Dynamic Hedging With a Deterministic Local Volatility Function Model

Thomas F. Coleman[†], Yohan Kim[†], Yuying Li[†], and Arun Verma[†]

October 20, 2000

Abstract

We compare the dynamic hedging performance of the deterministic local volatility function approach with the implied/constant volatility method. Using an example in which the underlying price follows an absolute diffusion process, we illustrate that hedge parameters computed from the implied/constant volatility method can have significant error even though the implied volatility method is able to calibrate the current option prices of different strikes and maturities. In particular the delta hedge parameter produced by the implied/constant volatility method is consistently significantly larger than that of the exact delta when the underlying price follows an absolute diffusion.

In order to compute a better hedge parameter, accurate estimation of the local volatility function in a region surrounding the current asset price is crucial. We illustrate that a suitably implemented volatility function method can estimate this local volatility function sufficiently accurately to generate more accurate hedge parameters. Hedging using this volatility function for the absolute diffusion example leads to a smaller average absolute hedging error when compared with using the implied/constant volatility rate.

When comparing the hedging performance in the S&P 500 index option market as well as the futures option market, we similarly observe that the delta hedge parameter from the implied/constant volatility method is typically greater than that using the volatility function approach. Examination of the hedging error reveals that the using a larger delta factor greater than that of the true volatility yields more positive average hedging error, assuming the underlying follows a deterministic volatility model. We observe that, in both the S&P 500 index option market and futures option market, the average absolute hedging error using the volatility function approach is smaller than that of the implied/constant volatility method for a sufficiently long hedging horizon, approximately 17 days for the S&P 500 index options and 6 days for the S&P 500 futures options. In addition, the average hedging error using the volatility function approach is always smaller than that of the implied/constant volatility method.

[†]Computer Science Department and Cornell Theory Center, Cornell University, Ithaca, NY 14850. [‡]Center for Applied Mathematics, Cornell University, Ithaca, NY 14850.

^{*}Research partially supported by the National Science Fundation NSF through grant DMS-9805602 and ONR through grant N00014-96-1-0050. This research is also partially supported by the Cornell Theory Center which receives funding from Cornell University, New York State, the National Center for Research Resources at the National Institutes of Health, the National Science Foundation, the Department of Defense Modernization Program, the United States Department of Agriculture, and members of the Corporate Partnership Program and the Financial Industry Solutions Center, an SGI/Cornell joint venture in New York.

1. Introduction

Option hedging error can arise from two different sources. First, the option value is a nonlinear function of the underlying and the Black-Scholes hedging is instantaneous; however, only discrete hedging, possibly with infrequent rebalancing due to transaction costs, can be done in practice. Second, errors in the underlying price model can lead to poor hedging performance.

Many studies have shown that the classical Black-Scholes constant volatility model does not adequately describe the stock price dynamics, see e.g., [17]. Implied volatility typically exhibits a dependence on both the option strike and maturity, referred to as the volatility smile. The *constant volatility* method, which assumes that the volatility rate is constant for all the options on the same underlying, can lead to a significant model specification error. To reduce this error, the *implied volatility* method which applies different volatility rates for options with different maturities and strikes is frequently used in practice for pricing and hedging. However, this method is ad hoc and inherently inconsistent. Although the implied volatility method yields accurate option price calibration, it does not address the issue of the appropriate model specification. One objective of this paper is to illustrate the importance of the accurate model specification in hedging.

A natural extension to the constant Black-Scholes pricing formula is to allow the volatility rate to be a deterministic function of the underlying price and time. In [12], Dupire shows that, under some regularity assumptions, this local volatility function can be uniquely determined if the prices of European options of all strikes and maturities are available. In this paper we refer a *volatility function method* as the method which computes a local deterministic volatility function from the market option prices. The computed volatility function can then be used for pricing options and computing hedge parameters for risk management. Various computational methods have been proposed [2, 3, 7, 9, 10, 14, 15, 16, 17] to compute a local volatility function calibrating a finite set of market option data. In the volatility function method [7], the volatility function is represented as a spline in an attempt to estimate the unknown volatility model accurately; the use of spline offers the flexibility in potentially modeling a complex volatility function and at the same time has the capacity to limit overfitting by judiciously choosing the number of spline knots.

More recently, there have been empirical studies comparing the local volatility function approach with the implied/constant volatility rate method in pricing and hedging [5, 4, 6, 11, 15]. In particular, Dumas, Fleming and Whaley [11] compare a few parametric volatility function models, including constant and quadratic forms, in option price predicting and hedging for S&P 500 index options using the market data from June 1988 to December 1993. They conclude that a parsimonious model is better for predicting the future option price [11]. In addition, a deterministic volatility function method does not appear to be an improvement over the implied/constant volatility rate model when delta hedging for a 1-week hedging period with continuous rebalancing [11].

In this paper, we provide empirical evidence illustrating that a suitably implemented volatility function method can be advantageous in discrete delta hedging for a longer time hedging horizon, compared with an implied/constant volatility rate method. The purpose of delta hedging is to eliminate the first order sensitivity of an option price with respect to the underlying price. We first quantify the dependence of this hedging error on the accuracy of the delta hedge parameter in §2. Then we illustrate in §3 that

the ad hoc implied volatility rate method, as well as the constant volatility method, generates erroneous hedge parameters, assuming the underlying price does not follow a geometric Brownian motion. In particular, we demonstrate that delta computed from the implied/constant volatility method tend to be larger than the true hedge parameters. In §4, we describe a spline volatility function method [7] and demonstrate, using a simple example, that it can produce significantly more accurate hedge parameters and smaller hedging error. More importantly, we provide empirical evidence in §5 and §6, using the S&P 500 index and futures option market data, that a suitably implemented volatility function method performs significantly better in dynamic hedging against the underlying price movement, in comparison with the implied/constant volatility method. We observe that the delta hedge parameter computed from the volatility method is on average less than that of the implied/constant volatility method. Moreover, dynamic hedging error for a sufficiently long hedge horizon, approximately 17 days for the S&P 500 index options and 6 days for the S&P 500 futures options.

2. Delta and Hedging Error

In dynamic hedging, hedge parameters are computed at each rebalancing time to adjust the position of the underlying holding in a hedge portfolio. Thus it seems intuitive that hedging with an accurate hedge parameter yields smaller hedging error.

For example, delta hedging error is directly related to the accuracy of the delta hedge parameter. To see this mathematically, let us assume that the underlying price S_t is described by a general deterministic local volatility model,

$$\frac{dS_t}{S_t} = (\mu - q)dt + \sigma(S_t, t)dW_t, \tag{1}$$

where W_t is a standard Brownian motion, μ is the drift, q is the dividend yield, and $\sigma(S, t)$ is a deterministic local volatility function satisfying suitable regularity conditions so that (1) admits a unique solution. In this paper we denote $V_t \equiv V(S_t, t)$ as the exact option value when the underlying has price S_t at time t. Hedging error from using an approximate delta hedge factor Δ_t for an infinitesimal time dt can be measured by the value Π_{t+dt} at t+dt of the self financing portfolio $\{\Delta_t S_t, B_t, -V_t\}$ formed at time t where $B(t) = V_t - \Delta_t S_t$ is the riskless bond. The value of the portfolio Π_{t+dt} equals $d\Pi_t$ where

$$d\Pi_t = \Delta_t dS_t + q\Delta_t S_t dt + r(V_t - \Delta_t S_t) dt - dV_t.$$

Using Ito's Lemma,

$$dV_t = \sigma S_t \frac{\partial V}{\partial S} dW_t + \left((\mu - q) S_t \frac{\partial V}{\partial S} + \frac{1}{2} \sigma^2 S_t^2 \frac{\partial^2 V}{\partial S^2} + \frac{\partial V}{\partial t} \right) dt,$$

and the Black-Scholes equation

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} = rV - (r-q)S \frac{\partial V}{\partial S}$$

the instantaneous hedging error is

$$d\Pi_t = \sigma S_t (\Delta_t - \frac{\partial V}{\partial S}) dW_t + (\mu - r) S_t (\Delta_t - \frac{\partial V}{\partial S}) dt.$$
⁽²⁾

Thus, assuming that the underlying price satisfies a stochastic equation (1) with $\sigma > 0$ and $\mu \neq r$, the instantaneous hedging error is determined by the accuracy of delta: the mean and standard deviation of the relative instantaneous hedging error $\frac{d\Pi_t}{S_t}$ is zero if and only if Δ_t equals the exact delta $\frac{\partial V}{\partial S}$.

3. Delta from the Implied Volatility

It is well recognized that the price of an exotic option, e.g., a barrier option, can be very sensitive to the volatility model specification. Less emphasized, however, is the sensitivity of the hedging error to the model specification error. In a market exhibiting a volatility smile, the constant volatility model can yield a poor approximation to the price movement. In addition, the implied volatility method is an inherently inconsistent way of modeling the underlying price in this context. They both can generate poor hedge parameter approximation.

To examine this in detail, let us consider a European call option with a fixed time to maturity T. Assume that the underlying price is modeled by a deterministic volatility model (1). For a call option with value $C(S_0, K, T, r, q, \sigma)$, the implied volatility is the volatility rate σ_{imp}^c such that

$$C_{\rm BS}(S_0, K, T, r, q, \sigma_{\rm imp}^c) = C(S_0, K, T, r, q, \sigma).$$
(3)

Here $C_{BS}(S_0, K, T, r, q, \sigma_{imp}^c)$ denotes the corresponding Black-Scholes formula. Holding K, T, r, and q constant, the implied volatility σ_{imp}^c is a function of the initial price S_0 unless the local volatility function σ is a constant. In general, we have $\sigma_{imp}^c = \sigma_{imp}^c(S_0, K, T, r, q)$. Thus

$$\frac{\partial C}{\partial S_0} = \frac{\partial C_{\rm BS}}{\partial S_0} + \frac{\partial C_{\rm BS}}{\partial \sigma_{\rm imp}^c} \frac{\partial \sigma_{\rm imp}^c}{\partial S_0}.$$
(4)

When the local volatility σ is not a constant, the implied volatility rate $\sigma_{imp}^{c}(S_0, K, T, r, q)$ can vary with the initial price S_0 .

Let us illustrate this with a simple example in which the underlying price is assumed to follow an absolute diffusion process,

$$\frac{dS_t}{S_t} = \mu \ dt + \frac{C}{S_t} dW_t,\tag{5}$$

where μ is the drift, W_t is a standard Brownian motion, and C > 0 is a constant. Here the volatility is a monotonically decreasing function of the underlying price but does not depend on time.

Consider specifically an example when the initial underlying price $S_0 = 100$, risk free interest rate r = .06, the dividend rate q = 0, and C = 40. Figure 1 graphs how the implied volatility rate σ_{imp}^c varies with the strike and time to maturity in relation to the true local volatility function C/K (note that the local volatility function is considered as a function of strike and time to maturity $\sigma(K, T)$ in the context of the forward Black-Scholes equation).

This simple volatility model is of interest since the computed volatility functions from the S&P 500 index options and futures options bear a strong resemblance to this model. Analysis of the hedging results for this simple example can potentially shed light on the performance of the hedging methods in the S&P 500 index and futures options



Figure 1: Implied Volatility and the Exact Volatility C/K

markets. When the underlying follows the absolute diffusion process (5), a closed-form solution for the European option price exists [8]; this closed form solution can be used to compute the exact option prices as well as hedge parameters.

For the absolute diffusion example, we can clearly observe the dependence of the implied volatility σ_{imp}^c on S_0 ; Figure 2 graphs $\sigma_{imp}^c(S_0)$ for a call option with K = 100 and T = 1.



Figure 2: Implied Volatility $\sigma_{imp}^{c}(S_0)$ and the Exact Volatility C/S0

Since here $\sigma_{imp}^c(S_0)$ is monotonically decreasing with the initial price S_0 and the Black-Scholes vega $\frac{\partial C_{BS}}{\partial \sigma_{imp}^c}$ is positive, equation (4) suggests that the implied volatility method overestimates the delta hedge parameter. Figure 3 confirms this is indeed the case: delta from the implied/constant volatility is significantly larger than the true delta. In addition, error in delta increases as time to maturity increases. The delta hedge parameters from the implied volatility rate and constant volatility rate are close for this example (we have $K = S_0 = 100$ here).

Sensitivity $\frac{\partial \hat{\sigma}_{imp}^c}{\partial S_0}$ provides useful information on the error in the implied volatility delta. Indeed, if this sensitivity is known, it can be used to adjust the Black-Scholes delta $\frac{\partial C}{\partial S_0} = \frac{\partial C_{BS}}{\partial S_0}$ according to (4) so that it becomes more accurate. Unfortunately, sensitivity $\frac{\partial \hat{\sigma}_{imp}^c}{\partial S_0}$ is not observable in practice. In [10], Derman, Kani and Zou assume that the local volatility function is a linear function of the underlying price S only and



Figure 3: Delta Comparisons: Using the Implied/Constant Volatility Rate

heuristically derive three rules of thumb for the S&P 500 index options. We use the following arguments suggesting that the implied volatility σ_{imp}^c , at least in the S&P 500 index option market, is typically a monotonically decreasing function of the initial price S_0 .

When the underlying price satisfies the stochastic equation (1), the European put and call prices are related through the reversal of K and S, and q and r via the following equation

$$C(S, K, T, r, q, \sigma) = P(K, S, T, q, r, \sigma),$$
(6)

where $P(S, K, T, q, r, \sigma)$ denotes the put price with strike K and time to maturity T. Therefore we have

$$P_{\rm BS}(K, S_0, T, q, r, \sigma_{\rm imp}^c(S_0, K, T, r, q)) = C_{\rm BS}(S_0, K, T, r, q, \sigma_{\rm imp}^c(S_0, K, T, r, q)).$$

But

$$C_{\rm BS}(S_0, K, T, r, q, \sigma_{\rm imp}^c(S_0, K, T, r, q)) = C(S_0, K, T, r, q, \sigma).$$

Using (6) again, we obtain

$$C(S_0, K, T, r, q, \sigma) = P(K, S_0, T, q, r, \sigma).$$

Let $\sigma_{imp}^p(K, S_0, T, q, r)$ be the implied volatility of the put option $P(K, S_0, T, q, r, \sigma)$, i.e.,

$$P(K, S_0, T, q, r, \sigma) = P_{BS}(K, S_0, T, q, r, \sigma_{imp}^p(K, S_0, T, q, r))$$

Therefore

$$P_{\rm BS}(K, S_0, T, q, r, \sigma_{\rm imp}^c(S_0, K, T, r, q)) = P_{\rm BS}(K, S_0, T, q, r, \sigma_{\rm imp}^p(K, S_0, T, q, r)).$$

Assuming that the implied volatility is unique, we have

$$\sigma_{\rm imp}^c(S_0, K, T, r, q) = \sigma_{\rm imp}^p(K, S_0, T, q, r).$$

Therefore

$$\frac{\partial \sigma_{\rm imp}^c(S_0, K, T, r, q)}{\partial S_0} = \frac{\partial \sigma_{\rm imp}^p(K, S_0, T, q, r)}{\partial S_0}$$

Note that the derivative on the left of the above equation is with respect to the initial price whereas the derivative on the right is with respect to the strike price. For futures options, we can regard q = r in the pricing function and thus $\frac{\partial \sigma_{imp}^c(S_0, K, T, r, q)}{\partial S_0} = \frac{\partial \sigma_{imp}^p(K, S_0, T, r, q)}{\partial S_0}$. When $q \approx r$, we can use $\frac{\partial \sigma_{imp}^p(K, S_0, T, q, r)}{\partial S_0}$ as an approximation to $\frac{\partial \sigma_{imp}^c(S_0, K, T, r, q)}{\partial S_0}$. It has been noted in [17] that, after the 1987 crash, the implied volatility of the S&P 500 index options exhibits a sneer: the implied volatility rate decreases monotonically as the option strike increases, i.e., $\frac{\partial \sigma_{imp}^p(K, S_0, T, q, r)}{\partial S_0} < 0$, with the rate of decrease increasing for options with short maturities. Therefore it is likely that $\frac{\partial \sigma_{imp}^c(S_0, K, T, r, q)}{\partial S_0} < 0$ when $q \approx r$.

4. A Spline Volatility Function Method

In order to compute accurate hedge parameters, it is important to estimate the local volatility function sufficiently accurately. More precisely, the local volatility function in a region $\mathcal{B} \in \Re^+ \times [0,T]$ around $(S_0,0)$ is crucial; the option price is most sensitive to the volatility function $\sigma(S,t)$ in the region $(S,t) \in \mathcal{B}$. This is depicted, in Figure 4, by the contours of the sensitivity matrices $\left[\frac{\partial V}{\partial \sigma_{i,j}}\right]$ for two European calls with (K,T) = (90,0.5) and (K,T) = (110,1) at a constant volatility level of .15; here $\frac{\partial V}{\partial \sigma_{i,j}}$ is the option sensitivity with respect to the volatility at the grid point $(S_i, t_j) = ((i-1)\frac{2S_0}{200}, (j-1)\frac{1}{50})$.



Figure 4: The Significant Region \mathcal{B}

We now summarize the main idea behind the volatility function method [7]. In addition, we illustrate, using the previous example in which the underlying price follows an absolute diffusion, that the local volatility function method produces more accurate hedge parameters than the implied/constant volatility method.

Assuming that the underlying price satisfies the stochastic equation (1), a local volatility function can be determined uniquely if the prices of European options for all strikes and maturities are available. Unfortunately, only a limited finite set of options are traded in a derivative market. This makes the problem of the local volatility function determination an ill-posed problem. This ill-conditioning is typically dealt with by smoothness regularization.

To achieve smoothness explicitly, a local volatility function is represented in [7] by a

bicubic spline which is computed by solving an inverse constrained nonlinear optimization problem as follows. Let $\{\bar{V}_j\}_{j=1}^m$ denote the *m* given market option prices and the number of spline knots $p \leq m$. Given $\{(\bar{S}_i, \bar{t}_i)\}_{i=1}^p$ spline knots with corresponding local volatility values $\bar{\sigma}_i \stackrel{\text{def}}{=} \sigma(\bar{S}_i, \bar{t}_i)$, an interpolating cubic spline $c(S, t; \bar{\sigma})$ with a fixed end condition, e.g., the natural spline end condition, is uniquely defined by setting $c(\bar{S}_i, \bar{t}_i) = \bar{\sigma}_i, i =$ $1, \dots, p$. Let the value $V_j(c(S, t; \bar{\sigma}))$ denote the theoretical option price with the same maturity and strike as that of the given option \bar{V}_j when the local volatility $\sigma(S, t) =$ $c(S, t; \bar{\sigma})$. The local volatility values $\{\bar{\sigma}_i\}_{i=1}^p$ at spline knots (hence the spline) can be determined by calibrating the market option prices:

$$\min_{\bar{\sigma}\in\Re^p} f(\bar{\sigma}) \stackrel{\text{def}}{=} \frac{1}{2} \sum_{j=1}^m [V_j(c(S,t;\bar{\sigma})) - \bar{V}_j]^2$$

subject to $l \le \bar{\sigma} \le u.$ (7)

Additional smoothness regularization can be added in the objective function; we choose to explicitly control the smoothness of the spline with a minimal number of spline knots (which also leads to a better computational efficiency). Additional weights can be introduced to take account of different accuracies of \bar{V}_j . The lower and upper bounds, $l, u \in \Re^p, l < u$ can be used to incorporate information concerning σ . More detailed discussion on computational issues pertaining to solving (7) can be found in [7].

Let us consider again the afore mentioned example in which the underlying price follows an absolute diffusion (5) with the parameters $S_0 = 100$, r = 0.06, q = 0, and C =40. Using the closed form price formula for (5), a set of 22 European call option prices with maturities T = [0.5, 1] (either half year or one year) and strikes $K = [90, 92, \dots, 110]$ are generated. Figure 5 compares the computed delta and gamma from these option prices using the spline volatility function method [7] with those from the exact local volatility function and the implied/constant volatility rate. For the plots on the left, the option maturity is fixed at T = 1. For the constant volatility method, the volatility rate is set to $\frac{C}{S_0} = 0.4$.

The hedge parameters graphed by the dash curves in Figure 5 are computed using the volatility function method [7] as follows. Applying the finite difference method for solving the generalized Black-Scholes partial differential equation, the volatility function method [7] computes a volatility function represented with 8 spline knots below

$$[.4S_0, .8S_0, 1.2S_0, 1.6S_0] \times [0, 1].$$

Once the volatility function is computed, hedge parameters are calculated using the finite difference approximation. From Figure 5, it is clear that both the delta and gamma computed using the volatility function method [7] are significantly more accurate compared to those from the implied/constant volatility rate.

In practice, hedging is done discretely. Using the known volatility function in the absolute diffusion option example as a benchmark, we now compare different volatility methods in the context of the discrete dynamic hedging. We consider three different rebalancing frequencies n = 52 (weekly), n = 104 (biweekly), and n = 365 (daily). For the constant volatility method, the volatility rate is fixed at $C/S_0 = .4$ over the entire hedge period, where S_0 is the initial underlying price. For the implied volatility method, implied volatility rates are computed at every rebalancing time. For the volatility function method [7], a single volatility surface is computed using the 8 spline knot locations



Figure 5: Hedge Parameter Comparisons: Using the Volatility Function Method

at the beginning of a hedge horizon; this surface is then used for the entire hedge period since the local volatility function here is independent of time.

We perform dynamic hedge simulation similar to that described in [13] to illustrate hedge effectiveness. Assuming that the expected return μ equals to 8% in the diffusion equation (5), underlying price paths are simulated using the Euler approximation.

To compare hedge effectiveness, hedging error of an option V needs to be quantified. Let $\{t_i\}_{i=0}^n$, $t_{i+1} = t_i + \delta t$, denote the discrete rebalancing times in the hedge horizon, $[0, \tau]$, $0 < \tau \leq T$, where T is the time to maturity of the option. For delta hedging, we consider the *delta hedge portfolio* $\{\Delta S(t), B(t), -V(t)\}$ with the value,

$$\Pi_{\Delta}(t) = \Delta(t)S(t) + B(t) - V(t)$$

where $\Pi_{\Delta}(t)$, V(t), S(t), and B(t) denote the prices of the portfolio, option, underlying, and money market account respectively. At the beginning of the hedge horizon, $\Pi_{\Delta}(0) = 0$, $B(0) = V(0) - S(0)\Delta(0)$ where $\Delta(0)$ is an approximate delta hedge parameter of the option V at t = 0. At each rebalancing time t_i , the hedge parameter $\Delta(t_i)$ is recomputed and the money market account is adjusted:

$$B(t_i) = e^{r \delta t} B(t_{i-1}) + S(t_i) (\Delta(t_{i-1}) - \Delta(t_i)).$$

Thus the portfolio is self-financed. The delta hedging error is defined as $\Pi_{\Delta}(\tau)$, the value of the portfolio Π_{Δ} at the end of the hedge horizon τ .

Similarly, we can consider gamma hedging which takes into account the curvature in the relationship between the option price and the underlying price, and thus produces better hedging results. To perform this second order hedging, one additional traded option on the same underlying is needed to make a portfolio both gamma and delta neutral. To measure gamma hedging error, a gamma hedge portfolio $\{m_1(t)S(t), m_2(t)I(t), B(t), -V_t\}$ is formed with

$$\Pi_{\Gamma}(t) = m_1(t)S(t) + m_2(t)I(t) + B(t) - V(t),$$

where $\Pi_{\Gamma}(t)$ denotes the dollar value of the portfolio, I(t) is the price of an additional instrument, which depends on the same underlying S(t), with the maturity greater than τ . The numbers of holdings $m_1(t)$ and $m_2(t)$ are chosen such that the portfolio Π_{Γ} is both delta and gamma neutral. For example, if the delta and gamma factors for the option V, $\frac{\partial V}{\partial S}$ and $\frac{\partial^2 V}{\partial S^2}$, and the delta and gamma for the option I, $\frac{\partial I}{\partial S}$ and $\frac{\partial^2 I}{\partial S^2}$ are used, then $m_1(t)$ and $m_2(t)$ are chosen to satisfy

$$m_1(t) + m_2(t)\frac{\partial I}{\partial S} - \frac{\partial V}{\partial S} = 0$$

$$m_2(t)\frac{\partial^2 I}{\partial S^2} - \frac{\partial^2 V}{\partial S^2} = 0.$$

At the beginning of the hedge horizon, $\Pi_{\Gamma}(0) = 0$, $B(0) = (V(0) - m_1(0)S(0) - m_2(0)I(0))$. At each rebalancing time t_i , both delta and gamma hedge parameters are recomputed and the money market account is adjusted:

$$B(t_i) = e^{r\delta t}B(t_{i-1}) + S(t_i)(m_1(t_{i-1}) - m_1(t_i)) + I(t_i)(m_2(t_{i-1}) - m_2(t_i)).$$

The gamma hedging error is then defined as $\Pi_{\Gamma}(\tau)$, the value of the gamma hedge portfolio Π_{Γ} at the end of the hedge horizon τ .

Rebalance Frequency		n=52	n = 104	n = 365
Delta Hedging	Constant Volatility	.1126	.1000	.0907
	Implied Volatility	.1228	.1119	.1008
	Volatility Function	.0737	.0551	.0303
	Exact Volatility	.0738	.0547	.0300
Gamma Hedging	Constant Volatility	.0240	.0226	.0228
	Implied Volatility	.0253	.0236	.0235
	Volatility Function	.0075	.0039	.0026
	Exact Volatility	.0074	.0038	.0021

Table 1: Average Absolute Relative Hedging Errors for One Year Hedge Period

Table 1 displays the average absolute relative hedging errors over 200 price paths at the maturity for the absolute diffusion European call option example with the strike K = 100, maturity T = 1 and $\tau = T$ in the described dynamic hedge simulation with $\mu = 0.08$. The average absolute relative hedging error here is defined as the expected absolute hedging errors at the maturity over the 200 price simulation paths divided by the initial option price V(0) = \$18.58. i.e., $\frac{E(|\Pi_{\Delta}(\tau)|)}{V(0)}$ and $\frac{E(|\Pi_{\Gamma}(\tau)|)}{V(0)}$. For gamma hedging, the put option with the strike K = \$98 and maturity T = 1.1 is used as the additional instrument.

The following observations can be made from Table 1. First, the performance of delta hedging using the volatility function method [7] is almost identical to the performance achieved with the true volatility function; delta hedging with the implied/constant volatility rate is significantly inferior. Regarding gamma hedging, again both the true volatility surface and the reconstructed surface from the volatility function method [7] significantly outperform the use of the implied/constant volatility rate. In this case the

use of the true surface is marginally better than using the reconstructed surface. In addition, delta hedging error decreases with more frequent rebalancing using the exact volatility function and the volatility function method [7]; the hedging error using the implied/constant volatility method does not decrease as quickly when the hedge portfolios are rebalanced more often.

In Figure 6, the plot on the left graphs the average absolute hedging error, with daily rebalancing, as a function of the length of the hedging period; the average hedging error is displayed on the right. As noted previously, delta hedge parameter from the implied/constant volatility rate is greater than the true delta for this absolute diffusion example; thus the average absolute hedging error increases as the length of the hedging period increases. In addition, the average absolute hedging error from using the exact delta and the spline volatility function method. Moreover, the hedging performance difference between using the exact delta and the spline volatility rate of $\frac{C}{S_0}$ is slightly better than the implied volatility rate for this particular example.



Figure 6: Average Hedging Errors over 200 Paths

5. Dynamically Hedging S&P 500 Index Options

The absolute diffusion example in §2 demonstrates that both delta and gamma hedging errors using the volatility function method [7] are significantly smaller than those from using the implied/constant volatility rate; delta hedging error using the volatility function method is close to that of the true volatility function. However, the encouraging performance on this example does not immediately imply that hedging with the volatility function method is better in a real market since a deterministic local volatility model (1) may not describe the underlying price dynamics exactly. By calibrating the market option prices and following the market price movement, we now provide similar evidence of the advantages of using the volatility function approach in dynamic hedging for S&P 500 index options. In §6, we provide comparison for the S&P 500 futures options.

data description

First we consider the S&P 500 index options market. In this study, we use the data used by Aït-Sahalia and Lo for nonparametric estimation of the state price density using

historical option prices [1]; this data set is derived from the S&P 500 index option market prices from January 4, 1993 to December 31, 1993. The S&P 500 index options are European style with the S&P 500 index as the underlying. During 1993, the mean and standard deviation of continuously compounded daily return of the S&P 500 index were 7.95% and 10.28% respectively. The short term interest rate ranged from 2.85% to 3.21%.

To circumvent the unobservability of the dividend rate and the difficulty in synchronizing the index price and the option price, Aït-Sahalia and Lo convert the S&P 500 index options equivalently to options on the futures using the spot-futures parity [1],

$$F_{t,T} = S_t e^{(r_{t,T} - q_{t,T})T},$$

where S_t is the spot price, $r_{t,T}$, $q_{t,T}$ are the constant interest rate and dividend rate in the interval [t, t + T], and $F_{t,T}$ is the futures price with the time to maturity T in years. At each time t, the futures price is implied from the at-the-money put and call option prices using the put-call parity. In addition, in-the-money call option prices are replaced by the implied put prices of the out-of-the-money options via the put-call parity; thus only the call option prices are in the data set. We refer a reader to [1] for more details on this data set.

To enhance the reliability of our study, similar to Dumas, Fleming and Whaley [11], we exclude the extremely short-term, deep in- and out- of the money options since they have a small time premia and are not actively traded [11]. We consider only options with $\frac{8}{252} \leq T \leq \frac{99}{252}$, assuming 252 trading days per year. For a given time to maturity T, we consider the options with the magnitude of moneyness $M \stackrel{\text{def}}{=} K/F - 1$ less than .1 where K is the strike price, and F is the current implied futures price, i.e., the exercise price of an option cannot be larger or smaller than the underlying futures price by 10 percent. In addition, we exclude option prices that violate the no-arbitrage conditions

$$\frac{\partial V}{\partial K} < 0 \qquad \text{and} \qquad \frac{\partial^2 V}{\partial K^2} > 0,$$

where V is the market call option price and K is the exercise price (the derivatives are approximated using finite difference). A summary of the remaining data set is given in Table 2. On average, there are 36.2 options each day.

	$T < \frac{40}{252}$	$\frac{40}{252} \le T < \frac{70}{252}$	$T \ge \frac{70}{252}$	row sum
M < -2.5%	1,511	1,263	994	3,768
$-2.5\% \le M < 2.5\%$	1,221	1,030	852	$3,\!103$
$2.5\% \leq M$	558	900	784	2,242
column sum	3,290	$3,\!193$	2,630	9,113

Table 2: The S&P 500 Option Data

volatility models

Now we compare the local volatility function approach with the implied/constant volatility rate method in dynamically hedging the S&P 500 options in the described data set. For the local volatility function approach, we consider both the parametric quadratic models used by Dumas, Fleming and Whaley in [11], as well as the nonparametric spline function method [7]. The potential advantage of the spline volatility function method [7], compared to the parametric volatility function method, lies in its capacity to model a complex volatility function with a judicious choice of the spline knots. The parametric quadratic models considered here are:

QUAD 1:
$$\max(0.01, a_0 + a_1F + a_2F^2),$$

QUAD 2: $\max(0.01, a_0 + a_1F + a_2F^2 + a_3T + a_4FT),$
QUAD 3: $\max(0.01, a_0 + a_1F + a_2F^2 + a_3T + a_4FT + a_5T^2).$

These models are used in [11]. The spline function models considered include both a 1-dimensional spline function and a 2-dimensional spline function with knots placed at

SP 1:
$$[.6F, .8F, F, 1.2F, 1.4F],$$

SP 2: $[.6F, .8F, F, 1.2F, 1.4F] \times [0 T],$

respectively, where F is the implied futures price and T here is the longest maturity date on any given day.

At each rebalancing time, for the implied volatility method, an implied volatility rate is computed for each option and the delta hedge parameter is calculated using this volatility rate. For the constant volatility method, we compute, on each rebalancing time, a single volatility rate by fitting the Black-Scholes model to all call option prices in the least squares sense.

The volatility surfaces computed using both the parametric 2-dimensional quadratic volatility function and the 2-dimensional spline function are typically predominantly a function of the underlying price within the region \mathcal{B} in which the volatility significantly influences the option value; the computed volatility functions typically show less dependence on time in this region. Figure 7 graphs the 1-dimensional spline function (SP 1) and a quadratic polynomial function (QUAD 1) implied from the S&P 500 index option prices on August 9, 1993. We observe that these two functions are close to each other around the the interval [420, 476] surrounding the current implied futures price \$450.8; this is the region around which volatility function [7] is stable with suitably chosen spline knots. In addition, around the interval [420, 476], similar to the absolute diffusion model (5), the volatility function appears to be a decreasing function of the underlying price.

hedging performance

Once the volatility in a model is estimated, it is used to compute the option delta. This delta factor is then used for dynamic hedging.

To evaluate dynamic hedging performance of different volatility models for S&P 500 index options, we similarly quantify the hedging error as follows; the difference is that here the underlying is the futures price F rather than the spot price since we treat S&P 500 index options as European futures options.

Consider a specified rebalancing frequency n and a hedge horizon τ : assume that $t_0 > 0$ denotes the starting hedge time and $t_1 < \ldots < t_n$ with $t_1 > t_0$ denote the subsequent rebalance times and $t_n - t_0 = \tau$. At t_i , $i = 0, 1, \ldots, n$, let V_i be an option price, Δ_i a corresponding approximate option delta, F_i the futures price, and Π_i the



Figure 7: Volatility Functions: 1-dimensional Models

value of the hedge portfolio. Let r_i denote the short term interest rate for the period $[t_i, t_{i+1}]$ where i = 0, 1, ..., n - 1.

At the initial time t_0 , for each option, we set up a self finance hedge portfolio Π_0 with one unit short position in the call option, B_0 units of money market account, and Δ_0 units of the futures contract; actually, $B_0 = V_0$ and the initial value of the portfolio Π is 0 since the futures contract has zero value. At time t_1 , we adjust the units of the futures contract held in the hedge portfolio so that there is Δ_1 units of the futures contract. The value of B_1 is adjusted to $B_1 = e^{r_0(t_1-t_0)}B_0 + \Delta_0(F_0 - F_1)$; the portfolio is self financed. At time t_2 , $B_2 = e^{r_1(t_2-t_1)}B_1 + \Delta_1(F_1 - F_2)$. Following this procedure, delta hedging error at t_n equals the value Π_n of the hedge portfolio at t_n . For example, for a hedge period of one day, we set up the hedge portfolio at t_0 and compute the hedging error on the next business day. For a hedging period of twelve days with daily rebalance, we set up a hedge portfolio at t_0 and rebalance the portfolio every business day and the hedging error is computed 12 business days later.

For a specified rebalance frequency Δt and a hedge horizon τ , the average hedging error is the average of the hedging errors of all possible option prices $\{V_{t_0}, V_{t_1}, \dots, V_{t_n}\}$ available in the data set, with $t_{i+1} - t_i = \Delta t$ and $t_n - t_0 = \tau$. In evaluating hedge performance, we consider average hedging error, average absolute hedging error and root mean squared error. The root mean squared error is the square root of the average squared error.

We now compare hedge performance of the volatility function approach with the implied/constant volatility rate method. For the volatility function approach, unlike the absolute diffusion example, the volatility function is estimated each time the hedge portfolio is rebalanced to compute delta hedge parameter. For the implied volatility method, the implied volatility rate at each rebalance time is used to compute the delta hedge parameters.

First we note that, similar to the absolute diffusion example in §2, the delta hedge parameter computed from the volatility function approach is usually smaller than the delta from the implied volatility/constant volatility method. A typical hedge parameter comparison is depicted in Figure 8. In Table 3, the average of the implied delta subtracted by the spline local volatility delta, standard deviation of the difference, and the number of options in each category are displayed. The number without any parenthesis is the average difference. The number in a bracket $\langle \cdot \rangle$ is the standard deviation and the number in a bracket (\cdot) is the number of options in each category. This table illustrates that the implied delta is greater than the spline local volatility delta on average. Moreover, the difference is greatest for at-the-money options with the longest maturities.



Figure 8: A Typical Delta Comparison

	M < -0.03	$-0.03 \le M < 0.03$	$M \ge 0.03$	all M
	0.0320	0.0579	0.0151	0.0418
T < 0.15	< 0.0348 >	< 0.0294 >	< 0.0123 >	< 0.0342 >
	(1, 262)	(1, 361)	(360)	(2, 983)
$.15 \le T < 0.25$	0.0851	0.0943	0.0211	0.0735
	< 0.0280 >	< 0.0280 >	< 0.0167 >	< 0.0392 >
	(981)	(1,053)	(636)	(2, 670)
$T \ge 0.25$	0.1128	0.1122	0.0358	0.0926
	< 0.0256 >	< 0.0268 >	< 0.0238 >	< 0.0423 >
	(1, 201)	(1, 362)	(897)	(3, 460)
all T	0.0753	0.0876	0.0269	0.0704
	< 0.0458 >	< 0.0366 >	< 0.0216 >	< 0.0444 >
	(3, 444)	(3,776)	(1, 893)	(9, 113)

Table 3: Statistics of the Implied Delta Minus the Spline Volatility Function Delta

In addition, we note that there is no significant difference in delta hedging errors for the 2-dimensional model and 1-dimensional model for this S&P 500 index options data set; both the 2-dimensional spline and quadratic models give slightly smaller average hedging errors. In Figure 9 we graph the delta hedging errors of the volatility function approach (SP 2 and QUAD 2) and the implied/constant volatility method against the length of the hedge horizon with daily rebalancing.

The curves in Figure 9 demonstrate a clear trend. For a short-term hedge horizon, the implied/constant volatility model has slightly smaller average absolute hedging error than the volatility function approach. This is consistent with the 1-week hedging results reported by Dumas, Fleming and Whaley [11] (although continuous rebalancing is assumed in their study). However, for a hedge horizon longer than 17 days, dynamic hedging with the volatility function approach (both the spline and quadratic models)

generates increasingly smaller average absolute hedging error, compared with the implied/constant volatility method.



Figure 9: Average Absolute and Root Mean Squared Hedging Error: 1-day Rebalance

The superior hedge performance of the volatility function approach compared to the implied/constant volatility method for a longer hedge horizon is very interesting; hedging comparison in [11] has been based on a 1-week time horizon. We believe that average absolute hedging error of the local volatility function is smaller than that of the implied/constant volatility method for a longer hedge period due to the following reason. Assume that the dynamics of the S&P500 index price is sufficiently captured by a deterministic volatility model

$$dS_t = (\mu - q)S_t dt + \sigma(S, t)S_t dW_t.$$

Let us form a self financing hedge portfolio $\Pi_0 = \{\Delta_0 F_0, B_0, -V_0\}$. When an approximate hedge parameter Δ_t is used, it can be similarly shown, using the spot and futures relationship (assuming constant interest rate), that the instantaneous hedging error is $d\Pi_t$ and

$$d\Pi_t = \sigma F_t (\Delta_t - \frac{\partial V}{\partial F}) dW_t + (\mu - r) F_t (\Delta_t - \frac{\partial V}{\partial F}) dt.$$

For our S&P 500 index option data set, the expected return μ of the S&P 500 index is 7.95% and the risk free interest r is in the range 2.85% to 3.25%. Hence the *average* hedging error is determined by the difference $\Delta_t - \frac{\partial V}{\partial F}$; a larger delta factor produces more positive average hedge error. If a hedging method consistently yields larger positive $\Delta_t - \frac{\partial V}{\partial F}$, the hedging error for a *longer* hedge horizon will become increasingly more positive due to the *averaging* effect of errors for different rebalancing subperiods.

In addition, if $\Delta_t - \frac{\partial V}{\partial F}$ for options with different strikes and maturities, are consistently more positive, the average hedge error (not the average absolute hedging error) will be more positive; indeed we observe, from the left plot in Figure 10, that the average hedging error of the volatility function approach is smaller than that of the implied/constant volatility method even for one day hedge horizon.

The same pattern of delta hedging performance comparison emerges for hedging with 5-day rebalancing, as illustrated in Figure 11. The volatility function approach has increasingly smaller long term average absolute and root mean squared hedging errors, whereas the implied/constant volatility method generates slightly smaller average absolute hedging error for a short hedge horizon. The implied volatility method is again



Figure 10: Average Hedging Error: 1-day Rebalance (Left) and 5-day Rebalance (Right)

outperformed by the volatility function methods for a hedging horizon of 17 days. It is interesting that the hedge performance of the volatility function approach surpasses the implied/constant volatility method around approximately the same hedge horizon regardless of the rebalancing frequency. Moreover, the average hedge error of the volatility function method is always smaller than that of the implied/constant volatility method as in the case of daily rebalance, see the right subplot in Figure 10.



Figure 11: Average Absolute and Root Mean Squared Hedging Errors: 5-day Rebalance

6. Dynamically Hedging S&P 500 Futures Options

Is the better hedge performance of the volatility function approach reported in §5 a mere coincidence? To further investigate this, we perform similar empirical hedge performance evaluation for a different option market: the S&P 500 futures options. The S&P 500 futures and futures options are fairly liquid and are traded in Chicago Mercantile Exchange (CME). Unlike the S&P 500 index options, S&P 500 futures option is American style. In addition, it is unclear whether the local volatility function can be determined from the option prices of all strikes and maturities for the American option contract. Nonetheless, as will be illustrated subsequently, we observe strikingly similar hedge performance comparisons among the volatility function approach and the implied/constant volatility method.

data description

We consider the S&P 500 futures option prices from June 11, 1997 to June 10, 1998. During this period, there are four different S&P 500 futures contracts with different maturity dates; the futures contracts mature in September and December of 1997 and in March and June of 1998. Both futures and options prices are daily closing prices supplied by the Futures Industry Institute.

Unlike S&P 500 index options, S&P 500 futures options mature on the same day as the underlying S&P 500 futures contract, i.e., Thursday before the third Friday of of the delivery month; all options on the same S&P 500 futures contract have the same maturity date.

To ensure the prices are reliable and the results are meaningful, we apply similar data selection criteria as before. First, options with time to maturity shorter than 8 days and longer than 99 days are excluded. Second, we only select options with moneyness satisfying |K/F - 1| < .1, where K is the exercise price and F is the underlying futures price. There are 8,714 calls and 7,553 puts satisfying the above criteria during this period. On average, there are 34.2 calls and 29.6 puts everyday. Table 4 provides some additional information about the futures and futures options data such as the maturity date, the period covered, the number of trading days for each futures and futures options in the data set, and the number of calls and puts for each futures contract. In addition, we obtain interest rates from the price quotes of the Treasury bill in the Wall Street Journal; interpolation is typically required to extract the relevant interest rate since the maturity date of the S&P futures rarely matches that of the Treasury bill exactly.

Maturity date	Starting day	Ending day	Trading Days	Calls	Puts
Sep. 18, 1997	Jun. 11, 1997	Sep. 10, 1997	64	2,118	1,708
Dec. 18, 1997	Sep. 10, 1997	Dec. 10, 1997	65	$2,\!178$	2,065
Mar. 19, 1998	Dec. 10, 1997	Mar. 11, 1998	62	2,011	2,859
Jun. 18, 1998	Mar. 11, 1998	Jun. 10, 1998	64	$2,\!407$	1,921

Table 4: Futures and Futures Option Data

volatility estimation

We assume that, for each maturity date, the S&P 500 futures price follows an independent stochastic process. Since the hedge performance of 2-dimensional volatility models does not differ significantly from the 1-dimensional volatility models, we report here only the results for the 1-dimensional volatility models: the volatility function is assumed to be time independent.

For the volatility function approach, we consider both the nonparametric 1-dimensional spline model [7] as well as the parametric 1-dimensional quadratic model. For the spline volatility model, the local volatility function $\sigma(F)$ is represented by a 1-dimensional spline with five knots placed at [.7F, .9F, F, 1.1F, 1.3F], where F is the current underlying futures price on a give day. The quadratic model specifies the local volatility function as $\max(0.01, a_0 + a_1F + a_2F^2)$. The 1-dimensional spline and the quadratic polynomial are estimated by fitting the market option prices in the least squares sense.

For the implied volatility method, the volatility rate is computed for each option

by inverting the Black-Scholes option formula. For the constant volatility model, we allow different rates for put and call; hence the constant model here does not completely conform to the constant volatility Black-Scholes assumption since different volatility rates are allowed for call and put options. We compute call volatility rate by fitting, on each rebalancing time, the Black-Scholes model to all call option prices in the least squares sense. Similarly, put volatility rate is estimated by fitting to all put option prices.



Figure 12: Constant and Implied Volatility

Figure 12 shows the constant and implied volatility rates on the same day. Implied volatility rate is a decreasing function of the strike price except near the really out-of-themoney option. Constant volatility rate lies approximately in the middle of the implied volatility rates.

As an example, Figure 13 displays the 1-dimensional spline and quadratic polynomial volatility functions estimated from the market data on January 21, 1998. We observe that two local volatility functions are close to each other around the current futures price F = \$975.40, the region around which the volatility is most important in determining option prices. The closeness of the estimated spline volatility function and the quadratic volatility function around the current futures price is typical for this S&P 500 futures option data set.



Figure 13: Local Volatility Functions: 1-dimensional Models

hedging result

Figure 14 graphs the average absolute hedging error with daily rebalancing against the length of the hedging horizon. We observe that, similar to Figure 9 for S&P 500 index options, the volatility function (both the 1-dimensional spline and the 1-dimensional quadratic model) has smaller average absolute hedging error, compared with the implied/constant volatility method, for a hedging horizon of more than six days. The performance of the 1-dimensional spline and 1-dimensional quadratic volatility models are comparable. In addition we note that the volatility function approach begins to outperform the implied/constant volatility method at a shorter hedge horizon, compared to the S&P 500 index options experiment. Moreover, the hedging error for the S&P 500 futures options seems to be greater than that of the index option for the same length of the hedge horizon: note that the S&P500 futures price here is greater than the implied futures price of the S&P 500 index options in §5. The implied volatility method is better for up to six days. In terms of the root mean squared hedging error, the volatility method.



Figure 14: Average Absolute and Root Mean Squared Hedging Error: 1-day Rebalance

For hedging with 5-day rebalancing, similar observations can be made. Figure 15 shows that volatility function approach outperforms the implied/constant volatility method for a hedge period of more than ten days. The 1-dimensional spline and quadratic polynomial volatility functions show little difference in terms of the hedging error.



Figure 15: Average Absolute and Root Mean Squared Hedging Error: 5-day Rebalance

Figure 16 graphs the *average* hedging errors against the length of the hedge horizon. These plots show that the volatility function approach always produces smaller mean hedging error, even when the hedging period is one day. The constant volatility method yields the largest mean hedging error.



Figure 16: Mean Hedging Error: 1-day Rebalance (Left) and 5-day Rebalance (Right)

Despite the striking similarity of the hedging performance comparison in both the S&P 500 index option market and S&P 500 futures option market, there are some differences. The most noticeable difference is that, for the S&P 500 futures options, the average hedging error is no longer monotonically increasing as the hedging period increases and it can become negative as well. Currently we do not have a clear explanation for this except noting that it may be attributed to the early exercise feature of the American option contract.

7. Concluding Remarks

It has been well recognized that the index option markets typically exhibit a volatility smile; therefore, direct use of the constant volatility Black-Scholes option pricing formula is questionable. Although the implied volatility method with different volatility rates applied to options of different strikes and maturities is able to price the vanilla options accurately, it has been demonstrated [7] that the hedge parameters computed in this fashion can be erroneous since the constant volatility model does not adequately describe the underlying price dynamics.

In this paper, we compare the dynamic hedging performance of the deterministic local volatility function approach and the implied/constant volatility method. We illustrate that the mean of the instantaneous hedging error is proportional to $(\mu - r)(\Delta - \frac{\partial V}{\partial S})$, where Δ is the delta used in dynamic hedging and $\frac{\partial V}{\partial S}$ is the true delta. This implies that the mean hedging error using a larger delta factor greater than that of the true volatility results in more positive mean hedging error, assuming $\mu > r$. We illustrate that when the underlying price follows an absolute diffusion model with $\mu > r$, the hedge parameter computed from the implied/constant volatility method is indeed significantly larger than the true delta hedge parameter. This leads to increasingly positive average absolute hedging error as the length of the hedging horizon increases.

To estimate a local volatility function from a finite set of option prices, a spline volatility method has been proposed [7]. We illustrate here, with the absolute diffusion

example, that the spline volatility function [7] yields smaller average absolute hedging error compared with the implied/constant volatility method due to more accurate hedge parameters.

When comparing the hedge performance in the S&P 500 index option as well as the futures option markets, we similarly observe that the delta from the implied/constant volatility method is typically greater than that of the local deterministic volatility function approach. We observe, in both the S&P 500 index option and futures option markets, the mean hedging error using the volatility function approach is always smaller than that of the implied/constant volatility error. Moreover, the average absolute hedging error using the volatility function is smaller than that of the implied/constant volatility method for a sufficiently long hedge horizon, approximately 17 days for the S&P 500 index options and 6 days for the S&P 500 futures options.

References

- Y. Aït-Sahalia and A. Lo. Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance*, 53(2):499–545, 1998.
- [2] Leif Andersen and Rupert Brotherton-Ratcliffe. The equity option volatility smile: an implicit finite-difference approach. The Journal of Computational Finance, 1(2):5–32, 1998.
- [3] M. Avellaneda, C. Friedman, R. Holemes, and D. Samperi. Calibrating volatility surfaces via relative entropy minimization. *Applied Mathematical Finance*, 4:37–64, 1997.
- [4] G. Bakshi, C. Cao, and Z. Chen. Empirical performance of alternative option pricing models. Journal of Finance, 52(5):2003–2049, 1997.
- [5] D. Bates. Post-'87 crash fears in the S&P 500 futures option market. Technical Report Working Paper, University of Iowa, 1995.
- [6] A. Buraschi and J. Jackwerth. Is volatility risk priced in the option market? Technical Report Working Paper, London Business School and University of Madison, 1999.
- [7] Thomas F. Coleman, Yuying Li, and Arun Verma. Reconstructing the unknown local volatility function. *The Journal of Computational Finance*, 2(3):77–102, 1999.
- [8] J. C. Cox and S. A. Ross. The valuation of options for alternative stochastic processes. *Journal of Financial Economics*, 3:145–166, 1976.
- [9] E. Derman and I. Kani. Riding on a smile. Risk, 7:32–39, 1994.
- [10] E. Derman, I. Kani, and J. Zou. The local volatility surface: Unlocking the information in index option prices. *Financial Analysts Journal*, pages 25–36, 1996.
- [11] B. Dumas, J. Fleming, and R. Whaley. Implied volatility function: empirical tests. *Journal of Finance*, 53:2059–2106, 1998.
- [12] B. Dupire. Pricing with a smile. *Risk*, 7:18–20, 1994.
- [13] J. Hull. Options, Futures, and Other Derivatives. Prentice Hull, 1997.
- [14] N. Jackson, E. Süli, and S. Howison. Computation of deterministic volatility surfaces. The Journal of Computational Finance, 2(2):5–32, 1999.
- [15] Jens Jackwerth and Mark Rubinstein. Recovering probability distributions from option prices. The Journal of Finance, 51(5):1611–1631, 1996.
- [16] R. Lagnado and S. Osher. Reconciling differences. Risk, 10:79–83, 1997.
- [17] Mark Rubinstein. Implied binomial trees. The Journal of Finance, 49(3):771–818, 1994.