps with larger mass on the right hump. The immediate implication of these densities is a negatively (positively) sloped implied volatility smirk for the high (low) belief case and a frown (or flat implied volatility curve) for the medium belief case. Some testable time-series properties of the densities will be examined in detail in the next subsection. Here we note however that based on the stationary density in Figure 2, the calibrated model implies a negative smirk for about 70% of the time, a frown or flat implied volatility curve for about 15% of the time and a positive smirk for the remaining 15% of the time.

The intuition underlying these findings is quite straightforward. The two state model implies that earnings growth is generated by a *mixture of distributions*, where the exact mixture at a point in time is determined by the beliefs of the agent. Returns, as characterized in

Proposition 1, are completely determined by the beliefs concerning earnings growth and hence inherit the properties of the mixture. The unconditional mixture of the two distributions for the fitted model implies a density which is both negatively skewed (because of the larger proportion of time spent in the high growth state) and has excess-kurtosis. The power of filtering theory becomes evident in agents making a finer distinction of the state of expected earnings growth and hence assigning dynamically evolving mixtures to the two possible distributions. For high beliefs, there is a large chance of a high density and a low chance of a low density creating in sum a negatively skewed density as seen in the top-left panel of Figure 3. Similar logic implies the two densities in the middle and bottom panel.

3.4 Tests for Time Varying Conditional Moments

The results of the previous section suggests that the time-variation in the filtered probability of a high earnings growth should be correlated with the shape of the risk-neutral return density. This section provides evidence that this is the case. We do so by using the standard Black-Scholes implied volatility, and two other popular measures of market sentiment, namely the conditional skewness premium and the conditional kurtosis premium, to elicit additional moments of the risk-neutral return density from option prices.

More specifically, from the options data described in section 3.1 we construct the time series of:

- (i) At-the-money implied volatility, computed as: $\mathcal{V}_t = \frac{1}{2} [C \cdot \text{Vol}(K/S_t = 1) + P \cdot \text{Vol}(K/S_t = 1)]$.
- (ii) Conditional Skewness Premium, computed as the ratio of 3% OTM put prices on 3% OTM call prices: $P_t(K/S_t = 1/1.03) / C_t(K/S_t = 1.03)$.⁹
- (iii) Conditional Kurtosis Premium, computed as:

$$Y_t = [P \cdot \text{Vol}(K/S_t = .97) + P \cdot \text{Vol}(K/S_t = 1.03) - 2 \cdot P \cdot \text{Vol}(K/S_t = 1)]. \quad (24)$$

The ATM implied volatility in point (i) is standard and needs no additional comment but the fact that our model predicts \mathcal{V}_t to be high when $\hat{\pi}_t \cdot (1 - \hat{\pi}_t)$ is high, as shown in the

⁹We follow Bates (1991) and set $K_{\text{put}} = S_t e^{(r-\delta)\tau} / 1.03$ and $K_{\text{call}} = S_t e^{(r-\delta)\tau} \times 1.03$ so that in Black-Scholes model, for instance, the P/C ratio would just equal 0.97.

bottom right panel of Figure 2. At the time-series level, an association between these two quantities is apparent in Panel C of Figure 4.

The quantity described in point (ii), the conditional skewness premium, is instead high when OTM puts trade at a premium relative to OTM calls. This occurs when the risk-neutral return density is negatively skewed, a situation that occurs when $\hat{\pi}_t$ is high (see Figure 3). Intuitively, when investors are confident that the fundamentals are strong, and hence the price-earnings ratio is high, they value highly the possibility that bad news will be disappointing and cause a sharp decline in the stock price. The top panel of Figure 5 plots the time-series of the skewness premium P_t/C_t from 1986 to 1996, along with its “theoretical” level $P_t^U(\hat{\pi}_t)/C_t^U(\hat{\pi}_t)$ implied jointly by the filtered beliefs $\{\hat{\pi}_t\}$ and the option pricing formula (17). We note that the empirical ratio fits our model fairly well except in the 1990-1991 period, when in the depths of a recession and negative earnings growth, our model implies a low $\hat{\pi}_t$ and hence a positive correlation between returns and volatility. Also as seen in the plot, the ratio trended up between October 1987 and 1994, in line with the observation by Jackwerth and Rubinstein (1995), that market participants seemed to be suffering from post-crash fears. Nonetheless, after peaking in 1994, the period of most rapid earnings growth, it then fell back in 1995 and 1996 as growth weakened, suggesting that factors other than post-crash fears played a role in determining the skewness premium. While our data runs until mid-1996, similarly constructed series on *Bloomberg* reveal that the put-call ratio also fell sharply in 1999 as earnings growth fell (see also Figure 4, Panel (B)). We finally note, that the behavioral finance explanation of the put-call ratio, (c.f. for e.g. Shefrin 1999) provides the opposite prediction to our model: following a sequence of strong earnings, individual investors become overconfident of future growth and bid up the relative prices of calls to puts.

Finally, the quantity in part (iii), the conditional kurtosis premium in option prices, also has a natural interpretation: During periods of high uncertainty (high $\hat{\pi}_t(1 - \hat{\pi}_t)$), investors are very uncertain about the growth of fundamentals, and hence are not likely to be surprised by very large or very small realizations of growth, and hence the revisions to their beliefs and stock prices, although rapid, are unlikely to be very large. Investors would therefore put higher odds on observing small rather than large price changes bidding up the implied volatility of ATM options relative to OTM puts and calls. Hence, Y_t is low when the kurtosis is low, and it

is decreasing with $\hat{\pi}_t(1 - \hat{\pi}_t)$.¹⁰ To see their positive association, in the bottom panel of Figure 5 we plot the time-series of the conditional kurtosis premium Y_t from 1986 to 1996, along with its “theoretical” level $Y_t^U(\hat{\pi}_t)$ implied jointly by the filtered beliefs $\{\hat{\pi}_t\}$ and the option pricing formula (17). Other authors have also found evidence of time-varying kurtosis in financial market prices (c.f., for e.g. Hansen 1994), even though alternative economic explanations of variation in kurtosis are hard to come by.

Panels A - C of Table 3 report the results of a time-series regression of the quantities in (i) to (iii) onto their theoretical counterparts constructed from the filtered beliefs $\{\hat{\pi}_t\}$ and the option pricing formula (17), as well as the filtered uncertainty level $\hat{\pi}_t \cdot (1 - \hat{\pi}_t)$ (or just $\hat{\pi}_t$ in case (ii)). The results are reported for the whole sample (first four columns) and for the ex-crash sample (last four columns), as described in Section 3.1.¹¹ In the interest of space, we only report results for the 3 months options. Very similar results were obtained for the 1 month and 6 months options.

A few comments are in order: First of all, taken collectively, the regression results offer strong supporting evidence that the filtered beliefs $\hat{\pi}_t$ are correlated with the time-varying fluctuations in the shape of the risk-neutral return density. In all panels and for both samples, we find the correct signs of the regression coefficients and a significance level of at least 5%. Second, when we control for lagged values of the dependent variable, we find that in all cases our “theoretical” explanatory variable retains its significance at the 5% level. In other words, serial autocorrelation does not seem to be the source of the regression results. This is especially true for the kurtosis premium, for which regression 11 shows only a small amount of variation from serial dependence.¹² Third, the R^2 's for the option price variables are between 8.2%

¹⁰Another way to look at the measure $Y_t(k)$, is that it be written as

$$Y_t = [\text{Vol.}(K/S_t = 1.03) - \text{Vol.}(K/S_t = 1)] - [\text{Vol.}(K/S_t = 1) - \text{Vol.}(K/S_t = .97)].$$

Therefore, Y_t is the second difference of volatilities across 3% strikes, and is therefore the volatility of volatility across strikes. The convexity measure can also be viewed as the implied volatility of a straddle, or a *butterfly spread* strategy. A butterfly spread strategy involves going long on a call option with a low and a high strike, and short on two call options with the a mid-way exercise price, and is therefore formally identical to the options used in constructing Y_t . The strategy leads to a profit if the stock price stays close to the mid-way exercise price, but gives rise to a small loss (due to the cost of the options) if there is a significant move in either direction (c.f., for example, Hull and White 1993).

¹¹As discussed, we report both sets of results, as the 87 crash generated an extremely high implied volatility (above 33%) which our model could not account for (see the top-right panel in Figure 2).

¹²We point out that since both π and the model-generated option prices depend on the first stage of the estimation in Table 2, the second stage regression coefficients in this table may be biased *downwards*, due to an error-in-variables problem (c.f., for e.g. Maddala and Nimalendran 1996). Therefore, correcting for the

(ATM volatility) and 12% (conditional skewness premium) for the whole sample, reaching up to 26% for ATM volatility in the ex-crash sample. We deem these R^2 's high in light of the fact that $\hat{\pi}_t$ is a macro-economic variable obtained solely from the history of earnings growth.

We conclude this section with two additional remarks about the relationship between earnings uncertainty and implied volatility in Table 3 (Panel A). First, while the improvement in explanatory power by including the model implied volatility is small, by conducting formal tests of Granger Causality (c.f. Section 11.2 Hamilton 1994) we find the interesting result that the model implied volatility Granger causes empirical implied volatility, while the converse is not true. Second, one may argue for a somewhat more direct economic mechanism that does not depend on investors' learning about fundamentals: stock market volatility is higher in periods of greater earnings volatility. We explored this channel by fitting various GARCH models to earnings growth, and found no support for it. In fact, the regression coefficient on the earnings volatility was insignificant and with a negative sign (see also David and Veronesi (2000) for similar evidence.)

3.5 Testing the Nine Conditional Moments of Option Prices

The previous section provided statistical evidence that the time-variation in the shape of the risk-neutral return density can be partly accounted for by changes in investors' beliefs about the economic prospects of a the economy. However, we have not yet formally tested whether the model-implied option prices, conditional on investors' beliefs, are "close" to the market traded option prices in the 1986-1996 period. This section provides such evidence. More specifically, given the set of parameters $\hat{\Theta}^U$ for the regime-switching model estimated in Table 2, let $C_\ell(S_t, K_\ell, t, T)$ and $C_\ell^U(S_t, K_\ell, t, T; \hat{\Theta}^U, \hat{\pi}_t, r_t, \delta_t)$, $\ell = 1, \dots, L$, be the observed call option prices, and the theoretical option prices from formula (17) conditional on $\hat{\Theta}^U$, S_t , $\hat{\pi}_t$, and r_t

bias would only strengthen the economic significance of model-based variables. Nonetheless, we perform two robustness checks: First, using the lagged independent variable as an instrument, we re-run each regression, and find that the model-based coefficients move up closer to one, as predicted by our theoretical model. Second, in the spirit of *robust estimation*, we sample 10,000 values of Θ^U from its asymptotic distribution based on the MLE, and for each value, we recompute the filtered belief series, π_t . Re-running regressions 1, 5, and 9, we find that regression 1 has a positive coefficient 98.1% of the time, regression 5, 100% of the time, and, regression 9, 95.5% of the time.

and δ_t . We define the pricing errors for calls (and analogously for puts) at time t for the option in class ℓ as:¹³

$$\epsilon_\ell(S_t, K_t, t, T; \hat{\pi}_t, r_t, \delta_t) = \frac{C_\ell^U(S_t, K_t, t, T; \hat{\Theta}^U, \hat{\pi}_t, r_t, \delta_t)}{C_\ell(S_t, K_t, t, T)} - 1, \quad (25)$$

We then combine the nine errors at each time t in the vector ϵ_t . To minimize notation, we suppress the arguments of the error. If $E[\epsilon_t | \hat{\pi}_t] = 0$, then by the law of large numbers, $\frac{1}{T} \sum_{t=1}^T \epsilon_t$ tends in probability to the zero vector. We therefore construct the test statistic

$$c = \left(\frac{1}{T} \sum_{t=1}^T \epsilon_t \right)' \cdot \Omega^{-1} \cdot \left(\sum_{t=1}^T \frac{1}{T} \cdot \epsilon_t \right) \quad (26)$$

where the variance-covariance matrix Ω is estimated using the Newey-West correction for heteroskedasticity and serial correlation (c.f. Equation 14.1.19 Hamilton 1994). Because the parameters of the fundamentals' processes, $\hat{\Theta}$, and the filtered belief process, $\{\hat{\pi}_t\}$, are *not* chosen to minimize c , under the null hypothesis of zero mean errors, c is distributed according to a $\chi^2(9)$. Similar statistics are constructed for subsets of the errors.

The chi-square statistics and their p-values for the 9 moments and their subgroupings by moneyness and maturity are shown in Table 4. We report numbers using 12 lags to adjust for autocorrelation; similar results are obtained using 6 or more lags. Of the individual moments, the three ATM put moments all have p-values in excess of 16%. 3-month OTM calls have a p-value of about 11%, but the other two maturities are considerably higher. OTM puts are the hardest to fit. As several other authors have noted (see e.g. Bates (1999)), OTM puts were expensive relative to most theoretical models. We find a p-value of only 8% for 3-month puts, but higher than 16% for the other two maturities. Combining all three maturities, gives a put p-value of 8%. The p-values of ATM and OTM calls are much higher. Across maturities, 1-month options have an overall p-value of a little over 10%, while the other two maturities have values of 39% and 75%. The overall p-value for all 9 moments is 52%. Therefore, we fail to reject the hypothesis that the option prices obtained from the filtered belief process with fixed earnings parameters prices the nine options considered with no conditional mean error.

¹³As discussed in Jacquier and Jarrow (2000), proportional (or relative) pricing errors ensure that low price, out-of-the-money options are not ignored in the diagnostic. In addition, they argue that for analyzing portfolio strategies, relative pricing error is the relevant measure of model risk.

4 Options' Implied Belief of an Earnings Expansion

Section 3 showed that filtered beliefs obtained from fitting a regime switching model to real earnings is able to partly explain the time-variation in the risk-neutral return density, as predicted by our model in Section 1. Implicit in this exercise, however, is the assumption that investors only use past earnings to update the probability π_t . In this section we relax this assumption and use the option pricing formula developed in proposition 2 to back-out investors' beliefs from traded options. The “distance” between the filtered beliefs and the options' implied beliefs would then reflect the additional sources of information that agents have while assessing the probabilities of a boom.

More specifically, given the set of parameters $\hat{\Theta}^U$ for the regime-switching model estimated in Table 2, at each date t we imply the belief level $\check{\pi}_t$ by non-linear least squares, to minimize the Proportional Root Mean Square Error (PRMSE) given by:

$$PRMSE_t = \sqrt{\frac{1}{L} \sum_{\ell=1}^L \epsilon_{\ell}(S_t, K_t, t, T; \check{\pi}_t, r_t, \delta_t)^2} \quad (27)$$

where $\epsilon_{\ell}(S_t, K_t, t, T; \check{\pi}_t, r_t, \delta_t)$ is defined in equation (25).

The resulting series $\{\check{\pi}_t\}$ for the period 1986:04 — 1996:05, extracted by minimizing (27) using all approximately 3 month option prices on the 15th of each month, is plotted in the top panel of Figure 6.¹⁴ For comparison, the filtered probability $\{\hat{\pi}_t\}$ is also plotted. The bottom panel of the same figure plots the options' implied uncertainty $\check{\pi}_t \cdot (1 - \check{\pi}_t)$. The two panels of this figure lend themselves to a nice interpretation of the events from 1986 to 1996. First, in the months preceding October 1987, the option-implied beliefs were highly volatile, showing swift changes in market participants' opinions. Notice that although there was no “recession” in 1986 as formally defined by e.g. the NBER, the filtered beliefs reveal that real earnings growth was at the bottom of its cycle in that year. It is likely that contradictory signals were received by market participants in those years. Second, the stock market crash of October '87 had a big impact on investors beliefs. From the bottom panel in Figure 6 we see that immediately after the crash, investors' options-implied uncertainty $\check{\pi}_t \cdot (1 - \check{\pi}_t)$ increased dramatically, reflecting the fear that the economy could be led to a recession by the

¹⁴Similar plots have been obtained when we use 1 months, 6 months and the whole set of options.

stock market crash. The uncertainty steadily declined as investors became more and more confident that the economic boom would persist. Third, the top panel in Figure 6 also show that the option-implied beliefs declined almost exactly at the start of the recession in 1990-91, and then recovered quickly at the beginning of 1991. This is consistent with the view that investors believed that the 1990-91 recession was not going to be very harsh, and that they soon received favorable signals about future growth, even though realized earnings growth in the period was still weak. We finally note that in the period between 1992 and 1996, a period of rapid earnings growth, both filtered and option-implied beliefs have hovered near almost perfect certainty of earnings being in the expansion state.

4.1 Specification Test from Options' Implied Beliefs

The time series of options' implied beliefs suggests a natural specification test of the uncertainty model described in section 1. In fact, we can test whether the options' implied beliefs $\{\tilde{\pi}_t\}$ retain the statistical properties that have to be satisfied by a belief process if the regime switching model in Assumptions 1-5 is true. More specifically, recall from the previous section that when we backed out these beliefs from traded options, we held the structural parameters $\hat{\Theta}^U$ constant as fitted in the regime-switching model in Table 2. If the model is correctly specified, Lemma 1 entails that the options' implied belief process, $\{\tilde{\pi}_t\}$, should satisfy the stochastic differential equation (3), that is, $d\pi_t = \mu_\pi(\pi_t; \Theta^U) dt + \sigma_\pi(\pi_t; \Theta^U) d\tilde{W}_t$.

In order to perform the test, we follow Diebold, Gunther, and Tay (1998), and compute the Kolmogorov's D_n statistic (c.f., for example Stuart and Ord 1991) from the transition CDF $F(\pi_{t+1}|\pi_t; \Theta^U)$ of the process (3) for π_t . We compute the latter by Monte-Carlo simulation.¹⁵ We then test if changes of the option' implied belief process, $\{\tilde{\pi}_t\}$ can be well approximated by $F(\pi_{t+1}|\pi_t; \Theta^U) = \int_0^{\pi_{t+1}} f(y|\pi_t; \Theta^U) dy$.

By the *probability-integral transformation*, if transitions of the implied beliefs process $\{\tilde{\pi}_t\}$ are generated by the CDF $F(\cdot)$, then the variable $z_t = F(\tilde{\pi}_{t+1}|\tilde{\pi}_t; \Theta^U)$ is distributed i.i.d. uniform (0,1). From the implied belief process, we write the sample CDF of the $\{z_t\}$ process as $S_n(z) = \frac{r}{n}$ for $z \in [z_{(r)}, z_{(r+1)})$, where $z_{(r)}$ are the order statistics, i.e. the observations are arranged so that $z_{(1)} < \dots < z_{(r)} < \dots < z_{(n)}$. The uniform CDF is simply $F_0(z) = z$, for

¹⁵For each initial belief in a grid of 200 values, we sampled 40,000 paths of eighty $N(0, \sqrt{dt})$ draws every month, where $dt = .001$.

$z \in [0, 1]$. Then, Kolmogorov's D_n statistic is given by

$$D_n = \sup_z |S_n(z) - F_0(z)|. \quad (28)$$

Rejection regions and p-values for sample sizes larger than 80 are well approximated by the formula:

$$\text{Prob}(D_n > zn^{-1/2}) = 2 \sum_{r=1}^n (-1)^{r-1} \exp(-2r^2 z^2). \quad (29)$$

Results for the D_n statistic are provided in Panel B of Table 5. As can be seen, we fail to reject with 95% confidence that the 1-month and 6-month implied belief series are generated by the hypothesized CDF with parameters $\hat{\Theta}^U$. The p-value for the 3-month beliefs is a little above 4%, and therefore leaves room to believe that it is generated by a different set of parameters. Diebold, Gunther, and Tay (1998) suggest that since the test is stringent in fitting the entire range of the CDF, its constructive to test separately for independence of the $\{z_t\}$ series, and its range. If the transitions are consistently generated by some other process in sub-samples, then one would expect to find that the powers of $\{z_t\}$ are auto-correlated. We test the first two powers using conventional Ljung-Box statistics, and find the p-values for no autocorrelation of almost 1. Therefore, the implied-beliefs do not seem to be conditionally inconsistent with our model in any sub-periods. The second diagnostic is to look at histograms of $\{z_t\}$ to see if all parts of the distribution are sufficiently covered. While we do not show these histograms, we find that the observations are fairly well distributed on the unit interval with the possible exception of the very lowest sub-interval. That is, the implied belief process in the sample did not have many much-smaller-than-expected outcomes. This is consistent with the time-series of the filtered and implied belief process: particularly in the 1991-1992 period, the filtered beliefs of an earnings expansion were close to zero, while the option implied beliefs moved quite rapidly closer to the one range. Over the remainder of the sample, there were few periods of very weak earnings growth, and therefore, sharp transitions to very low beliefs were hard to come by in the sample.

4.2 Fitting Alternative Models to Option Prices

In the previous section we obtained the time series of option-implied beliefs $\{\tilde{\pi}_t\}$ from the minimization of the PRMSE in equation (27). An additional output of this exercise is the

value of the minimal $PRMSE_t$ for each month t in the sample and for each option’s maturity $T = 1$ months, 3 months and 6 months. In order to gauge whether these pricing errors are “large” under some metric, we now compare them to those obtained from more standard models, such as the Black-Scholes (BS) model or the stochastic volatility (SV) model of Heston (1993). In addition, we will also look at the ability of each model to produce low hedging errors and to forecast future realized volatility. The latter exercise is particularly important in our context as it is a measure of dynamic consistency for the model.

We remark that Corollary 3 implies that in a single-state world of expected earnings growth the BS option pricing formula would obtain. Therefore, comparing the pricing errors of the U model to those of BS amounts to the comparison of a one-state model versus two-state model of expected earnings growth. Table 2 provides evidence that based on the time series of earnings growth, a one-state model can be rejected against the two-state model. In this section we will see that based on the errors of fitted option prices, we will reach the same conclusion.

For completeness, Appendix C contains the description of the SV model as well, a brief detail of the estimation method and estimate results. In a nutshell, let Θ^{SV} denote the structural parameters of the SV model. At each date t we obtain simultaneously an estimate Θ^{SV} and the volatility level V_t by non-linear least squares, through the minimization of the cross-sectional PRMSE. To fit the BS model we similarly find the level of volatility at that date that provides the best fit for all traded options.

4.2.1 In- and Out-of-Sample Fits

Table 6 contains the results of the relative empirical performance of the three models (U, BS, and SV) for two sample periods (“Full Sample” and “Ex- Crash”). We separate out the months following the crash because the calibrated U model permits a maximum implied volatility of 24 percent (see Figure 2) while actual implied volatility hovered above 30 percent in this period. The two competing models — BS and SV — are unconstrained in the level of implied volatility. Bold-faced items in each category denote ‘minimal’ PRMSE when the lowest mean item is different statistically with 95% confidence from the next (non-bold) lowest mean item under the t-statistic for difference in means: $t = \sqrt{T} \cdot (\bar{x}_1 - \bar{x}_2) / \sqrt{s_1^2 + s_2^2}$. In-sample fits are for the 15th (first following trading day) of each month, while out-of-sample fitting performance

of the three models is evaluated by implying parameters on the 15th (first following trading day) and computing the PRMSE on the next trading day.

Panels A and B show that the U-model's in-sample and out-of-sample fits are close to that of the benchmark SV model and better than those of the BS model. More specifically, for the full sample, the U-model has the lowest errors for 3-month options, SV has the lowest errors for 1-months options, and the two models are in a statistical dead heat for 6-month options. For the ex-crash sub-sample, the U-model has the better fits for 3-month options, the only cases where the two models have statistically different errors. The U-model has always a better in-sample fit than BS.

We remark that the monthly refit of the parameters Θ^{SV} of the SV model produces time-inconsistent and highly unstable parameters, as shown in Table 7 in Appendix C. Indeed, an important distinction between the stochastic volatility model of Heston (1993) and the U-model is that the former assumes a constant correlation between returns and volatility, a fact that is at odds with the evidence (see Introduction and the bottom panel in Figure 1). In contrast, the U-model obtains a time varying covariance between returns and volatility because of changes in the probability of being in an expansionary phase of earnings growth.

4.2.2 Hedging Performance

We now turn to another measure often used to compare option pricing models. Following the analysis in e.g. Bakshi, Cao, and Chen (1997), consider a short position in a call option at time t with τ periods to maturity which is hedged using a position $X_{S,t}$ in stocks and $X_{0,t}$ cash. This leads to the hedging error at $t + \Delta t$ is given by

$$H(t + \Delta t) = X_{S,t}^* S(t + \Delta t) + X_{0,t} e^{r\Delta t} - C(t + \Delta t, \tau - \Delta t). \quad (30)$$

The variance of the portfolio at $t + \Delta t$ is minimized by setting $X_{S,t}^* = \text{Cov}_t[dS_t, dC(t, \tau)] / \text{Var}_t[dS_t]$ and its actual value depends on the pricing model. Given the results in Bakshi, Cao, and Chen (1997) and in Corollary 4 we have the following:

- (i) Black Scholes: $X_{S,t}^{BS} = \Delta_{S,t}^{BS} = \frac{\partial C^{BS}}{\partial S}$ where $\Delta_{S,t}^{BS}$ is the BS delta;
- (ii) Heston (1993) SV model: $X_{S,t}^{SV} = \Delta_{S,t}^{SV} + \frac{\rho\sigma_v C_{V,t}^{SV}}{S_t}$ where $\Delta_{S,t}^{SV} = \partial C^{SV} / \partial S$ and $C_{V,t}^{SV} = \partial C^{SV} / \partial V$ are in Bakshi, Cao, and Chen (1997);

(iii) U-model: $X_{S,t}^U = \Delta_{S,t}^U + \frac{C_{\pi_t}^U}{S_t} \cdot \frac{A_3(\pi_t)}{A_2(\pi_t)}$ where $\Delta_{S,t}^U = \partial C^U / \partial S$ and $C_{\pi_t}^U = \partial C^U / \partial \pi$ are obtained in Corollary 4 and $A_2(\pi_t)$ and $A_3(\pi_t)$ are in eqs. (40) and (41) in Appendix 2.

In all cases we have that the position in bonds is $X_{0,t} = C(t, \tau) - X_{S,t}$. As can be seen the hedge ratios for the SV and U models have similar forms: each incorporates a direct effect due to the change in the stock price, Δ_S , and an indirect effect due to the change in the other state variable in the model. The indirect effect has the sign determined by the constant ρ_v in the SV model, but, as discussed earlier, the sign of the correlation between the stock price and beliefs is negative when π_t is larger than 0.5 and positive for π_t below 0.5.

To compare the empirical hedging performance of the three models, we use the implied parameters for each model at date t to use for computation of optimal hedge ratios and compute the errors at $t + \Delta t$, $H(t + \Delta t)$. The hedging errors are provided in panel C of Table 6. As can be seen, none of the models clearly dominates under this metric. For the full-sample, U has the minimal hedging error for 3-months options, while SV has the minimal errors for 1-month and 6-month. For the ex-crash sample, again U has the minimal error for 3-months, while the errors of all three models are virtually the same for 6-month options, and U and SV “tie” for 1-month options.

4.2.3 Variance Forecasting

The one day out-of-sample tests performed in subsection 4.2.1 may hide longer-term parameter instability in these models. Indeed, several authors have reported inconsistencies between parameters for the SV model that fit options data well and those estimated from historical returns (c.f. Bakshi, Cao, and Chen 1997, Bates 1999). In this section, we use the option-implied parameters of each model and formulate volatility forecasts. To compare the relative performance of the models, following Anderson and Bollerslev (1998), we run predictive regressions first suggested in Mincer and Zarnowitz (1969) of the form:

$$\frac{1}{T-t} \int_t^T r_s^2 ds = \beta_0 + \beta_1 \cdot E \left[\int_t^T \sigma_s^2 ds | \mathcal{F}_t \right], \quad (31)$$

and report $1 - \bar{R}^2$, a measure of the unexplained variation in future volatility.

The predictor on the right-hand side, $E \left[\int_t^T \sigma_s^2 ds | \mathcal{F}_t \right]$, can be computed explicitly for each option’s pricing model. For the BS model, $E \left[\int_t^T \sigma_s^2 ds | \mathcal{F}_t \right] = \bar{V}^{BS}(t, T; \sigma_t) = \sigma_t$. For the

SV model, Bates (1999) (Eq. (19)) shows that

$$\bar{V}^{SV}(t, T; \Theta^{SV}, V_t) = E_t \left[\int_t^T V_s ds \right] = \frac{\theta_v}{(\kappa_v - \lambda_V^{SV})} + \left(V_t - \frac{\theta_v}{(\kappa_v - \lambda_V^{SV})} \right) \cdot \frac{1 - e^{-(\kappa_v - \lambda_V^{SV})(T-t)}}{(\kappa_v - \lambda_V^{SV})(T-t)}. \quad (32)$$

The speed of mean reversion of volatility under the risk-neutral and objective measures differ by the market price of volatility risk. We note that because the measurement yardstick is the regression \bar{R}^2 , and not the distance from the predicted variance, the information content of the volatility forecasts is quite similar for a wide range of λ_V^{SV} parameters. For example, we obtained almost the same results by using λ_V^{SV} in the range $(-1, 1)$, and report results only for the case $\lambda_V^{SV} = 0$.¹⁶

Finally, for the U-option model the variance forecast is given by equation (22) in Proposition 4. Variance forecasts for horizons between 1-month and 6-months for the calibrated set of structural parameters, $\hat{\Theta}^U$ (Table 2), are provided in Figure 7. Because beliefs tend to be pulled towards the two humps of the stationary distribution, thus passing through the point of maximum uncertainty, variance forecasts are mean-reverting. When $\tilde{\pi}_t = 0.5$, forecasted variance is downward sloping. For $\tilde{\pi}_t = 0.1, 0.9$, the forecasts are upward sloping, although the level is higher at $\tilde{\pi}_t = 0.1$, as discussed in subsection 3.3.

The predictive regressions (31) are implemented as follows: First, as in e.g. Schwert (1989), the left hand side quantity $\int_0^T r_t^2 dt$ is replaced by its proxy obtained from daily returns, $\sum_{t=1}^T r_t^2$. Second, for each date we imply the parameters of the relevant model, for example $\{\Theta^{SV}, V_t\}$ for the SV model and $\{\pi_t\}$ for the U-model, and therefore obtain a time series of variance forecasts $E \left[\int_t^T \sigma_s^2 ds | \mathcal{F}_t \right]$. The regression (31) is then performed and the resulting $1 - \bar{R}^2$ reported.

The relative performance of the models is provided in panel D of Table 6. For the full sample, for each of the maturities considered, the variance forecast is minimal for the U model, and the highest for the SV model. For the ex-crash sample, the U model has the minimal error for 3-month and 6-month options, while the BS-model has the minimal error for 1-month options. Under each sub-category, the variance forecasting error of the U model is statistically lower than that of the SV model.

¹⁶We find that the correlation between monthly real consumption growth and realized S&P 500 variance between 1960 and 1998:9 was -0.10 , suggesting $\lambda_V^{SV} < 0$.

For 1-month options, as suggested by Anderson and Bollerslev (1998), the \bar{R}^2 are as high as 30 percent in the ex-crash sample. The fits deteriorate over longer horizons, but for the U model we find \bar{R}^2 as high as 14 percent for 3-month options and 6-month options for the full sample. The fluctuating uncertainty about earnings before the crash (see Figure 6) was quite successful in providing a high volatility forecast over these longer horizons. The forecast accuracy of the U model at these maturities is actually lower in the ex-crash sample.

The relatively weak performance of the SV model under this metric confirms the findings of others (c.f., for e.g. Bakshi, Cao, and Chen 1997, Bates 1999). In addition, we see in Table 7 in Appendix C that the speed of mean-reversion, κ_v , must be quite volatile to match the time-variation in the slope of the implied volatility curve, thus making the volatility forecasts more volatile [see eq. (32)]. The U model is able to fit deeper OTM options prices because of the hidden jump in the earnings drift rate, which translates under the fast learning set of parameters to a stock price process with almost jump-like moments for short horizons. Yet because these jumps are calibrated to be infrequent, the volatility of the process is persistent enough to provide superior volatility forecasts.

5 Conclusions

This paper explicitly links the dynamics of option prices to the dynamics of market participants' beliefs about real earnings growth. We develop a model for fundamentals and an option pricing formula that suggests a strong link between the moments of the risk-neutral return distribution, which affects option prices, and economic conditions, summarized by market participants' beliefs on future earnings growth. More specifically, our model shows that during good times, bad earnings news are likely to increase the fundamental uncertainty on economic growth, thereby endogenously producing a negative correlation between returns and volatility of returns. Hence, good times are characterized by negatively sloped implied volatilities, or, more generally, a negatively skewed and leptokurtic risk-neutral return distribution. In contrast, periods of high economic uncertainty, which may follow sequences of bad earnings news, are characterized by high implied volatilities, but essentially flat implied volatility curves. When uncertainty is high, large surprises on either direction are unlikely, generating a symmetric risk-neutral return density with low kurtosis. Finally, periods of weak earnings growth are characterized by a positive implied volatility curve. The model is then able to

provide an economic justification to the documented high variability of the risk-neutral return distribution.

When calibrated to real-earnings data, the model also produces economically significant effects. The fitted model produces average option prices that are in line with the ones traded in the market. In addition, also the dynamics of earnings and hence of filtered probabilities is consistent with the dynamics of the risk-neutral return distribution, as summarized by option prices. More specifically, we show that the calibrated model produces a time series of options' implied volatilities, skewness premia (ratios of OTM put-to-call) and kurtosis premia (implied volatilities of butterfly spreads) that are significantly correlated with the ones obtained from traded options.

We finally use the model to back-out market participants' beliefs about real earnings growth from options data. These "options' implied" beliefs reveal that investors' opinions vacillated rapidly before the crash of 1987, remained highly uncertain for a year afterwards, and, with the exception of a few months prior to the 1990-91 recession, were quite confident of rapid growth until 1996. We test whether these options' implied beliefs are consistent with our regime switching model and fail to reject. This time-variation pattern of beliefs is largely consistent with the implied parameters from Heston's (1993) stochastic volatility model, producing similar in- and out-of-sample fits and hedging errors. But, because the structural parameters of our model are held constant over time, and transitions in beliefs are well-specified, the model performs better in forecasting future realized volatility.

This paper focused on a two-state, regime switching model, which is appropriate for the investigation of the relationship between option prices and the phases in the cycle of earnings growth. A challenge for future research is to obtain a manageable option pricing formula for a number of states greater than two. This is likely to yield interesting results from an empirical standpoint, because *(i)* it makes it possible to model joint processes for real earnings and other important variables, such as inflation, and hence obtain a richer set of implications on option prices (for example: how does changes in expected inflation affect option prices?); *(ii)* it allows for the study of the possible occurrence of "Peso Problem" biases in the dynamics of option prices and *(iii)* in the same spirit as Section 4.2 it makes it possible to "back out" from traded option prices a full distribution of investors' beliefs about future real earnings growth and additional important state variables, such as future inflation.

Appendix A

Proof of Proposition 1: (a) From $m_t = e^{-\phi t} C_t^{-\gamma}$ and Ito's lemma, we find immediately that

$$\frac{dm_t}{m_t} = - \left(\phi + \gamma g - \frac{1}{2} \gamma (\gamma + 1) \sigma_C \sigma'_C \right) dt + \sigma_m d\widetilde{\mathbf{W}}_t.$$

where $\sigma_m = -\gamma \sigma_C$. Hence, we obtain immediately that the real interest rate $r = -E_t [dm_t/m_t] = \phi + \gamma g - \frac{1}{2} \gamma (\gamma + 1) \sigma_C \sigma'_C$.

(b) Let $\pi_{1,t} = (1 - \pi_t)$ and $\pi_{2,t} = \pi_t$. Rewrite (6) as

$$S_t = D_t E_t \left[\int_t^\infty \frac{m_s D_s}{m_t D_t} ds \right] = D_t \sum_{i=1}^2 E_t \left[\int_t^\infty \frac{m_s D_s}{m_t D_t} ds | \theta_t = \theta_i \right] \pi_{it}, \quad (33)$$

Define $M_s = m_s D_s$ to find

$$dM_t = M_t [-r + \theta_t + \sigma_m \sigma_D] dt + M_t (\sigma_D + \sigma_m) dW_t.$$

For convenience, define $\widehat{\theta}_t = r - \theta_t - \sigma_D \sigma'_m = \mu_S - \theta_t$ and $\sigma_M = (\sigma_D + \sigma_m)$. Then,

$$\frac{M_s}{M_t} = \exp \left(\int_t^s -\widehat{\theta}_u du - \frac{1}{2} \sigma_M \sigma'_M (s - t) + \sigma_M (W_s - W_t) \right). \quad (34)$$

By substituting (34) into (33), we obtain the same set up as in the proof of proposition 3.1 in Veronesi (2000).

By following the same steps, we find that we can define a matrix $\mathbf{A} = -\mathbf{\Lambda} + \text{diag}(\widehat{\theta}_1, \widehat{\theta}_2)$ to obtain the two constants C_1 and C_2 satisfying

$$C_i = E_t \left[\int_t^\infty \frac{M_s}{M_t} ds | \theta_t = \theta_i \right] \quad (35)$$

by

$$\mathbf{C} = \mathbf{A}^{-1} \mathbf{1}_n. \quad (36)$$

Notice finally that if the condition of proposition 1 (b) in the text is satisfied, then $\widehat{\theta}_t$ is always positive. From (34), this implies that C_i are well defined and positive. Finally, by solving explicitly for C_1 and C_2 we find that for $i, j = 1, 2$ with $i \neq j$

$$C_i = \frac{\mu_S - \theta_j + \lambda_{12} + \lambda_{21}}{(\mu_S - \theta_1 + \lambda_{12})(\mu_S - \theta_2 + \lambda_{21}) - \lambda_{12}\lambda_{21}}.$$

Hence, A and B in the text are given by $A = C_2 - C_1$ and $B = C_1$.

(c) By Ito's lemma

$$dS_t = S_t \left(\bar{\theta}_t + \frac{Ak(\bar{\pi} - \pi_t)}{\pi_t A + B} + \frac{A\pi_t(1 - \pi_t) \mathbf{b}\sigma'_D}{\pi_t A + B} \right) dt + S_t \left(\sigma_D + \frac{A\pi_t(1 - \pi_t) \mathbf{b}}{\pi_t A + B} \right) d\widetilde{\mathbf{W}}_t.$$

Recall that $\mathbf{b}\sigma'_D = ((\theta_2 - \theta_1), 0) (\sigma'_D, \sigma'_C)^{-1} \sigma'_D = ((\theta_2 - \theta_1), 0) (\sigma'_D, \sigma'_C)^{-1} (\sigma'_D, \sigma'_C) (1, 0) = \theta_2 - \theta_1$, so that little algebra yields

$$E_t \left[\frac{dS_t}{S_t} \right] = \frac{1}{\pi_t A + B} [\theta_1 B + B\pi_t(\theta_2 - \theta_1) + Ak(\bar{\pi} - \pi_t) + A\pi_t\theta_2]. \quad (37)$$

Notice now the following two equalities: (I) by recalling that $k = \lambda_{12} + \lambda_{21}$ and $\bar{\pi} = \lambda_{12}/k$

$$(1 - \pi_t, \pi_t) \mathbf{A}\mathbf{C} = Ak(\bar{\pi} - \pi_t),$$

and (II) using the equality implicit in (36) $(1 - \pi_t, \pi_t) \mathbf{A}\mathbf{C} = 1$ where $\mathbf{A} = -\mathbf{\Lambda} + \text{diag}(\hat{\theta}_1, \hat{\theta}_2)$ we find

$$(1 - \pi_t, \pi_t) \mathbf{A}\mathbf{C} = -\pi_t C_2 \theta_2 - (1 - \pi_t) C_1 \theta_1 + \mu_S \pi_t C_2 + \mu_S (1 - \pi_t) C_1 - 1.$$

Hence, by substituting $Ak(\bar{\pi} - \pi_t) = -\pi_t C_2 \theta_2 - (1 - \pi_t) C_1 \theta_1 + \mu_S \pi_t C_2 + \mu_S (1 - \pi_t) C_1 - 1$ into (37) and using the notation $A = C_2 - C_1$ and $B = C_1$ we find that most elements cancel leaving

$$\begin{aligned} E_t \left[\frac{dS_t}{S_t} \right] &= \frac{1}{\pi_t A + B} [\mu_S ((C_2 - C_1) \pi_t + C_1) - 1] \\ &= \mu_S - \frac{1}{\pi_t A + B} = \mu_S - \delta_t. \end{aligned}$$

As an equilibrium restriction, it is also immediate to check that

$$\mathbf{b}\sigma'_C = ((\theta_2 - \theta_1), 0) (\sigma'_D, \sigma'_C)^{-1} (\sigma'_D, \sigma'_C) (0, 1) = 0,$$

and hence,

$$\gamma \text{Cov}(dS_t/S_t, dC_t/C_t) = \gamma \sigma_{S,t} \sigma'_C = \gamma \sigma_D \sigma'_C = -\sigma_D \sigma'_m = \mu_{S_t} - r.$$

■

Proof of Corollary 1: The result follows from the fact that since $\sigma_\pi(\pi_t) = \pi_t(1 - \pi_t)(\theta_2 - \theta_1, 0) \cdot$

$(\sigma'_D, \sigma'_C)^{-1}$ and $\sigma_C = (0, 1) \cdot \begin{pmatrix} \sigma_D \\ \sigma_C \end{pmatrix}$, we have

$$\text{Cov}_t(dC_t/C_t, d\pi_t) = \sigma_C \cdot \sigma'_\pi(\pi_t)$$

$$= \pi_t(1 - p_{i_t})(0, 1) \cdot \begin{pmatrix} \sigma_D \\ \sigma_C \end{pmatrix} \cdot \begin{pmatrix} \sigma_D \\ \sigma_C \end{pmatrix}^{-1} \cdot (\theta_1 - \theta_2, 0) = 0$$

■

Proof of Corollary 2: The volatility function V can be written as $V = \sigma_D \sigma'_D + K_1 X + K_2 X^2$ where $X = g(\pi) = \frac{\pi(1-\pi)}{\pi A + B}$, $K_1 = 2A(\theta_2 - \theta_1)$ and $K_2 = A^2(\theta_2 - \theta_1)^2 \mathbf{b}\mathbf{b}'$. Since $0 < \pi < 1$, we have that the domain of X is the interval $[0, \bar{X}]$ for some \bar{X} . Since both K_1 and K_2 are positive, V is increasing on $[0, \bar{X}]$ and hence it has a unique maximum at \bar{X} . In addition, we see that the function $X = g(\pi)$ has negative second derivative, implying that it has a unique maximum as well. Hence, we can define $\hat{\pi}$ such that $\bar{X} = \max_{\pi} g(\pi) = g(\hat{\pi})$. In conclusion, the function V has a unique maximum at $\hat{\pi}$. Since the function is continuous, we must have $V_{\pi} < 0$ if and only if $\pi > \hat{\pi}$. ■

Proof of Proposition 2: The proposition follows from standard no-arbitrage arguments. See e.g. Duffie (1996). ■

Appendix B

Proof of Proposition 3 The Fourier Transform of the risk-neutral return density function over primitive states can be characterized by the usual Fokker-Planck backward equation (c.f., for example Karlin and Taylor 1982), which is the fundamental PDE (16) along with the initial condition $f(z_t, \pi_t, 0) = \delta_{z-z_t}$, where δ stands for the Dirac function, and the last argument of $f(\cdot, \cdot)$ denotes time to maturity. Substituting $z = \log S$ into (16), and letting $\tau \equiv T - t$ be the time to maturity, implies that the fundamental PDE can be written as

$$\begin{aligned} r(\pi)f &= -\frac{\partial f}{\partial \tau} + \frac{\partial f}{\partial z} \left[r - \delta(\pi) - \frac{1}{2} \left(\sigma_D \sigma'_D + \frac{A^2 \pi^2 (1-\pi)^2 b b'}{(\pi A + B)^2} + 2 \frac{A \pi (1-\pi)(\theta_2 - \theta_1)}{\pi A + B} \right) \right] \\ &+ \frac{\partial f}{\partial \pi} ((\lambda_{12} + \lambda_{21})(\pi^* - \pi)) \\ &+ \frac{1}{2} \frac{\partial^2 f}{\partial z^2} \left(\sigma_D \sigma'_D + \frac{A^2 \pi^2 (1-\pi)^2 b b'}{(\pi A + B)^2} + 2 \frac{A \pi (1-\pi)(\theta_2 - \theta_1)}{\pi A + B} \right) \\ &+ \frac{\partial^2 f}{\partial z \partial \pi} \left(\frac{\pi^2 (1-\pi)^2 A b b'}{\pi A + B} + \pi (1-\pi) \sigma_D b' \right) + \frac{1}{2} \frac{\partial^2 f}{\partial \pi^2} \pi^2 (1-\pi)^2 b b', \end{aligned} \quad (38)$$

where the dividend yield, $\delta(\pi) = 1/(\pi A + B)$.¹⁷ Define

$$A_1(\pi) = (\lambda_{12} + \lambda_{21}) \cdot (\pi^* - \pi), \quad (39)$$

$$A_2(\pi) = \frac{A^2 bb' (1 - \pi)^2 \pi^2}{(\pi A + B)^2} + \sigma_D \sigma'_D + \frac{2 A (1 - \pi) \pi (\theta_2 - \theta_1)}{\pi A + B}, \quad (40)$$

$$A_3(\pi) = \frac{A bb' (1 - \pi)^2 \pi^2}{\pi A + B} + (1 - \pi) \pi (\theta_1 - \theta_2), \quad \text{and}, \quad (41)$$

$$A_4(\pi) = bb' \pi^2 (1 - \pi)^2. \quad (42)$$

Taking a Fourier Transform on both sides of the PDE (38) with respect to z and using the frequency variable ω_1 and $i = \sqrt{-1}$, implies

$$\begin{aligned} r\bar{f} &= -\frac{\partial \bar{f}}{\partial \tau} - i\omega_1 \bar{f} \left(r - \delta(\pi) - \frac{1}{2} A_2(\pi) \right) + \frac{\partial \bar{f}}{\partial \pi} A_1(\pi) \\ &\quad - \frac{\omega_1^2}{2} \bar{f} A_2(\pi) - i\omega_1 \frac{\partial \bar{f}}{\partial \pi} A_3(\pi) + \frac{\partial^2 \bar{f}}{\partial \pi^2} \frac{1}{2} A_4(\pi), \end{aligned} \quad (43)$$

with initial condition $\bar{f}(\omega_1, \pi, 0) = e^{i\omega_1 z_t}$. Now taking another Fourier Transform on both sides of the PDE (43) with respect to the time variable and using ω_2 as the frequency variable implies that

$$-e^{i\omega_1 z_t} = \hat{f} \left(-r + i\omega_2 - i\omega_1 \left(r - \delta(\pi) - \frac{1}{2} A_2(\pi) \right) - \frac{\omega_1^2}{2} A_2(\pi) \right) + \hat{f}_\pi (A_1(\pi) - i\omega_1 A_3(\pi)) + \frac{1}{2} \hat{f}_{\pi\pi} A_4(\pi). \quad (44)$$

Using the frequency variables, ω_1 and ω_2 , as parameters, we further define

$$\begin{aligned} h_1(\pi) &= \frac{1}{2} A_4(\pi), \\ h_2(\pi) &= A_1(\pi) - i\omega_1 A_3(\pi), \quad \text{and}, \\ h_3(\pi) &= -r + i\omega_2 - i\omega_1 \left(r - \delta(\pi) - \frac{1}{2} \omega_1^2 A_2(\pi) \right) - \frac{\omega_1^2}{2} A_2(\pi). \end{aligned}$$

(44) can also be written concisely as $-e^{i\omega_1 z_t} = h_1(\pi) \hat{f}_{\pi\pi} + h_2(\pi) \hat{f}_\pi + h_3(\pi) \hat{f}$, the BVP in the statement of the proposition. The function $\bar{f}(z_t, \omega_1, \pi_t, T - t)$ can be obtained from the solution of this BVP by using the formula

$$\bar{f}(z_t, \omega_1, \pi_t, T - t) = \frac{1}{\Pi} \int_0^\infty \text{Re} \left[e^{-i\omega_2(T-t)} \hat{f}(z_t, \omega_1, \pi_t, \omega_2) \right] d\omega_2.$$

¹⁷To fit option prices in Section 4.2, we will instead use a constant dividend yield, δ_t to match the dividend yield observed on the S&P 500 at the date of the traded option.

Finally, using Theorem 1 in Bakshi and Madan (2000), by a second numerical Fourier inversion (shown below Proposition 3), we obtain the decomposition of the option price into primitive securities. This is completely analogous to the solutions for option pricing with stochastic jumps and stochastic volatility such as in Heston (1993), Bates (1996), Bakshi, Cao, and Chen (1997), and Scott (1997). ■

An M Term Polynomial Approximate Solution to BVP (18).

Looking ahead, it is evident that the solution of the system of difference equations that will solve the BVP (provided at the end of this proof) is homogeneous of degree 1 in the forcing term h_0 . Anticipating a future Fourier inversion with respect to ω_2 , we first solve the slightly simpler equation $-1 = h_1(\pi)\hat{f}_{\pi\pi} + h_2(\pi)\hat{f}_\pi + h_3(\pi)\hat{f}$. Due to homogeneity, the inverse Fourier transform of the solution of the first equation with respect to ω_2 , will coincide with $-e^{i\omega_1 z_t}$ times the Fourier inversion with respect to ω_2 of the solution of the latter equation (the logic is simply that $-e^{i\omega_1 z_t}$ will pass out of the integral for the Fourier inversion with respect to ω_2).

We solve $-1 = h_1(\pi)\hat{f}_{\pi\pi} + h_2(\pi)\hat{f}_\pi + h_3(\pi)\hat{f}$ with a series solution about an ordinary point of the BVP (c.f., for example, Rainville 1964, Boyce and DiPrima 1997). Because $A_4(0) = A_4(1) = 0$, 0 and 1 are singular points of the BVP. However, 0.5 is an ordinary point of the BVP, and the solution can be written in the form

$$\hat{f}(\omega_1, \pi, \omega_2) = \sum_{n=0}^{\infty} a_n \cdot (\pi - .5)^n. \tag{45}$$

Define the functions $p(\pi) = h_2(\pi)/h_1(\pi)$, and $q(\pi) = h_3(\pi)/h_1(\pi)$. By Theorem 5.3.1 in Boyce and DiPrima (1997), the radius of convergence of the series solutions is at least as large as the minimum of the radii of convergence of the series for p and q .¹⁸ It is well known that when $p(\pi)$ and $q(\pi)$ are ratios of complex polynomials (as is the case), then convergent series exist for these functions at π_0 when the denominator series of p and q does not vanish, i.e., $h_1(\pi_0) = \frac{1}{2}A_4(\pi_0) \neq 0$. Further, the radius of convergence of p and q is at least as large as the distance from π_0 to the nearest zero of the denominator series. Therefore, a series solution of the form in (45) converges at least for all $\pi \in (0, 1)$. The difference equation to be satisfied by the a_n , for $n = 2, \dots, \infty$, is provided below.

David (1997) (Property 1) showed that the updating process $\{\pi_t\}$ has ‘entrance’ boundaries at 0 and 1; neither boundary can be reached from the interior of the state space, but it is possible to consider processes that begin there. In such cases, it is impossible to impose arbitrary boundary conditions on the BVP (44) at 0 and 1. The boundary conditions are determined implicitly by formulating the same BVP as $\pi \rightarrow 0, 1$. Since $h_1(0) = h_1(1) = 0$, to satisfy the entrance boundary conditions at 0 and 1,

¹⁸This theorem has been attributed to Immanuel Lazarus Fuchs by several authors. See also Theorem 20 in Rainville (1964).

the coefficients a_n , $n = 1, \dots, M$, must also satisfy the equations:

$$h_2(0) \cdot \hat{f}_\pi(0) + h_3(0) \cdot \hat{f}(0) = -1, \quad (46)$$

$$h_2(1) \cdot \hat{f}_\pi(1) + h_3(1) \cdot \hat{f}(1) = -1. \quad (47)$$

Because we approximate the series solution (45) by a finite expansion $\hat{f}^M = \sum_{n=0}^M a_n(\omega_1, \omega_2) \cdot (\pi - .5)^n$, substituting $\hat{f}^M(0) = \sum_{n=0}^M a_n \cdot (-.5)^n$, and $\hat{f}^M(1) = \sum_{n=0}^M a_n \cdot (.5)^n$ into (46) and (47) provide two further equations in a_n , $n = 0, \dots, M$. Therefore, overall there are M equations in m unknowns, and since none of the equations is redundant, they permit a unique solution of (a_0, a_1, \dots, a_M) .

The coefficients a_n , $n = 0, \dots, M$, satisfying the difference equation shown below, are each polynomials in ω_1 and ω_2 , and we can write these as functions $a_n(\omega_1, \omega_2)$. Performing a term-by-term numerical Fourier inversion (the terms $(\pi - .5)^n$ pass out of the integration) of the solution (45) with respect to ω_2 using the Fourier inversion formula provides

$$b_n(\omega_1, T - t) = \frac{1}{\Pi} \int_0^\infty \text{Re} \left[e^{-i\omega_2(T-t)} a_n(\omega_1, \omega_2) \right] d\omega_2. \quad (48)$$

Overall, the M -term approximate solution to the Fourier Transform of the state return density function is given by¹⁹

$$\bar{f}^M(\omega_1, \pi, T - t; z_t) = e^{i\omega_1 z_t} \sum_{n=0}^M b_n(\omega_1, T - t) (\pi - .5)^n. \quad (49)$$

Coefficients for the Difference Equation

We explicitly provide the difference equation that must be satisfied by the coefficients a_n , for $n = 2, \dots, M$. in the finite approximation of the series solution (45). We start by making the finite Taylor Series expansions of the coefficient functions of the BVP (44). We define (h_{iq}) for $q = 1$ to $q = Q$ as the vector of coefficients of $h_i \cdot (\pi A + B)^2$ for $i = 1, 2, 3$. It can be easily verified that Q equals 7, 6, and 5 respectively for h_1 , h_2 , and h_3 respectively. Similarly h_{0q} , $q = 1, 2, 3$ are the coefficients of the expansion of $-(\pi A + B)^2$. To satisfy the BVP (44) standard series solutions methods (c.f., for example, Rainville 1964, Boyce and DiPrima 1997) imply that the coefficients must satisfy the following recursive system of equations.

$$h_{01} = h_{11} \cdot 2 \cdot a_2 + h_{21} \cdot a_1 + h_{31} \cdot a_0,$$

¹⁹Writing the Fourier Transform in this form rather than performing a double Fourier inversion on the solution (45) will be advantageous in the formulation of the partial derivatives of the option price with respect to the stock price and beliefs.

$$\begin{aligned}
h_{02} &= h_{11} \cdot 2 \cdot 3 \cdot a_3 + h_{12} \cdot 2 \cdot a_2 + h_{21} \cdot 2 \cdot a_2 + h_{22} \cdot a_1 + h_{31} \cdot a_1 + h_{32} \cdot a_0, \\
h_{03} &= h_{11} \cdot 3 \cdot 4 \cdot a_4 + h_{12} \cdot 2 \cdot 3 \cdot a_3 + h_{13} \cdot 1 \cdot 2 \cdot a_2 + h_{21} \cdot 3 \cdot a_3 + h_{22} \cdot 2 \cdot a_2 + h_{23} \cdot a_1 + h_{31} \cdot a_2 \\
&\quad + h_{32} \cdot a_1 + h_{33} \cdot a_0, \\
0 &= h_{11} \cdot 4 \cdot 5 \cdot a_5 + h_{12} \cdot 4 \cdot 3 \cdot a_4 + h_{13} \cdot 3 \cdot 2 \cdot a_3 + h_{14} \cdot 2 \cdot a_2 + h_{21} \cdot 4 \cdot a_4 + h_{22} \cdot 3 \cdot a_3 \\
&\quad + h_{23} \cdot 2 \cdot a_2 + h_{24} \cdot a_1 + h_{31} \cdot a_3 + h_{32} \cdot a_2 + h_{33} \cdot a_1 + h_{34} \cdot a_0, \\
0 &= h_{11} \cdot 5 \cdot 6 \cdot a_6 + h_{12} \cdot 4 \cdot 5 \cdot a_5 + h_{13} \cdot 3 \cdot 4 \cdot a_4 + h_{14} \cdot 2 \cdot 3 \cdot a_3 + h_{15} \cdot 2 \cdot a_2 + h_{21} \cdot 5 \cdot a_5 \\
&\quad + h_{22} \cdot 4 \cdot a_4 + h_{23} \cdot 3 \cdot a_3 + h_{24} \cdot 2 \cdot a_2 + h_{25} \cdot a_1 + h_{31} \cdot a_4 + h_{32} \cdot a_3 + h_{33} \cdot a_2 + h_{34} \cdot a_1 + h_{35} \cdot a_0, \\
0 &= h_{11} \cdot 6 \cdot 7 \cdot a_7 + h_{12} \cdot 5 \cdot 6 \cdot a_6 + h_{13} \cdot 4 \cdot 5 \cdot a_5 + h_{14} \cdot 3 \cdot 4 \cdot a_4 + h_{15} \cdot 2 \cdot 3 \cdot a_3 + h_{16} \cdot 1 \cdot 2 \cdot a_2 \\
&\quad + h_{21} \cdot 6 \cdot a_6 + h_{22} \cdot 5 \cdot a_5 + h_{23} \cdot 4 \cdot a_4 + h_{24} \cdot 3 \cdot a_3 + h_{25} \cdot 2 \cdot a_2 + h_{26} \cdot a_1 + h_{31} \cdot a_5 \\
&\quad + h_{32} \cdot a_4 + h_{33} \cdot a_3 + h_{34} \cdot a_2 + h_{35} \cdot a_1, \\
0 &= h_{11} \cdot (m+2) \cdot (m+1) \cdot a_{m+2} + h_{12} \cdot (m+1) \cdot (m) \cdot a_{m+1} + h_{13} \cdot (m) \cdot (m-1) \cdot a_m \\
&\quad + h_{14} \cdot (m-1) \cdot (m-2) \cdot a_{m-1} + h_{15} \cdot (m-2) \cdot (m-3) \cdot a_{m-2} + h_{16} \cdot (m-3) \cdot (m-4) \cdot a_{m-3} \\
&\quad + h_{17} \cdot (m-4) \cdot (m-5) \cdot a_{m-4} + h_{21} \cdot (m+1) \cdot a_{m+1} + h_{22} \cdot (m) \cdot a_m + h_{23} \cdot (m-1) \cdot a_{m-1} \\
&\quad + h_{24} \cdot (m-2) \cdot a_{m-2} + h_{25} \cdot (m-3) \cdot a_{m-3} + h_{26} \cdot (m-4) \cdot a_{m-4} + h_{31} \cdot a_m + h_{32} \cdot a_{m-1} \\
&\quad + h_{33} \cdot a_{m-2} + h_{34} \cdot a_{m-3} + h_{35} \cdot a_{m-4}, \quad \text{for } m \geq 6.
\end{aligned}$$

To approximate the solution by a series of length M , the set of $M-2$ equations are recursively solved for a_n , $n = 2, \dots, M$, as functions of a_0 , and a_1 . The boundary conditions provide the two remaining equations to solve the BVP.

Proof of Corollary 4. Using the M term polynomial approximation of the Fourier Transform $\bar{f}^M(\omega_1, \pi, T-t; z_t)$ in (49) and the characterization in Proposition 3, $G(\pi_t, T-t; z_t) = e^{z_t} \sum_{n=0}^M b_n(1/i, T-t)(\pi - .5)^n$, $G_S(\pi_t, T-t; z_t) = G/S$, and $G_\pi(\pi_t, T-t; z_t) = e^{z_t} \sum_{n=1}^M n b_n(1/i, T-t)(\pi - .5)^{n-1}$. The functions g_1 and g_2 in (19) and (20) can be written as series whose partial derivatives are evident:

$$\begin{aligned}
g_1(\omega_1, \pi, T-t; z_t) &= \frac{e^{i\omega_1 z_t} \sum_{n=0}^M b_n(1/i + \omega_1, T-t)(\pi - .5)^n}{\sum_{n=0}^M b_n(1/i, T-t)(\pi - .5)^n}, \quad \text{and,} \\
g_2(\omega_1, \pi, T-t; z_t) &= \frac{e^{i\omega_1 z_t} \sum_{n=0}^M b_n(\omega_1, T-t)(\pi - .5)^n}{\sum_{n=0}^M b_n(0, T-t)(\pi - .5)^n}.
\end{aligned}$$

Now by passing partial derivatives under the integral of the Fourier inversion, one obtains the desired functions

$$\Pi_{j\pi}(z_t, K, \pi_t, T-t) = \frac{1}{2} + \frac{1}{\Pi} \int_0^\infty \operatorname{Re} \left[\frac{e^{-i\omega_1 \log K} g_{j\pi}(\omega_1, \pi_t, T-t; z_t)}{i\omega_1} \right] d\omega_1.$$

Finally, using the product and chain rules of differentiation provide the stated expressions. ■

Proof of Proposition 4. The proof is very similar to the proof of Proposition 3, so we shall be brief. Applying Ito's Lemma to the variance forecast function $\bar{V}^U(t, T, \pi_t) = \int_0^T V(\pi_t) ds$ immediately implies the PDE in eq. (23). Taking a Fourier Transform of the PDE with respect to the time variable and using the frequency variable ω implies that the solution to the Fourier Transform of the PDE is the BVP

$$0 = -i w \bar{V}^U(\omega, T; \hat{\Theta}, \check{\pi}_t) + V(\pi_t) + \bar{V}_\pi^U(t, T; \hat{\Theta}, \check{\pi}_t) \mu_\pi(\pi_t) + \frac{1}{2} \bar{V}_{\pi\pi}^U(\omega, T; \hat{\Theta}, \check{\pi}_t) \sigma_\pi(\pi_t)^2. \quad (50)$$

Now as in Proposition 3, we formulate a series approximation of the form

$$\bar{V}^U(\omega, T; \hat{\Theta}, \check{\pi}_t) = \sum_{m=1}^M \hat{a}_m(\omega) (\check{\pi}_t - .5)^m,$$

where the coefficients $\hat{a}_m(\omega)$ satisfy a recursive relationship and the natural boundary conditions

$$\begin{aligned} 0 &= -i w \bar{V}^U(\omega, T; \hat{\Theta}, 0) + V(0) + \bar{V}_\pi^U(\omega, T; \hat{\Theta}, 0) \mu_\pi(0), \\ 0 &= -i w \bar{V}^U(\omega, T; \hat{\Theta}, 1) + V(1) + \bar{V}_\pi^U(\omega, T; \hat{\Theta}, 1) \mu_\pi(1). \end{aligned}$$

Finally, taking the Inverse Fourier Transforms of the coefficients as in eq. (48) provides the solution to the variance forecast function. ■

Appendix C

This appendix briefly describes Heston (1993) Stochastic Volatility model and the methodologies used to estimate its parameters. The SV model assumes that under the risk-neutral measure the processes for stock returns and volatility are given by

$$dz_t = [r - \delta] dt + \sqrt{V_t} dW_{s_t}, \quad \text{and}, \quad (51)$$

$$dV_t = [\theta_v - \kappa_v V_t] dt + \sigma_v \sqrt{V_t} dW_{v_t}, \quad (52)$$

where $z_t = \log(S_t)$, the interest rate, r , and dividend-yield, δ are assumed constant, and $\text{Cov}_t[dW_s, dW_v] = \rho_v dt$. The price of a call option is given by

$$C(t, \tau) = E_t^Q \left[e^{-r(T-t)} \max(S_T - K_t, 0) \right] \quad (53)$$

$$= S_t e^{-\delta t} \Pi_1(t, \tau, S, V) - K e^{-r\tau} \Pi_2(t, \tau, S, V), \quad (54)$$

where $E^Q[\cdot]$ indicates the expectation under the risk-neutral measure. The risk-neutral probabilities $\Pi_1(\cdot)$, and $\Pi_2(\cdot)$ are recovered by inverting the respective characteristic functions as in Eq. (21), and the characteristic functions $f_j(x, V, t, \tau, \omega)$, for $j = 1, 2$, are shown in Eq. (17) in Heston (1993) (we adjust the formulas for the dividend-yield). The equation for puts is similarly obtained using put-call parity.

Following Bates (1996) and Bakshi, Cao, and Chen (1997), we fit the parameters in (51) and (52) by the method of *implicit parameter estimation*. Let the set of parameters of the SV model be given by the vector $\Theta^{SV} = \{\theta_v, \kappa_v, \sigma_v, \rho_v\}$. Further, let $C_\ell(S_t, K_\ell, t, T)$, and $C_\ell^{SV}(S_t, K_\ell, t, T; \Theta_t^{SV}, V_t, r_t, \delta_t)$, for $\ell = 1, \dots, L$ be the prices of traded options and the ones implied by the SV model, respectively. For each date t we obtain simultaneously an estimate of Θ^{SV} and the volatility level V_t by non-linear least squares, through the minimization of the proportional root mean square error:²⁰

$$PRMSE_t = \sqrt{\frac{1}{L} \sum_{\ell=1}^L \frac{C_\ell(S_t, K_\ell, t, T) - C_\ell^{SV}(S_t, K_\ell, t, T; \Theta_t^{SV}, V_t, r_t, \delta_t)}{C_\ell(S_t, K_\ell, t, T)}} \quad (55)$$

Results from the empirical exercise above are in Table 7. We make the following summary observations on the estimations: (i) The average parameter estimates using Method 1 are quite similar to that in (Bakshi, Cao, and Chen 1997, Table III) who performed a similar exercise, even though the grouping of options by maturities and strikes is somewhat different. (ii) The implied correlation between returns and volatility, ρ_v , has been on average around -0.7 reflecting the fact that in the sample the correlation has mainly been negative. Nonetheless, as seen in Figure 1, there have been months in 1986 and 1988, when the implicit correlation was clearly positive, and months in 1991 when the correlation was only slightly negative. (iii) Using method 1, we found that the most unstable parameter, as in Bakshi, Cao, and Chen (1997), is the speed of mean reversion of volatility, κ_v . The level of κ_v decreases for higher maturities as does the instability. Nonetheless, the speed of mean reversion varies substantially from month to month. This may also explain the high value and the high standard error we obtained using the method of simulated moments. (iv) The volatility of variance is decreasing in maturity, which is due to the lower implicit kurtosis of implied return distributions. Nonetheless, as Bakshi, Cao, and Chen (1997) reported, the average value of σ_v is much larger than can be justified by its time-series estimate of 0.14 from stock returns.

²⁰Because we desire at least five contracts for a meaningful implicit parameter estimation of Heston's SV model, and some of the months have fewer three-month and six-month option prices, we increase the cutoff to 15 percent OTM options for these maturities. Despite this, we were forced to drop 14 months of data for six-month options.

References

- Ait-Sahalia, Y., 1998, "Maximum-Likelihood Estimation of Discretely-Sampled Diffusions: A Closed-Form Approach," working paper, Princeton University.
- Ait-Sahalia, Y., and A. W. Lo, 1998, "Nonparameteric Estimation of State-Price Densities Implicit in Financial Asset Prices," *The Journal of Finance*, 53(2), 499–548.
- Anderson, T. G., and T. Bollerslev, 1998, "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts," *International Economic Review*, (4), 885–905.
- Bakshi, G., C. Cao, and Z. Chen, 1997, "Empirical Performance of Alternative Option Pricing Models," *The Journal of Finance*, 52(5), 2003–2051.
- Bakshi, G., and D. Madan, 2000, "Spanning and Derivative Security Valuation," *The Journal of Financial Economics*, 55(2), 205–238.
- Barberis, N., M. Huang, and J. Santos, 2001, "Prospect Theory and Asset Prices," *Quarterly Journal of Economics* (forthcoming).
- Bates, D. S., 1991, "The Crash of '87: Was it Expected? The Evidence from Options Markets," *The Journal of Finance*, 46(3), 1009–1044.
- , 1996, "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in PHLX Deutsche Mark Options," *The Review of Financial Studies*, 9(1), 69–107.
- , 1999, "Post-'87 Crash Fears in the S&P 500 Futures Option Market," *Journal of Econometrics* (Forthcoming).
- Berk, J. B., R. C. Green, and V. Naik, 1999, "Optimal Investment, Growth Options, and Security Returns," *The Journal of Finance*, 54, 1553–1608.
- Black, F., and M. S. Scholes, 1973, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81(3), 637–654.
- Boyce, W. E., and R. C. DiPrima, 1997, *Elementary Differential Equations and Boundary Value Problems*. John Wiley and Sons, New York.
- Brennan, M. J., 1998, "The Role of Learning in Dynamic Portfolio Decisions," *European Finance Review*.
- Brennan, M. J., and Y. Xia, 2001, "Stock Price Volatility and the Equity Premium Puzzle," *The Journal of Monetary Economics* (forthcoming).
- Cagetti, M., L. P. Hansen, T. Sargent, and N. Williams, 2002, "Robustness and Pricing with Uncertain Growth," *Review of Financial Studies*.

- Campbell, J. Y., 1999, "Asset Prices, Consumption, and the Business Cycle," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford. North-Holland, Amsterdam, Netherlands.
- Campbell, J. Y., and J. H. Cochrane, 1999, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251.
- Campbell, S. D., and C. Li, 1999, "Option Prices with Unobserved and Regime-Switching Volatility," working paper, Department of Economics, University of Pennsylvania.
- Cecchetti, S. G., P.-S. Lam, and N. C. Mark, 1993, "The Equity Premium and the Risk-Free Rate: Matching the Moments," *Journal of Monetary Economics*, 31(1), 21–45.
- Cox, J. C., J. E. Ingersoll, and S. A. Ross, 1985, "An Intertemporal General Equilibrium Model of Asset Prices," *Econometrica*, 53, 363–384.
- David, A., 1993, "Fluctuating Confidence and Stock Market Returns," International Finance Discussion Papers 461, Board of Governors of the Federal Reserve System.
- , 1997, "Fluctuating Confidence in Stock Markets: Implications for Returns and Volatility," *Journal of Financial and Quantitative Analysis*, 32(4), 427–462.
- David, A., and P. Veronesi, 2000, "Inflation and Earnings Uncertainty and the Volatility of Asset Prices: An Empirical Investigation," Mimeo, Graduate School of Business, University of Chicago.
- Detemple, J. B., 1986, "Asset Pricing in a Production Economy with Incomplete Information," *The Journal of Finance*, 41, 383–391.
- Diebold, F., T. Gunther, and A. Tay, 1998, "Evaluating Density Forecasts, with Applications to Financial Risk Management," *International Economic Review*, 39, 863–883.
- Dothan, M. U., and D. Feldman, 1986, "A Theory of Asset Prices and the Term Structure of Interest Rates in a Partially Observable Economy," *The Journal of Finance*, 41, 369–382.
- Duffie, D., 1996, *Dynamic Asset Pricing Theory*. Princeton University Press, Princeton, New Jersey.
- Garcia, R., R. Luger, and E. Renault, 1999, "Asymmetric Smiles, Leverage Effects and Structural Parameters," Mimeo, CIRANO.
- Garica, R., 1998, "Asymptotic Null Distribution of the Likelihood Ratio Test in Markov Switching Models," *International Economic Review*, 39(3), 763–788.
- Gennotte, G., 1986, "Optimal Portfolio Choice Under Incomplete Information," *The Journal of Finance*, 41, 733–749.
- Guidolin, M., and A. Timmermann, 2000, "Option Prices and Implied Volatility Dynamics under Bayesian Learning," mimeo, University of California, San Diego.

- Hamilton, J. D., 1994, *Time Series Analysis*. Princeton University Press, Princeton, New Jersey.
- Hansen, B. E., 1994, "Autoregressive Conditional Density Estimation," *International Economic Review*.
- Heston, S. L., 1993, "A Closed-Form Solution for Options with Stochastic Volatility and Applications to Bond and Currency Options," *The Review of Financial Studies*, 6(2), 326–343.
- Hull, J. C., 1993, *Options, Futures, and Other Derivative Securities*. Prentice Hall, Englewood Cliffs, New Jersey.
- Hull, J. C., and A. White, 1987, "The Pricing of Options on Assets with Stochastic Volatility," *Journal of Finance*, 42(2), 281–300.
- Jackwerth, J. C., and M. Rubinstein, 1995, "Recovering Probability Distributions from Option Prices," *Journal of Finance*, pp. 1611–1631.
- Jacquier, E., and R. Jarrow, 2000, "Bayesian Analysis of Contingent Claim Model Error," *Journal of Econometrics*, 94(2), 145–180.
- Jones, R. A., and J. M. Ostroy, 1984, "Flexibility and Uncertainty," *Review of Economic Studies*, 51, 13–32.
- Karlin, S., and H. M. Taylor, 1982, *A Second Course in Stochastic Processes*. Academic Press, New York.
- Kendall, M., and A. Stuart, 1977, *The Advanced Theory of Statistics, 4th edition, Volume 1*. Macmillan Publishing Company, New York.
- Kim, C.-J., and C. R. Nelson, 1999, *State-Space Models with Regime-Switching*. The MIT Press, Cambridge, Massachusetts.
- Liptser, R. N., and A. N. Shiryaev, 1977, *Statistics of Random Processes I*. Springer Verlag, New York.
- Lo, A. W., 1988, "Maximum Likelihood Estimation of Generalized Ito Processes with Discretely-Sampled Data," *Econometric Theory*, 4, 231–247.
- Maddala, G. S., and M. Nimalendran, 1996, "Errors-in-Variables Problems in Financial Markets," in *Handbook of Statistics, Volume 14*, ed. by G. S. Maddala, and C. R. Rao. Elsevier Science, Amsterdam.
- Mincer, J., and V. Zarnowitz, 1969, "The Evaluation of Economic Forecasts," in *Economic Forecasts and Expectations*, ed. by J. Mincer.
- Rainville, E. D., 1964, *Intermediate Differential Equations*. The Macmillan Company, New York.

- 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, 1976 through August 31, 1978," *Journal of Finance*, 40, 455–480.
- Schwert, W. G., 1989, "Why Does Stock Market Volatility Change Over Time," *Journal of Finance*, pp. 1115–1153.
- Scott, L. O., 1997, "Pricing Stock Options in a Jump-Diffusion Model with Stochastic Volatility and Interest Rates: Applications of Fourier Inversion Methods," *Mathematical Finance*, 7(4), 413–426.
- Shefrin, H., 1999, "Irrational Exuberance and Option Smiles," *Financial Analysts Journal*.
- Stein, J., and E. M. Stein, 1991, "Stock Price Distributions with Stochastic Volatility," *The Review of Financial Studies*, 4, 727–752.
- Stuart, A., and J. K. Ord, 1991, *Kendall's Advanced Theory of Statistics, 5th edition, Volume 2*. Oxford University Press, New York.
- Veronesi, P., 1999, "Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model," *The Review of Financial Studies*, 12, 975–1007.
- , 2000, "How Does Information Quality Affect Stock Returns?," *The Journal of Finance*, 55(2).
- Yan, H., 1999, "Uncertain Growth Prospects, Learning, and Asset Prices," Job Market Paper 1, University of California at Berkeley.

Table 1: CBOE S&P 500 Index OTM Options Data

Maturity	Period Span	Full Sample		Ex-Crash	
		No.	No. Per Month	No.	No. Per Month
1-Month	1986:4 – 1996:5	2,780	22.78	2,650	22.84
3-Month	1986:5 – 1996:4	2,944	24.74	2,780	23.36
6-Month	1987:2 – 1996:1	1,692	16.75	1,618	15.41

Options have been selected as described in Section 3.1. 'Ex-crash' means that observations from October 1987 to March 1988 have been dropped.

Table 2: 2-State Model Calibrated to Real Earnings and Consumption Growth: 1960-1998

θ_1	θ_2	λ_{12}	λ_{21}	σ_x	g	σ_c	ρ_{xc}
-0.154 (0.048)*	0.109 (0.023)*	0.813 (0.366)*	0.406 (0.156)*	0.073 (0.004)*	0.0242 (0.004)*	0.0367 (0.003)*	0.0336 (0.039)

Log Likelihood: -1719.96

Log Likelihood Ratio (Garcia (1998)): 39.54 p-value: 0.000

Parameter estimates are based on maximum likelihood using the regime-switching estimation method as described in Hamilton (1994). Heteroskedastic and autocorrelation consistent asymptotic standard errors are in parenthesis. A “*” denotes statistical significance at the 5% level. The Log Likelihood Ratio statistics is for the null hypothesis of only one state θ , and the p-value is computed using the critical values in Table 1A in Garcia (1998). For consistency with Garcia (1998), the regime-switching model was re-estimated using only earnings data. The point estimates were essentially identical to the one reported.

Table 3: Time-Varying Moments of the Risk-Neutral Return Density

Full Sample				Ex Crash				
Panel A: ATM Implied Volatility \mathcal{V}_t								
No.	$\hat{\pi}_t(1 - \hat{\pi}_t)$	\mathcal{V}_t^U	\mathcal{V}_{t-1}	\bar{R}^2	$\hat{\pi}_t(1 - \hat{\pi}_t)$	\mathcal{V}_t^U	\mathcal{V}_{t-1}	\bar{R}^2
1.	0.177 (0.075)*			0.061	0.256 (0.063)*			0.196
2.		0.329 (0.125)*		0.082		0.472 (0.099)*		0.256
3.			0.810 (0.070)*	0.654			0.750 (0.080)*	0.554
4.		0.129 (0.061)*	0.785 (0.077)*	0.663		0.217 (0.061)*	0.641 (0.090)*	0.595
Panel B: 3 %OTM Put-to-Call Price P_t/C_t								
No.	$\hat{\pi}_t$	P_t^U/C_t^U	P_{t-1}/C_{t-1}	\bar{R}^2	$\hat{\pi}_t$	P_t^U/C_t^U	P_{t-1}/C_{t-1}	\bar{R}^2
5.	0.233 (0.075)*			0.080	0.265 (0.077)*			0.104
6.		0.306 (0.080)*		0.120		0.342 (0.102)*		0.148
7.			0.466 (0.098)*	0.219			0.451 (0.102)*	0.204
8.		.208 (0.064)*	0.406 (0.101)*	0.264		.241 (0.068)*	0.370 (0.106)*	0.264
Panel C: 3 %Kurtosis Premium $Y_t(k)$								
No.	$\hat{\pi}_t(1 - \hat{\pi}_t)$	Y_t^U	Y_{t-1}	\bar{R}^2	$\hat{\pi}_t(1 - \hat{\pi}_t)$	Y_t^U	Y_{t-1}	\bar{R}^2
9.	-0.739 (0.221)*			0.050	-0.750 (0.228)*			0.050
10.		0.259 (0.061)*		0.093		0.261 (0.064)*		0.092
11.			0.213 (0.067)*	0.044			0.215 (0.068)*	0.045
12.		0.231 (0.058)*	0.152 (0.048)*	0.107		0.235 (0.061)*	0.155 (0.050)*	0.107

Standard errors, corrected for heteroskedasticity and autocorrelation, are in parenthesis. The symbol * indicates statistical significance at least at the 5% level. 'Ex-crash' means that observations from October 1987 to March 1988 have been dropped. In Panel A, the dependent variable is the ATM implied volatility $\mathcal{V}_t = \frac{1}{2}[\text{C. Vol.}(K/S_t = 1) + \text{P. Vol.}(K/S_t = 1)]$. The right-hand variable \mathcal{V}_t^U is its theoretical counterpart from U option prices using Proposition 3 and the belief process filtered earnings data. In Panel B, the dependent variable is the 3 % put-to-call ratio P_t/C_t , i.e. the ratio of 3 % OTM puts over 3 % OTM calls (where OTM is defined with respect to the forward price, as in e.g. Bates (1991)). The right-hand variable P_t^U/C_t^U is its theoretical counterpart conditional on investors' filtered belief. In Panel C, the dependent variable is a convexity measure defined as $Y_t = (\text{P. Vol.}(K/S = 0.97) + \text{P. Vol.}(K/S = 1.03) - 2\text{P. Vol.}(K/S = 1))$. The right-hand side variable Y_t^U is the theoretical level conditional on investors' filtered belief. In all cases, implied volatility or prices are interpolated from available prices on the 15th (or closest to 15th) day of each month.

Table 4: Tests of Option Price Moments

	ATM	OTM Calls	OTM Puts	All
1-Month	0.0742 (0.7853)	0.4452 (0.8328)	1.8944 (0.1687)	6.2373 (0.1001)
3-Month	0.8744 (0.3497)	0.6938 (0.7922)	3.0506 (0.0807)	2.9601 (0.3978)
6-Month	0.4053 (0.5243)	0.1098 (0.7404)	1.0613 (0.3029)	1.1996 (0.7531)
All	5.1554 (0.1608)	4.0962 (0.2512)	6.7178 (0.0815)	8.0649 (0.5276)

Reported numbers are of the chi-square statistics of the nine option pricing moments:

$$c = \left(\frac{1}{T} \sum_{t=1}^T \epsilon_t \right)' \cdot \Omega^{-1} \cdot \left(\sum_{t=1}^T \frac{1}{T} \cdot \epsilon_t \right).$$

All standard deviations have been adjusted for contemporaneous and auto correlations using 12 lags of the Newey-West. Under the null hypothesis of zero means, c is distributed $\chi^2(k)$ where k moments are tested. P-values in parenthesis.

Table 5: Specification Test of Implied Belief Processes

Maturity	D_n	$\text{LBox}_{z_t - \bar{z}}(6)$	$\text{LBox}_{(z_t - \bar{z})^2}(6)$
1-Month	0.1176 (0.0967)	0.0715 (0.9999)	0.1064 (0.9999)
3-Month	0.1278 (0.0416)	0.0039 (0.9999)	0.0160 (0.9999)
6-Month	0.1189 (0.0750)	0.0339 (0.9999)	0.2110 (0.9998)

This table reports the results of the Kolmogorov's test for the diffusion $d\pi_t = k(\pi^* - \pi_t)dt + \pi_t(1 - \pi_t) \mathbf{b} dW_t$. P-values in parenthesis. Kolmogorov's D_n statistic is computed as follows: Let $z_t = F(\tilde{\pi}_{t+1} | \tilde{\pi}_t; \Theta^U)$. By the probability integral transformation, $\{z_t\}$ is distributed i.i.d. $U(0,1)$. Let $S_n(z_t)$ be the sample CDF of $\{z_t\}$ and $F_0(z) = z$ be the uniform CDF. Then, $D_n = \sup_z |S_n(z) - F_0(z)|$. P-values are obtained using the formula $\text{Prob}(D_n > zn^{-1/2}) = 2 \sum_{r=1}^n (-1)^{r-1} \exp(-2r^2 z^2)$. n equals 121, 120, and 117 for the three maturities respectively. Tests for the independence of z_t are constructed using Ljung-Box statistics: $\text{LBox}_X(J) = T \cdot \sum_{j=1}^J (T+2)/(T-j) \cdot \rho_j$, where ρ_j is its j th autocorrelation of X_t , and has a distribution of $\chi^2(J)$.

Table 6: The Relative Empirical Performance of U-model

	Full Sample			Ex-Crash		
Panel A: In-Sample Fit						
Maturity	U	BS	SV	U	BS	SV
1 Month	0.317	0.441	0.302	0.307	0.444	0.287
3 Month	0.267	0.411	0.312	0.245	0.405	0.308
6 Month	0.187	0.304	0.118	0.156	0.308	0.120
Panel B: Out-of-Sample Fit						
Maturity	U	BS	SV	U	BS	SV
1 Month	0.399	0.488	0.334	0.322	0.492	0.322
3 Month	0.341	0.533	0.402	0.328	0.510	0.411
6 Month	0.213	0.395	0.177	0.191	0.312	0.190
Panel C: Hedging Errors						
Maturity	U	BS	SV	U	BS	SV
1 Month	0.289	0.28	0.239	0.254	0.285	0.233
3 Month	0.164	0.189	0.245	0.148	0.234	0.191
6 Month	0.114	0.116	0.102	0.112	0.116	0.110
Panel D: Variance Forecast						
Maturity	U	BS	SV	U	BS	SV
1 Month	0.930	0.947	0.945	0.728	0.703	0.762
3 Month	0.867	0.906	0.963	0.889	0.924	0.985
6 Month	0.859	0.898	0.978	0.891	0.909	0.979

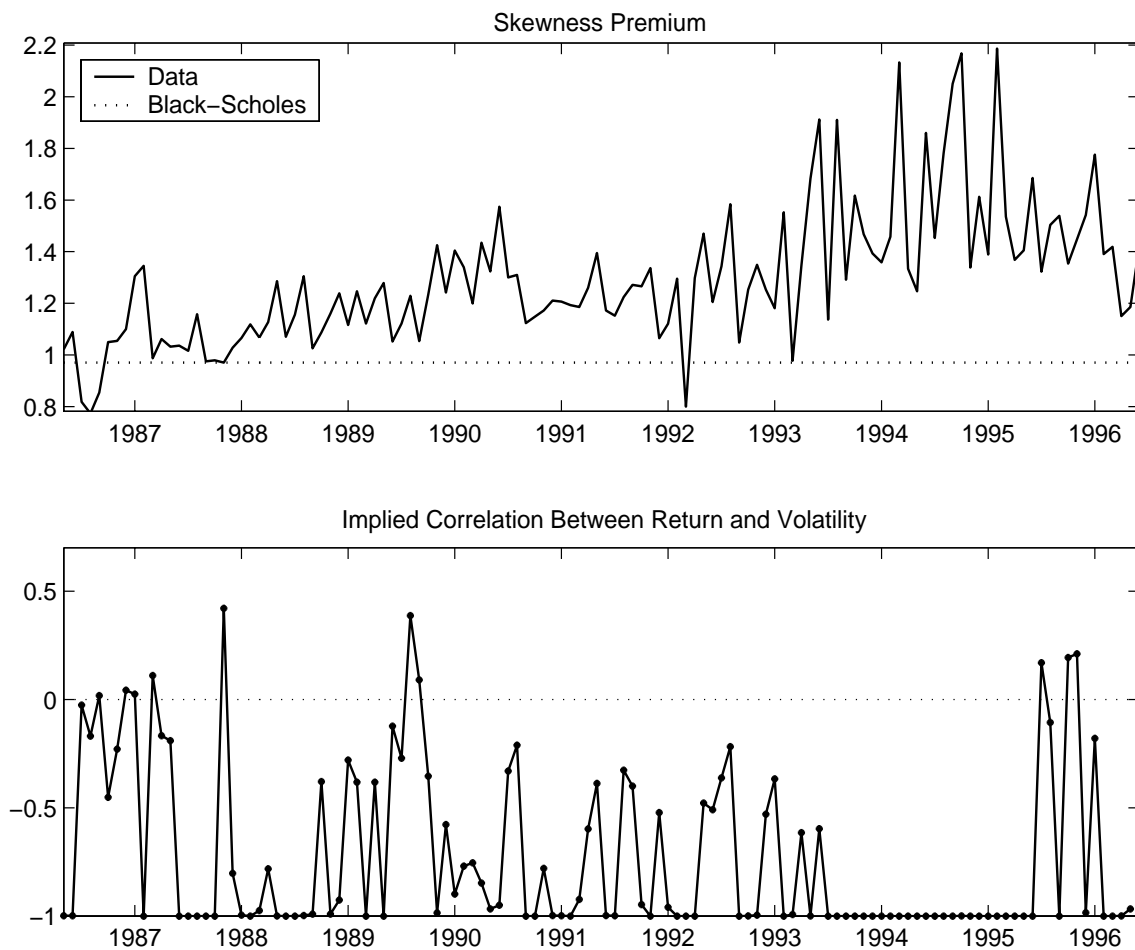
The models are fitted to each month's options data to minimize PRMSE, eq. (27). One parameter each, $\check{\sigma}_t$, for the B-S model, \check{V}_{t_j} for the SVF model, and $\check{\pi}_t$ for the U model, is implied and five parameters, \check{V}_t , $\check{\theta}_{v_t}$, $\check{\kappa}_{v_t}$, $\check{\sigma}_{v_t}$, and $\check{\rho}_t$ are implied for the SV model. 'Ex-crash' means that observations from October 1987 to March 1988 have been dropped. Single-instrument hedging errors for the three models are computed using eq. (30). Var. Forecast reports $1 - \bar{R}^2$ of the regression $\frac{1}{T-t} \int_t^T r_s^2 ds = \beta_0 + \beta_1 \cdot E[\int_t^T \sigma_s^2 ds | \mathcal{F}_t]$ for each of the models, when parameters are implied under each of the two objectives. **Bold Faced** items in each category denote 'minimal' PRMSE (Panel A and B), Hedging Errors (Panel C), or $1 - R^2$ (Panel D), when the lowest mean item is different statistically with 95% confidence from the next (non-bold) lowest mean item under the t-statistic for difference in means: $t = \sqrt{(T)} \cdot (\bar{x}_1 - \bar{x}_2) / \sqrt{s_1^2 + s_2^2}$. Sample standard deviations. s_i , are not shown for brevity.

Table 7: Parameter Estimates for Heston's Stochastic Volatility Model (1986:4 to 1996:5)

Parameter	1-Month	3-Month	6-Month
V	0.027 (0.023)	0.032 (0.019)	0.0282 (0.023)
θ_v	0.072 (.071)	0.034 (0.015)	0.036 (0.024)
κ_v	1.887 (3.927)	0.998 (0.582)	0.873 (0.482)
σ_v	0.509 (0.292)	0.289 (0.228)	0.271 (0.217)
ρ	-0.653 (0.401)	-0.718 (0.395)	-0.780 (0.321)

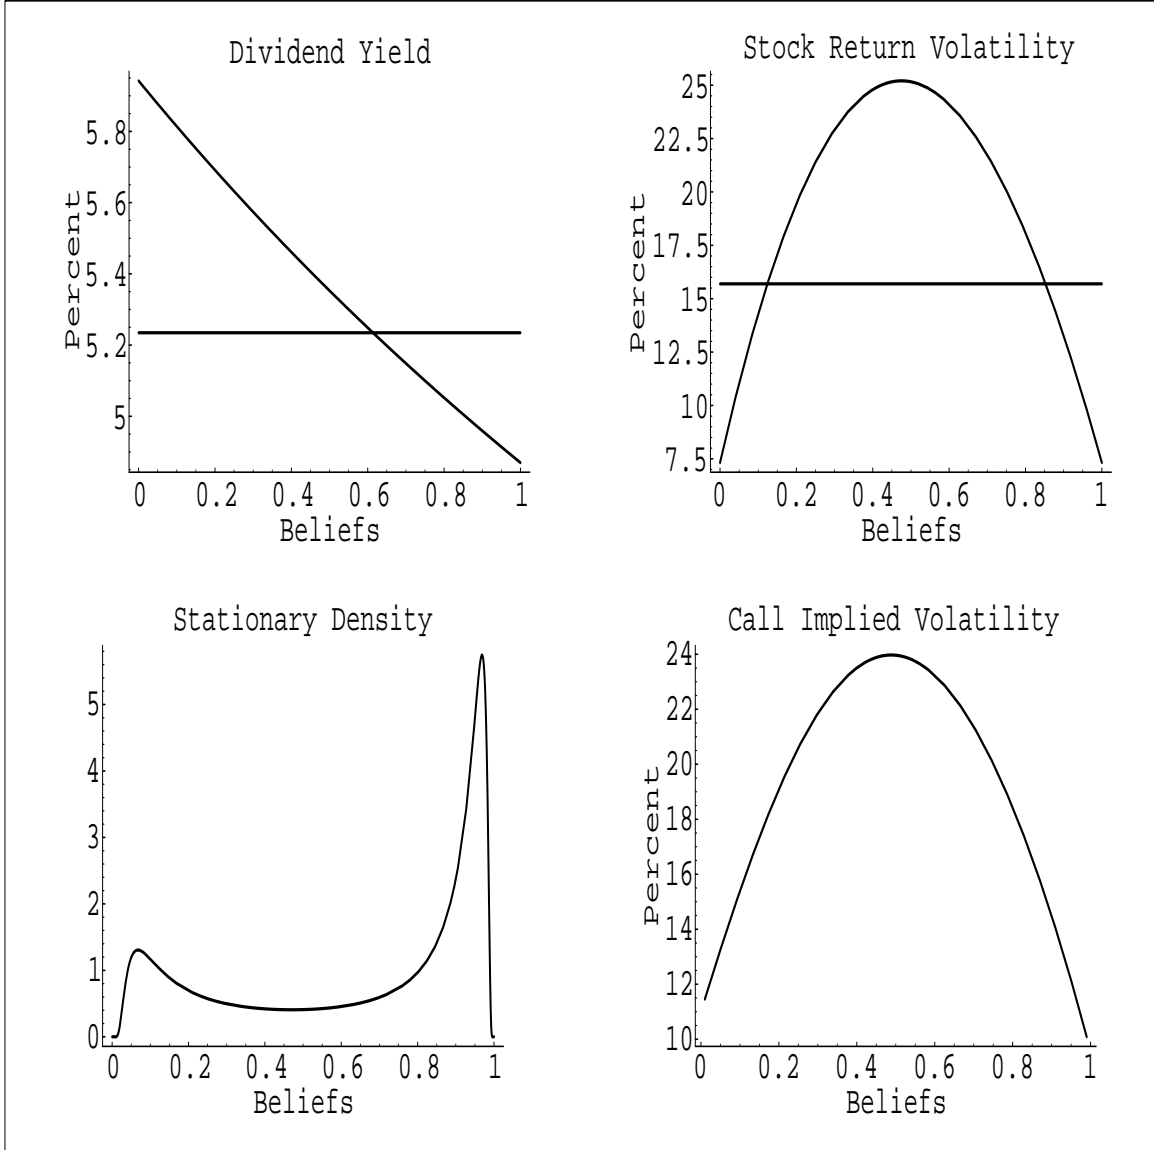
SV model fitted using implicit parameter estimation to options of each of the specified maturities to minimize PRMSE. Reported numbers are means of parameter estimates fitted on the 15th (first trading day) of each month. Sample standard deviation of parameter estimates in parenthesis.

Figure 1: Skewness Premium and Implied Return-Volatility Correlation (1986:04-1996:05)



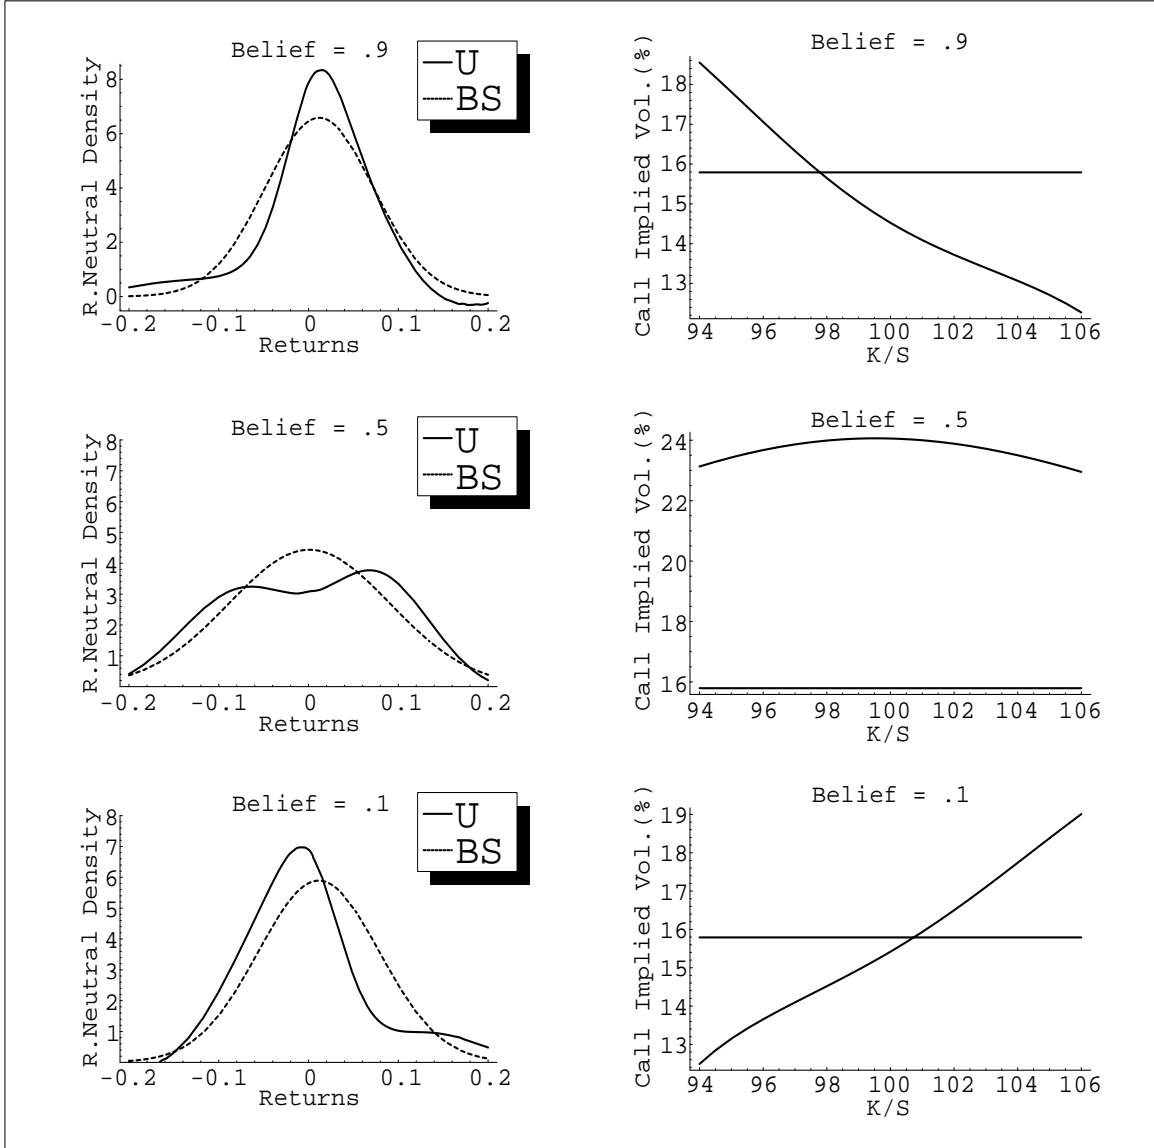
The top panel reports the ratio between 3% out-of-the-money put prices over 3% out-of-the-money call prices, in the data and as implied by the Black-Scholes model. In the bottom panel, the line shows the time-series of the implied ρ_v parameter in Heston's Stochastic Volatility Model, Eqs. (51) and (52), estimated by fitting all 3-month OTM option prices on the 15th (first following trading day) of each month by non-linear least squares to minimize the PRMSE shown in eq. (27).

Figure 2: Calibrated Model (1960:1-1999:9)



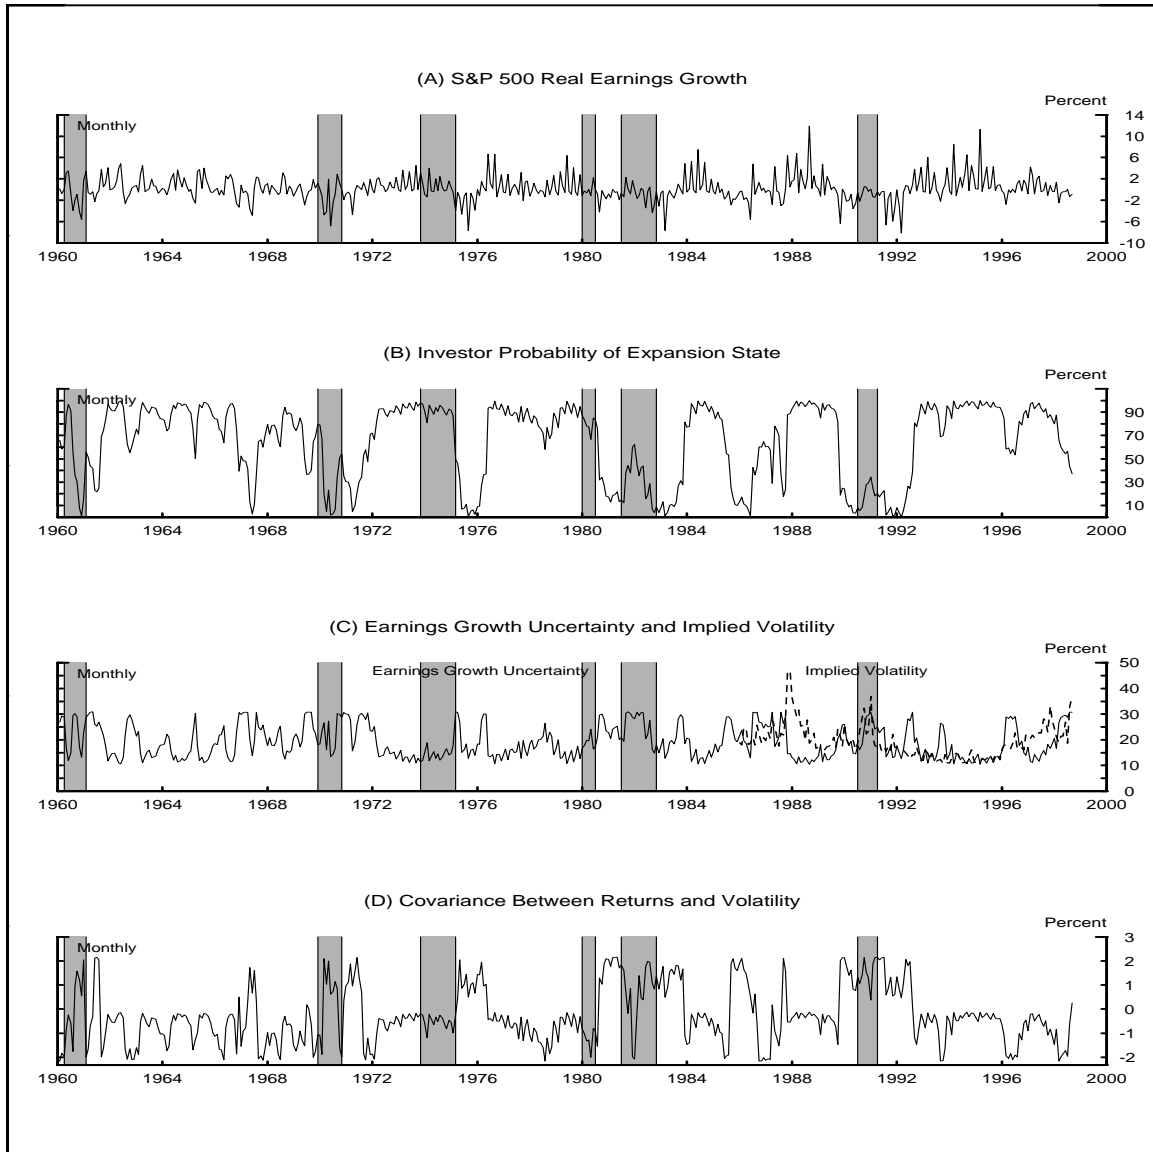
Top panels: The horizontal lines are the unconditional means of the variables shown computed by using the stationary density. Left panel plots the dividend yield, that equals $1/(\pi A + B)$, where $A = C_2 - C_1$, and $B = C_1$, and the C_i are computed using Proposition 1. The right panel plots the volatility function \sqrt{V} , computed using eq. (10). **Bottom panels:** Left panel plots the stationary density computed using eq. (5). The right panel plots the implied volatility for one-month at-the-money call options, obtained by equating the BS option price with the $C^U(S_t, K; \hat{\pi}_t)$ price as given in eq. (17). The other parameters used are those for the 2-state model calibrated to real earnings growth on the S&P 500, as given in Table 2.

Figure 3: Smirks and Frowns in the Two-State Calibrated Model



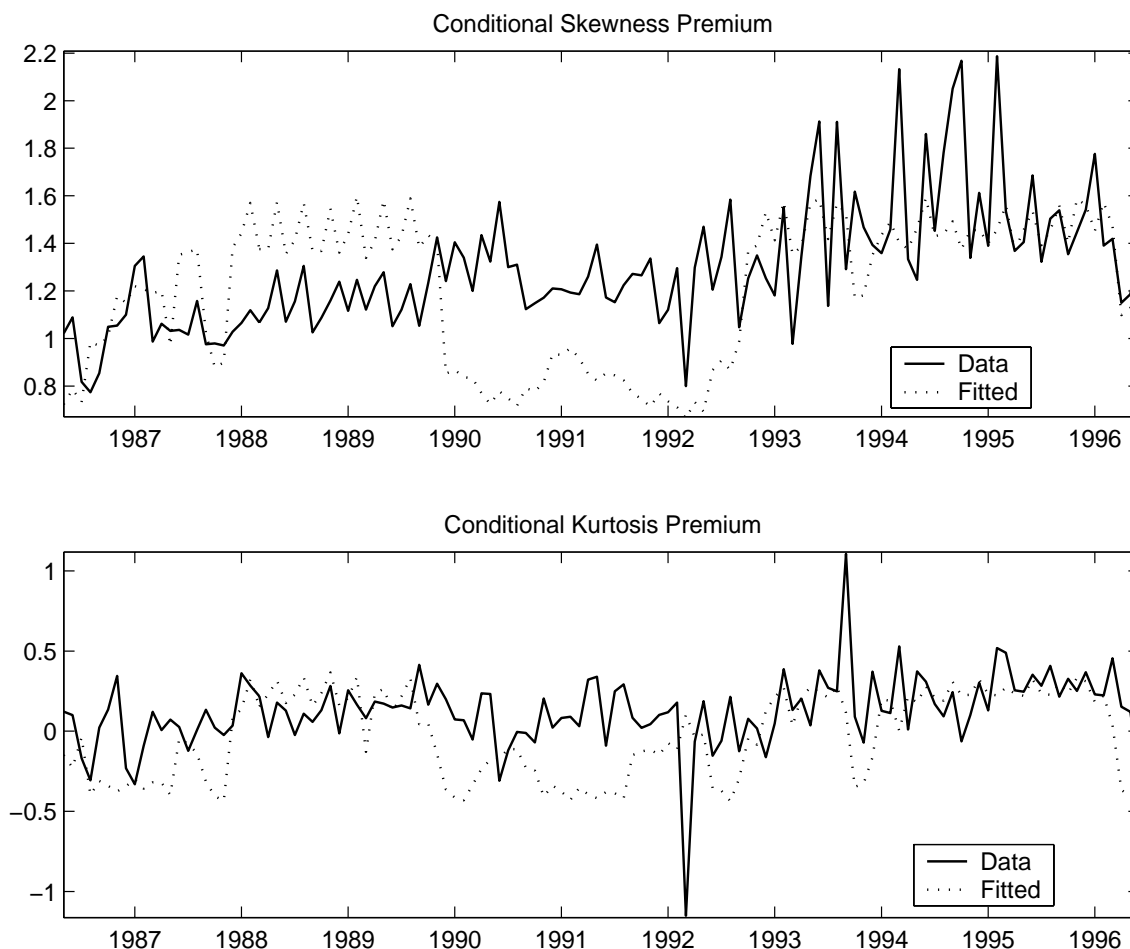
Left Panels: The risk-neutral return density functions conditional on three different initial beliefs $\pi = 0.9$ (top panel), $\pi = 0.5$ (center panel) and $\pi = 0.1$ (bottom panel) with two months to maturity. The density is the inverse of the Fourier Transform of the risk-neutral return density function given by $\bar{f}_r^M = \sum_{n=0}^M b_n(\omega_1, .167) \cdot (\pi - .5)^n$, where b_n are in (48) in Appendix 2. **Right Panels:** The corresponding implied volatility curves across different strike prices for call option prices with 2 months to maturity obtained by equating the BS option price with with the $C^U(S_t, K; \hat{\pi}_t)$ price as given in (17). The parameters used are those for the 2-state model calibrated to real earnings growth on the S&P 500, as given in Table 2.

Figure 4: Two-State Filter on S&P 500 Real Earnings Growth



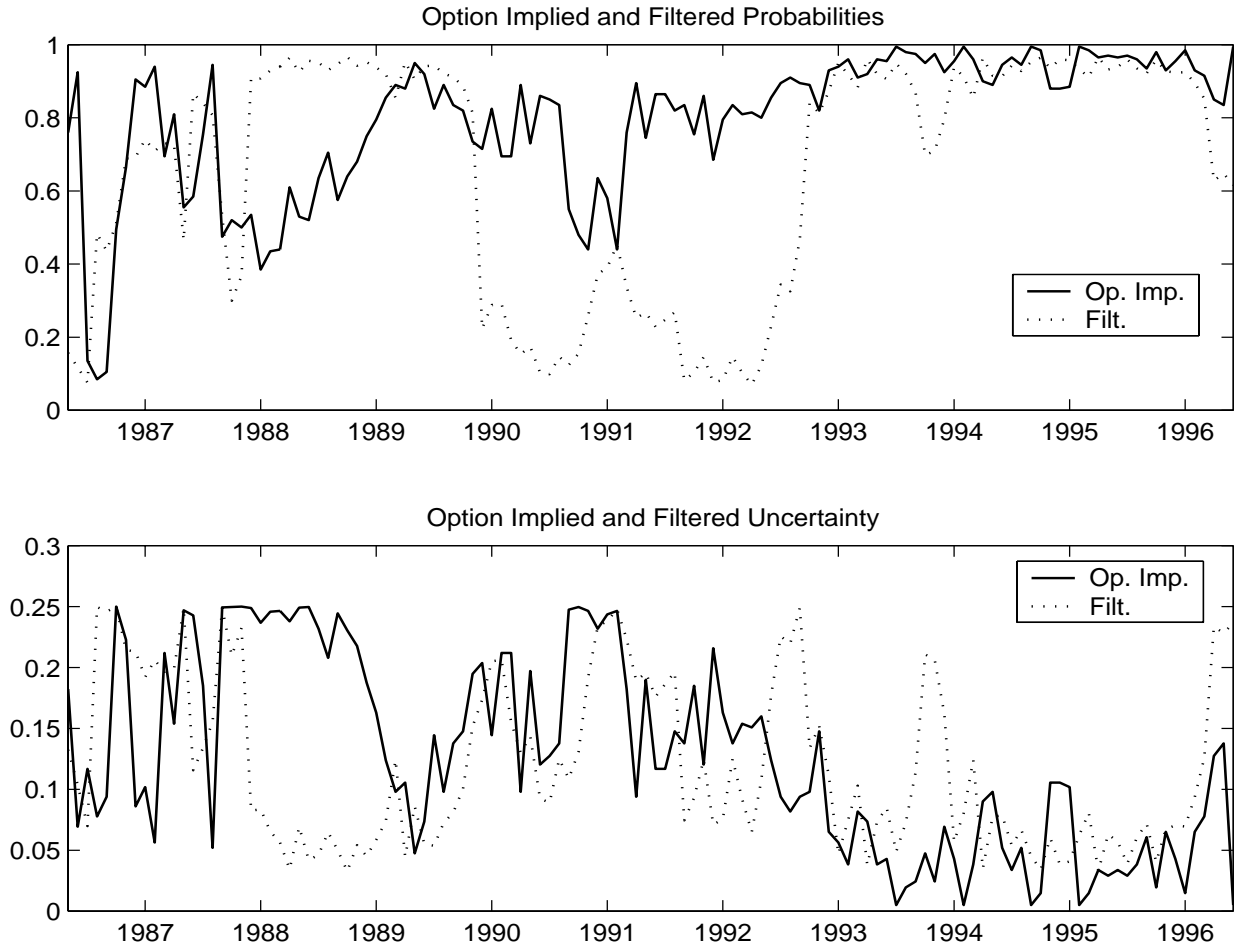
Panel (A) plots historical real earnings growth for S&P 500 firms from 1960:01 to 1998:10. **Panel (B)** plots the filtered posterior probability to be in the high state. **Panel (C)** plots the ATM implied volatility for 1-month options that equates the $C^U(z_t, K, \pi, T - t)$ price as given in (17) with the BS option price, using the belief process in (B) and the S&P 100 implied volatility index (VIX) from 1986:1 to 1998:9. **Panel (D)** plots the covariance between returns and volatility $\text{Cov}(dV, dR) = \sigma_S \sigma'_V$ given by equation (12) in the text. Shaded regions indicate NBER dated recessions of the US economy.

Figure 5: Realized and Fitted 3-Month Conditional Skewness and Kurtosis Premia (1986:4-1996:5)



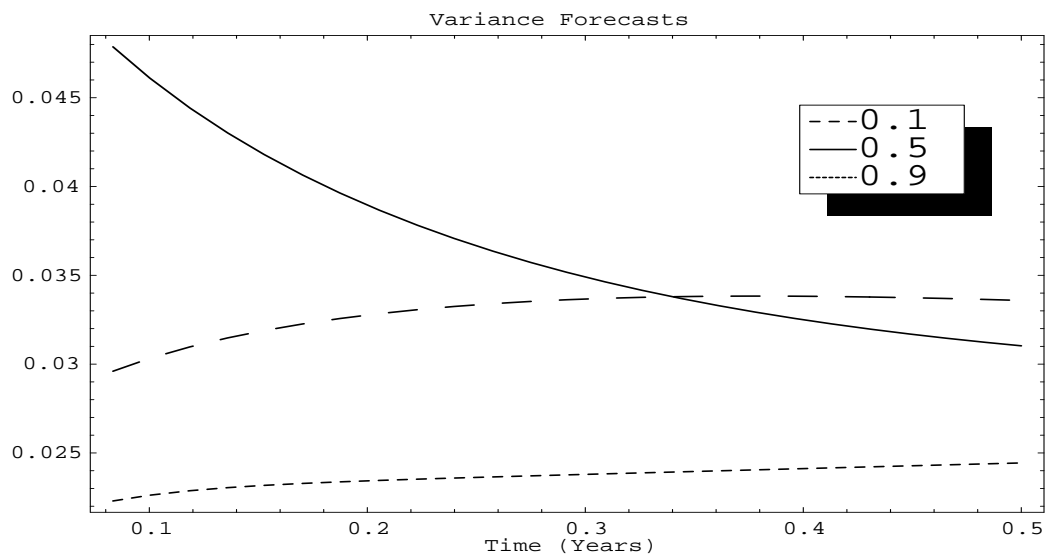
Series are based on data on the 15th (first following trading day) of each month. The realized skewness premium is defined as the ratio of 3% OTM put-to-call prices, $P(K/S_t = 1/1.03)/C(K/S_t = 1.03)$, and similarly, the fitted premium is the theoretical ratio, $P^U(K/S_t = 0.97; \hat{\pi}_t)/C^U(K/S_t = 1.03; \hat{\pi}_t)$, conditional on the filtered belief process, shown in Figure 4, Panel (B). The realized kurtosis premium is defined as $Y_t(3) = [\text{Vol.}(K/S_t = 0.97) + \text{Vol.}(K/S_t = 1.03) - 2 \cdot \text{Vol.}(K/S_t = 1)]$, where, Vol. stands for the BS implied volatility of the option. The fitted ratio is defined analogously, $Y_t^U(3) = [\text{Vol.}^U(K/S_t = 0.97; \hat{\pi}_t) + \text{Vol.}^U(K/S_t = 1.03; \hat{\pi}_t) - 2 \cdot \text{Vol.}^U(K/S_t = 1; \hat{\pi}_t)]$

Figure 6: Estimated Time Series of Filtered and Option-Implied Beliefs



The **Top panel** plots the monthly time series of option-implied beliefs (solid line) and filtered beliefs (dotted line). The **Bottom panel** plots the time series of option-implied “uncertainty” (solid line) and the filtered uncertainty (dotted line). Option-implied beliefs are obtained from the option pricing formula in Proposition 3 in order to minimize the proportional PRMSE as shown in (27). Uncertainty is defined as $\pi_t(1 - \pi_t)$.

Figure 7: Variance Forecasts for 1-month to 6-month Horizons for Alternative Option-Implied Beliefs



Variance forecasts for each option-implied belief, $\tilde{\pi}_t$, are calculated using Proposition 4, using the calibrated set of structural parameters, $\hat{\Theta}^U$ (Table 2).