Chapter 15

Equity Option Pricing Models

According to the original Black-Scholes option pricing model (Black and Scholes [1973]), the implied volatilities on European options on common stocks should be the same for options of all maturities and all strike prices. This is an implication of their assumption that the logarithm of a stock price follows a log-normal diffusion with constant volatility. In fact, especially subsequent to the "crash" of October, 1987, implied volatilities have exhibited a pronounced *smile* or *smirk*. A typical pattern of implied volatilities is displayed in Figure 15.1 for November 2, 1993. We see that call options that are deep in the money (put options that are deep out of the money) have higher implied volatilities than those that are nearer the money. Moreover, for a given degree of "out-of-the-moneyness," options with longer maturities tend to have lower implied volatilities. Neither of these patterns is consistent with the Black-Scholes model.

In this chapter we explore some of the models that have been put forth to explain these departures from the assumptions of the Black-Scholes model. Many of these models stay within the arbitrage-free framework of Black and Scholes and relax their strong assumptions about the distributions of stock prices. Others examine these empirical anomalies in an equilibrium setting starting with specifications of agents' preferences. Furthermore, while much of the empirical work on equity option pricing has focused on $S\&P500$ index options, recently several researchers have examined the the pricing of options on individual common stocks. In both the study of options on individual stocks and stock indices, a key question that has been addressed is which risks are priced in the markets and what are the properties of the associated "market prices of risk."

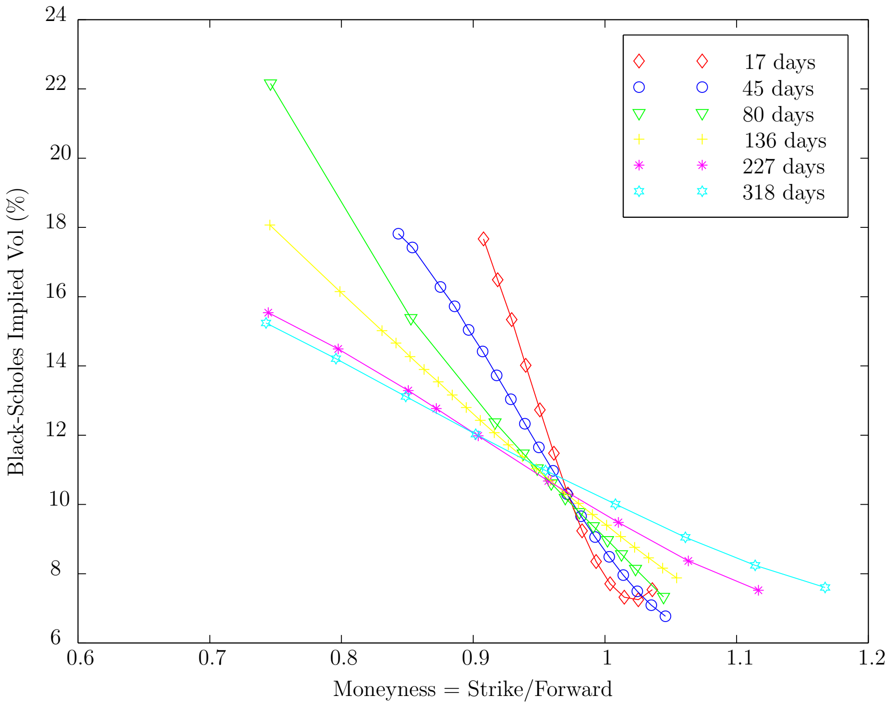

Figure 15.1: "Smile curves" implied by $S\&P$ 500 Index options of 6 different maturities. Option prices are obtained from market data of November 2, 1993.

We begin with an overview of no-arbitrage models that directly posit a pricing kernel with associated market prices of risk. This literature has focused on the relative contributions of priced volatility and jump risks in generating volatility smirks. Subsequently, we discuss the literature on preference-based models for option pricing.

15.1 No-Arbitrage Option Pricing Models

In an influential paper in the option-pricing literature, Heston [1993] showed that the risk-neutral exercise probabilities appearing in the call option-pricing formulas for bonds, currencies, and equities can be computed by Fourier inversion of a conditional characteristic function, which is known in closed form for his particular affine, stochastic volatility model. Building on this insight,[1] a variety of option-pricing models have been developed that can potentially explain these systematic deviations from the Black-Scholes model in the options data.

Initially, following the literature, we focus on the data generating process:

$$dS_t = \mu_S^{\mathbb{P}}(S_t, v_t) S_t \, dt + \sqrt{v_t} S_t \, dW_{St} + dZ_{St} \tag{15.1}$$

$$dv_t = \kappa_v^{\mathbb{P}}(\bar{v}^{\mathbb{P}} - v_t) \, dt + \sigma_v \sqrt{v_t} \left(\rho dW_{St} + \sqrt{1-\rho^2} dW_{vt} \right) + dZ_{vt}. \tag{15.2}$$

where $W = [W_S, W_v]'$ is a vector of independent Brownian motions in $\mathbb{R}^2$, $\rho \in (0,1)$ is a constant coefficient controlling correlation between the Brownian shocks to S and v, and $\mu_S^{\mathbb{P}}(S_t, v_t)$ is a stock price and volatility dependent component of the drift of dS_t/S_t. The processes (Z_S, Z_v) are jumps with intensities (mean arrival rates) ζ_{St} and ζ_{vt}, respectively. The amplitudes of the jumps, when they occur, may be random

Heston [1993] considered the special case of (15.1) - (15.2) in which $(\zeta_{St}, \zeta_{vt}) = 0$, so there were no jumps in either prices or volatility. Due to his focus on pricing (as contrasted with estimation using time series data), Heston assumed that the risk-neutral drift of S_t was $\mu_{St}^{\mathbb{Q}} = rS_t$ and $\mu_S^{\mathbb{P}}(S_t, v_t)$ was left unspecified. (The former is required by the assumption of no arbitrage opportunities– traded assets earn the riskfree rate when agents are risk-neutral.) He allowed for a volatility risk premium that was proportional to $\sqrt{v(t)}$, giving

$$dv_t = \left[\kappa_v^{\mathbb{P}}(\bar{v}^{\mathbb{P}} - v_t) + \eta_v v_t \right] \, dt + \sigma_v \sqrt{v_t} \, dW_{vt}^{\mathbb{Q}} \tag{15.3}$$

under $\mathbb{Q}$. Consequently, the $\mathbb{P}$ and $\mathbb{Q}$ drifts of v_t were both affine functions of v. We will refer to his model as model SV (with or without priced volatility risk).

Two recent studies that estimate models with stochastic volatility and a jump in stock returns (model SVJ) are Pan [2002] and Chernov and Ghysels [2000]. To interpret their assumptions about the market prices of price and volatility risks, we let $\mu_t^{\mathbb{P}} = (\mu_{St}^{\mathbb{P}}, \mu_{vt}^{\mathbb{P}})'$ denote the drift of $(d\,S_t/S_t, d\,v_t)$ under the physical measure, and $\mu_t^{\mathbb{Q}}$ its counterpart under the risk-neutral measure.

[1] Among the many recent papers examining option prices for the case of state variables following square-root diffusions are Bakshi, Cao, and Chen [1997], Bakshi and Madan [2000], Bates [1996], Bates [1997], Chen and Scott [1993], Chernov and Ghysels [2000], Pan [2002], Scott [1996, 1997], among others.

Recall that the relation between these two drifts is (within this affine setting) is

$$\mu_t^{\mathbb{Q}} = \mu_t^{\mathbb{P}} - \Sigma\sqrt{\Omega_t}\Lambda_t, \tag{15.4}$$

where Λ_t is the vector of market prices of risk associated with return and volatility risk,

$$\Sigma = \begin{pmatrix} 1 & 0 \\ \sigma_v\rho & \sigma_v\sqrt{1-\rho^2} \end{pmatrix}, \tag{15.5}$$

and Ω_t is the 2×2 diagonal matrix with v_t along the diagonal.

Chernov and Ghysels [2000]'s model abstracts from jumps (Z_t is omitted and $m_{SJ}^{\mathbb{Q}} = 0$) and replace $\mu_{St}^{\mathbb{P}}$ in (15.14) with $\bar{\mu}_S^{\mathbb{P}}$, a constant. They also treat the riskless interest rate as a constant r. Straight-forward calculation then shows that, for the drift of S_t to be r under the risk-neutral measure,

$$\Lambda_t^S = \frac{\bar{\mu}_s^{\mathbb{P}} - r}{\sqrt{v_t}}. \tag{15.6}$$

For the volatility process, Chernov and Ghysels assume that the drift under $\mathbb{Q}$ is $\kappa_v^{\mathbb{Q}}(\bar{v}^{\mathbb{Q}} - v_t)$, from which it follows that

$$\kappa_v^{\mathbb{P}}(\bar{v}^{\mathbb{P}} - v_t) - \sigma_v\rho\sqrt{v_t}\Lambda_t^S - \sigma_v\sqrt{1-\rho^2}\sqrt{v_t}\Lambda_t^v = \kappa_v^{\mathbb{Q}}(\bar{v}^{\mathbb{Q}} - v_t). \tag{15.7}$$

Substituting in (15.6) and solving for Λ_t^v gives

$$\Lambda_t^v = \frac{C_1}{\sqrt{v_t}} - C_2\sqrt{v_t}, \tag{15.8}$$

where

$$C_1 = \frac{\kappa_v^{\mathbb{P}}\bar{v}^{\mathbb{P}} - \kappa_v^{\mathbb{Q}}\bar{v}^{\mathbb{Q}} - (\bar{\mu}_S^{\mathbb{P}} - r)\sigma_v\rho}{\sigma_v\sqrt{1-\rho^2}}, \tag{15.9}$$

$$C_2 = \frac{\kappa_v^{\mathbb{P}} - \kappa_v^{\mathbb{Q}}}{\sigma_v\sqrt{1-\rho^2}}. \tag{15.10}$$

Notice that neither component of Λ_t in the Chernov-Ghysels model has a standard form for affine asset pricing models, which would have the risk premiums proportional to $\sqrt{v_t}$. Instead, they involves terms in the ratio of one over $\sqrt{v_t}$.

An implication of their formulation is that, as v_t approaches its lower bound of zero, Λ_t approaches infinity. As such, their model potentially admits arbitrage opportunities. This is precisely the same issue that arose in our discussions of the risk premiums in "essentially" affine *DTSM*s. There we constrained the parameters of the market prices of risk to rule out such arbitrage opportunities (see Chapter 12). Perhaps similar constraints on (C_1, C_2) could be derived to rule out arbitrage opportunities in this setting.

Pan [2002] uses a variant of the model in Bates [1997] in which there are no jumps in volatility ($\zeta_{vt} = 0$) and the intensity of the jump process Z_S is an affine function of volatility $\{\zeta_{vt} = \zeta_1 v_t \,:\, t \geq 0\}$, for a non-negative constant ζ_1. At the i-th jump time τ_i, the stock price was assumed to jump from $S(\tau_i-)$ to $S(\tau_i-)\exp(U_i^s)$, where U_i^s is normally distributed with mean $\mu_J^{\mathbb{P}}$ and variance δ_J^2, independent of W, of inter-jump times, and of U_j^s for $j \neq i$. The *mean relative jump size* is $m_{SJ}^{\mathbb{P}} = E(\exp(U^s) - 1) = \exp(\mu_J^{\mathbb{P}} + \delta_J^2/2) - 1$.

She assumes that the representation of (S, v) under $\mathbb{Q}$ is

$$dS_t = \left[r_t - q_t - \zeta_1 v_t m_{SJ}^{\mathbb{Q}}\right] S_t\, dt + \sqrt{v_t}\, S_t\, dW_{St}^{\mathbb{Q}} + dZ_{St}^{\mathbb{Q}}\,, \tag{15.11}$$

$$dv_t = \left[\kappa_v^{\mathbb{P}}\left(\bar{v}^{\mathbb{P}} - v_t\right) + \eta_v v_t\right] dt + \sigma_v \sqrt{v_t}\left(\rho\, dW_{St}^{\mathbb{Q}} + \sqrt{1-\rho^2}\, dW_{vt}^{\mathbb{Q}}\right), \tag{15.12}$$

where r is the riskless interest rate, q is the dividend payout rate, and $m_{SJ}^{\mathbb{Q}}$ is the risk-neutral mean of the relative jump size of the stock price. $W^{\mathbb{Q}} = \left(W_S^{\mathbb{Q}}, W_v^{\mathbb{Q}}\right)$ is a standard Brownian motion under $\mathbb{Q}$ defined by

$$W_t^{\mathbb{Q}} = W_t + \int_0^t \Lambda_u\, du\,, \quad 0 \leq t \leq T\,. \tag{15.13}$$

with $\Lambda_t = (\Lambda_t^S, \Lambda_t^v)'$ being the market prices of risk associated with the Brownian motions W_S and W_v. Additionally, she assumes that

$$\mu_{St}^{\mathbb{P}} = \begin{pmatrix} r_t - q_t + \eta_s v_t - \zeta_1 v_t m_{SJ}^{\mathbb{Q}} \\ \kappa_v^{\mathbb{P}}(\bar{v}^{\mathbb{P}} - v_t) \end{pmatrix}. \tag{15.14}$$

From these expressions we infer that Pan assumes the market price of risk Λ_t is given by

$$\Lambda_t^S = \eta_s \sqrt{v_t}, \tag{15.15}$$

$$\Lambda_t^v = -\frac{1}{\sqrt{1-\rho^2}}\left(\rho\eta_s + \frac{\eta_v}{\sigma_v}\right)\sqrt{v_t}\,, \tag{15.16}$$

with η_s and η_v being constant coefficients.

The risk-neutral mean relative jump size $m^{\mathbb{Q}}_{SJ}$ differs from its physical counterpart $m^{\mathbb{P}}_{SJ}$ in order to accommodate a premium for jump-size risk. Pan assumed that there is no risk premium associated with jump-timing risk (there is no distinction between ζ_{St} under $\mathbb{P}$ and $\mathbb{Q}$. Such a premium could be added by allowing the coefficient $\zeta^{\mathbb{Q}}_1$ for the risk-neutral jump-arrival intensity to be different from $\zeta^{\mathbb{P}}_1$. The motivation for assuming no premium on jump timing risk seemed to largely practical: it might be difficult to econometrically identify separate timing and amplitude premiums. Of course, by making this assumption, any risk premium is jump timing that is in fact present is being absorbed into one of the risk premiums that is allowed.

Two models with jumps in both returns and volatility where introduced by Duffie, Pan, and Singleton [2000]. In Chapter 7 we discussed the empirical properties of these models as descriptions of the historical behavior of returns. In this chapter we explore some of their implications for the pricing of options. Recall that model *SVIJ* assumes that the jump amplitude processes J_S and J_v are independent with respective amplitude distributions

$$J_{vt} \sim \exp(m_{Jv}) \text{ and } J_{St} \sim N(m_{JS}, \delta^2_{JS}). \tag{15.17}$$

Since J_v follows an exponential distribution, volatility can only jump up. Further, since J_v and J_S are independent, any "leverage" effects must be induced by the correlation among the diffusive shocks. A critical ingredient added by the jumps in volatility is the clustering of large return movements. Following an upward jump in v_t, the higher level of volatility persists due to its slow reversion to its mean.

Alternatively, model *SVCJ* has the jumps in returns and volatility driven by the same jump process (simultaneous jumps, $Z_v = Z_S$) and their amplitudes are correlated:

$$J_{vt} \sim \exp(m_{Jv}) \text{ and } J_{St}|J_{vt} \sim N(m_{JS} + \rho_J J_{vt}, \delta^2_{JS}). \tag{15.18}$$

This formulation introduces an additional leverage effect due to jumps when $\rho_J < 0$. A jump in volatility with a large amplitude J_{vt} (with $\rho_J < 0$), lowers the mean of the price jump amplitude thereby amplifying the leverage effect. At the same time, the positive jumps in volatility contribute to the right skewness of the distribution of volatility. Both features of this model tend to fatten the tails of the return distribution.

These models are easily extended to allow for state-dependent jump intensities as in Bates [1997] and Pan [2002]. Such a model *SVSCJ* with correlated jumps and stochastic arrival rate for jumps in stock prices,

$$\zeta_{St} = \zeta_0 + \zeta_1 v_t, \tag{15.19}$$

is explored in Eraker [2004]. With $\zeta_1 > 0$, jumps will tend to occur more frequently in high-volatility regimes. Furthermore, as with model *SVJ*, one can introduce risk premiums associated with both the jump amplitudes and the timing of jumps. As in Pan, Eraker assumed the former, but not the latter, risk was priced.

15.2 Option Pricing

Letting C_t denote the time-t price of a European-style call option on S, struck at K_t and expiring at T, and $X_t' = (r_t, q_t, v_t)$,

$$C_t = E^{\mathbb{Q}}\left[\exp\left(-\int_t^T r_u\, du\right)(S_T - K)^+ \,\middle|\, S_t, X_t\right]. \tag{15.20}$$

We can price this option using the time-t conditional transform of $\ln S_T$

$$\psi(u, X_t\,, T-t) = E^{\mathbb{Q}}\left[\exp\left(-\int_t^T r_s\, ds\right) e^{u\, \ln S_T} \,\middle|\, X_t, S_t\right], \tag{15.21}$$

for any $u \in \mathbb{C}$, introduced in Chapter 6. Letting $k_t = K_t/S_t$ be the time-t "strike-to-spot" ratio,

$$C_t = S_t\, O(X_t, T-t, k_t)\,, \tag{15.22}$$

where $O : \mathbb{R}_+^3 \times \mathbb{R}_+ \times \mathbb{R}_+ \to [0,1]$ is defined by

$$O(x, T-t, k) = \Pi_1 - k\,\Pi_2\,, \tag{15.23}$$

with

$$\begin{aligned} \Pi_1 &= \frac{\psi(1, x, T-t)}{2} - \frac{1}{\pi}\int_0^\infty \frac{\mathrm{Im}\left(\psi(1-iu, x, T-t)e^{i\,u(\ln k)}\right)}{u}\, du \\ \Pi_2 &= \frac{\psi(0, x, T-t)}{2} - \frac{1}{\pi}\int_0^\infty \frac{\mathrm{Im}\left(\psi(-iu, x, T-t)e^{i\,u(\ln k)}\right)}{u}\, du\,, \end{aligned} \tag{15.24}$$

and where $Im(\cdot)$ denotes the imaginary component of a complex number. The transform ψ is a known exponential-affine function of X_t.

15.3 Estimation of Option Pricing Models

Having described the models stock prices and reviewed the pricing of European options on stocks, we turn next to the estimation of the unknown parameters.

One, easy to implement, strategy is to minimize the squared deviations between the market and model-implied option prices. That is, letting

$$\epsilon_{it} \equiv \frac{C_{it}}{S_t} - O(X_t, T-t, k_{it}), \tag{15.25}$$

where i indexes options of possibly different strike prices and maturities, one minimizes the squared deviations ϵ_{it}^2 by choice of the parameter values of the model. Bakshi, Cao, and Chen [1997] estimate the parameters using a cross-section of strikes and maturities for a given day. (In this case, S_t cancels from the optimization problem.) Conceptually, this approach amounts to changing the model every period. As discussed in Chapter 12, such recalibration may well introduce dynamic arbitrage opportunities, viewed through the lens of the correct pricing model. Further, though it is common when using this approach to report the average values of the parameters across the days in a sample, such averages not are often not clearly linked to a specification of a "true" pricing model.

In contrast, Bates [1997] holds the parameters fixed over time and adopts an error components structure. Bates groups options according to their moneyness and maturity, and allows the group pricing errors to be serially correlated with normally distributed, group-specific shocks. In addition, he allows for an idiosyncratic shock for each option with a variance that is common to the group. That is, he assumed that for group I

$$\epsilon_{it} = \epsilon_{It} + \sigma_I \eta_{it}, \quad \text{for } i \in G(I,t), \tag{15.26}$$

$$\epsilon_{It} = \rho_I \epsilon_{I,t-1} + \nu_{It}, \tag{15.27}$$

where $G(I,t)$ is the set of observations for group I at date t, ν_{It} is a mean zero, Normally distributed shock that is common to all options in group I and the ν_{It} may be correlated across groups, and $\eta_{it} \sim N(0,1)$ and is uncorrelated with ν_{It}. He uses Kalman filtering methods to estimate the error components and a generalized least squares fitting criterion to estimate the parameters.

Chernov and Ghysels [2000] use the *SME* approach to estimation with an auxiliary model of the type suggested by Gallant and Tauchen [1996] (see

Chapter 4). They focus on short-term *ATM* call prices, with *ATM* defined as $k_t \in [.97, 1.03]$, for the sample period November, 1985 until October, 1994. The resulting series of call prices reflects variation over time in both the strike price, as *ATM* changes with market levels, and maturity, as the maturity of the short-term option that is closest to being *ATM* changes.

In constructing simulated series, with the simulation length $\mathcal{T}$ larger than the sample size T, one has to make assumptions about the changing nature of the strike prices and contract maturities of the *ATM* calls. Chernov and Ghysels address this issue by cycling through both the set of option maturities and the degree of moneyness in the data set. The latter was chosen, instead of strike prices, because the simulated cash prices may be very different than what was experienced in the historical sample.

A third estimation strategy, implied-state method of moments (*IS-GMM*), was pursued by Pan [2002].[2] Letting $y_t = \ln S_t - \ln S_{t-1} - (r - q)$, Pan constructs $M \geq K$ moment conditions of the form

$$E\left[h\left(\vec{y}_t^{\ell}, \vec{v}_t^{\ell}, \theta_0\right) \mid X_{t-1}\right] = 0\,, \tag{15.28}$$

where $\theta_0 \in \mathbb{R}^K$ is the population parameter for her model and $h : \mathbb{R}^{\ell} \times \mathbb{R}_+^{\ell} \times \Theta \rightarrow \mathbb{R}^M$ is the function defining the moments to be used in estimation. The moments may depend on an ℓ-history of the excess return and volatility. What distinguishes Pan's estimation strategy from conventional *GMM* is the construction of these moment functions using model-implied volatilities in place of v_t. Specifically, let $c_t = C_t/S_t$ denote the price-to-spot ratio of the option observed on date t, with time τ_t to expiration and strike-to-spot ratio k_t. Given this option price, S_t, and a value of the parameter vector θ, we invert the option pricing model (15.25) for v_t^{θ}. That is, for the domain of invertibility $\Xi \subset [0,1] \times \Theta \times \mathbb{R}_+ \times \mathbb{R}_+ \times \mathbb{R}_+ \times \mathbb{R}_+$ (with respect to volatility) of the option-pricing function O, the mapping $g : \Xi \rightarrow \mathbb{R}_+$ is uniquely defined by

$$O(g(c, \vartheta, \tau, k), \theta, \tau, k) = c\,, \tag{15.29}$$

[2]Pan allows for stochastic interest rates and dividend yields in her empirical analysis. The processes for interest rates and dividends are estimated first and then the parameters from these processes were input into the option pricing analysis as if they were the population parameter values (no adjustment to the standard errors of the option pricing model were made for two-stage estimation). We will abstract from this first stage and discuss estimation as if r and q are constant or follow known processes.

for all $(c, \theta, \tau, k) \in \Xi$. Thus, for any $\theta \in \Theta$, the date-t option-implied volatility is defined by

$$v_t^\theta = g(c_t, \theta, \tau_n, k_n)\,. \tag{15.30}$$

At $\theta = \theta_0$, $v_t^{\theta_0}$ is the true market-observed (according to this model) volatility.

Using these model-implied volatilities, and the associated sample moments

$$G_T(\theta) = \frac{1}{T}\sum_{t \le T} h\left(\vec{y}_t^{\ell}, \vec{v}_t^{\ell\theta}, \theta\right), \tag{15.31}$$

the *IS-GMM* estimator is defined as

$$\theta_T = \arg\min_{\theta \in \Theta}\; G_T(\theta)'\, \mathcal{W}_T\, G_T(\theta)\,, \tag{15.32}$$

where $\{\mathcal{W}_T\}$ is a $M \times M$ positive semi-definite distance matrix.

The asymptotic theory from Chapter 3 is often not directly applicable to *IS-GMM* estimation of option pricing models because, for exchange-traded options in particular, a time series of fixed-maturity options is not generally available. Therefore, it has become common practice to choose the maturity τ_t of the option whose price is observed at date t to be as close as possible to a given maturity (say 30 days), subject to the requirement that the option price is not too far out of the money. Typically, this selection rule implies a repetitive, nearly deterministic time path for $\{\tau_t\}$. A representative example of this problem is displayed in Figure 15.2 from Pan [2002].

To accommodate this problem, following Pan, we let $X_t^{\theta\ell}$ denote the ℓ-history of $X_t = (y_t, v_t^\theta, k_t)$ and $\vec{Y}_t^{\ell}$ denote the ℓ-history of τ_t. We then assume that

Assumption 15.1 (Time Stationarity of Y) *$\{Y_t\}$ has finitely many outcomes, denoted $\{1, 2, \ldots, I\}$. For each outcome i and each positive integer T, let $A_T^{(i)} = \{t \le T : Y_t = i\}$ be the dates, up to T, on which Y has outcome i. For each i, there is some $w_i \in [0,1]$, such that*

$$\lim_T \frac{\#A_T^{(i)}}{T} = w_i \quad a.s.\,, \tag{15.33}$$

where $\#(\,\cdot\,)$ denotes cardinality.

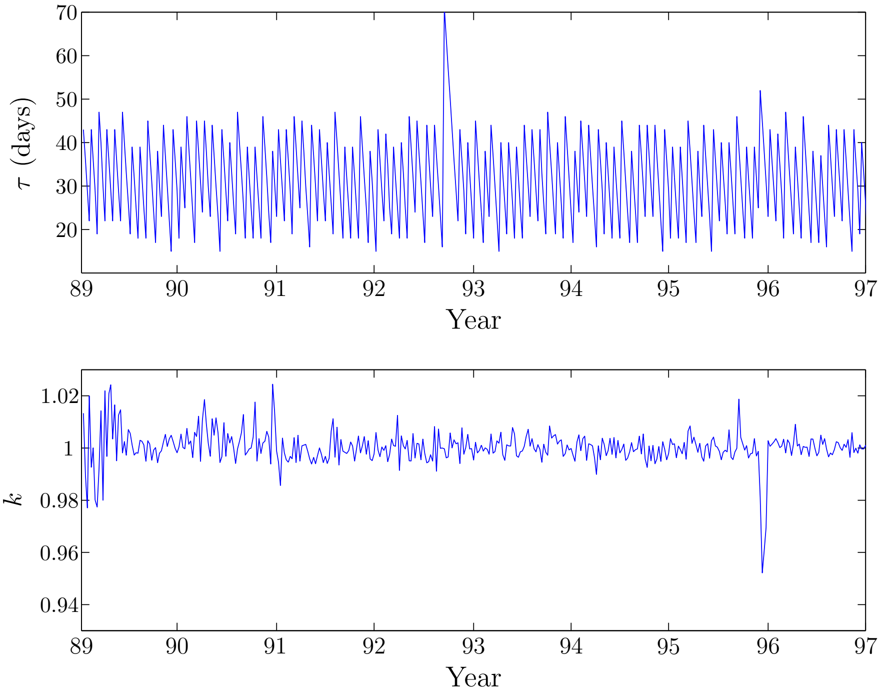

Figure 15.2: Time series of contract variables: time-to-expiration τ and strike-to-spot ratios k.

Assuming that X_t^θ is geometrically ergodic for given $\theta \in \Theta$, we known that functions of $X_t^{\theta\ell}$ satisfy a strong law of large numbers. For consistency we assume in addition that a uniform strong law of large numbers is satisfied under the sampling Assumption 15.1:

Assumption 15.2 (USLLN of $A^{(i)}$-Sampling) *For each outcome i of Y, letting*

$$G_T^{(i)}(\theta) = \frac{1}{\# A_T^{(i)}} \sum_{t \in A_T^{(i)}} h(\vec{X}_t^{\ell\theta}, \theta, i)\,,$$

$G_\infty^{(i)}(\theta) = \lim_T G_T^{(i)}(\theta)$ *exists (pointwise SLLN), and*

$$\sup_{\theta\in\Theta} |G_T^{(i)}(\theta) - G_\infty^{(i)}(\theta)| \to 0 \quad a.s.\,. \tag{15.34}$$

To develop more primitive assumptions that would imply Assumption 15.2 we would need to impose some type of Lipschitz condition on $h(x, \theta, i)$ as a function of θ as in Chapter 4. See Pan [2000] for further discussion of these issues.

It remains to discuss the selection of the moment function h. The primary motivation for pursuing the *IS-GMM* estimator was that it exploited the structure of the affine state process. As discussed in Chapter 6, the conditional characteristic function of an affine jump-diffusion is known and, hence, so are the conditional moments of the state vector. Pan [2002] chose the moments $E^\theta[\epsilon_t \mid \mathcal{F}_{t-1}] = 0$, where $\mathcal{F}_t$ was the information set generated by the history of (S_t, v_t) and the components of ϵ_t were

$$\begin{gathered} \epsilon_t^{y1} = y_t - E_{t-1}^\theta\left(y_t\right), \quad \epsilon_t^{y2} = y_t^2 - E_{t-1}^\theta\left(y_t^2\right) \\ \epsilon_t^{y3} = y_t^3 - E_{t-1}^\theta\left(y_t^3\right), \quad \epsilon_t^{y4} = y_t^4 - E_{t-1}^\theta\left(y_t^4\right) \\ \epsilon_t^{v1} = v_t - E_{t-1}^\theta\left(v_t\right), \quad \epsilon_t^{v2} = v_t^2 - E_{t-1}^\theta\left(v_t^2\right) \\ \epsilon_t^{yv} = y_t\, v_t - E_{t-1}^\theta\left(y_t v_t\right), \end{gathered} \tag{15.35}$$

where all of the conditional expectations are known functions of θ and v_{t-1}.[3]

For this set of conditional moment restrictions based on ϵ_t, the optimal set of instruments can be constructed following Hansen [1985] and the discussions in Chapters 3 and 6. In the case of $IS - GMM$ estimation, one cannot literally use Hansen's optimal instruments, as they involve terms of the form $E\left[\partial\epsilon_t/\partial\theta|\mathcal{F}_{t-1}\right]$. The error term ϵ_t depends on θ in two ways: directly through the functional dependence of the expectations $E[\cdot|\mathcal{F}_{t-1}]$ on θ and indirectly through the parameter dependence of v_t^θ. The contribution of the former term resides in $\mathcal{F}_{t-1}$ so the conditioning can be dropped. However, in general, the contribution of the second term is not in $\mathcal{F}_{t-1}$ and the conditional expectation of this term enters the optimal instruments. This expectation is unknown and therefore implementation of the optimal instruments would, at a minimum, be computationally demanding. Pan circumvented this problem by omitting the second term; that is, she computed the optimal instruments as if v_t^θ was known (did not depend on θ). The resulting computational simplicity is traded off against the loss in efficiency from omitting a component of the truly optimal instruments.

Finally, *MCMC* estimation (see Section 6.8) is a potentially attractive approach when one or more of the state variables are latent. This situation

[3]The fact that the conditional characteristic function of y_t depends only on v_{t-1} was seen in Chapters 6 and 7. This is a feature of this particular formulation of the stochastic volatility model.

arises in the analysis of option pricing models when the parameters of the SVJ model are fit to stock return data alone (volatility is latent) or, more generally, volatility follows a multifactor process and a smaller number of asset prices than state variables are used in estimation.

15.4 Econometric Analysis of No-Arbitrage Option Pricing Models

Comparing the results from the estimation of no-arbitrage option pricing models is complicated by fact that some authors have estimated their models using both options and spot market data, while others have used only options data;[4] various estimation methods have been used; and different sample periods for the data are often used. Eraker [2004] used both returns and options data (daily) on S&P 500 index over the period January 1987 to December 1990, and then out-of-sample tests were based on the period January 1991 to December 1996. Pan [2002] used both returns and option data (weekly) on S&P 500 index from January 1989 to December 1996. Bakshi, Cao, and Chen [1997] use only options data on the S&P 500 index from June 1988, to May 1991. Finally, Bates [2000] uses S&P 500 index options data from January 1988 to December 1993.

Comparing across these studies, several consistent patterns emerge. The estimate of the long-run mean of the stochastic volatility process, $\bar{v}$, tends to drop significantly with the introduction of jumps in returns (i.e., in moving from model SV to model SVJ). The same is true of the volatility of volatility, σ_v, though its magnitude varies considerably with the sample period. Together, these findings suggest that a substantial portion of the unconditional volatility of returns is due to the jump component.

Representative estimates of the parameters of models with jumps, estimated with using options or both options and return data, are displayed in Tables 15.1 and 15.2.[5] Focusing first on the parameters of the diffusion part

[4] In this chapter we focus primarily on studies that have used options prices directly in the estimation. Studies of the distribution of the underlying stock price processes were reviewed in Chapter 7.

[5] We follow the convention in Eraker [2004] and present the results for a daily time interval. Annualized numbers referenced in the subsequent discussions are obtained by multiplying by 252 (approximate number of trading days per year). We are grateful to Bjorn Eraker for assistance with the conversions of the annualized estimates reported in

		Diffusion Parameter Estimates				
Author	Model	$\kappa_v^{\mathbb{Q}}$	$\bar{v}$	σ_v	ρ	η_v
BCC	SVJ	0.008	1.60	0.15	-0.57	NA
Pan	SVJ	0.013	0.61	0.12	-0.53	0.012
Eraker	SVJ	0.011	1.65	0.20	-0.59	0.009
	SVCJ	0.011	1.35	0.16	-0.58	0.013
	SVSCJ	0.006	0.94	0.14	-0.54	0.017

Table 15.1: Estimates of the diffusion parameters using data on S&P500 index options prices.

of the models (Table 15.1), there is a notable similarity across these studies even though they used different sample periods and very different methods of estimation. The rate of mean reversion of the volatility, κ_v, is low under both $\mathbb{Q}$ and $\mathbb{P}$ consistent with a high degree of persistence in volatility that was documented in Chapter 7.[6] These studies also consistently find evidence for a substantial "leverage" effect in that the estimates of ρ are all roughly -0.5. The most notable difference is Pan's relatively lower estimate of the long-run mean of the volatility process, $\bar{v}$.

All three studies also show a positive volatility risk premium ($\eta_v > 0$)– investors tend to be adverse to increases in volatility. So when volatility is high, option prices are higher than what would be obtained by using historical volatility in pricing. Precise estimation of the risk premium parameters is difficult, partly because of the limited history of options data available for estimation and (in some cases) limited availability or use of prices on options far out of the money. It is the out-of-the-money puts and calls (in particular, deep in the money calls or deep out of the money puts) that are likely to be most revealing of risk premiums. Eraker reports that η_v is statistically significant (at conventional levels) in model *SVCJ*, but not in model *SVJ*.

Where these studies differ is in their implications for the role of jumps in option pricing. The model-implied average arrival rates of jumps $E[\zeta_S]$ (which is just the estimated stock price jump intensity when ζ_S is state-independent) suggest that jumps in returns occur about 0.002 times per day, or once every two years. Interestingly, this is a different picture than what

other papers.

[6]The mean reversion under $\mathbb{P}$, $\kappa_v^{\mathbb{P}}$, is obtained as $\kappa^{\mathbb{P}} = \kappa^{\mathbb{Q}} + \eta_v$. What is new here, relative to the insights learned from the study of stock returns alone (see Chapter 7), is that volatility remains persistent under $\mathbb{Q}$, after adjusting for the volatility risk premium.

		Jump Parameter Estimates						
Author	Model	ζ_v	$E[\zeta_S]$	$m^{\mathbb{Q}}_{SJ}(\%)$	$m^{\mathbb{P}}_{SJ}(\%)$	δ_J (%)	ρ_J	m_{vJ}
BCC	SVJ	NA	0.002	-5.00	NA	7.00	NA	NA
Pan	SVJ	NA	0.001	-19.2	-0.80	3.87	NA	NA
Eraker	SVJ	NA	0.002	-2.00	-0.39	6.63	NA	NA
	SVCJ	$= \zeta_S$	0.002	-7.51	-6.06	3.36	-0.69	1.64
	SVSCJ	$= \zeta_S$	0.002	-7.90	-1.54	2.07	-2.24	1.50

Table 15.2: Estimates of the jump parameters using data on the S&P500 index.

emerged from the student of stock return data alone. For example, Eraker, Johannes, and Polson [2003], in their study of jump-diffusions estimated with stock returns alone, obtained an arrival rate closer to 2 jumps per year. Eraker re-estimated model *SVCJ* for the same sample period underlying Table 15.2 using return data alone and obtained an arrival rate of about 1 jump per year. While the estimates clearly differ with sample period, these patterns do suggest that inclusion of the options data tends to reduce the predicted frequency of jumps over what is obtained from return data alone. One possible explanation of this pattern is that, when return data alone is studied, the jump parameters are *used* by the likelihood function to provide a better fit to the time-varying volatility of returns. An extreme version of this was seen in Chapter 7 where a pure jump diffusion model (with no stochastic volatility) gave rise to nearly 200 jumps per year! With the inclusion of options data, much more weight in the likelihood function is given to the role of jumps matching the volatility skew in this market. The estimated jump frequency will also depend on the actual incidence of jumps-like behavior during the sample period.

Table 15.2 reveals differences across studies in the average size of jumps in returns when they do occur, particularly under $\mathbb{Q}$. Where as the mean jumps sizes range between -5% to -10% for most studies, Pan obtains a mean of -19%. So her analysis gives somewhat less frequent jumps, with much more severe (under $\mathbb{Q}$) amplitudes when they do occur. Using a similar model, but a different data set, Bates [2000] also finds evidence for somewhat more severe jump amplitudes. These larger (in absolute value) jump amplitudes will generate a more pronounced volatility skew in model-implied Black-Scholes volatilities. Perhaps Pan's inclusion of an in-the-money call in her empirical analysis of model *SVJ* underlies her larger estimate of $m^{\mathbb{Q}}_{JS}$.

Notice that, comparing Pan's results for model *SVJ* with Eraker's results for model *SVSCJ* (which share a stochastic arrival rate for return jumps), quite similar estimates of $m^{\mathbb{P}}_{SJ}$ are obtained. As stressed by Eraker, $m^{\mathbb{P}}_{SJ}$ is difficult to estimate precisely, because it affects the distribution of stock returns but not option prices. Nevertheless, these findings suggest that differences in $m^{\mathbb{Q}}_{SJ}$ across studies arise due to very different estimates of the jump-amplitude risk premiums.

In order to preserve positivity of the volatility process, jumps in volatility in models *SVCJ* and *SVSCJ* are constrained to have (constant) positive amplitudes. As such, these models are by construction incapable of capturing abrupt declines in volatility, such as occurred after the period of high volatility during the 1987 crash. Similar observations have led some authors to examine multi-factor models for volatility; see, for example, Bates [2000] and Chernov, Gallant, Ghysels, and Tauchen [2000]. Pan [2002] and Eraker [2004] find evidence supportive of such an extension, but we are not aware of any systematic studies of option prices using a two-factor model for volatility.

How much does the enrichment of these models through additions of jumps to returns and volatility improve their fits to options prices and volatility skews? Bakshi, Cao, and Chen [1997] find a substantial improvement in model *SV*'s usefulness in hedging options positions, over the Black-Scholes model with constant volatility. However, adding jumps in returns (extending to model *SVJ*) does not lead to a further improvement in hedging performance. Though the pricing errors are on average smaller with than without jumps. Bates [2000] and Pan [2002] also find substantial improvements in fit for model *SVJ* over model *SV* to the volatility smirks, over a wide range of market conditions. These findings are in contrast to those of Eraker [2004], who finds that jumps in returns or volatility lead to quite small improvements in pricing errors. Similar results were obtained in his out-of-sample analysis.

The differences across these students may be partly attributable to the different estimation strategies use. As we have seen, some studies choose the parameter estimates to minimize squared pricing errors, while other use estimation methods that do not minimize pricing errors (e.g., *MCMC*). Differences in conclusions could also arise due to different ways of measuring pricing errors. However, most studies have adopted the convention of measuring errors in terms of dollars and cents, as apposed to say pricing errors as a percentage of the price of the underlying option. Alternative measures, like percentage errors, would effectively change the weighting given contracts by maturity and degree of moneyness.

Overall it seems that to explain the presence and temporal behavior of implied-volatility smirks, jumps in returns or volatility, or some other non-diffusing behavior of returns, will be necessary. We also found in Chapter 7 that such extensions of the basic log-normal diffusion model with stochastic volatility were also necessary to model to conditional distributions of stock returns.

15.5 Options and Revealed Preferences

Two complementary approaches to linking preferences and option prices have been pursued in the literature. Ait-Sahalia and Lo [2000], Jackwerth [2000], and Rosenberg and Engle [2002] use non-parametric and semi-parametric estimates of the marginal rate of substitution (or alternatively the pricing kernel) to back out the implied risk-aversion of the representative agent. Alternatively, Garcia, Luger, and Renault [2003] and Liu, Pan, and Wang [2004] explore specific parametric specifications of preferences and state variables with the goal of generating the empirical characteristics observed in option pricing data.

15.5.1 Non-parametric / Semi-parametric Approaches

The objective of this literature is to back out from option prices information about he risk aversion of market participants. In Chapter 8 we saw that the pricing kernel transforms the historical to the risk-neutral distribution of the state vector Y underlying risk in an economy. Suppose that we are in an economic environment where the pricing kernel q^* is a representative agent's marginal rate of substitution,

$$q^*_{t+1} = m_{t+1} \equiv \frac{U'_{t+1}}{U'_t}, \tag{15.36}$$

where U denotes the period utility function including agents' subjective discount factor. Then, using (8.25), we can write

$$m_{t+1} = e^{-r_t} \frac{f^{\mathbb{Q}}(Y_{t+1}|Y_t)}{f^{\mathbb{P}}(Y_{t+1}|Y_t)}. \tag{15.37}$$

It follows that, from knowledge of the historical and risk-neutral distributions of the state, we can infer agents' marginal rate of substitution.

Using results from Leland [1980], we can use this observation to construct a measure of agents' coefficient of absolute risk aversion (CRA). Specifically, assume markets are economically complete, there is one traded risky asset, and that the representative agent's consumption (C_t) is equal to her wealth that, in turn, is equal to the value of the risky asset (S_t), $C_t = S_t$. In this setting, agents' coefficient of relative risk aversion,

$$CRA_{t+1} = -\frac{U''_{t+1}}{U'_{t+1}} = -\frac{\partial m_{t+1}/\partial C_{t+1}}{m_{t+1}}, \tag{15.38}$$

can be expressed in terms of the densities $f^{\mathbb{P}}$ and $f^{\mathbb{Q}}$:

$$CRA_{t+1} = \frac{\partial f^{\mathbb{P}}(S_{t+1}|S_t)/\partial S_{t+1}}{f^{\mathbb{P}}(S_{t+1}|S_t)} - \frac{\partial f^{\mathbb{Q}}(S_{t+1}|S_t)/\partial S_{t+1}}{f^{\mathbb{Q}}(S_{t+1}|S_t)}. \tag{15.39}$$

It follows that, if the density functions $f^{\mathbb{P}}$ and $f^{\mathbb{Q}}$ can be estimated nonparametérically, then the risk aversion implicit in security prices can be computed without having to specify $U(\cdot)$ explicitly.

Ait-Sahalia and Lo [2000] and Jackwerth [2000] pursue this idea by using kernel density estimators of $f^{\mathbb{P}}$ constructed from historical return data. To compute $f^{\mathbb{Q}}$, Ait-Sahalia and Lo [2000] exploit the insight of Banz and Miller [1978] and Breeden and Litzenberger [1978] that the state-price density can be computed from the second derivative of an option price with respect to the strike price of the option. That is, letting $\mathcal{O}(S_t, K)$ denote the price of an option on the market index S struck at K and maturing at date $t+1$,

$$f^{\mathbb{Q}}(S_{t+1} = K|S_t) = e^{r_{t+1}} \frac{\partial^2 \mathcal{O}(S_t, K)}{\partial K^2}. \tag{15.40}$$

Following Ait-Sahalia and Lo [1998], market option prices and kernel regression techniques are used to estimate the function $\mathcal{O}(S_t, K)$ nonparametrically. This estimate is then differentiated with respect to K to estimate $f^{\mathbb{Q}}(S_{t+1}|S_t)$.

Alternatively, Jackwerth [2000] follows the approach of Jackwerth and Rubinstein [1996] in estimating $f^{\mathbb{Q}}$. The risk-neutral density if chosen nonparametrically to fit a cross-section of option implied volatilities, subject to exogenously given smoothness constraints on the shape of the conditional density. Rosenberg and Engle [2002] pursue a third approach by working with a parametric model of the pricing kernel and modeling the equity index return process using an asymmetric $GARCH$ process (see Chapter 7) with empirical innovation density.

Depending on the model and choice of nonparametric estimator, the resulting pricing kernel is not always monotonically decreasing in the wealth level of investors. Consequently, in the studies of Jackwerth [2000] and Rosenberg and Engle [2002] there are regions of negative risk risk aversion. (Ait-Sahalia and Lo 2000) find that relative risk aversion is positive, but non-monotonic and varies greatly (from 1 to 60 in the specified range, with the weighted average being 12.7). These authors also find that the marginal rate of substitution is non-monotonic in wealth levels. Together these studies suggest that the *CRRA* model of preferences (see Chapter 10) is not consistent with the options data. Jackwerth [2000] finds that, before the 1987 crash, the physical and the risk neutral distribution for the S&P 500 were more 'log-normal'-like than after the crash. Additionally, after the crash, the risk neutral distribution is more left-skewed and more peaked (leptokurtic).

A natural question is: how robust are these findings to assumptions being made about the underlying economic environment and the dimensionality of the state vector? These studies presume that stock prices follow a univariate process and, in particular, rule out stochastic volatility. Preferences are also defined over this single state variable, wealth. Garcia, Ghysels, and Renault [2004] illustrate how the approach taken by Rosenberg and Engle [2002] can lead to incorrect inferences about risk aversion, essentially because of a missing factor in the pricing kernel. It seems plausible that their concerns are nonparametric estimates of risk aversion are equally applicable to the other fitting methods applied to date. If agents' pricing kernels are state-dependent– depend on more than the price of security underlying the option being priced– then conclusions drawn about agents' preferences may well be misleading.

15.5.2 Preference-Based Models of Option Prices

With these cautionary observations in mind, we turn next to the literature that has used parametric specifications of preferences in an attempt to generate volatility smirks like those observed historically. Two specifications that are natural candidates for exploration in this setting are the recursive preferences of Epstein and Zin [1989] and preferences that accommodate habit formation.

Garcia, Luger, and Renault [2003] explore the properties of options prices implied by a model in which preferences are given by the recursive form (8.20) - (8.21). Consumption and dividend growths are assumed to follow a joint *i.i.d.* process conditional on a latent state variable that follows a Markov switching process. Three different models are compared: their most general option pricing model implied by the Epstein-Zin style preferences, the special case of expected utility (the *CRRA* is equal to the inverse of the intertemporal elasticity of substitution), and a preference-free model that amounts to a discrete-time version of the model examined in Hull and White [1987]. The latter model adjusts the basic Black-Scholes model to accommodate stochastic volatility that is *not priced*. They find that their most general model has the smallest pricing errors on average, and the preference-free model has the largest errors.

In a complementary study, these authors also examine the implications of their model for volatility smiles. Several features of their results suggest Epstein-Zin preferences, along with their Markov switching model for the state, cannot generate the smirk-like patterns observed in index options markets. Most notably, from Figure 10 in Garcia, Lugar, and Renault [2001], one sees that, holding the *CRRA* constant across models, the added flexibility of Epstein-Zin preferences over the nested expected utility model shifts the Black-Scholes implied volatility surface to horizontally to the right. This is approximately true in both of their regimes. Given their quoting convention, moving to the right highlights call options that are more in-the-money and put options that are deeper out-of-the money. The pronounced volatility smirk in actual markets implies that these options should be more expensive (have higher implied volatilities). However, for the parameters chosen, the graphs suggest that the model with recursive Epstein-Zin preferences generates cheaper, rather than more expensive, out-of-the money put option prices than the nested expected utility model. In other words, it appears that Epstein-Zin preferences will generate even less of a volatility smirk than the expected utility model.

Interestingly, relatively less attention has been given to investigating the properties of option prices in models of habit formation. Bansal, Gallant, and Tauchen [2004] price options on an aggregate consumption claim using a model with habit formation that is calibrated to U.S. aggregate data. However, they do not explore in depth the implications of their model for volatility smiles.

The premise of the analysis of Liu, Pan, and Wang [2004] is that standard recursive formulations of preferences are unlikely to be able to generate volatility smirks like those observed in the data, because of "rare event" premia implicit in the options prices. These authors consider a representative agent model in which the representative agent's aggregate endowment is affected by a diffusion component and a jump component, with the latter being the source of rare and unpredictable events. The agent is risk averse over both components.

A special feature of the jump component is that, since jumps occur infrequently, the agent is not fully informed about the parameters of the distribution of this component. As such, the agent's investment decisions reflect this uncertainty over and above the usual risk aversion. More precisely, the agent's endowment process under $\mathbb{P}$ is given by the following special case of (15.1):

$$dY_t = \mu Y_t dt + \sigma Y_t dW_t + (e^{J_t} - 1)Y_{t-} dZ_t, \tag{15.41}$$

where the Poisson jump process Z (with intensity λ) has been scaled by its random amplitude and $J_t \sim N(m_J, \delta_J^2)$. The capture uncertainty is about the jump component, the agent is presumed to consider alternative models for the jump process that are characterized by their Radon-Nikodym derivative with respect to the reference model $\xi_T = d\mathbb{P}(\xi)/d\mathbb{P}$:

$$d\xi_t = \left(e^{a+bJ_t-bm_J-\frac{1}{2}b^2\delta_J^2} - 1\right)\xi_{t-}dZ_t - (e^a - 1)\lambda\xi_t dt, \xi_0 = 1. \tag{15.42}$$

The authors show that this setup amounts to the agent examining models with arrival intensities and mean jump sizes (k) in the set

$$\lambda^\xi = \lambda e^a \text{ and } 1 + k^\xi = (1+k)e^{b\sigma_J^2}. \tag{15.43}$$

Finally, the representative agent is assumed to solve a "robust control" problem (see Anderson, Hansen, and Sargent [2000]) with preferences defined over the admissible a and b as:

$$U_t = \inf_{a,b} E_t^\xi \left[\int_t^T e^{-\rho(s-t)} \left(\frac{1}{\phi}\psi(U_s)H(a_s, b_s) + \frac{c_s^{1-\gamma}}{1-\gamma}\right) ds\right], \tag{15.44}$$

where H is the cost from deviating from the reference model:

$$H(a,b) = \lambda\left[1 + (a + \frac{1}{2}b^2\delta_J^2 - 1)e^a + \beta(1 + e^a(e^{a+b^2\delta_J^2} - 2))\right]. \tag{15.45}$$

Since out-of-the money options are more sensitive to large movements (jumps), they are able to separately identify the risk premia associated with risk aversion and uncertainty aversion. Introducing uncertainty aversion over the jump component increases the implied volatility for at-the-money options and also adds flexibility in fitting the out-of-the money prices so as to allow them to match the implied volatility smile / smirk patterns observed in the data. Absent this uncertainty aversion, their model (with what amounts to standard *CRRA* preferences) generates a relatively small smile. The authors argue further that extending the benchmark model to allow for habit formation (without uncertainty aversion) would not likely produce the required smirk in implied volatilities.

15.6 Options on Individual Common Stocks

Bakshi and Kapadia [2003] and Bakshi, Kapadia, and Madan [2003] document a number of empirical observations about individual stock options, especially in comparison to index options, offer intuition for these differences, and discuss some of their implications for pricing. Among the empirical observations they highlight are: (i)index volatility smiles are more negatively sloped than individual volatility smiles; (ii) individual stocks are mildly negatively skewed (and sometimes positively skewed) and are generally less negatively skewed than the index (which is never observed to be positively skewed); (iii) implied volatilities for individual options are higher than the corresponding historical return volatilities, but the differences are smaller than for index returns; and (iv) volatility risk premiums are smaller for individual stock options than for the index options.

Underlying these empirical observations is the proposition that any payoff function with bounded expectation is spanned by a continuum of option prices (Bakshi and Madan [2000]). Drawing upon this result, Bakshi, Kapadia, and Madan [2003] introduce payoffs that are powers of the return $R(t,\tau) = \log(S(t+\tau)) - \log(S(t))$, and show that the particular prices $V(t,\tau) = E^Q[e^{-r\tau}R(t,\tau)^2]$, $W(t,\tau) = E^Q[e^{-r\tau}R(t,\tau)^3]$, and $X(t,\tau) = E^Q[e^{-r\tau}R(t,\tau)^4]$, determine $SKEW^{\mathbb{Q}}(t,\tau)$ and $KURT^{\mathbb{Q}}(t,\tau)$. Furthermore, the prices (V, W, X) can be computed directly from call and put prices.

These authors also show that, for a model with power utility with *CRRA* γ), up to first order in γ, the risk-neutral skewness and the physical moments of the index are related by:[7]

$$SKEW^{\mathbb{Q}}(t,\tau) \approx SKEW^{\mathbb{P}}(t,\tau) - \gamma(KURT^{\mathbb{P}}(t,\tau)-3)STD^{\mathbb{P}}(t,\tau). \quad (15.46)$$

From (15.46) it is seen that skewness under $\mathbb{P}$ induces skewness under $\mathbb{Q}$. At the same time, even if $SKEW^{\mathbb{P}} = 0$, the risk-neutral distribution will tend to be skewed if the stock return exhibits excess kurtosis under $\mathbb{P}$. In general, these expressions for risk-neutral skews will not aggregate linearly across time if the $\mathbb{P}$ distribution of stock returns exhibit serial correlation. Interestingly, a positive autocorrelation tends to induce a "U"-shaped term structure of skewness, while negative autocorrelation gives skews of short-term returns that are more negative than their counterparts for long-term returns.

Denoting moneyness by $y = S/K$, these calculations also imply that the Black-Scholes implied volatility $\sigma_i(y;t,\tau)$ is related to the higher order moments of the $\mathbb{Q}$ distribution according to

$$\sigma_i(y;t,\tau) \approx \alpha_i[y] + \beta_i[y]SKEW_i^{\mathbb{Q}}(t,\tau) + \theta_i[y]KURT_i^{\mathbb{Q}}(t,\tau). \quad (15.47)$$

This expression is useful for its direct linkage of $SKEW^{\mathbb{Q}}$ and $KURT^{\mathbb{Q}}$ to the shape of the volatility skew. Bakshi et. al. show that firms with more negative skewness have larger implied volilties at low levels of moneyness, and firms with larger kurtoses have larger implied volatilities for both out-of- and in-the-money put options. Further, whereas skewness is a first order effect on the shape of the volatility smile - making it steeper, kurtosis is a second order effect on the smile and it affects out-of-the money call and put prices symmetrically.

To explore the higher moment properties of stock returns implicit in options data, Bakshi, Kapadia, and Madan [2003] use daily on the spot and options prices of the largest 30 stocks (by market capitalization) and the S%P 500. Though the options on individual stocks are "American" options, the authors found that their findings were largely insensitive to ignoring the early exercise premium. So they treated these options as "European" in their analysis. They found that the slope of the volatility curve tends to be much steeper for the index than for the individual stocks. Additionally, the at-the-money implied volatility for the index is lower than that of most individual stocks. Consistent with (15.47), individual stocks that exhibit more negative risk-neutral skewness also exhibited steeper volatility smile.

[7]This approximation actually holds somewhat more generally, to first order, and in particular applies to certain models with time-varying γ.

In a complementary study, Bakshi and Kapadia [2003] estimate the gains or losses on delta-hedged positions in individual stock options in order to assess the magnitude of the volatility risk premium in the markets for options on individual common stocks. They find that the premiums are negative, but less so on average than for the index option market.

Chapter 16

Pricing Fixed-Income Derivatives

Two quite distinct approaches to the pricing of fixed-income derivatives have been pursued in the literature. One approach takes the yield curve as given–essentially the entire yield curve is the current state vector. Then, assuming no arbitrage opportunities, prices for derivative claims with payoffs that depend on the yield curve are derived. Examples of models in this first group are the widely studied "forward-rate" models of Heath, Jarrow, and Morton [1992], Brace, Gatarek, and Musiela [1997], and Miltersen, Sandmann, and Sondermann [1997a]. Since the yield curve is an input, there is typically no associated *DTSM*; the model used to price derivatives does not price the underlying bonds. The second approach starts with a *DTSM*, often in one of the families discussed in Chapter 12, and this *DTSM* is used to simultaneously price the underlying fixed-income securities and the derivatives written against these securities, all under the assumption that there are no arbitrage opportunities. With the growing availability of time-series data on the implied volatilities on fixed-income derivatives, both approaches have been pursued in exploring the fits of pricing models to the historical implied volatilities of fixed-income derivatives.

In discussing the pricing of fixed-income derivatives, particular emphasis will be given to the formulations of the pricing models underlying recent empirical studies of derivatives pricing models. We begin with a review of pricing with affine *DTSM*s.[1] This is followed by an introduction to pricing

[1] See Leippold and Wu [2002] for a discussion of the pricing of fixed-income derivatives

with forward-rate based models. Since these models are being introduced for the first time, we discuss in some depth the various pricing measures that have been used to price derivatives. We then turn to a discussion of some of the more striking empirical challenges that have been raised based on examination of the historical data on implied volatilities in the LIBOR-based derivatives markets. We conclude this chapter with a discussion of how well various models address these challenges.

16.1 Pricing with Affine *DTSM*s

Particular attention has been given to the pricing of caps/floors and swaptions in the LIBOR/swap markets, no doubt in part because of the size and importance of these markets. To fix the notation for the tenor structure, let us suppose that, at time $t = 0$, there are N consecutive interest rate reset dates T_n, $n = 1, 2, \ldots, N$. The relevant rate for the time interval δ_n, $\delta_n = T_{n+1} - T_n$, is the Eurodollar deposit rate with tenor δ_n, $n = 1, 2, \ldots, N$ (with $T_{N+1} \equiv T_N + \delta_N$). For t greater than zero and less than or equal to T_N, we let $n(t) = \inf_{n \leq N} \{n : T_n \geq t\}$ denote the next delivery date on forward contracts.

Let $B(t,T)$ be the LIBOR discount factor at time t with maturity date T. Then, since LIBOR rates are set on a simple-interest basis,

$$B(T_n, T_{n+1}) = \frac{1}{1 + \delta_n \mathcal{R}(T_n)}, \tag{16.1}$$

where $\mathcal{R}$ is the quoted LIBOR rate for tenure δ_n at date T_n. The time-t forward LIBOR rate for a loan spanning the period $[T_n, T_{n+1}]$ is therefore given by

$$L_n(t) = \frac{1}{\delta_n} \left[\frac{B(t, T_n)}{B(t, T_{n+1})} - 1 \right]. \tag{16.2}$$

Note that $L_n(T_n) = \mathcal{R}(T_n)$, the LIBOR rate at date T_n.

A cap is a loan at a variable interest rate that is capped at some prespecified level. To price a cap, it is convenient to break up the cash flows into

in the class of quadratic-Gaussian *DTSM*s. Many of the solutions discussed subsequently to the pricing problems faced with affine *DTSM*s carry over, in suitably modified forms, to the *QG* class of models.

a series of "caplets" that capture the value of the interest rate cap in each period. Specifically, a caplet is a security with payoff $\delta_n\,[L_n(T_n)-k]^+$, determined at the reset date T_n and paid at the settlement date T_{n+1} (payment in arrears), where $L_n(T_n)$ is the spot LIBOR rate at T_n and k is the strike rate. The market value at time 0 of the caplet paying at date T_{n+1} is

$$\begin{aligned}\text{Caplet}_0(n) &= E^{\mathbb{Q}}\left[\exp\left(-\int_0^{T_{n+1}} r_u\,du\right)\delta_n\,(L_n(T_n)-k)^+\right]\\ &= E^{\mathbb{Q}}\left[\exp\left(-\int_0^{T_{n+1}} r_u\,du\right)\left(\frac{1}{B(T_n,T_{n+1})}-(1+\delta_n k)\right)^+\right]. \qquad (16.3)\end{aligned}$$

Within an affine *DTSM*, $1/B(T_n,T_{n+1}) = e^{-\gamma_0(\delta_n)-\gamma_Y(\delta_n)'Y_{T_n}}$ (see Chapter 12, equation (12.4)).[2] Therefore, valuing a caplet in this setting is equivalent to pricing a call option with contingent payoff $e^{-\gamma_0(\delta_n)-\gamma_Y(\delta_n)'Y_{T_n}}$ at date T_n and strike price $(1+\delta_n k)$. Referring back to the transform analysis for affine processes of Duffie, Pan, and Singleton [2000] discussed in Section 6.4, it is seem that the transform (6.39) can be applied directly to (16.3) to price the caplets and, hence, a cap.

Looking ahead to the case of coupon bonds, it is instructive to elaborate briefly on the pricing of an option on a zero-coupon bond. Using the $S-$forward measure $\mathbb{Q}_t^S$ induced on $\mathbb{Q}$ by the price of a zero-coupon bond issued at date t and maturing at time S, $B(t,S)$, the price $C(t,Y_t;S,T,K)$ of a call option with strike K and maturity S written on a zero-coupon bond with maturity T is

$$\begin{aligned}C(t,Y_t;S,T,K) &= E_t^{\mathbb{Q}}[e^{-\int_t^S r_u du}(B(S,T)-K)^+]\\ &= B(t,T)E_t^{\mathbb{Q}^T}[\mathbf{1}_{\{\mathbf{B(S,T)>K}\}}] - KB(t,S)E_t^{\mathbb{Q}^S}[\mathbf{1}_{\{\mathbf{B(S,T)>K}\}}]\\ &= B(t,T)Pr_t^T\{B(S,T)>K\} - KB(t,S)Pr_t^S\{B(S,T)>K\},\end{aligned}$$

where $Pr_t^S\{X>K\}$ is the conditional probability of the event $\{X>K\}$, based on the $S-$forward measure $\mathbb{Q}_t^S$. For the entire family of affine term structure models, these forward probabilities are easily computed using the

[2]See Duffie, Pan, and Singleton [2000] for a discussion of pricing caps when the payoff is expressed directly in terms of the floating rate $\mathcal{R}(T_n)$. Also, see Jarrow, Li, and Zhao [2004] for a discussion of pricing the payoff $(L_n(T_n)-k)^+$ using the transform methods in Duffie, Pan, and Singleton [2000] applied to the forward measure when $\log L_n(t)$ follows a square-root diffusion.

known conditional characteristic functions of affine diffusions and Levý inversion (Bakshi and Madan [2000], Duffie, Pan, and Singleton [2000]). That is, since $\{B(S,T) > K\} \equiv \{\gamma_0(T-S) + \gamma_Y(T-S)\cdot Y_S > \ln K\}$ and the characteristic function of $\gamma_Y(T-S)\cdot Y_S$ conditional on Y_t is known in closed form, two one-dimensional Fourier transforms give the requisite probabilities under the two forward measures. Note that only one-dimensional transforms are needed, even though the dimension of Y might be much larger.

The difficulty that arises in extending these ideas to the case of coupon bond options is that the exercise region is defined implicitly and, therefore, its probability is often difficult to compute. To illustrate the nature of the problem, let $V_t = V(t, Y_t; \{c_i\}_{i=1}^N, \{T_i\}_{i=1}^N)$ be the price of a fixed-income instrument with certain cashflows $c_1, c_2, \ldots, c_N$ payable at dates $T_1, T_2, \ldots, T_N$. Then price of a European option on this bond with strike K and maturity S is given by

$$\begin{aligned} C(t, Y_t; S, K, \{c_i\}_{i=1}^N, \{T_i\}_{i=1}^N) &= E_t^{\mathbb{Q}}\left[e^{-\int_t^S r_u du}(V_S - K)^+\right] \\ &= \sum_{i=0}^N c_i B(t, T_i) Pr_t^{T_i}\{V_S > K\} - KB(t,S)Pr_t^S\{V_S > K\}. \quad (16.4) \end{aligned}$$

The exercise region of this call option is

$$\{V_S > K\} \equiv \left\{\sum_{i=1}^N c_i B(S, T_i) > K\right\} \equiv \left\{\sum_{i=1}^N c_i e^{\gamma_0(T_i - S)} e^{\gamma_Y(T_i - S)\cdot Y_S} > K\right\},$$

where we are assuming that there are N remaining cashflows after the expiration date of the option.

If all the future cashflows c_i are positive, then this exercise boundary is a concave surface. Figure 16.1 illustrates these observations by plotting exercise boundaries for five-year at-the-money calls on thirty-year 10% coupon and discount bonds implied by the two-factor square-root model (an $A_2(2)$ model), with parameter values taken from Duffie and Singleton [1997] and the state variables evaluated at their long-run means.

Various approximations to the option price (16.4) have been proposed in the literature. Singleton and Umantsev [2002] exploit properties of the conditional distribution of the state variables in typical affine *DTSM*s to locally approximate the exercise boundaries in Figure 16.1 with straight line segments. This leads to a very accurate approximate pricing formula in terms of the

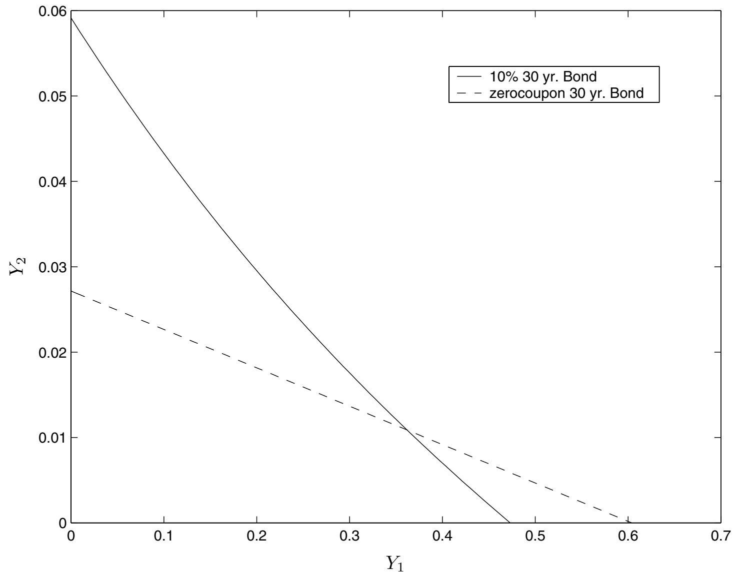

Figure 16.1: Exercise boundaries for 5-year a.t.m. Calls on 30-year 10% coupon and zero-coupon bonds implied by an $A_2(2)$ affine term structure model. Source: Singleton and Umantsev [2002].

values of options on zero-coupon bonds. Alternatively, Collin-Dufresne and Goldstein [2002b] propose another accurate approximation strategy based on an Edgeworth expansion of the probability distribution of the future value of the underlying bond.

Turning to the specific case of a swaption,[3] the value of a (settled in arrears) swap today (date t) that matures at date T_n is given by

$$V_t = c \sum_{i=i_t}^{N} B(t, T_i) + B(t, T_n) - \frac{B(t, T_{i_t})}{B(T_{i_t - 1}, T_{i_t})}, \tag{16.5}$$

where the T_i are the cashflow dates and i_t is the index of the next cashflow date at time t. The last term in (16.5) appears, because the LIBOR floating side of the contract is settled in arrears using the LIBOR rate at the preceding cashflow date. An important consequence of this settlement convention is that V_t depends not only on the current state, but also on the value of the

[3]This discussion follows Singleton and Umantsev [2002] where a more detailed discussion of the implementation of the pricing formulas can be found.

state on the previous cashflow date. Only on cashflow dates, when the last term simplifies to unity, does the direct parallel between a swap and a coupon bond emerge.

On cashflow dates, the floating side of swap is at par so the swaption price is equal to the price of a call of the same maturity and strike of one written on a coupon bond with maturity and coupon rate equal to those of the swap. Specifically, letting $T = T_N - S$, at the inception of a "T-in-S" swaption– the right to enter into a T-period swap at some future date S– the swaption price is

$$E_t^{\mathbb{Q}}\left[e^{-\int_t^S r_u du} \cdot \left(c\sum_{i=i_S}^{N} B(S,T_i) + B(S,T_n) - 1\right)^+\right], \tag{16.6}$$

where r_t is being set to the discount rate implicit in the pricing of swaps. We could easily extend this valuation approach to the case where the counterparties in the swaption contract had different ratings than those (say AA) underlying the pricing of generic swaps by introducing a different discount rate for pricing swaptions versus pricing swaps.

The pricing of both swaps and swaptions recognizes the two-sided nature of the credit risk of swaps. However, following Duffie and Singleton [1997], we are assuming that the counterparties have symmetric credit risks. As shown by Duffie and Huang [1996], asymmetry of credit quality has very little effect on the pricing of at-market interest rate swaps. Within this framework, the pricing of newly issued swaptions proceeds as in the case of a coupon-bond option.

16.2 Pricing using Forward Rate Models

A significant part of the literature on the pricing of fixed-income derivatives[4] has focused on forward-rate models in which the terminal payoff $Z(T)$ is assumed to be completely determined by the discount function $(B(t,T) : T \geq t)$ (as in Ho and Lee [1986]), or equivalently, the forward curve $(f(t,T) : T \geq t)$ (as in Heath, Jarrow, and Morton [1992]) defined by

$$f(t,T) = -\frac{\partial \log B(t,T)}{\partial T}, \quad \text{for any } T \geq t. \tag{16.7}$$

[4]This section draws extensively from Dai and Singleton [2003a].

The time t price of a fixed-income derivative with terminal payoff $Z(T) = Z(f(T, T+x) : x \geq 0)$ is then given by

$$Z(t) = E^{\mathbb{Q}}\left[e^{-\int_t^T f(u,u)du} Z(f(T, T+x) : x \geq 0) \middle| f(t, t+x) : x \geq 0 \right]. \quad (16.8)$$

For this model to be free of arbitrage opportunities, Heath, Jarrow, and Morton [1992] show that the risk-neutral dynamics of the forward curve must be given by

$$df(t, T) = \left[\sigma(t, T) \int_t^T \sigma(t, u)\, du \right] dt + \sigma(t, T)\, dW(t), \text{ for any } T \geq t, \quad (16.9)$$

and for a suitably chosen volatility function $\sigma(t, T)$. This forward-rate representation of prices is particularly convenient in practice, because the forward curve can be taken as an input for pricing derivatives and, once the functions $\sigma(t, T)$, for all $T \geq t$, are specified, then so are the processes $f(t, T)$ under $\mathbb{Q}$. This approach, as typically used in practice, allows the implied r_t and Λ_t to follow general Ito processes (up to mild regularity conditions); there is no presumption that the underlying state is Markov in this forward-rate formulation. Additionally, taking $(f(t, T) : T \geq t)$ as an input for pricing means that a forward-rate based model can be completely agnostic about the behavior of yields under the actual data generating process.

Building off of the original insights of Heath, Jarrow, and Morton, a variety of different forward-rate based models have been developed and used in practice. The finite dimensionality of $W^{\mathbb{Q}}$ was relaxed by Musiela [1994], who models the forward curve as a solution to an infinite-dimensional stochastic partial differential equation (*SPDE*) (see Da Prato [1992] and Pardoux [1993] for some mathematical characterizations of the *SPDE*). Specific formulations of infinite-dimensional *SPDE*s have been developed under the labels of "Brownian sheets" (Kennedy [1994]), "random fields" (Goldstein [1995]), and "stochastic string shocks" (Santa-Clara and Sornette [2001]). The high dimensionality of these models gives a better fit to the correlation structure, particularly at high frequencies. Since solutions to *SPDE*s can be expanded in terms of a countable basis (cylindrical Brownian motions – see, e.g., Da Prato [1992] and Cont [1999]), the *SPDE* models can also be viewed as infinite-dimensional factor models. Though these formulations are mathematically rich, in practice, they often add little generality beyond finite-state forward-rate models, because practical considerations often lead modelers to work with a finite-dimensional $W^{\mathbb{Q}}$.

Key to all of these formulations is the specification of the volatility function, since this determines the drift of the relevant forward rates under $\mathbb{Q}$ (as in Heath, Jarrow, and Morton [1992]). Amin and Morton [1994] examine a class of one-factor models with the volatility function given by

$$\sigma(t,T) = [\sigma_0 + \sigma_1(T-t)]\, e^{-\lambda(T-t)} f(t,T)^{\gamma}. \tag{16.10}$$

This specification nests many widely used volatility functions, including the continuous-time version of Ho and Lee [1986] ($\sigma(t,T) = \sigma_0$), the lognormal model ($\sigma(t,T) = \sigma_0 f(t,T)$), and the Gaussian model with time-dependent parameters as in Hull and White [1993]. When $\gamma \neq 0$, (16.10) is a special case of the "separable specification" $\sigma(t,T) = \xi(t,T)\eta(t)$ with $\xi(t,T)$ a deterministic function of time and $\eta(t)$ a possibly stochastic function of Y. The state vector may include the current spot rate $r(t)$ (see, e.g., Jeffrey [1995]), a set of forward rates with fixed time-to-maturity, or an autonomous Markovian vector of latent state variables (Cheyette [1994], Brace and Musiela [1994], and Andersen, Chung, and Sorensen [1999b]). In practice, the specification of $\eta(t)$ has been kept simple to preserve computational tractability, often simpler than the specifications of stochastic volatility in yield-based models. On the other hand, Y often has a large dimension (many forward rates are used) and $\xi(t,T)$ is given a flexible functional form. Thus, there is the risk with forward-rate models of mis-specifying the dynamics through restrictive specifications of $\eta(t)$, while "over-fitting" to current market information through the specification of $\xi(t,T)$.

More discipline, as well as added computational tractability, is obtained by imposing a Markovian structure on the forward rate processes. Two logically distinct approaches to deriving Markov *HJM* models have been explored in the literature. Ritchken and Sankarasubramanian [1995], Bhar and Chiarella [1997], and Inui and Kijima [1998] ask under what conditions, taking as given the current forward rate curve, the evolution of future forward rates can be described by a Markov process in an *HJM* model. These papers show that an N-factor *HJM* model can be represented, under certain restrictions, as a Markov system in $2N$ state variables. While these results lead to simplifications in the computation of the prices of fixed-income derivatives, they do not build a natural bridge to Markov, spot-rate based *DTSM*s. The distributions of both spot and forward rates depend on the date and shape of the initial forward rate curve.

Carverhill [1994], Jeffrey [1995], and Bjork and Svensson [2001] explore conditions under which an N-factor *HJM* model implies an N-factor Markov representation of the short rate r. In the case of $N = 1$, the question can be posed as: Under what conditions does a one-factor *HJM* model – that by construction matches the current forward curve – imply a diffusion model for r with drift and volatility functions that depend only on r and t? Under the assumption that the instantaneous variance of the T-period forward rate is a function only of (r, t, T), $\sigma_f^2(r, t, T)$, Jeffrey proved the remarkable result that $\sigma_f^2(r, t, T)$ must be an affine function of r (with time-dependent coefficients) in order for r to follow a Markov process. Put differently, his result essentially says that the only family of "internally consistent" one-factor *HJM* models (see also Bjork and Christensen [1999]) that match the current forward curve and imply a Markov model for r is the family of affine *DTSM*s with time-dependent coefficients. Bjork and Svensson discuss the multi-factor counterpart to Jeffrey's result.

An important recent development in the HJM modeling approach, based on the work of Miltersen, Sandmann, and Sondermann [1997b], Miltersen, Sandmann, and Sondermann [1997a], Brace, Gatarek, and Musiela [1997], Musiela and Rutkowski [1997a], and Jamshidian [1997], is the construction of arbitrage-free models for forward LIBOR rates at an observed discrete tenor structure. Besides the practical benefit of working with observable forward rates (in contrast to the unobservable instantaneous forward rates), this shift overcomes a significant conceptual limitation of continuous-rate formulations. Namely, as shown by Morton [1988] and Sandmann and Sondermann [1997], a lognormal volatility structure for $f(t, T)$ is inadmissible, because it may imply zero prices for positive-payoff claims and, hence, arbitrage opportunities. With the use of discrete-tenor forwards, the lognormal assumption becomes admissible. The resulting *LIBOR market model* (LMM) is consistent with the industry-standard Black model for pricing interest rate caps.

In addition to taking full account of the observed discrete-tenor structure, the LMM framework also facilitates tailoring the choice of "pricing measures" to the specific derivative products. (See Section 8.3.2 for a discussion of the pricing measures $m(P)$ based on the numeraire P.) The LIBOR market model is based on either one of the following two pricing measures: the *terminal (forward) measure* proposed by Musiela and Rutkowski [1997a] and the *spot LIBOR measure* proposed by Jamshidian [1997].

Letting $C_n(t)$ denote the price of the caplet, Brace, Gatarek, and Musiela [1997] show that, in the absence of arbitrage, both $\frac{B(t,T_n)}{B(t,T_{n+1})}$ (and hence $L_n(t)$) and $\frac{C_n(t)}{B(t,T_{n+1})}$ are Martingales under the forward measure, $\mathbb{Q}^{n+1} \equiv$

$m(B(t, T_{n+1}))$, induced by the LIBOR discount factor $B(t, T_{n+1})$. Furthermore, under the assumption that $L_n(t)$ is log-normally distributed,[5] the Black model for caplet pricing obtains:

$$C_n(t) = \delta_n B(t, T_{n+1}) \left[L_n(t) \, N(d_1) - k \, N(d_2) \right], \tag{16.11}$$

$$d_1 \equiv \frac{\log \frac{L_n(t)}{k} + \frac{v_n}{2}}{\sqrt{v_n}}, \quad d_2 \equiv \frac{\log \frac{L_n(t)}{k} - \frac{v_n}{2}}{\sqrt{v_n}}, \tag{16.12}$$

where $N(\cdot)$ is the cumulative normal distribution function and v_n is the cumulative volatility of the forward LIBOR rate from the trade date to the delivery date: $v_n \equiv \int_t^{T_n} \sigma_n(u)' \sigma_n(u) \, du$. The price of a cap is simply the sum of all un-settled caplet prices (including the value of the caplet paid at settlement date $T_{n(t)}$ which is known at t).

The Black-Scholes type pricing formula (16.11)–(16.12) for caps is commonly referred to as the *cap market model*. The simplicity of the cap market model derives from the facts that (a) each caplet with reset date T_n and payment date T_{n+1} is priced under its own forward measure $\mathbb{Q}^{n+1}$; (b) we can be completely agnostic about the exact nature of the forward measures and their relationship with each other; and (c) we can be completely agnostic about the factor structure: the caplet price C_n does not depend on how the total cumulative volatility v_n is distributed across different shocks W^n.

The simplicity of the cap market model does not immediately extend to the pricing of securities whose payoffs depend on two or more spot LIBOR rates with different maturities, or equivalently two or more forward LIBOR rates with different reset dates. A typical example is a European swaption with expiration date $n \geq n(t)$, final settlement date T_{N+1}, and strike k. Let

$$S_{n,N}(t) = \frac{B(t, T_n) - B(t, T_{N+1})}{\sum_{j=n}^{N} \delta_j B(t, T_{j+1})}$$

be the forward swap rate, with delivery date T_n and final settlement date T_{N+1}, the payoff of the payer swaption at T_n is a stream of cash flows paid at T_{j+1} and in the amount $\delta_j [S_{n,N}(T_n) - k]^+$, $n \leq j \leq N$, where the spot swap

[5] That is,

$$\frac{dL_n(t)}{L_n(t)} = \sigma_n(t)' \, dW^n(t),$$

where W^n is a vector of standard and independent Brownian motions under $\mathbb{Q}^n$, and $\sigma_n(t)$ is a deterministic vector commensurate with W^n.

rates $S_n(T_n)$ are completely determined by the forward LIBOR rates $L_j(T_n)$, $n \leq j \leq N$. The market value of these payments, as of T_n, is given by

$$\sum_{j=n}^{N} \delta_j B(T_n, T_{j+1}) \left[S_n(T_n) - k\right]^+ = \left[1 - B(T_n, T_{N+1}) - k \sum_{j=n}^{N} \delta_j B(T_n, T_{j+1})\right]^+ .$$

In order to price instruments of this kind, we need the joint distribution of the forward LIBOR rates $\{L_j(t) : n \leq j \leq N, 0 \leq t \leq T_n\}$, under a *single* measure. The *LIBOR market model* arises precisely in order to meet this requirement.

Musiela and Rutkowski [1997a] show that under the *terminal measure* $\mathbb{Q}^* \equiv \mathbb{Q}^{N+1}$, i.e., the probability measure induced by the LIBOR discount factor $B(t, T_{N+1})$, the forward LIBOR rates can be modeled as a joint solution to the following stochastic differential equations (SDEs): for $n(t) \leq \forall n \leq N$,

$$\frac{dL_n(t)}{L_n(t)} = \sigma_n(t)' \left[- \sum_{j=n+1}^{N} \frac{\delta_j L_j(t)}{1 + \delta_j L_j(t)} \sigma_j(t)\, dt + dW^*(t)\right], \tag{16.13}$$

where W^* is a vector of standard and independent Brownian motions under $\mathbb{Q}^*$. These SDEs have a recursive structure that can be exploited in simulating the LIBOR forward rates: first, the drift of $L_N(t)$ is identically zero, because it is a Martingale under $\mathbb{Q}^*$; second, for $n < N$, the drift of $L_n(t)$ is determined by $L_j(t)$, $n \leq j \leq N$.

Jamshidian [1997] proposes an alternative construction of the LIBOR market model based on the so-called the *spot LIBOR measure*, $\mathbb{Q}^B$, induced by the price of a "rolling zero-coupon bond" or "rolling CD" (rather than a continuously compounded bank deposit account which induces the risk-neutral measure):

$$B(t) \equiv \frac{B(t, T_{n(t)})}{B(0, T_1)} \prod_{j=1}^{n(t)-1} \left[1 + \delta_j L_j(T_j)\right].$$

He shows that, under this measure, the set of LIBOR forward rates can be modeled as a joint solution to the following set of SDEs: for $n(t) \leq \forall n \leq N$,

$$\frac{dL_n(t)}{L_{n(t)}(t)} = \sigma_n(L_{n(t)}(t), t)' \left[\sum_{j=n(t)}^{n} \frac{\delta_j L_j(t)}{1 + \delta_j L_j(t)} \sigma_i(L_i(t), t) dt + dW^B(t)\right],$$

(16.14)

where W^B is a vector of standard and independent Brownian motions under $\mathbb{Q}^B$ and the possible state-dependence of the volatility function is also made explicit. These SDEs also have a recursive structure: starting at $n = n(t)$, $L_{n(t)}(t)$ solves an autonomous SDE; for $n > n(t)$, the drift of $L_n(t)$ is determined by $L_j(t)$, $n(t) \leq j \leq n$.

Under the LIBOR market model, the time-t price of a security with payoff $g(\{L_j(T_n) : n \leq j \leq N\})$ at T_n is given by

$$\begin{aligned} P_t &= B(t, T_{N+1})\, E_t^* \left[\frac{g\left(\{L_j(T_n) : n \leq j \leq N\}\right)}{B(T_n, T_{N+1})} \right] \\ &= B(t, T_{n(t)})\, E_t^B \left[\frac{g\left(\{L_j(T_n) : n \leq j \leq N\}\right)}{\prod_{j=n(t)}^{n-1} (1 + \delta_j L_j(T_j))} \right], \end{aligned} \tag{16.15}$$

where $E_t^*[\cdot]$ denotes the conditional expectation operator under the terminal measure $\mathbb{Q}^*$ and $E_t^B[\cdot]$ denotes the conditional expectation operator under the spot LIBOR measure $\mathbb{Q}^B$. The Black model for caplet pricing or the cap market model is recovered under the assumption that the proportional volatility functions $\sigma_j(t)$ are deterministic.[6]

According to equation (16.15), the price of a payer swaption with expiration date T_n and final maturity date T_{N+1} is given by

$$\begin{aligned} P_{n,N}(t) &= B(t, T_{N+1})\, E_t^* \left[\frac{\left(1 - B(T_n, T_{N+1}) - k \sum_{j=n}^{N} \delta_j B(T_n, T_{j+1})\right)^+}{B(T_n, T_{N+1})} \right] \\ &= B(t, T_{n(t)})\, E_t^B \left[\frac{\left(1 - B(T_n, T_{N+1}) - k \sum_{j=n}^{N} \delta_j B(T_n, T_{j+1})\right)^+}{\prod_{j=n(t)}^{n-1} (1 + \delta_j L_j(T_j))} \right]. \end{aligned}$$

Under the assumption of deterministic proportional volatility for forward LIBOR rates, the above expression can not be evaluated analytically. In order to calibrate theoretical swaption prices directly to market quoted Black

[6]The pricing equation (16.15) holds even when the proportional volatility of the forward LIBOR rates are stochastic. Narrowly defined, the LIBOR market model refers to the pricing model based on the assumption that the proportional volatilities of the forward LIBOR rates are deterministic. Broadly defined, the LIBOR market model refers to the pricing model based on any specification of state-dependent proportional volatilities (as long as appropriate Lipschitz and growth conditions are satisfied).

volatilities for swaptions, a more tractable model for pricing European swaptions is desirable. Jamshidian [1997] shows that such a model can be obtained by assuming that the proportional volatilities of forward swap rates, rather than those of forward LIBOR rates, are deterministic. The resulting model is referred to as the *swaption market model*.

The swap market model is based on the *forward swap measure*, $\mathbb{Q}^{n,N}$, induced by the price of a set of fixed cash flows paid at T_{j+1}, $n \leq j \leq N$, namely,

$$B_{n,N}(t) \equiv \sum_{j=n}^{N} \delta_j B(t, T_{j+1}), \;\; t \leq T_{n+1}.$$

Under $\mathbb{Q}^{n,N}$, the forward swap rate $S_{n,N}(t)$ is a Martingale:

$$\frac{dS_{n,N}(t)}{S_{n,N}(t)} = \sigma_{n,N}(t)' dW^{n,N},$$

where $W^{n,N}$ is a vector of standard and independent Brownian motions under $\mathbb{Q}^{n,N}$. Thus, the price of a European payer swaption with expiration date T_n and final settlement date T_{N+1} is given by

$$P_{n,N}(t) = B_{n,N}(t)\, E_t^{n,N}\left[(S_{n,N}(T_n) - k)^+\right], \;\; t \leq T_n. \tag{16.16}$$

Under the assumption that the proportional volatility of the forward swap rate is deterministic, the swaption is priced by a Black-Scholes type formula:

$$\begin{aligned} P_{n,N}(t) &= B_{n,N}(t)\, \left[S_{n,N} N(d_1) - kN(d_2)\right], \\ d_1 &\equiv \frac{\log \frac{S_{n,N}}{k} + \frac{v_{n,N}}{2}}{\sqrt{v_{n,N}}}, \;\; d_2 \equiv \frac{\log \frac{S_{n,N}}{k} - \frac{v_{n,N}}{2}}{\sqrt{v_{n,N}}}, \end{aligned}$$

where $v_{n,N} \equiv \int_t^{T_n} \sigma_{n,N}(u)' \sigma_{n,N}(u)\, du$ is the cumulative volatility of the forward swap rate from the trade date to the expiration date of the swaption.

Several approaches have been taken to translate these ideas into econometrically tractable models for the analysis of time-series data on derivatives prices. Longstaff, Santa-Clara, and Schwartz [2001b] use a version of the *LMM* to price caps and swaptions.[7] They take the LIBOR Forward Rates

[7]Though these authors refer to their model as a "string" model, they construct their pricing model using a finite number of forward rates with discrete tenors. Therefore, the resulting framework is usefully thought of as an *LMM*. See Kerkhof and Pelsser [2002] for a formal discussion of the equivalence of the *LMM* and discrete string models.

$F_i = F(t, T_i, T_i + 1/2)$ to be the fundamental variables driving the term structure, and assume that their $\mathbb{Q}$ dynamics are:

$$dF_i = \alpha_i F_i dt + \sigma_i F_i dW_i^{\mathbb{Q}}. \tag{16.17}$$

The shocks $d^{\mathbb{Q}}W_i$ are correlated across forward rates with the time-homogenous covariance matrix Σ. To obtain the model-implied representation of zero-coupon bond prices, Longstaff et. al. note that $F_i = \frac{360}{a}(\frac{B(t,T_i)}{B(t,T_i+1/2)} - 1)$, and then they use Ito's lemma to obtain

$$dB = rBdt + J^{-1}\sigma F\, dW^{\mathbb{Q}}, \tag{16.18}$$

where $\sigma F dW^{\mathbb{Q}}$ is formed by stacking the $\sigma_i F_i dW_i$ and the Jacobian matrix J, obtained from the mapping from bond prices to forwards, has a banded diagonal form.

As noted above, the *LMM* does not lead to closed-form expressions for the prices of European swaptions. Longstaff, Santa-Clara, and Schwartz [2001b] proceed using simulation methods, based on their full characterization of the joint distribution of forward rates, to compute prices. Calibration then amounts to choosing Σ to match the market data on derivatives prices. They proceed from the spectral decomposition of the historical correlation matrix of changes in forward rates, $H = U\Lambda U'$. The relevant covariance matrix for the forward rates is assumed to be of the form $\Sigma = U\Phi U'$. In other words, the eigenvectors of H are assumed to be those of Σ in the pricing model, and all that remains is to select the eigenvalues of Σ, $diag[\Phi]$, to match the historical swaption prices. The assumption that H and Σ share the same eigenvectors amounts to imposing a special structure on the $\mathbb{P}$ drifts of the forward rates.

Han [2004] extends the simple *LMM* model in Longstaff, Santa-Clara, and Schwartz [2001b] to allow for stochastic volatility. Starting directly with bond prices, he assumes that

$$\frac{dB(t,T)}{B(t,T)} = r(t)dt - \sum_{k=1}^{N} \beta_k(T-t)\sqrt{v_k(t)}\, dZ_k^{\mathbb{Q}}(t), \tag{16.19}$$

where the volatility factors $v_k(t)$ are assumed to follow square-root diffusions ($A_1(1)$ processes). Rather than using simulation methods to price caps and swaptions, Han derives approximate prices. Specifically, for caplets, he combines the Black model (16.11) - (16.12) with the approach taken by Hull and

White [1987] in pricing equity options. Under the assumption that volatility risk is not priced, the deterministic cumulative variance v_n in (16.12) can can be replaced by the expected average variance under the $A_1(1)$ model for stochastic volatility. A similar approximation is used to price swaptions, exploiting the joint distribution of the forward rates implied by his *LMM*.

16.3 Risk-Factors and Derivatives Pricing

Much of the recent literature applying *DTSM*s to the pricing of derivatives has focused on two features of the distributions of swap rates and implied volatilities on LIBOR-based derivatives. First, a substantial portion of the variation in the prices of options on fixed-income securities is uncorrelated with the variation in the prices of the underlying bonds on which the option are based. Second, developing a model that prices various types of options written on the same underlying securities has been challenging, particularly for the case of caps and swaptions. We briefly review each of these puzzles prior to discussing the econometric studies of derivative pricing models.

16.3.1 Unspanned Stochastic Volatility

Heidari and Wu [2001] document that the common interest rate factors that explain over 99% of the variation in the yield curve can explain under 60% of the variation in swaption implied volatilities. These results come from examining yields on LIBOR contracts with maturities ranging between one and twelve monthws; interest rate swaps with maturities ranging between two and thirty years; and at-the-money (*ATM*) swaptions on one up to ten-year swap contracts with maturities of 1, 3 and 6 months. Their sample period was from October 1995 to July 2001.

The interest rate factors– the first three principal components of LIBOR and swap yields– have the familiar interpretation of "level," "slope," and "curvature" factors (see Chapter 12). To these three factors they add three volatility related factors extracted from the swaptions data. Adding these sequentially to the interest rate factors increases the average variation of implied volatilities explained to 85%, 96% and almost 98%, respectively. Thus, their findings suggest that a large fraction of the variation in implied volatilities on swaption contracts is largely uncorrelated with the sources of variation in the underlying swap rates.

Similarly, Collin-Dufresne and Goldstein [2002a] explore how much of the variation in straddles (of *ATM* caps and floors) can be explained by the variation in swap rates. These authors focus on straddles, because they are relatively insensitive to small changes in the level of yields while being highly sensitive to changes in bond-price volatility. The data for their analysis was the six-month LIBOR rate and swap rates for one through ten years to maturity from the U.S., U.K., and Japan; and straddles ranging in maturity from six months to ten year for the U.S., and six-month to eleven years for the U.K. and Japan. Their sample period was February 1995 through December 2000.

When they regressed the changes in straddle prices on changes in swap rates, they obtained relatively low (adjusted) R^2's: 0.085 - 0.391% for the U.S., -0.071 - 0.134% for the U.K., and 0.044 - 0.254% for Japan. To put these low numbers in perspective, the authors simulated data from an estimated $A_1(3)$ affine *DTSM* and used this data to rerun these regressions. They obtained R^2 of roughly 90%, clearly well above those observed in the historical data. Based on this evidence, the authors conclude that bonds do not span the fixed income markets, and specifically floors and caps seem to be sensitive to stochastic volatility which cannot be hedged by a position solely in bonds.

To better understand the nature of this unspanned stochastic volatility (*USV*), Collin-Dufresne and Goldstein [2002a] construct the principal components of the covariance matrix of the residuals from their historical regressions of straddle prices on swap yields. The first *PC* explained over 80% of this residual variation and the second explained an additional 10% or more for all three countries. Thus, the portions of the option prices that are not spanned by the bond yields appear to have a common factor that accounts for most of their variations. Of course, this *USV* puzzle, though striking, is not logically inconsistent with arbitrage-free pricing. We saw in Section 13.6 that factors can affect the distribution of the short rate r, but not affect the yields on bonds of all maturities.

Fan, Gupta, and Ritchken [2003] raise several caveats about interpreting the empirical evidence in the studies by Heidari and Wu [2001] and Collin-Dufresne and Goldstein [2002a] as indicative of *USV* or, more generally, economically incomplete markets. First, an *ATM* straddle has a highly convex payout structure while at the same time being nearly delta neutral with respect to changes in the underlying bond yields. The latter property implies that shocks (at least small ones) to the *PC*s of swap rates should have small effects on the prices of straddles. Large shocks to bond yields, on the other hand, will have an effect through the nonlinear (highly convex) dependence of straddles on their underlying risk factors. Second, in assessing the degree of economic incompleteness of a market, one is interested in how changes in the prices of straddles are related to changes in the prices of traded securities. Changes in swap rates do not correspond to the changes in prices of relevant replicating portfolios of bonds. It remains an empirical question then as to whether *DTSM*s with low-dimensional factor structures are capable of describing the time-series behavior of derivatives prices.

16.3.2 Relative Pricing of Caps and Swaptions

Though financial theory predicts a close link between the prices of caps and swaptions (as they are both LIBOR-based derivatives), developing a model that simultaneously prices both contracts has proved challenging. Explanations for this "swaption/cap puzzle" often focus on the nature of the model-implied factor volatilities and/or correlations and their roles in determining prices. For instance, Rebonato and Cooper [1997] and Longstaff, Santa-Clara, and Schwartz [2001a] compare the correlations among forward swap rates with those implied by low-dimensional factor models and find that the correlations implied by the models are much larger than those in the data. Brown and Schaefer [1999] and Carverhill [2002] find similar results using Treasury strip yields.

We can anticipate the difficulty standard *DTSM*s will have in matching yield correlations by comparing historical and model-implied correlations among weekly changes in the yield spreads for non-overlapping segments of the U.S. dollar swap yield curve. The correlation "3-2/4-3," for example, in Table 16.1 represents the correlation of daily changes in the 3yr-2yr swap spread with changes in the 4yr-3yr spread. The rows labeled "2 PC" and "4 PC" present the corresponding correlations for fitted spreads from projections onto the first two and four principal components (PCs), respectively. Notably, even using four PC's the segment correlations are larger than their sample counterparts, and the match is much worse using only two PC's (in the spirit of a two-factor *DTSM*).

Not surprisingly, when we compute model-implied segment correlations from the affine $A_M(N)$ models with $N \leq 3$, using swap data, they are all substantially larger than their historical counterparts. The same is true for the fitted, relative to the historical, treasury yields from Ahn, Dittmar, and

Segment	3-2/4-3	4-3/5-4	5-4/7-5	7-5/10-7
Historical	0.34	0.09	0.13	0.14
2 PC	0.99	0.99	0.99	0.99
4 PC	0.81	0.96	0.84	0.32

Table 16.1: Correlations of changes in swap yield spreads for various yield-curve segments. 10-7,3-2, for example, indicates the correlation between changes in the 10-7 year yield spread and the 3-2 year yield spread.

Gallant [2001]'s study of *QG* models.

Closely related to this "swaptions/caps" pricing puzzle, many have found that model-implied volatilities extracted from cap prices tend to be larger than those backed out from swaption prices. One interpretation of this puzzle is based on the observation that a cap can be viewed as a portfolio of options on forward LIBOR rates, whereas a swaption can be viewed as an option on a portfolio of forward LIBOR rates. As such, cap prices are relatively insensitive to the correlation structure of forward LIBOR rates, whereas the swaption prices depend crucially on the correlation structure. Indeed, a one-factor model can be calibrated exactly to all *ATM* cap prices, but it likely misprice swaptions because forward rates are perfectly correlated in such a model. If swaptions and caps have different sensitivities to a model's (in)ability to match yield curve segments or forward-rate correlations, then this could resolve the pricing puzzles. However, the literature is not fully in agreement about the relative responses of the prices of caps and swaptions to changes in factor volatilities or correlations.

16.4 Affine Models of Derivatives Prices

Empirical work addressing the fit of *DTSM*s to the joint distributions of swap and swaption prices has been limited. Upon fitting an $A_3(3)$ model (with independent factors) to historical swap yields, Jagannathan, Kaplan, and Sun [2001] find that their model is incapable of accurately pricing caps and swaptions.[8] However, in the light of the preceding discussion, reliable pricing of swaptions would seem to depend on using swaption data in estimation in order to "pick up" the effectively unspanned factors. That is, if one fits an

[8] To value the swaptions, Jagannathan, Kaplan, and Sun [2001] use a method developed by Chen and Scott [1995] that is specific to multi-factor $A_N(N)$ (CIR) models.

$A_M(3)$ model, for example, to swap rates alone, then the likelihood function will tend to select factors that (suitably rotated) are highly correlated with the first three *PC*s of swap rates. Including option prices directly will change the "weight" given by the likelihood function to matching the structure of volatility and will likely lead to a different factor structure.

This is confirmed by Umantsev [2001], who estimated $A_M(3)$ models using data on swap rates and swaption volatilities simultaneously over the sample period 1997 though 2001. He finds that $A_M(3)$ models, for $M = 1, 2$, fit the data notably better than an $A_3(3)$ model. Moreover, as anticipated by the descriptive findings of Heidari and Wu [2001], the third factor is related more to volatility in the swaption market than to "curvature" in the swap curve (the more typical third factor in *DTSM*s fit to yield data alone). Relating Umantsev's findings back to those of Collin-Dufresne, Goldstein, and Jones [2004], the latter study found that the third factor (beyond the first two *PC*s of bond yields) in $A_1(3)$ models was either the curvature factor or a proxy for the volatility of the short rate. However, with only three factors, an $A_1(3)$ model was not capable of matching both the curvature factor and the volatility factor. By explicitly including the implied volatilities of options into his *ML* estimation, Umantsev effectively forced selection of a volatility factor as the third factor. The results of Collin-Dufresne et. al. suggest that adding a fourth factor may allow a match to both the curvature factor and the common volatility factor that Heidari and Wu found in swaption volatilities.

16.5 Forward-Rate Based Pricing Models

Longstaff, Santa-Clara, and Schwartz [2001b] examined data on swap rates between July, 1992 and July 1999, and a cross section of thirty-four swaptions and cap prices from January, 1997 to July 1999. Working with an four ($N = 4$) factor model, the first three eigenvectors, constructed from the $\mathbb{P}$-covariance matrix of the forwards, were the familiar level, slope, and curvature factors, and the fourth factor affected the very short end of the yield curve.[9] In this regard, they extended the complementary analysis in Hull and White [2000] where a three-factor *LMM* was examined.

[9] See Chapter 13 for a discussion of the role of this factor in explaining the properties of the very short end (under one year) of the yield curve.

The probability model for forward rates was obtained by solving for the eigenvalues in Φ (from the decomposition $\Sigma = U\Phi U'$ of covariance matrix of forward rates) that minimized the pricing errors for the swaptions. This minimization was done cross-sectionally so Φ was updated every week. The temporal variation in the elements of Φ suggests the presence of time-varying conditional second moments of the forward rates, a feature of the data that was not formally taken into account in the pricing of swaptions or caps. The presence of such stochastic volatility/correlation motivates the analyses of Han [2004] and Collin-Dufresne and Goldstein [2001] (see below).

Longstaff et. al. found that their four-factor model did a quite good job of matching the cross section of thirty-four swaption prices, except during the fall of 1998 period when Russia defaulted on its domestic debt. However, when they used the model calibrated to swaption prices to price *ATM* caps, they found mean pricing errors ranging from 23% for two-year caps to 5% for five-year caps, with caps being undervalued relative to swaptions.

In a more systematic study of forward-rate and LIBOR market models Driessen, Klaassen, and Meleberg [2003] examine the effect of the number of factors on the pricing and hedging of caps and swaptions. Their *HJM*-style models have

$$df(t,T) = \mu_f(t,T,\omega)dt + \sum_{i=1}^{K} \sigma_{f,i}(t,T,\omega)\, dW_i(t). \tag{16.20}$$

For their *LMM*, LIBOR rates are assumed to follow the processes

$$dL_n(t) = ...dt + \sum_{i=1}^{K} \sigma_{L,i}(t,n,\omega) L_n(t) dW_i(t). \tag{16.21}$$

Up to three factor models are considered with the time-homogeneous volatility specifications: parametric *HJM* model: $\sigma_{f,i}(T-t) = \sigma_i e^{\kappa_i (T-t)}$; *LMM*: $\sigma_{L,i}(T-t) = h_i(T-t)$, where the h_i are deterministic functions. The data set for this analysis covers money market and swap rates, swaptions, and caps for the period January 1995 to June 1999. In both the *HJM* and LIBOR market models, the three factors resemble the standard *PC*s of swap yields, level, slope and curvature, which together explain over 96% of the variation in yields.

The smallest pricing errors are obtained for the models estimated over rolling windows, so that that parameters are allowed to change over time.

Further, incorporation of the options data directly into the estimation improves the fit of both the cap and swaption pricing models. Consistent with previous studies, their three-factor model (the largest number of factors considered) fits the best.

More recently, Jarrow, Li, and Zhao [2004] use cap price data from SwapPX to examine volatility smiles in fixed-income derivatives markets. The smile is asymmetric with *ITM* caps having a stronger skew than *OTM* caps. Furthermore, the smile is more pronounced after September, 2001. Particular attention is given to the relative performance of alternative *LMM*s in explaining these smiles. Toward this end, the authors used data on caps of maturities from one to ten years and ten different strike prices, and they divide their sample period into four sub-samples: September, 2000 to March, 2001; March, 2001 to August, 2001; November 2001 to May 2002; and May, 2002 to November, 2002.

The framework they consider has

$$\frac{dL_n(t)}{L_n(t)} = \alpha_n(t)dt + \sigma_n(t)dW^{\mathbb{P}}_{L,n+1}(t) + dZ^{\mathbb{P}}_n(t) \tag{16.22}$$

where $W^{\mathbb{P}}_{L,n+1}$ is a standard Brownian motion and $Z^{\mathbb{P}}$ is an independent jump process under $\mathbb{P}$. To reduce the dimensionality of the parameter space, Jarrow et. al. follow Longstaff, Santa-Clara, and Schwartz [2001b] and Han [2004] and assume that the instantaneous covariance matrix of changes in LIBOR rates takes the form $\Sigma_t = U\Phi_t U'$, where U is the $N \times 3$ (N is the number of forward LIBOR rates included) matrix of the first three eigenvectors of the historical covariance matrix of LIBOR rates. In other words, they assume that there are three factors underlying the temporal variation in the instantaneous variances and covariances of LIBOR rates. The i^{th} diagonal element of Φ_t, $v_i(t)$, is the instantaneous variance of the i^{th} common factor and it is assumed to follow a square root diffusion. The jump is assumed to take the same form as that in Pan [2002]'s study of equity options, including her assumption that the jump timing risk is not priced (see Chapter 15).

To complete their model, Jarrow, Li, and Zhao [2004] assume that the risk premiums associated with both the jump amplitude and stochastic volatilities are linear functions of time to maturity. A richer parametrization is not identified econometrically, because they are not studying both the cash and derivatives prices simultaneously. Finally, cap prices are obtained using the transforms in Duffie, Pan, and Singleton [2000] applied under various forward measures. The model is estimated using cap data and a variant of the implied-state *GMM* methods proposed by Pan [2002] (see Section 15.3).

Consistent with previous studies using similar models, the first three principal components underlying the construction of Σ_t corresponded to the level, slope, and curvature of the LIBOR yield curve. Based on cap data for the sample period June, 1997 through July, 2000, the level factor had the most volatile stochastic volatility (v_{it}), and the slope factor was the least volatile. Allowing all three factors to have stochastic volatilities led to notably smaller pricing errors for caps than when a subset of the three factors had constant volatilities. However, even in their most flexible model with stochastic volatility, they found significant underpricing of *ITM* and overpricing of *OTM* caps. These patterns suggest a potentially important role for jumps, and are indicative of misspecification of previous *LMM*s with constant or stochastic volatility and no jumps (e.g., Longstaff, Santa-Clara, and Schwartz [2001b] and Han [2004]).

The model in Jarrow et. al. with stochastic volatility is not capable of capturing a volatility smile, because the volatility factors underlying Φ_t and the Brownian motions $W_n^{\mathbb{P}}(t)$ driving forward rates are mutually independent. We saw in Sections 7.7 and 15.4 that negative correlation between volatility and price shocks is a potentially important contributor to skewness in returns a smile in equity options markets. However, in equity markets, this "leverage" effect does not generate sufficient skewness to match the observe smile in implied volatilities of options.

Introducing jumps in LIBOR rates (while preserving the independence between LIBOR and volatility shocks) substantially improves the fit of the *LMM*. Analogously to prior findings for equity markets, introducing jumps lowers volatilities of the factors v_i. At the same time, there is evidence of large negative jumps in LIBOR rates under the forward measure. The estimated arrival intensities of jumps were between 2 and 6% per year with very large mean relative jumps sizes (between -50% and -90%).

16.6 On Model-Based Hedging

An alternative, informative means of assessing model performance, besides the magnitudes of pricing errors, is the effectiveness of hedge positions based on an estimated model. This approach to model assessment seems particularly relevant in this literature, because of the ongoing debates about the importance of unspanned risk factors in the bond markets. The low correlations between the *PC*s of bond yields and implied option volatilities suggest, at first glance, that hedges of option position constructed from bond positions should be ineffective against key sources of risk in the options markets.

Longstaff, Santa-Clara, and Schwartz [2001b] used their four-factor model and the nested Black model to compute hedge ratios for each of the swaptions in their data base. From these ratios, they computed hedging errors, defined to be the change in swaption price less the change in value of the hedge portfolio. For the Black model, each swaption had its own hedge instrument (the underlying swap rate), while in the four-factor model there are only four common instruments. Nevertheless, the performance of the two models was almost indistinguishable (89.28% vs 89.35% of variability explained). Notably, a large percentage of the variation in changes in swaption prices was explained by the hedge portfolios.

The question of whether a small number of hedge instruments might lead to effective hedges for swaptions is the primary focus of the analysis in Fan, Gupta, and Ritchken [2003]. Table 16.2 displays the absolute hedging errors (in basis points)[10] for biweekly swaption data from March 1, 1998, to October 31, 2000. The number of factors N is the number of factors underlying the covariation in forward LIBOR rates, and it determines the number of instruments used to construct the hedge portfolios. The difference between the columns with and without "recalibration" is that in the former case the parameters are recalibrated every week, while in the latter case they are fixed for four weeks at the values used for the first week of hedging.

The results show a substantial decline in the root mean-squared errors of the hedged positions with the addition of factors out to $N = 3$. The improvement in hedging performance with the addition of one more factor ($N = 4$) is economically small and statistically insignificant. Furthermore, the (unadjusted) R^2 show that, for a one-week horizon, the hedges account for over 90% of the variation in the unhedged positions. Even over horizons of four weeks, the percentage of the variation in the unhedged positions explained by the hedges remains large.

Fan, Gupta, and Ritchken [2003] repeat their hedging analysis with straddles, motivated in large part by the earlier findings by Collin-Dufresne and Goldstein [2002a] that a significant fraction of the variation of changes the

[10] The authors multiply the root mean square of the hedging errors for a contract by 10,000 so that it becomes interpretable as a basis point error.

Expiry (years)	Swap Mat.	Unhedged Swaption	Number of Factors in the Model: With Recalibration 1	2	3	4	Without Recalibration 1	2	3	4
1	2	13.1	5.5	3.3	3.3	3.2	6.1	3.9	3.9	3.8
	3	18.9	6.6	4.4	4.1	4.6	7.1	4.8	4.7	4.8
	4	24.1	7.8	5.7	5.4	5.7	8.5	6.4	6.3	6.3
	5	28.9	8.4	6.5	6.3	6.2	8.8	7.0	6.8	6.7
5	2	10.8	5.5	4.0	3.9	4.5	5.9	4.3	4.3	4.5
	3	16.4	7.9	5.7	5.8	5.7	8.4	6.2	6.2	6.2
	4	21.5	10.1	7.2	7.2	7.1	10.7	7.9	7.9	7.9
	5	26.3	12.0	8.6	8.6	9.5	12.6	9.3	9.3	11.2
R2 -1week out-of-sample			0.83	0.92	0.92	0.91	0.80	0.90	0.90	0.90
R2 -4 weeks out-of-sample			0.67	0.86	0.92	0.91	0.62	0.80	0.86	0.86

Table 16.2: Root mean squared errors (in basis points) of the hedged and unhedged portfolios one week out-of-sample, with and without recalibration of the models. Source: Fan, Gupta, and Richkten [2003].

prices of straddles was unrelated to variation in swap rates. Again they find that large percentages (over 80%) of the variation of the unhedged straddle positions were explained by the bond-based hedge portfolios. Fan et. al. argue that a primary reason for the difference between their findings and those of Collin-Dufresne and Goldstein is that they are using actual traded bonds to construct hedges rather than running regressions of straddle prices on swap yields. Indeed, when Fan et. al. regressed the corresponding straddle volatilities (computed from swaption contracts as opposed to cap contracts) onto swap yields, they obtained low R^2's comparable to those reported in Collin-Dufresne and Goldstein [2002a] for their analysis of straddle prices. Overall, Fan, Gupta, and Ritchken [2003] conclude that the role of USV in understanding the prices of LIBOR-based derivatives is likely to be economically small.

Driessen, Klaassen, and Meleberg [2003] undertake a similar analysis of hedging, examining both caps and swaptions and a wider variety of models and methods for computing hedge ratios. They also find that a large percentage of the variation in prices of unhedged positions in derivatives is explained by changes in the prices of hedge portfolios constructed from bonds. However, these percentages are smaller than those documented by Fan et. al. As a result, the authors reach a more measured conclusion regarding the potential importance of USV, noting a potentially important role for stochastic volatility and jumps in understanding the behavior of cap and swaption prices.

16.7 Pricing Eurodollar Futures Options

A different perspective on the pricing of fixed income derivatives is offered by the study of Bikbov and Chernov [2004] of option on Eurodollar futures contracts. These authors examine $A_0(3)$, $A_1(3)$, and $A_1(3) - USV$ affine *DTSM*s estimated using weekly data on Eurodollar futures and options over the sample period January 1994 through June 2001. The model with USV is obtained by imposing the constraints on the canonical $A_1(3)$ model that lead to the volatility factor having no affect on bond prices for all maturities (see Collin-Dufresne and Goldstein [2002a]). The market prices of risk were of the form proposed by Cheridito, Filipovic, and Kimmel [2003] as an extention of Duffee [2002]'s essentially affine formulation (see Chapter 12). Call prices were computed using the transform results in Duffie, Pan, and Singleton [2000]. All of the futures and options contracts were allowed to be priced with errors (no contracts were priced perfectly by the models), and estimation was accomplished using *QML* and Kalman filtering.

When only futures data was used in estimation, signified by the superscript 'f', Bikbov and Chernov found that the volatility factor is highly correlated with the a "butterfly" Eurodollar futures position (long the six-month and ten-year futures and short two times the two-year futures) for model $A_1(3)^f$, while in model $A_1(3)^f - USV$ it was most highly correlated with the long end of the futures curve. Thus, imposing the USV constraint seems to effect a factor rotation, a feature of constraints that we discussed in Chapter 13. When both options and futures data were used in estimation, signified by the superscript 'fo,' the volatility factor was most highly correlated with the slope of the futures curve (ten-year minus six-month futures) in model $A_1(3)^{fo}$, while it was highly correlated with the implied variance from the options market in model $A_1(3)^{fo} - USV$. Thus, only in this final model did the authors find that the volatility factor is closely matched to the implied volatility in the options market. Perhaps the most surprising aspect of these results is the labelling of the factors for model $A_1(3)^{fo}$. One might have expected, as for example in Umantsev [2001], that inclusion of options prices in the estimation would lead to at least one of the factors matching up closely with the implied volatilities in this market.

Of particular interest to the issues raised in this chapter about risk factors in cash and options markets is Bikbov and Chernov's formal analysis of the *USV* restrictions on model $A_1(3)$. Comparing models $A_1(3)^{fo}$ to model $A_1(3)^{fo} - USV$, they find a notable deterioration in the fit (measured by pricing errors) for *both* the futures and options data. The percentage errors in options prices are more than twice as large at the six-month expiration and more than six times are large at the one-year expiration. Formal likelihood ratio tests of the *USV* constraints also indicate rejection at conventional significance levels.

Furthermore, a familiar tension arises: the models are able to fit certain moments at the expense of others. In the case of the $A_1(3)$ models 'fo,' without and without *USV*, the models match the kurtosis in the data quite well, but fail to match the historical volatilities. Imposing the *USV* constraints increases this tension relative to that in the canonical $A_1(3)$ model.