

Dynamic Hedging in a Volatile Market

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

May 27, 1999

1. Introduction

In financial markets, errors in option hedging can arise from two sources. First, the option value is a nonlinear function of the underlying; therefore, hedging is instantaneous and hedging with discrete rebalancing gives rise to error. Frequent rebalancing can be impractical due to transaction costs. Second, errors in specifying the model for the underlying price movement (model specification error) can lead to poor hedge performance. In this article, we compare the effectiveness of dynamic hedging using the constant volatility method, the implied volatility method, and the recent volatility function method [3]. We provide evidence that dynamic hedging using the volatility function method [3] produces smaller hedge error. We assume that there are no transaction costs, and both the risk-free interest rate r and the dividend rate q are constant.

Many studies have shown that the classical Black-Scholes constant volatility model does not adequately describe the stock price dynamics, see e.g., [13]. Implied volatility typically exhibits a dependence on both option strike and maturity, referred to as the volatility smile. Thus the *constant volatility* method, which assumes that the volatility is constant for all the options on the same underlying, can lead to a significant model specification error. To reduce this error in practice, the *implied volatility* method, which uses different constant volatilities for options with different maturities and strikes, is frequently used in pricing and hedging. Although the implied volatility method yields accurate valuation of vanilla options, it does not eliminate the model specification error.

A volatility function method attempts to accurately approximate the unknown local volatility function $\sigma(s, t)$ from the available market option prices in the context of a 1-factor continuous model. The computed volatility function can then be used for pricing options and determining hedge factors. This type of method follows the work of Dupire [8] who shows that the local volatility function can be uniquely determined if the prices of European options of all strikes and maturities are available. Various methods have been proposed [1, 2, 3, 5, 6, 10, 11, 12, 13] to compute a local volatility function calibrating a finite set of market option data. In the volatility function method [3], the volatility function is represented as a 2-dimensional spline; this method is distinguished from the

[†]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 Applied Mathematical Sciences Research Program (KC-04-02) of the Office of Energy Research of the U.S. Department of Energy under grant DE-FG02-90ER25013.A000 and NSF through grant DMS-9505155 and ONR through grant N00014-96-1-0050. The research was 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.

others in its emphasis on accurate approximation of the unknown volatility function, a key component in accurate hedging.

The main idea behind the volatility function method [3] is now briefly summarized. Let a general 1-factor diffusion process describe the underlying price movement,

$$\frac{dS_t}{S_t} = \mu(S_t, t)dt + \sigma(S_t, t)dB_t, \quad (1)$$

where B_t represents a standard Brownian motion, $\mu(s, t)$ is the drift, and $\sigma(s, t)$ is a deterministic local volatility function. The 1-factor model (1) is calibrated with a given set of option prices to approximate the local volatility function $\sigma(s, t)$ using a spline. In [3], the local volatility function is represented 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 $c(\bar{s}_i, \bar{t}_i) = \bar{\sigma}_i, i = 1, \dots, p$. The local volatility values $\bar{\sigma}_i$ at knots (hence the spline) are determined by calibrating the market option prices:

$$\begin{aligned} \min_{\bar{\sigma} \in \mathbb{R}^p} f(\bar{\sigma}) &\stackrel{\text{def}}{=} \frac{1}{2} \sum_{j=1}^m [v_j(c(s, t; \bar{\sigma})) - \bar{v}_j]^2 \\ \text{subject to} \quad &l \leq \bar{\sigma} \leq u. \end{aligned} \quad (2)$$

Additional weights can be added to take account of different accuracies of \bar{v}_j . The value $v_j(c(s, t; \bar{\sigma}))$ denotes the theoretical option price with the same maturity and strike as the given option \bar{v}_j when the local volatility $\sigma(s, t) = c(s, t; \bar{\sigma})$. The lower and upper bounds, $l, u \in \mathbb{R}^p, l < u$ can be used to incorporate additional information on σ . More detailed discussion on computational issues pertaining to solving (2) can be found in [3].

The synthetic European option example used in [3, 12] demonstrates that the hedge factors computed using the constant and implied volatility methods can be erroneous. In this example, it is assumed that the underlying follows a known absolute diffusion process,

$$\frac{dS_t}{S_t} = \mu dt + \frac{C}{S_t} dB_t, \quad (3)$$

where μ is the drift, B_t is a standard Brownian motion, and $C > 0$ is a constant. Here the volatility is a function of the underlying but does not depend on time. When the underlying follows the absolute diffusion process (3), a closed-form solution for the European options exists [4]; this closed form solution can be used to compute option prices as well as hedge factors. Let the initial underlying value be $S_0 = 100$. Assume the risk free interest rate $r = .06$, the dividend rate $q = 0$, and $C = 40$. Figure 1 displays the computed hedge factors for the European call options with the strike $K = 100$ using the exact volatility, the constant volatility method, the implied volatility method, and the volatility function method [3]. For the top subplots, the option maturity is $T = 1$. For the constant volatility method, the volatility constant equals $\frac{40}{100}$. The hedge factors graphed by the dash-dot curves in Figure 1 are computed as follows. Using the closed form formula for (3), 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. Applying the finite




Figure 1: The Delta and Gamma Hedge Factors of European Call Options

difference method for solving the generalized Black-Scholes partial differential equation, the volatility function method [3] described above computes a volatility function using 22 spline knots; more details can be found in [3]. It is clear that both the delta and gamma factors computed using the volatility function method [3] are significantly more accurate compared to the constant and implied volatility methods.

The focus of this article is dynamic hedging in which a hedge portfolio is rebalanced discretely. Delta and gamma hedge factors are computed at each rebalancing time to adjust positions of the instruments in the hedge portfolio. Gamma hedging takes into account the curvature in the relationship between the option price and the underlying price, and thus produces better hedge 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.

Subsequently we compare the effectiveness of the volatility function method [3] with the popular constant volatility and implied volatility methods in dynamic hedging. We first demonstrate, with a synthetic example, that delta hedging with the volatility function method [3] produces significantly smaller hedge errors than both the constant and implied volatility methods; the delta hedge errors also decrease faster as the rebalancing frequency increases. More importantly, we provide evidence, using the S&P 500 futures option market data, that the volatility method [3] performs significantly better in dynamic hedging against the market futures price movement, in comparison to the implied and constant volatility methods.

2. Hedging the Synthetic European Option

We consider the synthetic European option example described in §1 since the known volatility function can be used as the benchmark for comparison. Three different rebalancing frequencies n are considered, $n = 52$ (weekly), $n = 104$ (biweekly), and $n = 365$ (daily). For the constant volatility method, the volatility parameter is fixed over the entire hedge period; the constant volatility used is C/S_0 , where S_0 is the initial underlying price. For the implied volatility method, an implied volatility is obtained at every rebalancing time. For the volatility function method [3], a single volatility surface is computed at the beginning of a hedge horizon in order to avoid the computing time required to reconstruct a volatility surface at each rebalancing time for each simulation; this surface is then used for the entire hedge period.

We perform dynamic hedge simulation similar to that described in [9] to illustrate the hedge effectiveness. From the diffusion equation (3), paths of the underlying price movement are simulated using the Euler approximation.

To compare hedge effectiveness, hedge error of an option 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 maturity of the option. For delta hedging, we consider the *delta hedge portfolio* $\{V(t), S(t), B(t)\}$ with,

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

where $\Pi_{\Delta}(t)$, $V(t)$, $S(t)$, and $B(t)$ denote the dollar values of the portfolio, option, underlying, and money market account respectively. At the beginning of the hedge horizon, $\Pi_{\Delta}(0) = 0$, $B(0) = -(V(0) + S(0)\Delta_V(0))$ where $\Delta_V(0) = [\frac{\partial V}{\partial S}]_{t=0}$ is the delta hedge factor of the option V at $t = 0$. At each rebalancing time t_i , the hedge factor $\Delta_V(t_i)$ is recomputed and the money market account is adjusted:

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

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

For gamma hedging, the *gamma hedge portfolio* $\{V(t), I(t), S(t), B(t)\}$ is formed with

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

where $\Pi_{\Gamma}(t)$ denotes the dollar value of the portfolio, $I(t)$ is an additional instrument, which depends on the same underlying $S(t)$, with the maturity greater than T . 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, i.e.,

$$\begin{aligned} \Delta_V(t) + m_1(t) + m_2(t)\Delta_I(t) &= 0 \\ \Gamma_V(t) + m_2(t)\Gamma_I(t) &= 0. \end{aligned}$$

The values $\Delta_V(t), \Gamma_V(t)$ are the delta and gamma factors for the option $V(t)$; the values $\Delta_I(t), \Gamma_I(t)$ are delta and gamma for the option $I(t)$. At the beginning of the hedge horizon, $\Pi_{\Gamma}(0) = 0$, $B(0) = -(V(0) + S(0)m_1(0) + I(0)m_2(0))$. At each rebalancing time t_i , both delta and gamma hedge factors are recomputed and the money market account is adjusted:

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

Rebalancing Frequency		$n=52$	$n=104$	$n=365$
Delta Hedging	Constant Volatility	.1108	.1005	.0890
	Implied Volatility	.1188	.1112	.0969
	Volatility Function	.0803	.0610	.0392
	Exact Volatility	.0771	.0580	.0292
Gamma Hedging	Constant Volatility	.0251	.0233	.0224
	Implied Volatility	.0256	.0244	.0235
	Volatility Function	.0138	.0134	.0121
	Exact Volatility	.0074	.0041	.0015

Table 1: Average Relative Hedge Errors Over 200 Sample Paths: A Synthetic Example

The gamma hedge error is defined as $|\Pi_\Gamma(\tau)|$, the absolute value of the gamma hedge portfolio Π_Γ at the end of the hedge horizon of the option.

Table 1 displays the average relative hedge errors at the maturity for the synthetic European call option with the strike $K = 100$, maturity $T = 1$ and $\tau = T$ in the described dynamic hedge simulation. The average relative hedge error is defined as the average of the hedge errors at the maturity over the 200 price simulation paths divided by the initial option price $V(0) = \$18.58$. For gamma hedging, the put option with the strike $X = 98$ and maturity $T = 1.1$ is used as the additional instrument. To illustrate the change of the hedge portfolio values in the course of the hedge period, the relative values of the hedge portfolios are graphed in Figure 2 for a sample path in one year with the rebalancing frequency $n = 104$.




Figure 2: The Relative Values of the Hedge Portfolios Along a Sample Path

The following observations can be made from Table 1 and Figure 2. First, the performance of delta hedging using the volatility function [3] is almost identical to the performance achieved with the true volatility function. Delta hedging with constant volatility or implied volatility is significantly inferior. Regarding gamma hedging, again both the true volatility surface and the reconstructed surface [3] significantly outperform the use of constant or implied volatility. In this case the use of the true surface is marginally better than using the reconstructed surface – probably due to the fact that the reconstructed volatility surface is computed just once during the entire hedge horizon. Finally, delta hedge error decreases with more frequent rebalancing using the exact volatility function and the volatility function method [3]; the hedge errors using the

constant and implied volatility methods do not decrease as quickly when the hedge portfolios are rebalanced more often.

3. Hedging the S&P 500 Futures Options

The synthetic example in §2 demonstrates that both delta and gamma hedge errors using the volatility function method [3] are significantly smaller than those from using the implied and constant volatility methods; the delta hedge error using the volatility function method [3] is close to that from using the true volatility function. However, this encouraging performance on synthetic data does not immediately imply that hedging with the volatility function method [3] is better in a real market. Calibrating from the market option prices and following the market price movement, we now provide evidence illustrating the advantages of using the volatility function method [3] in dynamic hedging.

We consider dynamic hedging for the S&P 500 futures options traded in Chicago Mercantile Exchange. Here the market futures price movement is used as the path against which hedge performance is measured. Although these options are American, the spline inverse optimization formulation (2) remains a reasonable way to estimate the local volatility function from a given set of option prices; the American option values are computed using a partial differential equation approach as described in [9].

We consider the market futures option price data from May 1997 to March 1998. There are three index futures in this data set: the first index future matures on September 18, 1997, the second on December 18 1997, and the third on March 19 1998. The futures and options mature on the same day. Therefore, we correspondingly separate the option prices into three data sets. We choose, on each Wednesday, 12 calls and 12 puts whose strikes are nearest to the futures price; we only consider at-the-money and near-the-money options since their prices are more accurate than deep in-the-money or out-of-the-money options. Thus the first data set contains call and put options on the S&P 500 September 97 index futures from May 21 to September 10 in 1997. The second data set contains option prices on the December 97 futures from July 30 to November 19, 1997. The third data set includes options on the March 98 futures from January 7 to March 11, 1998. The third data set covers relatively shorter period than the first two since we do not have the option prices on the March 98 futures near the end of 1997. The hedge portfolios are rebalanced weekly; the volatility function, implied volatility, constant volatility parameter, and hedge factors are recomputed weekly. Figure 3 displays the futures prices and option strikes in the first data set for the September 18, 1997 futures; the solid line depicts the futures prices and the circles/squares display the strikes of the call/put options.

For the constant volatility method, we choose the constant which best fits for the 12 call option prices in the least squares sense. The volatility parameter for the put options is defined similarly. In the implied volatility method, each option has a different implied volatility parameter. For the volatility function method [3], a volatility surface is computed, at each rebalancing time, by solving (2) using the 24 call and put options with 9 spline knots placed on the mesh,

$$[.6S_0, S_0, 1.4S_0]^T \times [0, .25T, .75T],$$

where S_0 denotes the initial futures price and T is the maturity. We perform one week delta and gamma dynamic hedging on each option in a data set similar to §2. Delta and




Figure 3: Futures and the Option Strikes for the September 18, 1997 Data Set

gamma hedge factors of options are computed using each of the three methods; options are hedged for a 1-week period using these hedge factors.

Table 2–4 display the average weekly hedge errors and the standard deviations using the constant volatility, implied volatility, and the volatility function method [3]. The average weekly hedge error is the sum of the weekly hedge errors of all the options in the data set divided by the number of options. The numbers in the parenthesis are the standard deviations of the hedge errors.

		Call	Put
Delta Hedging	Constant Volatility	1.6474 (1.3170)	1.6310 (1.2966)
	Implied Volatility	1.6348 (1.3126)	1.6188 (1.2885)
	Volatility Function	1.4339 (1.2052)	1.4154 (1.1837)
Gamma Hedging	Constant Volatility	0.0400 (0.0453)	0.0468 (0.0704)
	Implied Volatility	0.0361 (0.0326)	0.0364 (0.0349)
	Volatility Function	0.0254 (0.0264)	0.0276 (0.0386)

Table 2: Hedge Error : Options On The Sep 97 Futures

		Call	Put
Delta Hedging	Constant Volatility	2.7216 (2.0647)	2.6844 (2.0912)
	Implied Volatility	2.6801 (1.9935)	2.6085 (1.9712)
	Volatility Function	2.0069 (1.3201)	1.9899 (1.3196)
Gamma Hedging	Constant Volatility	0.0862 (0.0834)	0.1206 (0.2202)
	Implied Volatility	0.0703 (0.0791)	0.1003 (0.1920)
	Volatility Function	0.0458 (0.0469)	0.0673 (0.0999)

Table 3: Hedge Error : Options On The Dec 97 Futures

From Tables 2–4, we observe that the volatility function method [3] is superior to both the implied and constant volatility methods with up to 33% less delta hedge

		Call	Put
Delta Hedging	Constant Volatility	1.7175 (0.8912)	1.6661 (0.8642)
	Implied Volatility	1.6687 (0.8480)	1.6204 (0.8238)
	Volatility Function	1.6167 (0.8187)	1.5664 (0.8161)
Gamma Hedging	Constant Volatility	0.0501 (0.0613)	0.0522 (0.0657)
	Implied Volatility	0.0486 (0.0511)	0.0485 (0.0499)
	Volatility Function	0.0456 (0.0657)	0.0392 (0.0309)

Table 4: Hedge Error : Options On The Mar 98 Futures

error; gamma hedging also leads to smaller errors. The most significant performance differences between the volatility function method [3] and the constant and implied volatility methods occur in the December 97 futures data set. We note that the hedge performance of the volatility function method [3] depends on the choice for the number of knots and their placement. These decisions should in general be made by some cross validation method.

The volatility smile exhibited in an option index market indicates that the price distribution is not lognormal; indeed the implied distribution from the volatility function method [3] is typically not lognormal. Figure 4 illustrates the implied risk neutral distribution of the December 97 futures as seen on August 6, 1997; the risk neutral distribution of the constant volatility model is graphed for comparison. Each risk neutral distribution is computed using the Fokker-Planck equation [7].




Figure 4: Comparison of the Implied Risk-Neutral Distributions

4. Concluding Remarks

It has been well recognized that the global index options markets typically exhibit a volatility smile; therefore, the direct use of the constant volatility Black-Scholes option pricing formula is questionable. Although the implied volatility approach with different constant volatility applied to options of different strikes and maturities is able to price the vanilla options adequately, it has been demonstrated [3] that the hedge factors computed in this fashion can be erroneous: the constant volatility model does not sufficiently

describe the underlying price dynamics. Assuming that the price of the underlying follows a 1-factor continuous diffusion process, it is important to accurately reconstruct the local volatility function for option hedging as well as pricing. In this paper, we compare the performance of dynamic hedging using the constant volatility method, the implied volatility method, and the volatility function method [3]. With a synthetic European option example, we demonstrate that the volatility function method [3] yields significantly more accurate hedge factors and smaller hedge errors. Using the S&P 500 futures option market data and hedging against the market futures price movement, the volatility function method [3] is shown to perform significantly better in dynamic hedging when compared with the constant and implied volatility methods.

References

- [1] 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.
- [2] M. Avellaneda, C. Friedman, R. Holes, and D. Samperi. Calibrating volatility surfaces via relative entropy minimization. *Applied Mathematical Finance*, 4:37–64, 1997.
- [3] Thomas F. Coleman, Yuying Li, and Arun Verma. Reconstructing the unknown local volatility function. *The Journal of Computational Finance*, 2(3):77–102, 1999.
- [4] J. C. Cox and S. A. Ross. The valuation of options for alternative stochastic processes. *Journal of Financial Economics*, 3:145–166, 1976.
- [5] E Derman and I Kani. Riding on a smile. *Risk*, 7:32–39, 1994.
- [6] E Derman and I Kani. The local volatility surface: Unlocking the information in index option prices. *Financial Analysts Journal*, pages 25–36, 1996.
- [7] Darrell Duffie. *Dynamic Asset Pricing Theory*. Princeton, 1996.
- [8] B. Dupire. Pricing with a smile. *Risk*, 7:18–20, 1994.
- [9] J. Hull. *Options, Futures, and Other Derivatives*. Prentice Hall, 1997.
- [10] N. Jackson, E. Süli, and S. Howison. Computation of deterministic volatility surfaces. *The Journal of Computational Finance*, 2(2):5–32, 1999.
- [11] Jens Jackwerth and Mark Rubinstein. Recovering probability distributions from option prices. *The Journal of Finance*, 51(5):1611–1631, 1996.
- [12] R. Lagnado and S. Osher. Reconciling differences. *Risk*, 10:79–83, 1997.
- [13] Mark Rubinstein. Implied binomial trees. *The Journal of Finance*, 49(3):771–818, 1994.