

A Nonlinear Factor Analysis of S&P 500 Index Option Returns

CHRISTOPHER S. JONES*

ABSTRACT

Growing evidence suggests that extraordinary average returns may be obtained by trading equity index options, and that at least part of this abnormal performance is attributable to volatility and jump risk premia. This paper asks whether such priced risk factors are alone sufficient to explain these average returns. To provide an answer in as general as possible a setting, I estimate a flexible class of nonlinear models using all S&P 500 Index futures options traded between 1986 and 2000. The results show that priced factors contribute to these expected returns but are insufficient to explain their magnitudes, particularly for short-term out-of-the-money puts.

IN A RECENT PAPER, COVAL AND SHUMWAY (2001) examine the returns on delta-hedged option positions, focusing on the at-the-money straddle, and find average returns close to minus 3% per week. Since over short time horizons these positions are approximately invariant to movements in the underlying index, Coval and Shumway argue that their CAPM betas should be near zero and their expected returns close to the risk-free rate. They conclude that large deviations from the CAPM are the norm, which they suggest is evidence that some other systematic factor, such as stochastic volatility, might be priced by the market. Similar conclusions are made by Jackwerth (2000), who finds risk-adjusted profitability from selling puts, and by Bakshi and Kapadia (2003), who find that the volatility risk premium contributes significantly to higher prices for calls and puts. Buraschi and Jackwerth (2001) concur, suggesting that more than one priced risk factor is necessary to explain option prices.

A less direct set of evidence is the now standard observation that the Black-Scholes implied volatilities of equity index options are upward-biased predictors of future realized volatility. Numerous studies, such as Fleming, Ostdiek, and Whaley (1995) and Christensen and Prabhala (1998), document that implied volatilities of short-term at-the-money equity index options are on average several percentage points higher than the volatilities realized over the option's

*Christopher S. Jones is at University of Southern California. I am grateful to Doron Avramov, Luca Benzoni, Josh Coval, Rick Green, Eric Ghysels, John Griffin, Larry Harris, John Long, Spencer Martin, Bruce Mizrach, Allen Potesman, Robert Stambaugh, Robert Tompkins, Harold Zhang, an anonymous referee and associate editor, and seminar participants at Arizona State University, the University of Rochester, the University of Southern California, the University of North Carolina, the 2001 European Finance Association summer meetings, the 2002 American Finance Association meetings, and the 2004 meetings of the Society for Nonlinear Dynamics and Econometrics for many useful comments. All errors are my own.

life. This bias is easily seen in the VIX index (an index of the implied volatility of a hypothetical 22-day at-the-money S&P 100 option), whose average over the period from January 1986 to December 2003 was 21.6%. The realized index volatility over the same period was just 19.2%. Under the Black–Scholes assumptions, at-the-money 1-month put and call options were on average overpriced by 10% to 15% over this period.

Finally, a related conclusion follows from estimating the risk-neutral parameters of stochastic volatility or jump-diffusion models used in option pricing. Both Bakshi, Cao, and Chen (1997) and Bates (2000) argue that the risk-neutral parameters required to fit the Heston (1993) model and more general stochastic volatility models to options prices are unrealistic. Benzoni (2002), Chernov and Ghysels (2000), Jones (2003), and Pan (2002) also find large volatility risk premia. In many of these papers, the price of volatility risk is sufficiently high to induce explosive volatility dynamics under the risk-neutral measure, implying volatility term structures that are too upward-sloping to be empirically plausible.

Recently, these findings have come under some criticism. Branger and Schlag (2004) argue that conclusions about volatility risk premia that are based on option hedging errors are likely biased toward finding risk premia where none exist. One of the problems they identify is the use of continuous-time hedge ratios in hedges over discrete intervals. Another problem is the likely misspecification of the affine stochastic variance model that is typically assumed, a realistic possibility given the deficiencies of this model that Andersen, Benzoni, and Lund (2002) and Jones (2003) identify.

Even more fundamental is the question of whether the risk premia estimated in any of these studies are *sufficient* to explain large negative option returns and the high option prices they imply. The main purpose of this paper is, therefore, to ask whether the extraordinary returns that appear to be possible by short selling put options can be explained by *any* set of priced factors, or whether options appear mispriced under very general assumptions about the nature of systematic risk.

In order to consider as large a class of models as possible, I adopt a semiparametric approach to option pricing and hedging. Specifically, I approximate option hedge ratios, expected returns, and other unknown quantities using flexible but parsimonious sets of orthogonal polynomials. This approach should approximate models that have been proposed already as well as many more models that have not. In addition, I make relatively weak assumptions about the sources of systematic risk in option returns, giving the approach an additional degree of flexibility beyond what the literature typically allows. In particular, I assume that factors contain both observable and unobservable components and that they may depend nonlinearly on observables such as the market return.

The semiparametric approach is related to the research by Hutchinson, Lo, and Poggio (1994) and Garcia and Gençay (2000) that seeks to identify an unknown option pricing formula. Related studies include Ait-Sahalia and Lo (1998), Broadie et al. (2000), and Christoffersen and Hahn (1999). Longstaff (1995) and Buraschi and Jackwerth (2001) use polynomial approximations of

the pricing kernel to test the martingale restriction on option prices. Also related are the semiparametric hedging methods introduced by Bossaerts and Hillion (1997, 2003), which, like this paper, examine factor structures in option returns. The aim of this paper is fundamentally different, however, since its focus, like that of Coval and Shumway (2001), is on modeling expected option returns, rather than realized returns or option prices, and decomposing these expected returns into components related to systematic risk and components attributed to mispricing.

The paper's focus on explaining the entire cross section of option prices also relates it to the burgeoning literature on implied volatility surfaces, which aims to develop tractable models that are consistent with the observed "smile" and "term structure" patterns in implied volatilities. This literature consists of both descriptive papers that document implied volatility dynamics (e.g., Skiadopoulos, Hodges, and Clewlow (1999), Cont and da Fonseca (2002)) and theoretical papers that embed those dynamics into no-arbitrage models of option pricing (e.g., Ledoit, Santa-Clara, and Yan (2002)). The key difference between these papers and this study is that their emphasis is on enforcing no-arbitrage relations between different options, while this paper aims to understand the link between average option returns and more fundamental risk factors. In addition, on a more technical level, volatility surface models typically specify the price process of the underlying asset as a diffusion, which may not be appropriate given growing evidence on the presence of jumps in stock returns (e.g., Eraker, Johannes, and Polson (2003)).

The flexibility of the polynomial approximations and the partially latent nature of the factors makes estimation accuracy particularly important. As such, I adopt a Bayesian approach with uninformative priors, which should provide an approximation to maximum likelihood and therefore result in high efficiency. I also use an extremely large sample of option returns, totaling about 34,000 observations, which should allow a large number of parameters to be estimated with accuracy.

The results suggest that two- or three-factor models are most successful in explaining both expected and realized option returns. The first two factors can usually be interpreted as market and volatility factors, respectively. Simpler one-factor specifications are clearly oversimplistic, and more complex multifactor models offer little improvement in fit to justify their out-of-sample unreliability. Moreover, while models specifically incorporating nonlinearity into the pricing kernel occasionally display some advantages, they are no more useful than similar linear models in explaining average option returns.

The analysis shows further that volatility risk and possibly jump risk are priced in the cross section of index options, but, more importantly, that these systematic risks are insufficient to explain average option returns. In all specifications I consider, short-term deep-out-of-the-money put options appear to be overpriced relative to other options, often generating negative abnormal returns in excess of 0.5% per day.

The outline of the paper is as follows. Section I introduces and motivates the model, and Section II discusses the econometric approach. I present estimation

results in Section III. Section IV contains some in-sample and out-of-sample hedging and portfolio allocation experiments. Section V concludes.

I. A Factor Model of Option Returns

A principal difficulty in estimating a factor model for option returns is the implausibility of constant betas. As a given option contract evolves over time, it changes from a long-run option to a short-run option in an obvious manner. The *moneyiness* of the option, defined as the strike price divided by the forward price, also changes randomly over time. The effects of these changes can be dramatic. Assuming an annualized volatility of 15% and an interest rate of 5%, the Black–Scholes delta of a 22-day call is about 0.99 for an option that is 10% in the money, but only 0.02 for an option that is 10% out of the money. The instantaneous betas of these two options, defined for the Black–Scholes model as

$$\text{B-S delta} \times \frac{\text{underlying price}}{\text{option price}}, \quad (1)$$

are, respectively, 9.6 and 64.4. Given that either of these options could have evolved, say, from a contract that was at the money and had 6 months to expiration (and a beta of 7.7) when it was issued, the magnitude of the time variation in betas has clear importance.

As opposed to time-varying beta approaches such as Shanken (1990) or Ferson and Harvey (1999), in which betas respond to changes in aggregate predictive variables, changing options betas in a Black–Scholes environment are driven by changes in the characteristics of the options themselves. In a more general environment in which volatility is stochastic, option betas will depend on this aggregate variable as well.

The difficulty in dealing with such extreme time variation is mitigated by the fact that while the betas of the different assets are in general unknown, they must be highly related since the assets are all options on the same underlying security. Given the same values for the state variables (such as stochastic volatility), two different options at two different points in time with the same moneyiness and the same maturity must have the same betas. Thus, there is a potential parsimony not enjoyed by studies that examine “unrelated” assets.

A. Motivation in Continuous Time

A diffusion setting provides more concrete motivation. Abstract from possible nonlinearity of the pricing kernel, I focus on asset pricing in an idealized setting.

Suppose that asset price uncertainty is driven by a K -dimensional zero-drift diffusion process S_t . We observe the prices of a variety of options on the same underlying security that differ on the basis of time to expiration and moneyiness. I assume that the price processes of all these options are time-homogeneous and that they obey the stochastic differential equation (SDE)

$$\frac{dP_{it}}{P_{it}} = [r_t + \mu(\tau_{it}, \kappa_{it}, v_t, T_i)] dt + \beta'(\tau_{it}, \kappa_{it}, v_t, T_i) dS_t, \quad (2)$$

where r_t is the instantaneous short rate, $\beta(\cdot)$ is a $K \times 1$ vector, τ_{it} is the time to expiration of option i at time t , κ_{it} is the moneyness (strike/forward price) of option i at time t , v_t is a vector of aggregate-level conditioning variables observed at time t , T_i is a variable indicating the security type (put or call), and where $\mu(\cdot)$ and $\beta(\cdot)$ are assumed continuous (in all arguments except T_i) and identical across all option contracts.

The absence of arbitrage implies the existence of a $K \times 1$ function $\lambda(v_t)$, the vector of market prices of risk, such that

$$\frac{dP_{it}}{P_{it}} = [r_t + \beta'(\tau_{it}, \kappa_{it}, v_t, T_i)\lambda(v_t)] dt + \beta'(\tau_{it}, \kappa_{it}, v_t, T_i) dS_t. \quad (3)$$

Equation (3) implies that instantaneous expected excess returns are linear in factor betas.

In the special case of Black and Scholes (1973), there are no conditioning variables ($v_t = \emptyset$) and the riskless rate is constant. We can therefore define the single noise process as the unexpected component of the market return, or

$$dS_t = dP_t^M / P_t^M - [r + \mu^M] dt, \quad (4)$$

where P_t^M is the value of the market index and

$$dP_t^M / P_t^M = (r + \mu^M) dt + \sigma^M dz_t.$$

Since the price of risk in this case is simply the expected excess return on the market ($\lambda = \mu^M$), the resulting excess returns process,

$$\frac{dP_{it}}{P_{it}} - r dt = \beta(\tau_{it}, \kappa_{it}, T_i) \left(\frac{dP_t^M}{P_t^M} - r dt \right), \quad (5)$$

is, by no arbitrage, consistent with the continuous-time CAPM.

B. Factor Pricing in Discrete Time

Turning to discrete time, it has been well known since Dybvig and Ingersoll (1982) that the CAPM cannot price securities with nonlinear payoffs. The problem is that the CAPM implies a pricing kernel that is linear in the market return, that is, the pricing kernel can become negative in states of high market returns, an arbitrage opportunity that could be exploited with options. As an alternative to standard linear asset pricing models, Bansal and Viswanathan (1993) and Bansal, Hsieh, and Viswanathan (1993) propose the “nonlinear APT,” which models the pricing kernel as a flexible nonnegative function of a small set of observables such as the market return. They find support for their nonlinear model in a set of returns on nonderivative assets and argue that nonlinearities are undoubtedly even more relevant for pricing options.

Dittmar (2002) shows that a nonlinear “one-factor” model outperforms not only linear one-factor models, but linear multifactor models as well.

As is standard, I assume the existence of a pricing kernel $\mathcal{M}_t$ that prices all assets. Defining R_{it} as the time- t excess return on asset i , we have

$$E_t[\mathcal{M}_{t+1}R_{it+1}] = 0. \quad (6)$$

Following Chapman (1997), I assume that the pricing kernel is well approximated by a finite-order polynomial expansion

$$\mathcal{M}_t^* = \frac{1}{1+r_t} + \sum_{k=1}^K h_k F_{kt}, \quad (7)$$

where F_{kt} is a polynomial of some underlying variables adjusted so that $E_{t-1}[F_{kt}] = 0$.

Thus if we write excess returns as

$$R_{it+1} = \mu_{it} + \sum_{k=1}^K \beta_{ikt} F_{kt+1} + \epsilon_{it+1}, \quad (8)$$

it is well known (see Cochrane (2001) for example) that the absence of arbitrage implies there exist λ_{kt} such that

$$\mu_{it} = \sum_{k=1}^K \beta_{ikt} \lambda_{kt}. \quad (9)$$

The only distinction between this result and the one in the previous section is the interpretation of the “factors” F_{kt} , which are now nonlinear functions of observables.

C. A Model of Option Returns

For tractability, I work with the discrete-time model

$$R_{it+1} = \mu(\tau_{it}, \kappa_{it}, v_t, T_i) + \beta'(\tau_{it}, \kappa_{it}, v_t, T_i) F_{t+1} + \epsilon_{it+1}, \quad (10)$$

$$\epsilon_{it+1} \sim N(0, \sigma(\tau_{it}, \kappa_{it}, v_t, T_i)^2). \quad (11)$$

In this equation, R_{it+1} measures the excess return on asset i , $\beta(\cdot)$ represents a $K \times 1$ vector of factor loadings, and F_{t+1} is a yet-unspecified $K \times 1$ vector of mean-zero “factors.” While residual errors are not called for in a theoretical continuous-time model, in practice the realities of options data require their use, with bid-ask spreads, infrequent trading, and nonsynchronicity between different quotes all contributing to this error. I therefore assume that all returns are subject to idiosyncratic conditionally Gaussian shocks. These shocks are assumed independent across assets and across time.

This paper’s focus on options on the market index necessitates that we estimate time-varying betas using an approach that is notably different from that of existing studies such as Shanken (1990) or Ferson and Harvey (1999). In these studies of equity returns, the time- t beta for a particular stock i might

be defined as $\beta_{0i} + \beta'_{1i}v_t$, where the pair (β_{0i}, β_{1i}) must be estimated for each firm. In working with highly related assets such as options on the same underlying security, however, we enjoy a parsimony that comes from the natural smoothness of the betas. For example, two call options that are close to each other both in moneyness and maturity should have betas that are close to one another, at least given the same values for relevant conditioning variables. This smoothness means that rather than estimating contract-specific parameters, it is more appropriate to estimate a *function* for β that applies to *all* contracts.

These considerations suggest the use of an identical approach in specifying $\mu(\cdot)$ and $\sigma(\cdot)$. Sufficient flexibility in the specification of $\mu(\cdot)$ allows both for priced risk factors and the possibility of option mispricing, while flexibility in $\sigma(\cdot)$ permits residual volatilities to systematically differ across contracts and in accordance with movements in economywide variables such as market volatility.

D. Specification of the Factors

Following Bauer and Tamayo (2000), I use economic variables other than returns to help infer, but not completely determine, the realizations of the factors. I therefore assume a nonlinear relation between the $K \times 1$ factor vector and a vector of observable “factor proxies” f_t :

$$F_{t+1} = \psi(f_{t+1}) - E_t[\psi(f_{t+1})] + \eta_{t+1}, \quad (12)$$

where η_t is a $K \times 1$ matrix of i.i.d. residuals that are assumed Gaussian for tractability.

Given that f_t includes variables such as the market return, it might be expected that most of the variation in F_t will be captured by the observable component $\psi(f_t)$. Still, some role for η_t is likely. At the very least, η_t should help correct for “proxy error,” which arises, for example, from using implied volatility as a proxy for actual market volatility.

Note that, for the construction of the factors it is necessary to be able to compute $E_t[\psi(f_{t+1})]$, which requires a model of the factor proxies, $f_{t+1} \sim p_f(\theta_f, v_t)$, where θ_f is a parameter vector. This model, which I describe in detail below, allows dynamics to change with the conditioning variable v_t and is flexible enough to capture the nonnormality apparent in factor proxy variables such as the market return.

Finally, for simplicity, the covariance matrix Ω of the factor equation residuals η_t is assumed to be diagonal. Thus, in addition to its role in producing nonnormality, f_t is also the only possible source of correlation between the factors.

E. Measuring Mispricing

Given $\mu(\cdot)$ and $\beta(\cdot)$ and a specification for factor risk premia, $\lambda(v_t)$, that depends only on aggregate conditioning variables, the pricing restriction in (9) may be rewritten as

$$\mu(\tau_{it}, \kappa_{it}, v_t, T_i) = \beta'(\tau_{it}, \kappa_{it}, v_t, T_i)\lambda(v_t). \quad (13)$$

Without sufficient restrictions on the specifications of any of these functions, the condition will not necessarily hold, giving the specification some testable content.

As in Geweke and Zhou (1996, hereafter GZ), the focus of this paper is on measuring the magnitude of the deviations from the pricing relation (13) rather than testing the point hypothesis of zero mispricing. GZ assume that both the regression parameters and the vector λ of risk prices are constant. Given this framework, GZ examine the following measure of mispricing:

$$Q^2 = \min_{\lambda} \frac{1}{N} (\mu - \beta' \lambda)' (\mu - \beta' \lambda). \quad (14)$$

This is a simple sum of the squared “alphas” of N assets, and λ is defined implicitly as the arg min of this expression.

I adapt the GZ approach to deal with time-varying regression parameters. First, rather than being equally weighted across assets, my measure of mispricing is instead equally weighted across the N return observations, that is,

$$Q^2 = \min_{\lambda(\cdot)} \frac{1}{N} \sum_{i,t} [\mu(\tau_{it}, \kappa_{it}, v_t, T_i) - \beta'(\tau_{it}, \kappa_{it}, v_t, T_i) \lambda(v_t)]^2. \quad (15)$$

Thus, option contracts that appear in the sample over longer periods receive more weight. Second, because the vector of risk prices in my model is not constant, but rather a function of the conditioning variable v_t , the minimization in (15) is actually over the parameters of this function.

II. The Econometric Approach

Since the true functional forms of $\mu(\cdot)$, $\beta(\cdot)$, $\sigma(\cdot)$, $\psi(\cdot)$, and $\lambda(\cdot)$ are unknown, flexibility in their specification is desirable. I use a flexible set of orthogonal polynomial basis functions that are able to generate very accurate approximations while retaining tractability. For $\mu(\cdot)$, $\beta(\cdot)$, and $\sigma(\cdot)$, I estimate functions separately for puts and calls.

A. The Aggregate Time Series Used

In the spirit of standard stochastic volatility models used in the literature, the conditioning variable v_t consists solely of a measure of market volatility. Although it is likely that other state variables influence the mean and covariance structure of option returns, none does so as clearly as market volatility. For instance, while a short-term interest rate may be another obvious candidate to include in v_t , studies such as Bakshi et al. (1997) argue that the impact of this variable is very small, so I omit it from v_t .

The market volatility measure I use is the historical version of the Chicago Board Options Exchange Market Volatility Index (VIX).¹ This carefully

¹ What I refer to in this paper as the VIX index is now called the VXO. Since 2004, the VIX symbol has referred to an index of implied volatilities on S&P 500 Index options. Because this latter series is only available starting in 1990, I do not use it here.

constructed index represents an average of eight Black–Scholes implied volatilities from options on the S&P 100 index; Whaley (1993) describes the construction of the VIX index in detail.

Blair, Poon, and Taylor (2001) find that the VIX is a reliable indicator of future stock market volatility, generally outperforming measures of volatility based on past daily returns. Use of the VIX could be criticized because it is a measure of the volatility of the S&P 100, not the S&P 500, and because, as a measure of implied volatility, it does not necessarily coincide with volatility under the true probability measure. However, as long as the relation between the VIX and the true S&P 500 volatility is approximately deterministic and not extremely nonlinear, the Legendre polynomials should be flexible enough to account for any bias.

In addition to the conditioning variable v_t , the factor proxies f_t must be specified. Since f_t should contain variables that are useful for describing realized option returns, I include the contemporaneous return on the S&P 500 (R_{SP}), the change in the logarithm of the VIX index (ΔVIX), and the change in the logarithm of the 3-month T-bill yield (ΔY_{3M}).² As Chen (2003) argues, volatility is likely to be priced because, in the framework of Merton's (1973) ICAPM, it represents a factor that drives the investment opportunity set. Thus, all three variables are likely to appear in the pricing kernel because they proxy for aggregate investment performance or for changes in future investment prospects.

B. The Approximation Scheme

As in Chapman (1997), I approximate all functions using Legendre polynomials. Betas, in particular, are approximated using third-order Legendre series in three variables, namely log maturity $\ln(\tau_i)$, standardized moneyness $(\kappa_i - 1)/(\sqrt{\tau_i}v_i)$, and log volatility $\ln(v_i)$. This polynomial consists of a constant and the first-, second-, and third-order Legendre polynomials in each argument. It also includes cross terms formed by taking products of the first- and second-order polynomials of different variables. Experiments suggest that this basis is superior to simply using the untransformed variables τ_i , κ_i , and v_i . Finally, since Legendre polynomials are orthogonal on $[-1, 1]$, I linearly rescale all three variables to lie on that interval.

To gain some insight into how this approximation scheme might perform, I explore how the method would fare in estimating put option betas under the Heston (1993) model.³ I generate 25,000 random volatilities, maturities, and moneyness levels by drawing $v_i (i = 1, \dots, 25,000)$ from a uniform distribution between 8% and 40% (annualized), τ_i from a uniform distribution between 10 and 250 days, and κ_i from a uniform distribution between 0.8 and 1.1. Given these three inputs, I calculate several Heston model “betas” similarly to (1). Specifically, I consider the option betas with respect to changes in underlying

² In a slight abuse of notation, ΔVIX and ΔY_{3M} both denote changes in logs rather than levels.

³ The model (with parameter values on a daily frequency) is given by $dS_t = (0.06/264)S_t dt + v_t S_t dB_{1t}$, $dv_t^2 = 0.01(0.0001 - v_t^2)dt + 0.001v_t dB_{2t}$, and $\text{Corr}(dB_{1t}, dB_{2t}) = -0.5$.

price (the Heston model delta times underlying price divided by put price), volatility (using vega), or the interest rate (using rho). The put values are computed using Heston's formula, while the "Greeks" are computed by taking numerical derivatives.⁴

Because the Legendre approximation is linear in parameters, it may be fit by running an ordinary least squares regression. I therefore regress the true Heston model betas on the approximating polynomial to assess whether that polynomial is capable of matching the true nonlinear relation among option characteristics (maturity and moneyness), market volatility, and betas. The results suggest that very tight approximations obtain using this approach. In particular, for the market beta (delta times underlying price divided by put price), the R^2 of the approximation is 0.994. The R^2 of the volatility beta is almost as high, at 0.987, while the goodness of fit for interest rate betas is a full 0.997. In all three cases, visual inspection of the approximation fit reveals no obvious deviations.

In addition to the third-order polynomials for each dimension of the two $\beta(\tau_{it}, \kappa_{it}, v_t)$ vectors (one vector for puts, one for calls), the $\psi(f_t)$, $\lambda(v_t)$, $\mu(\tau_{it}, \kappa_{it}, v_t)$, and $\sigma(\tau_{it}, \kappa_{it}, v_t)$ functions must also be approximated, where the latter two are also defined separately for puts and calls. Following Dittmar's (2002) analysis, I use up to third-order polynomials in the approximation of $\psi(f_t)$, but in the interest of parsimony only include cross terms that are products of first-order polynomials. Furthermore, I include no constants in the approximation of $\psi(f_t)$ since the polynomial is always demeaned according to (12). Each dimension of $\lambda(v_t)$ is specified as a third-order Legendre polynomial, while the $\mu(\tau_{it}, \kappa_{it}, v_t)$ functions are taken to be second-order expansions given that they are expected (like means in general) to be estimated with relatively low accuracy. Finally, I approximate $\ln \sigma(\tau_{it}, \kappa_{it}, v_t)$ identically to each dimension of $\beta(\tau_{it}, \kappa_{it}, v_t)$.

C. A Model of the Factor Proxies f_t

The pricing restriction (9) only holds when the factors are nonzero since only then can $\mu(\cdot)$ be interpreted as the conditional mean of excess returns. As in Chen, Roll, and Ross (1986), I construct factors by removing predictable components of the observable variables from which they are constructed, leading us to define the factor realizations as in (12), where $f_t = \{f_{1,t}, f_{2,t}, f_{3,t}\} = \{R_{SP,t}, \Delta VIX_t, \Delta Y_{3M,t}\}$, as described in Section II.A. We therefore require a model to compute $E_t[\psi(f_{t+1})]$.

Since $\psi(f_t)$ is potentially nonlinear in f_t , the model must be capable of matching not only the mean of f_t , but higher moments as well. It must also account for return heteroskedasticity and a possible relation between conditional means and volatility. To achieve these ends, I use the skewed t -distribution of

⁴ I drop draws resulting in options with prices below 0.001 cents per dollar of notional principal because of difficulties with numerical differentiation and the implausibility of observing these options in practice.

Fernández and Steel (1998) to represent the distribution of the error term. In the resulting model, the change in the log of the VIX index (ΔVIX_t) is

$$\Delta \text{VIX}_t = f_{1,t} = \phi_1(\text{VIX}_{t-1}) + \omega_1(\text{VIX}_{t-1})e_{1,t}, \quad (16)$$

where $e_{1,t} \sim f_{\text{FS}}(s_1, d_1)$ is Fernández and Steel's asymmetric t -distribution with asymmetry parameter s_1 and degrees of freedom parameter d_1 .⁵ Fernández and Steel show that when $s_1 > 1$, the distribution is positively skewed, whereas $0 < s_1 < 1$ generates negative skewness. The degrees-of-freedom parameter, d_1 , of the underlying symmetric t -distribution controls the thickness of both tails of the density of $e_{1,t}$. Combinations of the two can generate markedly non-Gaussian and asymmetric densities.

An important assumption here is that the VIX index is Markovian. However, I do not assume the precise form of the mean and volatility functions. Rather, I again use Legendre polynomials to approximate both $\phi_1(\text{VIX}_{t-1})$ and $\ln \omega_1(\text{VIX}_{t-1})$. Also as before, the polynomials are in the log, rather than level, of VIX_{t-1} . Now, however, the order of the polynomials is chosen according to the Bayes Information Criterion (BIC).

Because some terms in $\psi(f_{t+1})$ involve products of different elements of f_{t+1} , the model must incorporate dependence among the elements of f . I achieve this by "regressing" R_{SP} on the contemporaneous ΔVIX in addition to the lagged information variables. To allow for return predictability and mean reversion in the short rate, these information variables also include the log T-bill yield. Specifically, the models for R_{SP} and ΔY_{3M} are

$$R_{SP,t} = f_{2,t} = \phi_2(\Delta \text{VIX}_t, \text{VIX}_{t-1}, Y_{3M,t-1}) + \omega_2(\text{VIX}_{t-1}, Y_{3M,t-1})e_{2,t} \quad (17)$$

$$\Delta Y_{3M,t} = f_{3,t} = \phi_3(\Delta \text{VIX}_t, R_{SP,t}, \text{VIX}_{t-1}, Y_{3M,t-1}) + \omega_3(\text{VIX}_{t-1}, Y_{3M,t-1})e_{3,t}. \quad (18)$$

Note that $e_{2,t}$ and $e_{3,t}$ also have the skewed t -distributions proposed by Fernández and Steel, though with different parameter values. Again, mean and volatility functions are Legendre polynomials with order determined by the BIC.

The above model is flexible in a number of ways. First, the underlying skewed t -distributions allow the model to generate conditional distributions with potentially large levels of skewness and kurtosis. Second, the model permits the level of the VIX to influence the mean and volatility of some of the elements of f_t without requiring an equal effect on all. Thus, the model allows for the fact that the volatility of market returns should be roughly proportional to VIX_t , while the volatility of changes in the T-bill yield may be only weakly related to the level of the VIX. In addition, the model generates mean reversion in the

⁵ The exact form of the density follows from applying Fernández and Steel's equation (1) to a standard t -distribution with d_1 degrees of freedom. Fernández and Steel's γ parameter is s_1 in my notation.

VIX and the T-bill yield and allows for predictability in market returns. Finally, these effects may be nonlinear if called for by the data.

D. The Option Returns Data

The original data set consists of all transactions of put and call options on S&P 500 Index futures. Both the futures and the options trade on the Chicago Mercantile Exchange. The sample starts in January 1986 and ends in September 2000.

I construct a much smaller daily data set for each option contract by keeping the transaction closest to 3:00 p.m. Transactions more than 15 minutes from this time are discarded, as are transactions in which no corresponding futures price is observed within 1 minute of the option trade. To eliminate at least the largest data errors, I exclude observations that have Black–Scholes implied volatilities (BSIV) less than 1%. I also eliminate large reversals in BSIVs, which include any BSIVs that are three times greater, or one-third the level, of both the previous day's and next day's values for the same contract. To account for the extreme events surrounding the crash of 1987, I retain all observations between October 15 and 22 of that year. Manual inspection of the data over that period reveals no obvious errors. Finally, because of data problems that other studies identify, I only consider options with at least 10 trading days prior to expiration.

The analysis of S&P index futures options is complicated by the fact that they are American. That is, the strategy of buying an option at the close of trade one day, forgoing early exercise, and selling that option at the close of trade the following day may be suboptimal. The option might therefore appear overpriced, with the strategy having a negative mean risk-adjusted return, even though a superior strategy involving early exercise would perform better. Rather than attempting to solve the early exercise problem, I instead examine only out-of-the-money (OTM) options, for which immediate exercise is most likely irrational.⁶

From the above data, I calculate returns only from options that were observed on consecutive days. The resulting data set consists of 33,928 observations of 1-day option returns, of which 58% are puts and 42% are calls. The average number of returns observed per day is 9.4, with a minimum of 1 and a maximum of 40. Table I reports the composition of the sample, sorted by strike and time to expiration.

It is apparent that the greatest number of observations correspond to slightly OTM puts ($0.9 < \kappa \leq 1.0$). Returns are observed over a widest range of strike prices for moderately short maturities. Deep-OTM call returns are relatively

⁶ The sample includes the return from time t to $t + 1$ on a particular option if the option is OTM on day t . Thus, there is no look-ahead bias that would result from inclusion based on the time $t + 1$ price. It is possible, therefore, that the option might be in the money at time $t + 1$ and that early exercise some time during day $t + 1$ might have been optimal. Results from Whaley (1986) and others suggest that the early exercise premium is likely to be less than 1% of the value of the options in this paper's sample, implying that the value of forgoing early exercise for 1 day should be extremely small.

Table I
Sample Size by Time to Expiration and Moneyness

The table lists the number of 1-day option returns observed for a variety of maturity and moneyness categories. Maturity, denoted by τ , is equal to the number of trading days until option expiration. Moneyness, denoted by κ , is equal to the present value of the strike price divided by the current price of the underlying security, in this case the S&P 500 Index. Since all in-the-money options are discarded, all puts have $\kappa \leq 1$ and all calls have $\kappa \geq 1$.

	Puts			Calls		
	$\kappa \leq 0.8$	$0.8 < \kappa \leq 0.9$	$0.9 < \kappa \leq 1.0$	$1.0 < \kappa \leq 1.1$	$1.1 < \kappa \leq 1.2$	$1.2 < \kappa$
$\tau \leq 22$	92	739	7,655	6,042	52	3
$22 < \tau \leq 44$	228	889	5,912	4,811	161	86
$44 < \tau \leq 66$	96	397	2,231	1,862	146	3
$66 < \tau \leq 128$	56	280	1,018	999	63	4
$128 < \tau$	8	15	41	36	3	0

rare, as are all options with expirations more than 6 months (128 trading days) away.

E. The Estimation Procedure

I estimate the option returns model (10) in a Bayesian framework under flat priors (exactly proportional to a constant) using a Gibbs sampling approach similar to Geweke and Zhou (1996). This method alternates between a step involving draws of the parameter vector and a “data augmentation” step in which the latent factors $\mathbf{F}$ are drawn conditional on those parameters. The Appendix contains details. Under general conditions, a sequence of such draws converges to the unknown posterior.

A caveat is that the parameters θ_f that enter the model of the factor proxies in (16)–(18) are estimated without using information in the option returns; rather, they are estimated by maximum likelihood using only the structure imposed by those three equations. In principle, option returns can be helpful in estimating θ_f because these parameters determine the expectation of $\psi(f_{t+1})$, which appears in (12) as a component of F_{t+1} . Unfortunately, I find no feasible sampling algorithm that allows a draw of θ_f conditional on both the factor proxy variables and the option returns, so the resulting “pseudo-posteriors” must be interpreted with this qualification in mind.

III. Estimation Results

Altogether, 10 different specifications are investigated. These differ in three dimensions, the first of which is the definition of f_t . In the most general model, f_t includes the market return (R_{SP}), the log change in the VIX index (ΔVIX), and the log change in the 3-month T-bill yield (ΔY_{3M}). “Stochastic volatility” specifications include just the first two variables, while “market model” specifications include only the market return.

Models also differ according to whether they allow nonlinearity in the relation between the unobservable factors F_t and the factor proxies f_t . In “linear”

specifications, the function $\psi(f_t)$ from (12) is linear in f_t . “Nonlinear” specifications use the third-order expansions for $\psi(\cdot)$ discussed in Section II.B. Other forms of nonlinearity, such as in the relation between option characteristics and betas, are present in all specifications. The “linear” and “nonlinear” labels strictly refer to the assumed properties of the factors and hence the pricing kernel.

Finally, the number of factors, K , must be chosen. Only for linear specifications, where there is one factor for each factor proxy, is it the case that K must equal the dimension of f_t . For nonlinear specifications, K can be larger or smaller than the dimension of f_t . For example, a “nonlinear market model” specification might have multiple factors, though all the factors are different polynomials in the same market return. On the other hand, a flexible one-factor model could allow the single factor to depend nonlinearly on R_{SP} , ΔVIX , and ΔY_{3M} .

For each specification, I run a Gibbs sampler for 110,000 iterations, discarding the first 10,000 iterations to negate the effects of initial conditions. The remaining 100,000 draws of the parameter vector, which are assumed to come from the posterior distribution, form the basis for all inferences drawn in the remainder of the paper.

A. Parameter Estimates

Table II presents the estimated models used to compute expectations of $\psi(f_{t+1})$. As mentioned above, dependencies among the models are captured by first modeling ΔVIX , then R_{SP} conditional on ΔVIX , and then finally ΔY_{3M} conditional on both ΔVIX and R_{SP} . The polynomial orders of all terms are chosen according to the Bayesian Information Criterion, errors have the asymmetric student- t distribution of Fernández and Steel (1998), and models are estimated by maximum likelihood.

For the most part, the estimated model for ΔVIX conforms with the results of previous studies. Mean reversion is demonstrated by the significant negative coefficients on the first and second Legendre polynomials in the lagged logarithm of the VIX index, denoted by $\mathcal{L}_1(VIX_{t-1})$ and $\mathcal{L}_2(VIX_{t-1})$, where they appear in the mean of ΔVIX . Consistent with Jones (2003), nonlinear level-dependence in the volatility of ΔVIX is also apparent, as both $\mathcal{L}_1(VIX_{t-1})$ and $\mathcal{L}_2(VIX_{t-1})$ have positive coefficients where they appear in the volatility of ΔVIX . Finally, the estimated parameters of the asymmetric student- t distribution indicate positive skewness (since $\hat{s} = 1.08 > 1$) and fat tails (degrees of freedom $\hat{d} = 4.15 \ll 30$). These results are consistent with Eraker et al.’s (2003), who find that volatility is likely discontinuous and therefore conditionally non-Gaussian.

The estimated model for R_{SP} reaffirms the standard finding that market returns are negatively correlated with changes in volatility. I find that this dependence is essentially linear, as only the first Legendre polynomial in contemporaneous ΔVIX , denoted $\mathcal{L}_1(\Delta VIX_t)$, is significant. Other coefficients are somewhat more difficult to interpret since they serve to describe only the

Table II
Preferred Models of the Factor Proxy Variables

The following equations represent estimated models of ΔVIX , R_{SP} , and ΔY_{3M} . Each model is estimated by maximum likelihood using daily data from January 1986 to September 2000. In each equation, $\mathcal{L}_n(x)$ represents the n th-order Legendre polynomial in the variable x . Values in parentheses underneath the estimated parameter values are asymptotic standard errors. In each equation, the order of the Legendre polynomial in each explanatory variable is chosen according to the Bayesian Information Criterion. $f_{FS}(s, d)$ denotes the asymmetric t -distribution of Fernández and Steel (1998) with skewness parameter s and degrees-of-freedom parameter d .

VIX equation

$$\begin{aligned} \Delta VIX_t = & -0.011 - 0.016 \mathcal{L}_1(VIX_{t-1}) - 0.0063 \mathcal{L}_2(VIX_{t-1}) + \exp[-3.10 + 0.38 \mathcal{L}_1(VIX_{t-1}) \\ & (0.002) \quad (0.003) \quad (0.0027) \quad (0.02) \quad (0.03) \\ & + 0.15 \mathcal{L}_2(VIX_{t-1})] e_{1,t} \\ & (0.03) \\ e_{1,t} \sim & f_{FS}(1.08, 4.15) \\ & (0.03)(0.28) \end{aligned}$$

Return equation

$$\begin{aligned} R_{SP,t} = & -0.0026 - 0.015 \mathcal{L}_1(\Delta VIX_t) - 0.0013 \mathcal{L}_2(\Delta VIX_t) + 0.000051 \mathcal{L}_3(\Delta VIX_t) \\ & (0.0003) \quad (0.002) \quad (0.0002) \quad (0.000031) \\ & - 0.0018 \mathcal{L}_1(VIX_{t-1}) + 0.00055 \mathcal{L}_1(Y_{3M,t-1}) + \exp[-5.15 + 0.69 \mathcal{L}_1(VIX_{t-1})] e_{2,t} \\ & (0.0003) \quad (0.00019) \quad (0.02) \quad (0.04) \\ e_{2,t} \sim & f_{FS}(1.17, 4.73) \\ & (0.03)(0.42) \end{aligned}$$

T-bill equation

$$\begin{aligned} \Delta Y_{3M,t} = & -0.000020 - 0.0015 \mathcal{L}_1(\Delta VIX_t) - 0.0021 \mathcal{L}_1(R_{SP,t}) - 0.00022 \mathcal{L}_1(Y_{3M,t-1}) \\ & (0.000232) \quad (0.0004) \quad (0.0004) \quad (0.00021) \\ & + \exp[-4.99 + 0.42 \mathcal{L}_1(VIX_{t-1}) + 0.090 \mathcal{L}_2(VIX_{t-1}) - 0.22 \mathcal{L}_1(Y_{3M,t-1})] \\ & (0.03) \quad (0.04) \quad (0.034) \quad (0.04) \\ & + 0.078 \mathcal{L}_2(Y_{3M,t-1})] e_{3,t} \\ & (0.036) \\ e_{3,t} \sim & f_{FS}(0.99, 2.98) \\ & (0.02) \quad (0.19) \end{aligned}$$

portion of R_{SP} that is not explained by contemporaneous changes in ΔVIX . Nevertheless, I find negative mean dependence and positive volatility dependence on the lagged VIX index ($\mathcal{L}_1(VIX_{t-1})$) and positive dependence on the lagged T-bill rate ($\mathcal{L}_1(Y_{3M,t-1})$). Residuals in this model are again fat-tailed, with $\hat{d} = 4.73$, and somewhat surprisingly are positively skewed, as $\hat{s} > 1$. (Overall, however, R_{SP} is negatively skewed because of the asymmetry it inherits from its negative dependence on ΔVIX , which is positively skewed.)

The model for ΔY_{3M} displays contemporaneous relations with both ΔVIX and R_{SP} , though the lagged T-bill yield ($\mathcal{L}_1(Y_{3M,t-1})$) is insignificant, indicating a lack of mean reversion in the sample. The dependence of the volatility of ΔY_{3M} is nonlinear in both lagged VIX and lagged Y_{3M} , and the residuals are fat-tailed ($\hat{d} = 2.98$) though not significantly skewed.

Since the coefficients of the Legendre expansions of $\mu(\cdot)$, $\beta(\cdot)$, $\sigma(\cdot)$, and $\psi(\cdot)$ are difficult to interpret, I do not report them in the paper. Instead, Table III gives a summary of the number of parameters in each specification and the

Table III
Parameter Posterior Summary

The table presents, for each of the 10 specifications considered in the paper, a variety of summary statistics on the number of factor model parameters and their posterior precisions. The columns labeled “signed/total” report the number of parameters in a particular group for which 95% of the posterior distribution is on one side of zero, as well as the total number of parameters within that group. (Since parameters in Ω are all signed a priori, only their total number is reported.) Columns labeled “median t -ratio” report the median absolute value of the ratios of univariate posterior means to posterior standard deviations. In cases in which there are no results for the parameters of the ψ function, all the parameters of that function are implied by identification restrictions and are therefore not estimated. All models are estimated using daily data from January 1986 to September 2000.

	K	Signed/Total Number of Parameters	Median t -Ratios					
			All	μ	β	σ	ψ	Ω
Linear specifications								
R_{SP}	1	81/97	4.51	1.86	7.79	4.09		44.14
R_{SP} and ΔVIX	2	108/136	3.94	1.31	5.86	2.58		46.44
R_{SP} , ΔVIX , and ΔY_{3M}	3	121/175	2.98	1.13	3.46	2.53		35.00
Nonlinear specifications								
R_{SP}	1	80/99	5.19	2.21	9.27	3.75	40.21	54.23
R_{SP}	3	137/175	4.40	3.39	5.49	2.53		38.58
R_{SP} and ΔVIX	2	117/146	4.80	2.72	6.35	2.24	7.89	45.15
R_{SP} and ΔVIX	6	199/298	2.43	2.96	2.41	2.28	2.21	8.23
R_{SP} , ΔVIX , and ΔY_{3M}	1	86/108	4.45	2.21	8.31	3.37	3.73	51.61
R_{SP} , ΔVIX , and ΔY_{3M}	3	154/202	3.20	2.42	3.60	3.19	1.85	32.28
R_{SP} , ΔVIX , and ΔY_{3M}	6	133/328	1.08	2.46	1.01	2.42	0.75	2.17

overall accuracy with which they are estimated. In the linear market model specification, for example, there are 97 parameters, out of which 81 can be signed with at least 95% probability, that is, for each of the 81, at least 95% of the marginal posterior distribution is on one side of zero. Finally, the median absolute “ t -ratio” (posterior mean divided by standard deviation) is equal to 4.51, indicating that most posterior means are far from zero.

Generally, models with three or fewer factors produce relatively precise posterior distributions, with most parameters signed with high probability and most t -ratios above two. For the more restricted specifications, parameters of the $\mu(\cdot)$ function are less accurately estimated than are parameters of the $\beta(\cdot)$ or $\sigma(\cdot)$ functions. For the six-factor specifications, even the $\beta(\cdot)$ functions are difficult to estimate. For the most general model, which is shown on the last line of the table, the median t -ratio is just 1.01, suggesting that this model may imply an overly complex covariance structure.

B. Explained Variation in Realized Returns

The analysis of model performance begins by examining how well the models explain realized option returns. Since we are outside the linear regression

framework, the measurement of such variation is somewhat ambiguous. I settle on a version of the standard R -squared,

$$1 - \text{Var}(\epsilon_{it})/\text{Var}(R_{it}),$$

where ϵ_{it+1} is defined via (10). Because residuals are calculated conditional on a particular set of parameters and latent factors, we obtain a posterior distribution of each R^2 rather than a single value.

Table IV presents posterior means and standard deviations of these R -squares for different classes of options, classified both by maturity and moneyness. Short- and long-maturity classes refer to values of τ from 10 to 44 and 45 to 264 days, respectively. I also split puts into two different moneyness categories, those with $\kappa \leq 0.92$ and those with $0.92 < \kappa \leq 1$. Because of the scarcity of deep-OTM call option observations, all calls are grouped together within each maturity class.

The variances implicit in each calculation are those of observations pooled both across contracts and over time, and in a nonlinear model there is no guarantee that R -squares will increase with the number of factors or even that they will be positive in all cases. The first row of Table IV gives an example of such an anomaly, as the linear market model displays an “ R -squared” of -2.296 for short-term deep-OTM put options. The interpretation of R -squares closer to one is much more straightforward, and most of the models do not display the gross failings of the first specification.

In general, only two or three factors appear to be necessary to explain the vast majority of realized returns across all the different categories of moneyness and maturity, explaining close to 90% of the variation in all options and in calls and puts separately. One-factor specifications all display some difficulty in fitting at least one of the maturity/moneyness groups, though the nonlinear one-factor models are vastly superior to their linear counterparts. For two- or three-factor models, allowing nonlinearity in the factors yields little benefit. Finally, adding additional factors beyond the first two or three offers little improvements in fit, especially considering that the measure used is unadjusted for differences in the number of parameters or factors.

C. Mispricing

Similar to the stochastic volatility/jump-diffusion models common in the literature, risk premia in the model are functions only of the level of the volatility proxy, v_t . For given $\mu(\cdot)$ and $\beta(\cdot)$, I follow Geweke and Zhou (1996) and measure mispricing by choosing the coefficients of the Legendre polynomials in $\lambda(v_t)$ to minimize Q^2 , the mean squared pricing error defined in (15). With the solution in hand, we may examine the mispricing functions

$$\alpha(\tau_{it}, \kappa_{it}, v_t, T_i) \equiv \mu(\tau_{it}, \kappa_{it}, v_t, T_i) - \beta(\tau_{it}, \kappa_{it}, v_t, T_i)\lambda(v_t) \quad (19)$$

separately for puts and calls, which measure the amount by which option expected returns exceed their risk-adjusted benchmark levels. Of primary interest is whether these functions are reliably nonzero for any ranges of τ and κ .

Tables V and VI contain two measures of mispricing for a variety of classes of options. Table V reports the posterior means and standard deviations of various measures of root mean squared mispricing. For the full sample, root mean squared mispricing is equal to Q , from (15); for specific moneyness/maturity categories it is computed similarly over a subset of the returns data. Table VI reports the posterior means and standard deviations of average alphas, indicating whether particular option classes are overpriced or underpriced.

Several patterns emerge from the two tables. First, from Table V, the magnitude of mispricing across all OTM options is generally large. The most general six-factor model, for instance, has a root mean squared mispricing of 0.53, indicating that the typical “alpha” is around 0.5% (plus or minus) per day. One-factor specifications imply mispricing measures between two and three times as high. Average mispricing across all options, as reported in Table VI, is close to zero, though this result is expected given the way that the mispricing measures are constructed.

Second, mispricing is primarily reduced by increasing the number of factors. Nonlinearity, by itself, does not reduce root mean squared pricing errors even if it can better explain realized returns (as Table IV showed). For example, the best-performing model based on the results in Table V is the linear three-factor model, with performance that is actually superior to both nonlinear three-factor models. While Table VI is somewhat more ambiguous about which specification is superior, in many cases the linear three-factor model outperforms the nonlinear models in terms of pricing biases as well.

Perhaps the most important result in both tables is that no model explains the negative average returns associated with short-term deep-OTM put options. At a minimum, the “alpha” on deep-OTM puts is estimated at about -0.33% per day, and most specifications imply substantially more negative values. Thus, the table confirms the overpricing of short-term OTM puts that Jackwerth (2000) and others observe by looking at risk-neutral densities and that Coval and Shumway (2001) find by examining option return time series.

At first glance, Table VI provides an equally surprising result, namely, that longer-term puts and calls generally have *positive* alphas. This should be interpreted with caution, however, given that Blume and Stambaugh (1983) show that the use of transactions prices should induce a positive bias in average returns as a result of bid-ask bounce. More fundamentally, however, the prices of risk are chosen via a minimization process in (15) that implies, by construction, that alphas will approximately average out to zero. For some options to have negative alphas, others must therefore have positive alphas. Thus, the results are more accurately interpreted as stating that short-term OTM puts are overvalued *relative to* other option classes.

D. A Closer Look at Several Models

Bansal and Viswanathan (1993) and Dittmar (2002) both find substantial benefits in modeling the pricing kernel as a nonlinear function of the market return. While the results in Tables V and VI are somewhat more negative about

Table V
Root Mean Squared Mispricing

The table reports posterior means and standard deviations (in parentheses) of the root mean squared pricing error, or Q , defined as the square root of

$$\min_{\lambda(\cdot)} \frac{1}{N} \sum_{i,t} [\mu(\tau_{it}, \kappa_{it}, v_t, T_i) - \beta'(\tau_{it}, \kappa_{it}, v_t, T_i) \lambda(v_t)]^2.$$

These values are reported in the column labeled “ Q .” While the above minimization is performed using all N return observations, the $\lambda(\cdot)$ functions that result are then used to construct similar root mean squared mispricing measures using specific maturity and moneyness subsets of the return observations. Within the column headings, “short maturity” refers to all options with less than 2 months until expiration and “long maturity” refers to expirations longer than 2 months. Option moneyness, denoted by κ , is defined as the present value of the strike price divided by the current price of the underlying index. All values result from posteriors computed using daily data from January 1986 to September 2000.

	K	All OTM Options	All OTM Puts	All OTM Calls	Short Maturity			Long Maturity		
					$\kappa \leq 0.92$ Puts	$0.92 < \kappa \leq 1$		$\kappa \leq 0.92$ Puts	$0.92 < \kappa \leq 1$	
						Puts	Calls		Puts	Calls
Linear specifications										
R_{SP}	1	1.25 (0.12)	0.91 (0.11)	1.59 (0.20)	1.76 (0.28)	0.62 (0.07)	1.70 (0.22)	1.09 (0.20)	0.58 (0.10)	1.07 (0.20)
R_{SP} and ΔVIX	2	0.64 (0.08)	0.70 (0.10)	0.55 (0.10)	1.37 (0.24)	0.44 (0.06)	0.57 (0.11)	0.95 (0.18)	0.47 (0.09)	0.41 (0.09)
R_{SP} , ΔVIX , and ΔY_{3M}	3	0.55 (0.08)	0.62 (0.11)	0.40 (0.08)	1.29 (0.27)	0.35 (0.07)	0.43 (0.09)	0.69 (0.17)	0.45 (0.09)	0.28 (0.08)
R_{SP}	1	1.46 (0.13)	1.04 (0.09)	1.88 (0.22)	2.05 (0.24)	0.72 (0.07)	2.04 (0.24)	1.08 (0.19)	0.57 (0.09)	1.20 (0.21)
Nonlinear specifications										
R_{SP}	3	0.93 (0.10)	1.02 (0.11)	0.78 (0.11)	2.26 (0.29)	0.55 (0.07)	0.85 (0.12)	1.00 (0.14)	0.47 (0.08)	0.46 (0.07)
R_{SP} and ΔVIX	2	1.04 (0.09)	1.12 (0.11)	0.89 (0.14)	2.28 (0.25)	0.77 (0.08)	0.97 (0.15)	1.04 (0.12)	0.57 (0.08)	0.51 (0.08)
R_{SP} and ΔVIX	6	0.51 (0.06)	0.51 (0.08)	0.51 (0.07)	1.02 (0.21)	0.26 (0.04)	0.50 (0.07)	0.79 (0.16)	0.37 (0.08)	0.54 (0.10)
R_{SP} , ΔVIX , and ΔY_{3M}	1	1.46 (0.12)	1.06 (0.09)	1.87 (0.21)	2.04 (0.25)	0.75 (0.07)	2.02 (0.23)	1.10 (0.20)	0.63 (0.09)	1.17 (0.21)
R_{SP} , ΔVIX , and ΔY_{3M}	3	0.66 (0.08)	0.72 (0.11)	0.57 (0.09)	1.50 (0.27)	0.44 (0.06)	0.62 (0.10)	0.62 (0.12)	0.46 (0.09)	0.35 (0.07)
R_{SP} , ΔVIX , and ΔY_{3M}	6	0.53 (0.07)	0.55 (0.09)	0.50 (0.08)	1.07 (0.24)	0.30 (0.07)	0.48 (0.09)	0.82 (0.25)	0.39 (0.11)	0.55 (0.10)

Table VI
Mean Mispricing

The table reports posterior means and standard deviations (in parentheses) of the average pricing error, or average “alpha,” defined as

$$\frac{1}{N} \sum_{i,t} [\mu(\tau_{it}, \kappa_{it}, v_t, T_t) - \beta'(\tau_{it}, \kappa_{it}, v_t, T_t) \lambda(v_t)],$$

where $\lambda(\cdot)$ functions are defined as in Table VI. In addition to average alphas computed using all N return observations, similar average pricing errors are computed using specific maturity and moneyness subsets of the return observations. Within the column headings, “short maturity” refers to all options with less than 2 months until expiration, “medium maturity” refers to options expiring within 2 to 4 months, and “long maturity” refers to expirations longer than 4 months. Option moneyness, denoted by κ , is defined as the present value of the strike price divided by the current price of the underlying index. All values result from posteriors computed using daily data from January 1986 to September 2000.

	K	All OTM Options	All OTM Puts	All OTM Calls	Short Maturity			Long Maturity		
					$\kappa \leq 0.92$	$0.92 < \kappa \leq 1$	$1 < \kappa$	$\kappa \leq 0.92$	$0.92 < \kappa \leq 1$	$1 < \kappa$
					Puts	Puts	Calls	Puts	Puts	Calls
Linear specifications	1	-0.01 (0.10)	-0.02 (0.09)	0.00 (0.13)	-0.98 (0.24)	0.00 (0.10)	-0.20 (0.15)	0.86 (0.23)	0.45 (0.11)	0.71 (0.17)
	2	0.00 (0.03)	0.05 (0.02)	-0.06 (0.04)	-0.50 (0.16)	0.03 (0.05)	-0.10 (0.04)	0.67 (0.23)	0.39 (0.10)	0.10 (0.12)
	3	0.03 (0.03)	0.06 (0.03)	-0.02 (0.04)	-0.73 (0.18)	0.12 (0.05)	-0.04 (0.04)	0.50 (0.24)	0.40 (0.10)	0.07 (0.12)
	1	0.02 (0.10)	-0.05 (0.09)	0.12 (0.13)	-1.14 (0.24)	0.00 (0.10)	-0.05 (0.16)	0.91 (0.23)	0.39 (0.11)	0.72 (0.18)
	3	-0.05 (0.03)	-0.01 (0.03)	-0.10 (0.04)	-1.10 (0.16)	0.14 (0.04)	-0.12 (0.04)	0.21 (0.20)	0.32 (0.09)	-0.01 (0.11)
Nonlinear specifications	2	-0.06 (0.03)	-0.03 (0.02)	-0.12 (0.04)	-0.99 (0.18)	0.04 (0.04)	-0.17 (0.04)	0.55 (0.22)	0.37 (0.09)	0.06 (0.11)
	6	0.05 (0.02)	0.05 (0.02)	0.04 (0.03)	-0.37 (0.11)	0.06 (0.04)	0.00 (0.03)	0.53 (0.17)	0.22 (0.09)	0.22 (0.12)
	1	0.01 (0.10)	-0.06 (0.09)	0.11 (0.13)	-1.21 (0.23)	0.00 (0.10)	-0.05 (0.16)	0.93 (0.24)	0.43 (0.11)	0.69 (0.17)
	3	0.01 (0.02)	0.05 (0.02)	-0.04 (0.04)	-0.89 (0.16)	0.16 (0.04)	-0.06 (0.03)	0.33 (0.20)	0.38 (0.10)	0.01 (0.11)
	6	0.05 (0.02)	0.05 (0.03)	0.05 (0.04)	-0.33 (0.13)	0.06 (0.04)	0.01 (0.03)	0.49 (0.29)	0.22 (0.11)	0.19 (0.12)

this kernel’s ability to explain the cross section of average option returns, I briefly compare the linear and nonlinear one-factor “market models,” in which the return on the S&P 500 Index (R_{SP}) is the sole element of the factor proxy vector, f_t .

The pricing kernel and the single factor are related (see Cochrane (2001)) via

$$\mathcal{M}_t = \frac{1}{1 + r_t} + \frac{\lambda(v_t)}{\text{Var}_{t-1}(F_t)} F_t, \tag{20}$$

so the relation between f_t and F_t translates directly into a relation between f_t and $\mathcal{M}_t$. Figure 1 displays posterior distributions of the pricing kernel expected as a function of R_{SP} , the return on the S&P 500 Index, where the expectation comes from setting the error term η_t in the factor equation (12) to zero. The figure shows kernels calculated for both the linear and nonlinear market models under both normal and high volatility environments.




Figure 1. Shape of the pricing kernel in one-factor market models. Each panel reports posterior means and 90% confidence intervals of the expected pricing kernel as a function of R_{SP} , the return on the S&P 500 Index. Panels in the top row report results for the linear one-factor model (where $\psi(f_t) = R_{SP}$), while the bottom panel contains results for the corresponding nonlinear model (where $\psi(f_t) = \psi(R_{SP})$). Pricing kernels are computed conditional on two different values of the VIX index.

It is clear from Figure 1 that nonlinearities are apparent in the way that one would expect from examining models such as Black–Scholes. Unlike Black–Scholes, however, in which the shape of the pricing kernel remains constant over time, Figure 1 shows that there are significant differences in the slope of the kernel between the two volatility regimes, with the linear kernels flattening and the nonlinear kernels steepening when volatility is high. Least squares fits of the Black–Scholes pricing kernel to the posterior means displayed in the bottom two panels of the figure result in relative risk aversion estimates of about 7.4 for the normal volatility environment and 11.3 for high volatility.

Another model that is similar to those currently popular in the option pricing literature is the nonlinear three-factor model in which $f_t = \{R_{SP,t}, \Delta VIX_t, \Delta Y_{3M,t}\}$. As with traditional option pricing models, this model includes what one might easily term a “market factor.” While the first factor could have a virtually arbitrary relation with R_{SP} , ΔVIX , and ΔY_{3M} , it turns out that it is closely linked only with R_{SP} . Figure 2 plots posterior means of the first factor against the contemporaneous realizations of R_{SP} and reveals a very tight and approximately linear relation between the two. The interpretation of the first factor as a generic market factor is therefore clear.

The second factor in the same model is only slightly more difficult to interpret. Figure 3 plots the posterior mean of the second element of $\psi(f_t) - E_{t-1}[\psi(f_t)]$ as a function of R_{SP} and ΔVIX (as, again, ΔY_{3M} appears to be unimportant). The figure therefore tells us what we expect the second factor to be given f_t ,




Figure 2. Posterior means of the first factor in a three-factor nonlinear model. This figure displays the relation between R_{SP} , the realized return on the S&P 500 Index, and the posterior mean of the first factor in the most general three-factor nonlinear model (in which $f_t = \{R_{SP,t}, \Delta VIX_t, \Delta Y_{3M,t}\}$).




Figure 3. Interpretation of the second factor in a three-factor nonlinear model. This figure plots the posterior mean of the second dimension of $\psi(f_t) - E_{t-1}[\psi(f_t)]$ as a function of R_{SP} , the return on the S&P 500 Index, and ΔVIX , the change in the log of the VIX index. This represents the value of the second factor that would be expected given the observation of R_{SP} and ΔVIX . The plot is computed for the most general three-factor nonlinear model (in which $f_t = \{R_{SP,t}, \Delta VIX_t, \Delta Y_{3M,t}\}$).

and it is apparent that this expectation is positively related both to ΔVIX and to the square of R_{SP} . I therefore label this the “volatility” factor.

Unlike the first two factors, the third factor’s relation with the factor proxy variables is hard to characterize. Rather than plotting it against an observable series, Figure 4 simply provides a time-series plot of the posterior means of




Figure 4. Time series of the third factor in a three-factor nonlinear model. This figure plots the time series of posterior means of the third factor in the most general three-factor nonlinear model (in which $f_t = \{R_{SP,t}, \Delta VIX_t, \Delta Y_{3M,t}\}$).




Figure 5. Risk premia in a three-factor nonlinear model. Panels in the top row display the prices of risk, $\lambda(v_t)$, of the three factors in a nonlinear three-factor model. Posterior means and 90% confidence intervals are plotted as functions of the lagged value of the VIX index. The middle panel shows posterior means and 90% confidence intervals of the model-implied betas, as a function of the VIX index, for a 1-month put option with a strike price equal to 0.95 times the current forward price. The bottom row displays products of the first two rows, indicating the risk premia received for each of the factors by the same 1-month put.

the third factor. What is immediately evident is that this factor's time-series variation is dominated by the crash and subsequent turmoil of October 1987, and that in general the third factor is very close to zero. Other large spikes correspond to the stock market crash of 1989, the Russian financial crisis of 1998, and the dot-com crash of 2000. Because most of these events coincide with large 1-day returns, which are usually negative, I refer to this as the "jump" factor.

Each of these three factors is assumed to command a risk premium that is a function of the level of the VIX index. Figure 5 displays, in its top row, the price of risk functions $\lambda(v_t)$ computed according to (15), and depicts both the posterior means and 90% confidence intervals. The middle row of the figure gives the betas for a typical option in the data set, a slightly OTM 1-month

put option with a strike price of 0.95 times the current underlying price. The bottom row of Figure 5 displays the products of the first two rows, indicating the risk premia received for each factor by the same 1-month OTM put.

I find that all three factors command risk premia that appear to be larger for higher levels of the VIX, as are the risk premia assumed in affine stochastic volatility models such as those in Duffie, Pan, and Singleton (2000). For the first two factors, the prices of risk are reliably nonzero, with narrow confidence intervals that do not include zero. For the 1-month OTM put I consider, the typical risk premia for both the market and volatility factors are each sufficient to induce price declines of 1%–2% per day. In contrast, 90% confidence intervals for the price of “jump” factor risk often include zero and appear to contribute comparatively little to risk premia on the same put.

These results contrast markedly with Broadie, Chernov, and Johannes (2005), who find no premium for diffusive volatility risk and a significant price of jump risk. While there are many implementation details that differ between our papers, perhaps the most notable is that Broadie et al. (2005) assume a particular stochastic volatility/jump model specification. In contrast, this paper takes a much less parametric approach. One possible explanation for the fundamental differences between our conclusions is that the affine dynamics assumed in that paper are somehow inaccurate.

IV. Hedging and Portfolio Allocation

Given the large sample of option returns and the very high explanatory power of the factor model, specifications with as many parameters as those considered here may not be overly complex. Nevertheless, the flexibility of the models in this paper suggests that some type of out-of-sample exercise would be useful for assessing the usefulness of the proposed approach. I perform such an exercise by using the models both for hedging and asset allocation applications.

For the purpose of performing out-of-sample analysis, I reestimate all 10 specifications using only the first half of the sample (up to June 1993). I also reestimate, over the same period, the models summarized in Table II that are used to compute expectations of $\psi(f_t)$. This leaves July 1993 to September 2000 as an out-of-sample test period.

In both hedging and portfolio allocation exercises, some estimates of the conditional covariance matrices of option returns are required. Let β_{t-1} denote the $N_t \times K$ matrix of betas observed for the N_t options whose time- t returns are observed. Let σ_{t-1}^2 be the corresponding $N_t \times 1$ vector of idiosyncratic variances. The covariance matrix of time- t option returns is therefore

$$\beta_{t-1} \text{Cov}_{t-1}(F_t, F_t') \beta_{t-1}' + \text{diag}(\sigma_{t-1}^2),$$

where

$$\text{Cov}(F_t, F_t') = \text{Cov}(\psi(f_t), \psi(f_t)') + \Omega.$$

I compute these expressions using 10,000 Monte Carlo simulations and parameters estimated using data only through June 1993.

A. Hedging

I first consider a hedging exercise that is designed to mimic the problem faced by an option writer who has a book of option positions in a variety of strikes and maturities and who would like to hedge these options by trading a small number of more liquid puts. For each day I choose three put options of strikes and maturities that tend to see high volume. The first is the put nearest to a 1-month maturity that is also close to at-the-money. The second put has around a 3-month maturity and is as close as possible to being 10% OTM. The third is also close to at-the-money, but with a maturity as near as possible to 6 months.

The book that is being hedged is an equally weighted combination of long positions in all remaining puts and short positions in all available calls. This combination of long and short positions is designed to eliminate the natural delta-hedging that would arise in a portfolio that was long or short in both calls and puts. Since the long-short option portfolio will, to some degree, naturally hedge out volatility risk, I also consider a scenario in which the option book contains long puts and no calls.

For the estimated models, I choose hedges to minimize portfolio variance. For comparison, I also consider strategies based on the Black-Scholes “delta,” “gamma,” and “vega,” which are calculated using the options’ own implied volatilities. When hedging only one or two of the Greeks, I continue to use three hedge assets so that diversification of the idiosyncratic element of option returns can still be achieved. When hedging out delta only, all three puts in the hedge portfolio are held in equal weights. When hedging out delta in addition to either gamma or vega, the second and third puts are assumed to be held in equal quantities.

The top panel of Table VII displays the standard deviation and the mean absolute deviation of the put-only portfolio returns. I calculate “in-sample” results over the 1,275 days between January 1986 and June 1993 for which there are three puts available to form the hedge portfolio and at least one other put option to form the book. “Out-of-sample” results were computed from 1,229 days from July 1993 to September 2000. The bottom panel of Table VII examines the performance in hedging both long puts and short calls. Here, in-sample and out-of-sample results are based on 1,555 and 1,390 days of hedged returns, respectively, where I again require that three puts be available to form the hedge and at least one other put or call option to form the book.

As before, having two or three factors seems to represent a good balance between oversimplification and overfitting. For hedging the long put portfolios, the best performances both in- and out-of-sample come from the two nonlinear three-factor models and the one nonlinear two-factor model. For the long put and short call portfolio, which should naturally have lower volatility risk, the linear three-factor model does best, perhaps suggesting that market risk is more easily captured by a linear model. Both six-factor models display substantial evidence of an out-of-sample breakdown, consistent with earlier results suggesting that those specifications are overly complicated. One-factor models also display some failings, particularly for hedging long puts only.

Table VII
Hedging In and Out of Sample

In the top panel, an equal-weighted portfolio of puts is hedged using three other puts that are not among those in the original portfolio. The bottom panel examines portfolios that are long puts and short calls. Strategies that use the models proposed in this paper try to minimize the variance of the hedged portfolio, while Black–Scholes strategies construct hedged portfolios by imposing zero deltas and other Black–Scholes “Greeks.” All parameters are estimated using a daily subsample sample of the primary data set that starts in January 1986 and ends in June 1993. “In sample” results pertain to hedging performance over that period, while “out-of-sample” results are computed using daily data from July 1993 to September 2000.

		Standard Deviation		Mean Absolute Deviation	
		In Sample	Out of Sample	In Sample	Out of Sample
Panel A: Long Puts					
Unhedged returns		43.05	32.40	20.12	22.91
Black–Scholes strategies					
Delta		9.85	6.61	4.53	4.64
Delta and Gamma		12.43	9.49	6.46	6.53
Delta and Vega		15.01	9.11	5.92	6.12
Delta, Gamma, and Vega		21.50	29.65	10.84	15.61
Linear specifications					
R_{SP}	1	10.87	10.66	5.06	6.81
R_{SP} and ΔVIX	2	11.24	8.49	4.72	5.46
R_{SP} , ΔVIX , and ΔY_{3M}	3	10.86	7.94	4.78	5.31
Nonlinear specifications					
R_{SP}	1	10.32	8.48	4.75	5.69
R_{SP}	3	9.58	7.35	4.46	4.99
R_{SP} and ΔVIX	2	10.07	7.49	4.49	5.01
R_{SP} and ΔVIX	6	11.55	17.38	5.30	8.55
R_{SP} , ΔVIX , and ΔY_{3M}	1	10.01	9.17	4.67	5.87
R_{SP} , ΔVIX , and ΔY_{3M}	3	10.03	7.32	4.51	5.05
R_{SP} , ΔVIX , and ΔY_{3M}	6	11.47	14.95	5.21	7.71
Panel B: Long Puts and Short Calls					
Unhedged returns		28.63	32.04	18.81	23.00
Black–Scholes strategies					
Delta		22.83	16.36	9.01	10.48
Delta and Gamma		25.10	18.72	12.18	12.90
Delta and Vega		23.53	17.27	10.74	12.01
Delta, Gamma, and Vega		41.65	48.89	24.37	32.10
Linear specifications					
R_{SP}	1	12.14	14.26	7.51	8.69
R_{SP} and ΔVIX	2	11.57	13.20	7.03	8.27
R_{SP} , ΔVIX , and ΔY_{3M}	3	11.15	12.99	6.90	8.08
Nonlinear specifications					
R_{SP}	1	12.09	14.11	7.50	8.63
R_{SP}	3	12.48	14.56	7.55	9.27
R_{SP} and ΔVIX	2	15.63	13.61	7.56	8.77
R_{SP} and ΔVIX	6	14.05	16.62	7.88	10.12
R_{SP} , ΔVIX , and ΔY_{3M}	1	12.06	14.25	7.46	8.75
R_{SP} , ΔVIX , and ΔY_{3M}	3	15.05	14.51	7.63	9.23
R_{SP} , ΔVIX , and ΔY_{3M}	6	13.97	14.60	7.64	9.35

Given these results, it is somewhat surprising that the best Black–Scholes delta-hedging scheme, which is essentially a one-factor hedge, is the best-performing strategy of all for hedging puts. For hedging both puts and calls, however, all the Black–Scholes strategies fare relatively poorly.

Together, these findings indicate that even half the sample used to generate the main results of this paper is sufficient to produce fitted models that are useful in real applications. Even given the complexity of the models I estimate, this might not be surprising given the size of the data set. Even more importantly, the returns we are looking at are very tightly linked to one another by arbitrage restrictions. Thus, the option returns data are naturally more informative than, say, returns on different stocks, and it is likely that only because of that informativeness are such flexible models of much use.

B. Mean–Variance Portfolios

Given previous pricing results, it is worthwhile to explore how an investor might optimally form a portfolio of different put and call options to exploit their apparent mispricing. This section considers mean–variance optimization that might be performed by an investor attempting to capture these returns by trading solely in put and call options of different strikes and maturities. While it is natural to consider strategies that involve trading in the underlying as well, I do not analyze them here.

On each day in the sample, I form Sharpe ratio-maximizing portfolios of options. I construct conditional covariance matrices as described above, while I compute conditional means using current conditioning variables combined with posterior means of the parameters in the $\mu(\cdot)$ functions. I consider both unrestricted portfolios and portfolios that are constrained to be somewhat “crash resistant.” This constraint requires that the net position in put options be long. More precisely, if n_i is the number of contracts bought on put i , then I require that $\sum n_i \geq 0$. This implies that the total value of all puts will be bounded below at a level determined by the lowest strike price.⁷

Table VIII reports in-sample (Panel A) and out-of-sample (Panel B) results. For comparison with the optimal portfolios, equal-weighted portfolios of puts and calls are also considered, as is the S&P 500 Index. All standard errors are adjusted for heteroskedasticity.

In sample, average excess returns are positive under every model-based strategy. Portfolio weights are normalized so that the portfolio’s target standard deviation is 5% per day, and the realized volatilities are usually close to that value. In-sample average returns on unconstrained portfolios are between 0.5% and 1.09% per day, resulting in daily Sharpe ratios between 0.093 and 0.155. The portfolios are therefore vastly superior to either the S&P 500 Index, which has a Sharpe ratio of 0.016, or equally weighted portfolios of long calls or short

⁷ In the event of a market crash in which the underlying price falls below all strike prices, all options approach their intrinsic values. Ignoring discounting, the approximate value of a portfolio of puts, $\sum n_i \max\{K_i - P^M, 0\}$ is bounded below by $\sum n_i (K_i - \min K)$.

Table VIII
Option Portfolios

The table reports summary statistics on the results of in- and out-of-sample portfolio allocation exercises. For each day, a Sharpe ratio-maximizing portfolio is constructed given estimated parameter values and current information variables. This portfolio is held in a quantity that is consistent with a target return standard deviation of 5%. Crash-resistant portfolios require a net long position in puts. For each model, the mean, standard deviation, Sharpe ratio, skewness, and kurtosis are computed from the excess returns that result from the portfolio based on that model. The final column measures the modified alpha of Leland (1999). All *t*-statistics are computed using heteroskedasticity-adjusted standard errors. All models are estimated using data from January 1986 to June 1993. In Panel A, portfolios are formed over the same period. In Panel B, portfolios are formed over the out-of-sample period from January 1986 to June 1993.

	Mean Return/ <i>t</i> -Statistics	Standard Deviation	Sharpe Ratio	Skewness	Kurtosis	Market Model		Leland's Modified α	
						Alpha/ <i>t</i> -Statistics	Beta/ <i>t</i> -Statistics		
Panel A: In-Sample Performance									
Unconstrained Portfolios									
Linear specifications	1	1.06/6.47	7.03	0.1504	0.41	16.38	1.05/6.40	0.68/2.13	1.05
	2	0.76/5.84	5.61	0.1357	-3.72	43.34	0.74/5.69	1.26/2.70	0.74
	3	0.74/6.15	5.17	0.1429	-3.14	32.24	0.71/6.12	1.56/4.26	0.71
	1	1.09/6.68	7.02	0.1554	0.44	15.13	1.07/6.65	1.27/3.88	1.07
	3	0.68/6.32	4.66	0.1468	-0.99	8.28	0.66/6.46	1.50/7.05	0.66
	2	0.67/4.81	5.97	0.1118	-5.72	98.85	0.62/4.97	2.65/4.27	0.62
Nonlinear specifications	6	0.55/4.99	4.72	0.1160	-2.25	20.58	0.51/5.26	2.07/9.58	0.51
	1	1.06/6.24	7.33	0.1451	-1.02	29.02	1.03/6.22	1.99/3.22	1.03
	3	0.59/4.55	5.57	0.1057	-6.56	123.42	0.54/4.70	2.58/4.10	0.54
	6	0.50/4.00	5.42	0.0930	-6.95	128.90	0.46/4.18	2.74/4.61	0.45
	Crash-Resistant Portfolios								
	1	0.99/6.19	6.84	0.1443	0.67	16.47	0.98/6.12	0.75/2.38	0.98
Nonlinear specifications	2	0.65/5.31	5.21	0.1239	-2.43	22.19	0.63/5.28	1.25/5.17	0.63
	3	0.65/5.48	5.07	0.1278	-2.30	22.81	0.62/5.51	1.52/5.57	0.62
	1	1.03/6.45	6.83	0.1504	0.69	15.59	1.01/6.42	1.28/3.86	1.01
	3	0.59/5.46	4.61	0.1273	-0.91	7.97	0.57/5.48	1.27/5.12	0.57
	2	0.54/4.03	5.72	0.0939	-5.50	94.40	0.50/4.06	2.38/3.88	0.49
	6	0.48/4.23	4.88	0.0987	-3.05	38.17	0.45/4.46	2.17/6.59	0.45
Nonlinear specifications	1	0.99/5.93	7.12	0.1385	-0.69	27.49	0.95/5.90	1.86/3.13	0.95
	3	0.49/3.84	5.53	0.0895	-7.62	162.35	0.45/3.86	2.46/3.39	0.45
	6	0.42/2.98	6.02	0.0695	-11.25	280.91	0.37/2.94	3.01/3.22	0.36

Panel B: Out-of-Sample Performance									
Unconstrained Portfolios									
Linear specifications									
R_{SP}	1	0.71/3.22	9.24	0.0771	-1.61	25.36	0.59/2.72	2.59/5.58	0.58
R_{SP} and ΔVIX	2	0.32/2.25	5.92	0.0539	-2.31	17.29	0.23/1.66	1.90/5.85	0.22
R_{SP} , ΔVIX , and ΔY_{3M}	3	0.49/3.46	5.90	0.0830	-3.44	31.02	0.36/2.73	2.79/9.54	0.35
Nonlinear specifications									
R_{SP}	1	0.88/4.17	8.79	0.0999	-0.63	23.77	0.74/3.69	2.87/7.17	0.74
R_{SP}	3	0.36/2.73	5.53	0.0655	-1.93	15.08	0.26/2.10	2.22/7.76	0.25
R_{SP} and ΔVIX	2	0.39/2.56	6.35	0.0614	-4.62	61.94	0.26/1.82	2.73/8.44	0.25
R_{SP} and ΔVIX	6	0.40/3.54	4.72	0.0847	-2.19	24.50	0.29/2.85	2.37/11.92	0.28
R_{SP} , ΔVIX , and ΔY_{3M}	1	0.85/4.19	8.51	0.1003	-0.91	15.91	0.72/3.71	2.86/7.36	0.71
R_{SP} , ΔVIX , and ΔY_{3M}	3	0.49/3.82	5.35	0.0914	-2.01	17.94	0.37/3.19	2.50/10.18	0.37
R_{SP} , ΔVIX , and ΔY_{3M}	6	0.37/3.31	4.62	0.0794	-2.69	31.30	0.26/2.61	2.23/10.93	0.26
Crash-Resistant Portfolios									
Linear specifications									
R_{SP}	1	0.58/2.84	8.49	0.0685	-1.84	20.19	0.48/2.39	2.21/5.12	0.48
R_{SP} and ΔVIX	2	0.33/2.52	5.49	0.0606	-2.00	15.77	0.24/1.90	1.86/6.50	0.24
R_{SP} , ΔVIX , and ΔY_{3M}	3	0.38/2.58	6.04	0.0623	-4.32	48.11	0.26/1.86	2.72/8.79	0.25
Nonlinear specifications									
R_{SP}	1	0.78/3.91	8.24	0.0944	-0.91	16.37	0.65/3.44	2.74/7.43	0.65
R_{SP}	3	0.32/2.41	5.50	0.0582	-1.96	16.05	0.22/1.76	2.10/7.67	0.21
R_{SP} and ΔVIX	2	0.33/2.31	5.94	0.0557	-3.43	36.12	0.22/1.61	2.48/8.12	0.21
R_{SP} and ΔVIX	6	0.38/3.27	4.78	0.0788	-2.36	28.67	0.28/2.65	2.16/10.21	0.28
R_{SP} , ΔVIX , and ΔY_{3M}	1	0.74/3.82	8.06	0.0921	-0.84	13.29	0.63/3.39	2.54/6.81	0.62
R_{SP} , ΔVIX , and ΔY_{3M}	3	0.44/3.43	5.28	0.0828	-1.68	18.09	0.34/2.88	2.37/9.80	0.33
R_{SP} , ΔVIX , and ΔY_{3M}	6	0.32/2.87	4.65	0.0693	-2.65	31.12	0.23/2.22	2.00/9.10	0.23
Benchmark Strategies									
EW short call portfolio		-1.76/2.18	33.78	-0.0521	-2.10	13.73	-0.42/0.91	-28.58/24.04	-0.44
EW short put portfolio		2.56/3.74	28.54	0.0897	-1.92	13.04	1.35/3.82	25.76/32.42	1.33
S&P 500 index return		0.05/2.05	0.96	0.0491	-0.22	7.16			
Benchmark Strategies									
EW short call portfolio		-0.55/0.88	26.98	-0.0205	-1.80	11.50	-0.25/0.55	-16.91/5.96	-0.27
EW short put portfolio		1.08/1.28	36.45	0.0297	-14.35	376.43	0.60/1.14	27.08/5.82	0.58
S&P 500 index return		0.02/0.69	1.11	0.0160	-3.61	70.93			

puts, with Sharpe ratios of 0.02 and 0.03, respectively. Although the crash-resistant portfolios perform somewhat worse than the unconstrained versions, their performance is still strikingly good.

Somewhat surprisingly, strategies based on the three different one-factor models post the highest average returns. In part, this appears to be due to overly risky strategies; while all portfolios are normalized to have an anticipated daily volatility of 5% per day, the one-factor portfolios often have volatilities significantly above that target. Nevertheless, these models manage to produce the highest Sharpe ratios as well.

Market model regressions reveal significant alphas and negative betas, consistent with a put-writing strategy. Because alpha, as well as the Sharpe ratio, may be inappropriate for evaluating strategies that result in fat-tailed return distributions, I also report Leland's (1999) modified alpha, which takes higher moments into account. In many cases, skewness and kurtosis coefficients are very large. Not surprisingly, the unhedged equally weighted short put portfolio has the fattest tails, but even in this case Leland's adjustment does not amount to much.

Out-of-sample results result in roughly the same ordering of the models. With the exception of the one linear specification, one-factor models continue to generate high Sharpe ratios despite continuing to exceed the 5% volatility target. However, the magnitude of the abnormal performance of all models is lessened by roughly one third to one half, leading to Sharpe ratios between 0.056 and 0.1. Roughly half of the strategies produce statistically significant market model alphas. In addition, while all of the Sharpe ratios are superior to that of the S&P 500 Index, the simple strategy of writing naked puts now generates a very competitive Sharpe ratio of about 0.09. Given the virtually unbounded risks this strategy entails, it is not obvious that it would be preferred over one of the crash-resistant portfolios, even given their somewhat worse ex post performance.

V. Conclusion

The factor models I consider in this paper appear to offer no rational explanation for the large excess returns that appear possible in the market for S&P 500 futures options. The paper therefore complements the recent findings of Bondarenko (2003), who rejects a large class of models in which the pricing kernel depends only on the market return. Allowing for additional sources of priced risk reduces the degree of mispricing for most options, but additional factors are still insufficient to explain the anomalous average returns on a wide range of put and call options, in particular, short-term deep-OTM puts.

The best-performing models tend to be those with two to three factors, as one-factor models are obviously oversimplistic and more complex models are unstable out of sample. While allowing factors to be nonlinear occasionally leads to improvements in hedging performance, linear multifactor models are also useful in both hedging and asset allocation and provide competitive or even superior descriptions of expected and realized returns.

Overall, the failures of all the specifications I consider suggest several possible avenues for future work. As in other related work, here I assume that the level of volatility (proxied by the VIX index) is sufficient for describing time variation in covariances and risk premia. It is possible that the addition of some unknown state variable may resolve these puzzles, though it is difficult to speculate on what that state variable might be. The second possibility is that, more generally, the standard representative agent/perfect markets paradigm cannot provide an adequate explanation of expected option returns. Fortunately, many avenues exist for broadening this framework, such as heterogeneous preferences (Bates (2001)) or beliefs (Buraschi and Jiltsov (2005)), uncertainty about fundamentals (David and Veronsi (2002), Guidolin and Timmermann (2003)), behavioral biases (Potoshman (2001)), or proximate causes such as buying pressure (Bollen and Whaley (2004)). Discriminating among these effects and quantifying their ability to resolve the puzzles identified in this paper and elsewhere represents a promising direction for additional research.

Appendix: The Posterior Sampling Procedure

I compute the posterior using the Gibbs sampler with the following four blocks: (i) $\mu(\cdot)$ and $\beta(\cdot)$, (ii) $\sigma(\cdot)$, (iii) $\mathbf{F}$, and (iv) Ω and $\psi(\cdot)$. For simplicity, the derivations below are presented as though the sample consists of only calls or only puts. In practice, there are separate $\mu(\cdot)$, $\beta(\cdot)$, and $\sigma(\cdot)$ functions for both puts and calls.

Throughout this appendix, it is more convenient to replace the double subscript “ i t ” (for security i at time t) with a single subscript $i = 1, \dots, N$, where $N = 33,928$ is the size of the entire sample. Therefore τ_i, κ_i , and all other i -subscripted symbols are to be interpreted as variables that correspond to return i . In addition, F_i refers to the column of $\mathbf{F}$ (a $K \times T$ matrix) that is contemporaneous with the observation of R_i , so F_i is *not* the i th column of $\mathbf{F}$. Similarly, let v_i denote the value of the conditioning variable v that corresponds to observation i . Finally, in the derivations below, let $\mathcal{L}_i^\mu$ refer to the row vector of Legendre polynomials used in the expansion of the function $\mu(\cdot)$ and evaluated at the point $(\tau_i, \kappa_i, v_i, T_i)$. Other polynomial expansions are defined similarly.

A. Identification

As in Geweke and Zhou (1996) and other papers, the latent factor nature of the model implies that identification is not possible without placing restrictions on some model parameters. For example, multiplying the factor matrix $\mathbf{F}$ by any invertible $K \times K$ matrix, and the return betas by the inverse of that matrix, delivers the same likelihood. To prevent this situation, each factor is identified by prespecifying its loading on a particular observable variable. Since the Legendre expansion of the $K \times 1$ vector $\psi(f_i)$ is $G\mathcal{L}_i^\psi$, we have

$$F_t = G\mathcal{L}_t^\psi - \mathbf{E}_{t-1}[G\mathcal{L}_t^\psi] + \eta_t = G(\mathcal{L}_t^\psi - \mathbf{E}_{t-1}[\mathcal{L}_t^\psi]) + \eta_t, \quad (\text{A1})$$

where $\eta_t \sim \mathbf{N}(0, \Omega)$. For identification, I therefore set

$$G = \begin{bmatrix} G_1 \\ G_2 \end{bmatrix},$$

where $G_1 = I_K$ is the $K \times K$ identity matrix. Thus, each factor is linked one-for-one with one of the elements in $\mathcal{L}_t^\psi$, and that element does not affect any of the other factors. The result is that “rotation” of the factor matrix is impossible, since doing so would change the form of the G matrix. If the element of $\mathcal{L}_t^\psi$ that is “assigned” to a particular factor turns out to be irrelevant, its impact can be eliminated if the factor takes on large amounts of residual variance or dependence on the other elements of $\mathcal{L}_t^\psi$. Thus, it is important that G_2 and the diagonal elements of Ω be unrestricted.

While it is not theoretically necessary that the off-diagonal elements of Ω be fixed, I find virtually no ability to estimate these off-diagonal elements. Posteriors are extremely imprecise, and the Gibbs sampler that allows these elements to be unrestricted is very slow to converge. I therefore restrict Ω to be a diagonal matrix.

B. Drawing the Parameters of $\mu(\cdot)$ and $\beta(\cdot)$

Conditional on the factor matrix $\mathbf{F}$ and the parameters of the $\sigma(\cdot)$ function, the factor model may be written as

$$R_i = \mathcal{L}_i^\mu a + \mathcal{L}_i^\beta B F_i + \epsilon_i, \quad (\text{A2})$$

where a is a column vector and B is a matrix with K columns. Thus, $\mu(\tau_i, \kappa_i, v_i, T_i) = \mathcal{L}_i^\mu a$ is the scalar regression intercept and $\beta(\tau_i, \kappa_i, v_i, T_i) = \mathcal{L}_i^\beta B$ is a $1 \times K$ vector of factor loadings. Using Proposition 10.4 of Hamilton (1994), we rewrite the return equation as

$$R_i = \mathcal{L}_i^\mu a + (F_i' \otimes \mathcal{L}_i^\beta) \text{vec}(B) + \epsilon_i, \quad (\text{A3})$$

where $\otimes$ denotes the Kronecker product and $\text{vec}(B)$ is the column vector formed by stacking the K columns of B .

Defining b by stacking a on top of $\text{vec}(B)$, we obtain the standard linear regression

$$R_i = [\mathcal{L}_i^\mu \ F_i' \otimes \mathcal{L}_i^\beta] b + \epsilon_i \equiv X_i b + \epsilon_i. \quad (\text{A4})$$

Stacking all N observations and writing in vector form, we have $R = Xb + \epsilon$, where $\epsilon \sim N(0, \Sigma)$ and Σ is the diagonal $N \times N$ residual covariance matrix whose (i, i) element is equal to $\sigma^2(\tau_i, \kappa_i, v_i, T_i)$.

This setup implies the standard Bayesian generalized least squares (GLS) result that under flat priors,

$$b \sim N(\hat{b}, V_b), \quad (\text{A5})$$

where $\hat{b} = (X' \Sigma^{-1} X)^{-1} X' \Sigma^{-1} R$ and $V_b = (X' \Sigma^{-1} X)^{-1}$. The parameters of interest, a and B , are given by the appropriate partitions of b .

C. Drawing the Parameters of $\sigma(\cdot)$

We may rewrite the return equation as

$$R_i = \mathcal{L}_i^\mu a + X_i b + \exp(\mathcal{L}_i^\sigma s) e_i, \quad (\text{A6})$$

where e_i is an i.i.d. standard normal random variable. We seek to draw s , the coefficients of the Legendre expansion of $\sigma(\cdot)$, from its distribution given all other unknowns.

By squaring and taking logs, the equation approximately implies that

$$\log([R_i - X_i b]^2 + 0.001) = 2\mathcal{L}_i^\sigma s + e_i^*, \quad (\text{A7})$$

where $e_i^* = \log(e_i^2)$ and where b and X_i are defined as in (A4). Clearly, the approximation arises from the extra 0.001 term appearing in the left-hand side.⁸ Given b and $\mathbf{F}$, this is a standard linear regression, with the complication that e_i^* is not Gaussian or even mean zero, although it is i.i.d. More importantly, the density of e_i^* does not depend on any unknown parameters.

Kim, Shepard, and Chib (1998, hereafter KSC) propose an efficient and very accurate approximate solution for the posterior of s . Their procedure consists of approximating the distribution of $e_i^* = \log(e_i^2)$ by a mixture of seven different normal distributions. The normal means, variances, and mixture probabilities that KSC calculated can be found in table IV of their paper.

The procedure involves the introduction of a latent state variable π_i that takes on an integer value between one and seven, indexing the normal distribution from which e_i^* is drawn. The KSC procedure therefore requires the addition of an additional “block” in the Gibbs sampler, the draw of the π_i , although this block is by itself of little interest.

Following KSC, I draw π_i for $i = 1, \dots, N$ from the probability mass function

$$P(\pi_i = j | a, B, \mathbf{F}, s) \propto q(j) \phi(\log([R_i - X_i b]^2 + 0.001); 2\mathcal{L}_i^\sigma s + c(j), d(j)), \quad (\text{A8})$$

where $\phi(x; \mu, v)$ is the density function of a $N(\mu, v)$ random variable evaluated at x .

Conditional on π_i , we may view e_i^* as a normal random variable with mean $c(\pi_i)$ and variance $d(\pi_i)$. After subtracting this mean from each side of (A7), we may proceed to draw s using Bayesian GLS. Let $\mathcal{L}^\sigma$ denote the matrix formed by stacking the $\mathcal{L}_i^\sigma$ for all i . Let C denote the column matrix formed from the N values of $c(\pi_i)$ and D denote the diagonal matrix formed from the corresponding values of $d(\pi_i)$. Then

$$s \sim N(\hat{s}, V_s), \quad (\text{A9})$$

⁸ As Kim et al. (1998) discuss, adding the offset 0.001 before taking logs makes the transformation more robust since it reduces the “inlier” problem associated with taking logarithms of very small numbers.

where $V_s = ((2\mathcal{L}^\sigma)'D^{-1}(2\mathcal{L}^\sigma))^{-1}$ and $\hat{s} = V_s(2\mathcal{L}^\sigma)'D^{-1}(\log[(R - Xb)^2 + 0.001] - C)$, and the square denotes an element-by-element operation.

D. Drawing F_t

For this draw I follow the derivations of Geweke and Zhou (1996), which necessitates the introduction of some new notation. Specifically, let R_t denote the vector of all returns observed at time t . Let ϵ_t denote the idiosyncratic component of those returns, whose diagonal covariance matrix is $\Sigma_t = \text{diag}(\exp(\mathcal{L}_t^\sigma s))$, where $\mathcal{L}_t^\sigma$ denotes the rows of $\mathcal{L}^\sigma$ that correspond to all time- t observations.

Recall that

$$F_t = \psi(f_t) - \mathbf{E}_{t-1}[\psi(f_t)] + \eta_t,$$

where $\eta_t \sim N(0, \Omega)$. Conditional on f_t , the parameters of $\psi(\cdot)$, the parameters that determine the expectation $\mathbf{E}_{t-1}[\psi(f_t)]$, but *not* the contemporaneous returns R_t , this implies that $F_t \sim N(\hat{F}_t, \Omega)$, where $\hat{F}_t = \psi(f_t) - \mathbf{E}_{t-1}[\psi(f_t)]$.

Since $R_t = \mu_t + \beta_t' F_t + \epsilon_t$, we have

$$\begin{bmatrix} F_t \\ R_t \end{bmatrix} \sim N \left(\begin{bmatrix} \hat{F}_t \\ \beta_t' \hat{F}_t \end{bmatrix}, \begin{bmatrix} \Omega & \Omega \beta_t \\ \beta_t' \Omega & \beta_t' \Omega \beta_t + \Sigma_t \end{bmatrix} \right).$$

This implies the following distribution for F_t conditional on all other variables, including R_t :

$$F_t | R_t \sim N(\hat{F}_t + (\beta_t' \Omega \beta_t + \Sigma_t)^{-1} \beta_t' \Omega (R_t - \mu_t - \beta_t' \hat{F}_t), \Omega - \Omega \beta_t (\beta_t' \Omega \beta_t + \Sigma_t)^{-1} \beta_t' \Omega). \quad (\text{A10})$$

E. Drawing Ω and the Parameters of $\psi(\cdot)$

Conditional on $\mathbf{F}$ and all parameters except Ω , we have

$$F_t - \hat{F}_t = \eta_t \sim N(0, \Omega).$$

Since Ω is diagonal, each element may be drawn independently, and under flat priors we obtain the standard result that

$$\sqrt{\Omega_{k,k}} \sim \text{IG}(\sqrt{(1/T)\Sigma_{t=1}^T \eta_{kt}^2}, T). \quad (\text{A11})$$

From (A1), we see that the equation for the k th factor may be written as

$$F_{kt} = G_k(\mathcal{L}_t^\psi - \mathbf{E}_{t-1}[\mathcal{L}_t^\psi]) + \eta_{kt},$$

where G_k is the k th row of G . Since η_t has a diagonal covariance matrix, each of these K equations may be treated independently. Since the first K elements of G_k are given by identification restrictions, it is convenient to define Z_1 as the $T \times K$ matrix formed by stacking the first K elements of $\mathcal{L}_t^\psi - \mathbf{E}_{t-1}[\mathcal{L}_t^\psi]$

for all t . Let Z_2 denote the corresponding matrix of the remaining elements of $\mathcal{L}_t^\psi - E_{t-1}[\mathcal{L}_t^\psi]$. We may then rewrite the equation in matrix notation as

$$\mathbf{F}_k - Z_1 G'_{1k} = Z_2 G'_{2k} + \eta_k,$$

where G_{1k} is the k th row of G_1 , the identity matrix, and G_{2k} is the k th row of G_2 , which is unknown. Since η_{kt} is Gaussian and homoskedastic, the draw of G_2 under flat priors is the standard

$$G'_{2k} \sim N(\hat{G}_{2k}, V_{G_{2k}}), \quad (\text{A12})$$

where $\hat{G}_{2k} = (Z_2' Z_2)^{-1} Z_2' (\mathbf{F}_k - Z_1 G'_{1k})$ and $V_{G_{2k}} = \Omega_{kk} (Z_2' Z_2)^{-1}$.

REFERENCES

- Ait-Sahalia, Yacine, and Andrew Lo, 1998, Nonparametric estimation of state-price densities implicit in financial asset prices, *Journal of Finance* 53, 499–547.
- Andersen, Torben G., Luca Benzoni, and Jesper Lund, 2002, An empirical investigation of continuous-time equity return models, *Journal of Finance* 57, 1239–1284.
- Bakshi, Gurdip, Charles Cao, and Zhiwu Chen, 1997, Empirical performance of alternative option pricing models, *Journal of Finance* 52, 2003–2049.
- Bakshi, Gurdip, and Nikunj Kapadia, 2003, Delta-hedged gains and the negative market volatility risk premium, *Review of Financial Studies* 16, 527–566.
- Bansal, Ravi, David A. Hsieh, and S. Viswanathan, 1993, A new approach to international arbitrage pricing, *Journal of Finance* 48, 1719–1747.
- Bansal, Ravi, and S. Viswanathan, 1993, No arbitrage and arbitrage pricing: A new approach, *Journal of Finance* 48, 1231–1262.
- Bates, David S., 2000, Post-'87 crash fears in the S&P 500 futures option market, *Journal of Econometrics* 94, 181–238.
- Bates, David S., 2001, The market for crash risk, Working paper, University of Iowa.
- Bauer, Gregory, and Ane Tamayo, 2000, The cross-section of expected returns with imperfectly observed factors, Working paper, University of Rochester.
- Benzoni, Luca, 2002, Pricing options under stochastic volatility: An empirical investigation, Working paper, University of Minnesota.
- Black, Fischer, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–659.
- Blair, Bevan J., Ser-Huang Poon, and Stephen J. Taylor, 2001, Forecasting S&P 100 volatility: The incremental information content of implied volatilities and high frequency index returns, *Journal of Econometrics* 105, 5–26.
- Blume, Marshall E., and Robert F. Stambaugh, 1983, Biases in computed returns: An application to the size effect, *Journal of Financial Economics* 12, 387–404.
- Bollen, Nicolas P.B., and Robert E. Whaley, 2004, Does net buying pressure affect the shape of implied volatility functions? *Journal of Finance* 59, 711–753.
- Bondarenko, Oleg, 2003, Why are put options so expensive? Working paper, University of Illinois at Chicago.
- Bossaerts, Peter, and Pierre Hillion, 1997, Local parametric analysis of hedging in discrete time, *Journal of Econometrics* 81, 243–272.
- Bossaerts, Peter, and Pierre Hillion, 2003, Local parametric analysis of derivatives pricing and hedging, *Journal of Financial Markets* 6, 573–605.
- Branger, Nicole, and Christian Schlag, 2004, Can tests based on option hedging errors correctly identify volatility risk premia? Working paper, Goethe University.
- Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2005, Model specification and risk premiums: Evidence from futures options, *Journal of Finance*, forthcoming.

- Broadie, Mark, Jérôme Detemple, Eric Ghysels, and Olivier Torr s, 2000, American options with stochastic dividends and volatility: A nonparametric investigation, *Journal of Econometrics* 94, 53–92.
- Buraschi, Andrea, and Jens Jackwerth, 2001, The price of a smile: Hedging and spanning in option markets, *Review of Financial Studies* 14, 495–527.
- Buraschi, Andrea, and Alexei Jiltsov, 2006, Model uncertainty and option markets in heterogeneous economies, *Journal of Finance*, forthcoming.
- Chapman, David A., 1997, Approximating the asset pricing kernel, *Journal of Finance* 52, 1383–1410.
- Chen, Joseph, 2003, Intertemporal CAPM and the cross-section of stock returns, Working paper, University of Southern California.
- Chen, Nai-Fu, Richard Roll, and Stephen A. Ross, 1986, Economic forces and the stock market, *Journal of Business* 59, 383–403.
- Chernov, Mikhail, and Eric Ghysels, 2000, A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation, *Journal of Financial Economics* 56, 407–458.
- Christensen, B.J., and N.R. Prabhala, 1998, The relation between implied and realized volatility, *Journal of Financial Economics* 50, 125–150.
- Christoffersen, Peter, and Jinyong Hahn, 1999, Nonparametric testing of ARCH for option pricing, in Yaser S. Abu-Mostafa, Blake LeBaron, Andrew W. Lo, and Andreas S. Weigend, eds.: *Proceedings of the Sixth International Conference on Computational Finance* (MIT Press, Cambridge).
- Cochrane, John H., 2001, *Asset Pricing* (Princeton University Press, Princeton).
- Cont, Rama, and Jos  da Fonseca, 2002, Dynamics of implied volatility surfaces, *Quantitative Finance* 2, 45–60.
- Coval, Joshua D., and Tyler Shumway, 2001, Expected options returns, *Journal of Finance* 56, 983–1009.
- David, Alexander, and Pietro Veronesi, 2002, Option pricing with uncertain fundamentals: Theory and evidence on the dynamics of implied volatilities, Working paper, Washington University in St. Louis and University of Chicago.
- Dittmar, Robert F., 2002, Nonlinear pricing kernels, kurtosis preference, and evidence from the cross-section of equity returns, *Journal of Finance* 57, 369–403.
- Duffie, Darrell, Jun Pan, and Kenneth Singleton, 2000, Transform analysis and asset pricing for affine jump-diffusions, *Econometrica* 68, 1343–1376.
- Dybvig, Philip H., and Jonathan E. Ingersoll, Jr., 1982, Mean-variance theory in complete markets, *Journal of Business* 55, 233–251.
- Eraker, Bj rn, Michael Johannes, and Nicholas Polson, 2003, The impact of jumps in volatility and returns, *Journal of Finance* 58, 1269–1300.
- Fern ndez, Carmen, and Mark F.J. Steel, 1998, On Bayesian modeling of fat tails and skewness, *Journal of the American Statistical Association* 93, 359–371.
- Ferson, Wayne E., and Campbell R. Harvey, 1999, Conditioning variables and the cross section of stock returns, *Journal of Finance* 54, 1325–1360.
- Fleming, Jeff, Barbara Ostdiek, and Robert E. Whaley, 1995, Predicting stock market volatility: A new measure, *Journal of Futures Markets* 15, 265–302.
- Garcia, Ren , and Ramazan Gen ay, 2000, Pricing and hedging derivative securities with neural networks and a homogeneity hint, *Journal of Econometrics* 94, 93–115.
- Geweke, John, and Guofu Zhou, 1996, Measuring the pricing error of the arbitrage pricing theory, *Review of Financial Studies* 9, 557–587.
- Guidolin, Massimo, and Allan Timmermann, 2003, Option prices under Bayesian learning: Implied volatility dynamics and predictive densities, *Journal of Economic Dynamics and Control* 27, 717–769.
- Hamilton, James D., 1994, *Time Series Analysis* (Princeton University Press, Princeton).
- Han, Bing, 2004, Limits of arbitrage, sentiment, and the pricing kernel: Evidence from index options, Working paper, Ohio State University.
- Heston, Steven L., 1993, A closed-form solution for options with stochastic volatility with applications to bond and currency options, *Review of Financial Studies* 6, 327–343.

- Hutchinson, James M., Andrew W. Lo, and Tomaso Poggio, 1994, A nonparametric approach to pricing and hedging derivative securities via learning networks, *Journal of Finance* 49, 851–890.
- Jackwerth, Jens, 2000, Recovering risk aversion from option prices and realized returns, *Review of Financial Studies* 13, 433–451.
- Jones, Christopher S., 2003, The dynamics of stochastic volatility: Evidence from underlying and options markets, *Journal of Econometrics* 116, 181–224.
- Kim, Sangjoon, Neil Shepard, and Siddhartha Chib, 1998, Stochastic volatility: Likelihood inference and comparison with ARCH models, *Review of Economic Studies* 65, 361–393.
- Ledoit, Olivier, Pedro Santa-Clara, and Shu Yan, 2002, Relative pricing of options with stochastic volatility, Working paper, UCLA.
- Leland, Hayne E., 1999, Beyond mean-variance: Performance measurement in a nonsymmetrical world, *Financial Analysts Journal* 55, 27–36.
- Longstaff, Francis, 1995, Option pricing and the martingale restriction, *Review of Financial Studies* 8, 1091–1124.
- Merton, Robert C., 1973, An intertemporal capital asset pricing model, *Econometrica* 41, 867–887.
- Pan, Jun, 2002, The jump-risk premia implicit in options: Evidence from an integrated time-series study, *Journal of Financial Economics* 63, 3–50.
- Poteshman, Allen M., 2001, Underreaction, overreaction, and increasing misreaction to information in the options market, *Journal of Finance* 56, 851–876.
- Shanken, Jay, 1990, Intertemporal asset pricing: An empirical investigation, *Journal of Econometrics* 45, 99–120.
- Skiadopoulos, George, Stewart Hodges, and Les Clewlow, 1999, The dynamics of the S&P 500 implied volatility surface, *Review of Derivatives Research* 3, 263–282.
- Whaley, Robert E., 1986, Valuation of American futures options: Theory and empirical tests, *Journal of Finance* 41, 127–150.
- Whaley, Robert E., 1993, Derivatives on market volatility: Hedging tools long overdue, *Journal of Derivatives* 1, 71–84.

Copyright of Journal of Finance is the property of Blackwell Publishing Limited and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.