Gamma Hedging in Incomplete Markets Under Transactions Costs

Les Clewlow
Stewart D Hodges

Abstract

This paper explores the role of gamma hedging strategies in situations where there are either jumps or stochastic volatility, and it is costly to make options transactions in order to hedge. The motivation for this work stems from earlier empirical work we have completed on hedging options portfolios. A previous empirical analysis (Hodges, Clewlow *et al* (1993)) of hedging CME options on the S&P futures confirmed that gamma hedging provided a significant risk reduction compared to delta hedging, and did so even when rebalancing was done daily. The reason why the improvement is so marked appears to be because real asset processes involve jumps and stochastic volatility. However, the study also showed that this kind of standard gamma hedging involved high turnover of options positions - which with transactions costs would make the hedge prohibitively costly. We therefore have the following problem: What is the best way to manage gamma when options are costly to transact and we have jumps and/or stochastic volatility?

We have some insights into this problem from the literature on optimal delta hedging under transactions costs. The gamma hedging problem, though, is very much harder since, unlike delta, there are many different candidate instruments that could be used to modify gamma. We would have to have many state variables to be able to solve the problem using a dynamic programming formulation (as we did, for example, in Hodges and Neuberger (1989) and Hodges and Clewlow (1993)), though it could of course be written in that framework in a formal sense. Instead, we have therefore used a simulation approach to search for optimal methods of hedging based on heuristics which are consistent with myopic policies.

1. Introduction

This paper explores the role of gamma hedging strategies in situations where there are either jumps or stochastic volatility, and it is costly to make options transactions in order to hedge. The motivation for this work stems from earlier empirical work we have completed on hedging options portfolios. A previous empirical analysis (Hodges, Clewlow *et al* (1993)) of hedging CME options on the S&P futures confirmed that gamma hedging provided a significant risk reduction compared to delta hedging, and did so even when rebalancing was done daily. The reason why the improvement is so marked appears to be because real asset processes involve jumps and stochastic volatility. However, the study also showed that this kind of standard gamma hedging involved high turnover of options positions,- which with transactions costs would make the hedge prohibitively costly. We therefore have the following problem: What is the best way to manage gamma when options are costly to transact and we have jumps and/or stochastic volatility?

We have some insights into this problem from the literature on optimal delta hedging under transactions costs. The gamma hedging problem, though, is very much harder since, unlike delta, there are many different candidate instruments that could be used to modify gamma. We would have to have many state variables to be able to solve the problem using a dynamic programming formulation (as we did, for example, in Hodges and Neuberger (1989) and Hodges and Clewlow (1993)), though it could of course be written in that framework in a formal sense. Instead, we have therefore used a simulation approach to search for optimal methods of hedging based on heuristics which are consistent with myopic policies. The only other paper we are aware of which takes a similar approach is that of Kat (1993).

The structure of the paper is as follows. In the following section we review previous empirical work on hedging and also the literature on optimal delta hedging with transaction costs. Section 3 discusses some of the issues related to processes with either jumps or stochastic volatility. We first give some limited evidence from frequency distributions concerning the role of these effects. Next we discuss the relationship between jumps or stochastic volatility and volatility smile effects, and finally their relationship to delta hedging errors, in a situation without transactions costs. Section 4 describes the issue of moneyness as it applies to gamma hedges, and the notion of speed which provides a local measure of the robustness of a gamma hedge. We then describe our simulation analysis of alternative hedging approaches under a variety of assumptions. The final section draws some conclusions and sketches some ideas for further work in this area.

2. Related Earlier Work

In this section we review previous empirical work on hedging and also the literature on optimal delta hedging with transaction costs. Although there are quite a few theoretical studies of hedging errors (for example, Boyle and Emmanuel (1980)), and there are others which employ Monte Carlo simulation to study hedging effectiveness (eg. Figlewski (1989)), there are few published empirical studies of hedging effectiveness. The earliest study is the pre-CBOE work of Black and Scholes (1992). Other studies of interest are the Hull and White (1987) paper on hedging FX options, and the work of Dolbear (1992), which compares alternative methods of hedging lookback options.

Hodges, Clewlow *et al* (1993) describe the results of using delta, delta-gamma, and delta-vega hedges to hedge the CME S&P 500 index option contract. They hedged ten non-overlapping contracts between 1985 and 1992, and report on both the terminal replication errors and on the daily replication errors from marking to market. The main results are reproduced in Table 1. It is evident that delta hedging brings about a dramatic reduction in the risk exposure, and that delta-gamma hedging roughly halves again the standard deviation of the replication error. The vega hedge is particularly naive, as it is mismatched on expiry dates but yet pays no attention to the way in which long term volatilities move less than short term ones. Table 2 reproduces the levels of turnover involved from the various hedges. Note that daily rebalancing of the gamma or vega hedges implies prohibitively high turnover, both for the option contract and also for the delta adjustment made through the futures contract.

We are therefore led to consider work on optimal hedging under transactions costs and its implications for the current problem. The work in this area has so far only considered delta hedging. Leland (1985) provided important early insights into this problem. Hodges and Neuberger (1989) (building on earlier work by Davis and Norman (1990)), formulated the optimal delta hedging problem under transactions costs as a problem in stochastic optimal control. This has a natural dynamic programming structure (ie. similar to the binomial method) where we work backwards from the expiry date. By choosing an exponential utility function they reduced the problem to one with just two state variables, asset price and portfolio delta, and they were able to provide numerical solutions. Other related work on this problem includes papers by Bensaid *et al* (1992), Davis *et al* (1992, 1993), Edirisinghe *et al* (1991), Hodges and Clewlow (1993), and by Wilmott (1993).

The nature of the optimal solutions to this problem is that with proportional transactions costs (i.e. no fixed component) the optimal hedging strategy is to transact only as much as is necessary to stay within a region which can be computed. This region takes the form of a band around the preferred delta, which itself differs from the Black-Scholes delta (and indeed the Black-Scholes delta can even lie outside the optimal region). When we think of the gamma hedging problem, the difficulty of a computable solution is immediately apparent. Although the problem could be formally stated in a dynamic programming framework, at best we will end up with state variables for the holdings in every single available option, and this cannot be computationally tractable.

3. Jumps and Stochastic Volatility

In this section we discuss some of the issues related to processes with either jumps or stochastic volatility. We first give some limited evidence from frequency distributions concerning the role of these effects. A useful published paper of this topic is Turner and Weigel (1992). Using the CME data referred to earlier, we have computed our own summary statistics for holding period returns on the futures and changes in implied volatility of the options contracts, over 1 day, 5 day and 20 day periods. We find that both distributions exhibit a significant degree of fat-tailedness, whether measured by their kurtosis or by the ratio of the range to the standard deviation. The daily equity returns show a range of 14.3 standard deviations and an excess kurtosis of 8.6 in the period after the 1987 crash. This fat-tailedness declines, but more slowly than it would (under the central limit theorem) if successive returns were entirely independent. Twenty day returns exhibit a range of 8.6 standard deviations and excess kurtosis of 1.5. Evidence from GARCH models suggests that even after normalising by the predictable component of volatility, the distribution is far from normal. This suggests that neither a stationary jump process for returns, nor a stochastic volatility model based only on diffusions will be adequate to capture the realities of the market place. However, we will proceed by considering in turn the implications of each of these more general processes separately, before putting them together. Similar conclusions apply to the processes for other kinds of financial assets, eg - bonds and FX.

Option valuation models already exist for both jump-diffusion processes (eg. Merton (1976)) and for stochastic volatility (eg. Hull and White (1988) and see also Clewlow and Xu (1992, 1993)). We shall now contrast the differences between the kinds of implied volatility smiles to be expected under these two processes, and also the difficulties they pose for delta hedging in the absence of transactions costs. If we plot implied volatility against the log of the strike price we always expect to find that the shorter maturity options produce greater curvature. However, this is a misleading metric, for the range of likely outcomes for the underlying is much smaller for the shorter maturity instrument. To obtain a metric which reflects the underlying shape of the distribution (i.e. its skewness and kurtosis) we should normalise the x-axis by dividing by the square root of the time to expiry. Under a jump process with constant parameters we have initially high excess kurtosis, but this reduces rapidly for more distant horizons under the central limit theorem. The volatility smile under our preferred metric is curved for short expiries but flattens out rapidly for longer ones. Under the stochastic volatility model the opposite is the case. The instantaneous distribution is always conditionally normal and the degree of fat-tailedness depends on the coefficient of variation of the variance. In the short run the variance has not had time to move far and we see very little smile effect. It grows more marked for longer term options but may reduce again further out if we have mean-reversion in the volatility and through the effect of the central limit theorem.

Finally in this section we consider the nature of the expected replication error variance for delta hedging strategies under jumps and stochastic volatility but without transactions costs. Discrete-time delta hedging under Black-Scholes assumptions was first dealt with by Boyle and Emmanuel (1980). We can also provide similar first approximations to the rate at which hedging variance accumulates under both jump processes and diffusion processes exhibiting stochastic volatility. Formulae for the three processes are given in the table below:

Hedging Variance per unit time $\cong$

$\tfrac{1}{2}\Gamma\sigma_S^4 dt$	Black-Scholes
$\tfrac{1}{2}\Gamma\sigma_S^4\left(dt + K/2\right)$	Poisson with excess kurtosis K at one year
$Vega^2 \times Var\left(\Delta\sigma_{imp}\right) \cong \Gamma^2\sigma_S^4 CV(vol)\tau^2$ where $CV(vol) = SD\left(\sigma_{imp}\right) / \sigma_{imp}$	Stochastic volatility process (τ is time to expiry of option)

Note that while the rate of hedging variance in a Black-Scholes world is proportional to the time between revisions, for the jump process there is an additional component which depends on the kurtosis of the jump distribution and which cannot be reduced purely by delta hedging. Thus while delta hedging in a Black-Scholes world takes the standard deviation of the replication error to zero as the square root of the rebalancing interval, under a jump process it goes asymptotically to some fixed amount. Replication error under stochastic volatility behaves in a related way, and cannot be controlled by simply delta hedging. The entry in the table just shows a crude approximation to the rate of replication variance which stems from vega risk alone. Here, since vega is (at least in special cases) proportional to gamma times the time to expiry, the rate at which hedging variance accrues depends on time to expiry as well as gamma. Note that it is the coefficient of variation of the volatility which matters, and we can again interpret this as being related to the fat-tailedness of the asset distribution at the expiry date. The results confirm the general importance of jump risk. They also confirm that the risk due to stochastic volatility is much more important for long dated options than for short dated ones.

4. Some Issues Related to Gamma Hedges

In this section we describe the issue of moneyness as it applies to gamma hedges, and the notion of speed which provides a local measure of the robustness of a gamma hedge.

A delta hedge is a delta hedge, though we may look to gamma to warn us of how rapidly our delta neutral position can become unstuck. A gamma hedge is a little more complex. To adjust gamma we have to buy or sell some kind of option, and depending on what strike price it has we will obtain a different gamma profile. The gamma of an option peaks at spot prices close to the present value of the strike price. It is therefore usually advisable to hedge using an at the money option, for if we neutralise gamma by buying or selling options a long way from the money, we are likely to be introducing undesirable peaks of gamma at other levels of the asset price.

A lesser known "greek" called speed (see Garman (1992)) measures the change in gamma with respect to the underlying, and in constructing gamma hedges (and choosing the strike price as well as the option amounts) we will want to try to ensure that we manage speed as well as gamma and delta.

5. Structure of our Simulations

The simulations run delta and delta-gamma hedging portfolios under three different assumptions for the nature of "the world" and under a variety of methods of implementing the hedges.

The three "worlds" are:

- a Black-Scholes world, where the asset follows Geometric Brownian Motion,
- a Stochastic Volatility world, based on the Hull and White (1988) model,
- a Jump-Diffusion world, based on the Merton (1976) model, and where the jump is to a conditionally normal distribution with fixed standard deviation.

In each of these worlds we have a fairly low level of costs for transacting in the underlying (or future) at 0.5% but a much higher transaction cost for buying or selling options to adjust the gamma exposure. The latter cost is a fixed spread chosen so that it is 2.5% of the premium on a six month at the money option. The spirit of the analysis is to understand the hedging of aggregate portfolios of options. However, in order to have a "portfolio" that we really understand we have chosen initially to simply hedge a single one year written call option. We gamma hedge using options with an expiry of six months from our starting date. The hedges are rebalanced dynamically, and we look at the hedge performance at the end of a three month period.

Each of our worlds has a risk free interest rate of 5% and gives the underlying asset a volatility of either 0.30 or reverting to 0.30. In the jump diffusion world the intensity parameter is 5 (i.e. we expect 5 jumps a year) and it jumps to a distribution with a standard deviation of 0.05. The volatility of the diffusion is chosen so that the overall annualised standard deviation is 0.30. In the stochastic volatility model (which takes the CIR square root form for the variance) we have a mean reversion parameter of 1, and a volatility parameter of 0.30 as the volatility of volatility.

Within each of our three worlds we explore the performance of a number of hedging strategies. We first look at naive delta and delta-gamma strategies, implemented with a range of rebalancing periods from one day up to 25 days. We then look at heuristics designed to control the level of transactions incurred in a nearly optimal way. Specifically, we apply a myopic optimisation model which at each date minimises the sum of a quadratic function of gamma and speed plus the transactions costs incurred in rehedging. We also rebalance delta, but without considering the cost of this since it is relatively small. We simulate under a variety of parameter values for our optimisation in order to extract the rules which provide efficient cost/risk trade-offs.

6. Results of the Simulations

Figures 1 to 3 show the results of our simulations in each of the three "worlds". In each case we have plotted the expected level of transactions cost against the standard deviation of the replication error. The naive delta and gamma strategies have been run with a variety of rebalancing frequencies. Our myopic optimisation approach has been run with daily and weekly rebalancing and we have then computed its subsequent performance on an out of sample basis.

Given our low level of transactions cost of changing delta, the delta hedging approach consistently produces less replication variance as we rebalance more frequently, and the expected cost of this only increases fairly slowly. As expected, it performs less well once the Black-Scholes assumptions are violated by jumps, or to a lesser extent by stochastic volatility.

The naive gamma hedging strategy is rarely able to outperform the delta hedge, because of the high level of transactions costs it incurs. Note that the transactions costs worsen the replication variance as well as the expected cost.

Our simulations have enabled us to calibrate reasonable and robust parameter values for the myopic optimisation policy. This policy gives a significant advantage over the simple delta hedging approach under either jumps or stochastic volatility.

Finally, Figure 5 illustrates the nature of our preferred heuristic. The dots illustrate how revisions are made to the portfolio in gamma-speed space. It will be noted that they correspond to a kind of barrier strategy, but where the exact location of the optimal region varies slightly through time depending on the opportunities available in the option market.

7. Conclusions

Our simulations have confirmed that gamma hedging is particularly useful in realistic situations with stochastic volatility or jumps, but can be dangerous if applied in a naive manner. It is clear that the management of transactions costs is important, while the choice of model for calculating deltas and gammas seems to make little difference. It is also worth noting that our previous empirical study (Hodges, Clewlow *et al* (1993)) also found that the choice of volatility input used to construct hedges made little difference to hedge performance.

Further work remains to be done to provide more refined prescriptions for conserving transactions costs, and also to extend the analysis to a world with both stochastic volatility and jumps.

References

Bensaid B, J-P Lesne, H Pages and J Scheinkman, (1992), "Derivative Asset Pricing with Transactions Costs", *Mathematical Finance* 2, No 2, pp 63-86.

Black F and M Scholes, (1972), "The Valuation of Option Contracts and a Test of a Market Efficiency", *Journal of Finance*, 27, pp 399-418.

Black F and M Scholes, (1973), "The Pricing of Options and Corporate Liabilities", *Journal of Political Economy* 72, pp 637-659.

Boyle P and D Emanuel, (1980), "Discretely Adjusted Option Hedges", *Journal of Financial Economics* 8, pp 259-282.

Clewlow L J and X Xu, (1992), "A Review of Option Pricing with Stochastic Volatility", FORC Pre-print, Warwick Business School, University of Warwick.

Clewlow L J and X Xu, (1993), "The Dynamics of Stochastic Volatility", FORC Pre-print, Warwick Business School, University of Warwick.

Davis M H A and Norman A R, (1990), "Portfolio Selection with Transactions Costs", *Mathematics of Operations Research* 15, pp 676-713.

Davis M H A and V G Panas, (1991), "European Option Pricing with Transactions Costs", Proc 30th I E E E Conference on Decision and Control, pp 1299-1304.

Davis M H A , V G Panas and T Zariphopoulou, (1992), "European Option Pricing with Transactions Costs", Working Paper, Imperial College, London.

Dolbear S, (1992), "Hedging Lookback Options: An Empirical Analysis" in Hodges S D (Ed), *Options: Recent Advances in Theory and Practice, Volume 2*, Manchester University Press.

Edirisinghe C, V Naik and R Uppal, (1991), "Optimal Replication of Options with Transactions Costs", Working Paper, Faculty of Commerce, University of British Columbia.

Figlewski S, (1989), "Options Arbitrage in Imperfect Markets", *Journal of Finance*, 44, pp 1289-1311.

Garman M, (1992), "Charm School", *Risk*, July-August, pp 53-56.

Hodges S D and A Neuberger, (1989), "Optimal Replication of Contingent claims Under Transactions Costs, *The Review of Futures Markets*, 8, pp 222-239.

Hodges S D, L J Clewlow, R Martinez, M J P Selby, C Strickland and X Xu, (1993), "Hedging Option Position Risk: An Empirical Examination", FORC Preprint, Warwick Business School, University of Warwick.

Hodges S D and L J Clewlow, (1993), "Optimal Delta Hedging Under Transactions Costs", FORC Pre-print 93/36, Warwick Business School, University of Warwick.

Hull J and A White, (1987), "Hedging the Risks from Writing Foreign Currency Options", *Journal of International Money and Finance*, 6, pp 131-152.

Hull J and A White, (1988), "An Analysis of the Bias in Option Pricing Caused by a Stochastic Volatility", *Advances in Futures and Options Research*, 3, pp 29-61.

Kat H M, (1993), "Replicating Call Options in an Imperfect market: A Stochastic Simulation Study", Working Paper, University of Amsterdam.

Leland H E, (1985), "Option Pricing and Replication with Transactions Costs", *Journal of Finance* 40, pp 1283-1301.

Merton R C, (1976), "Option Pricing When Underlying Stock Returns are Discontinuous", *Journal of Financial Economics*, 3, pp 141-183.

Turner A L and E J K Weigel, (1992), "Daily Stock Market Volatility: 1928-1989", *Management Science*, 38, pp 1586-1609.

Wilmott P and E Whalley, (1993), "Counting the Costs", *Risk*, October, pp 59-65.

Table 1
Summary of Hedging Results: $100 Initial Premium

Contract Date	Expiry Hedging Errors Under: No Hedge	Delta	D-Gamma	D-Kappa
8509	-81.8	11.2	-35.4	-39.9
8606	-929.5	-28.3	16.9	15.1
8703	-77.8	-16.0	6.4	20.8
8712	104.8	-66.9	-5.6	64.6
8809	75.8	52.6	14.7	-49.2
8906	-141.3	32.5	2.5	-70.9
9003	104.9	-1.7	17.0	132.6
9012	104.9	4.6	4.5	-24.4
9112	-64.9	21.8	-20.0	-45.3
9209	31.8	35.0	-5.1	-65.8
Mean	-87.3	4.5	-0.4	-6.2
SD	294.0	33.2	16.0	61.7
SD from daily errors:				
underlying	207.0	28.1	14.7	87.9
volatility	58.5	58.5	45.1	34.3
Total	202.3	52.4	36.1	89.5

Table 2
Average Turnover

	Futures Turnover (#'s of contracts)		
Volatility	**Delta**	**D-Gamma**	**D-Kappa**
Implied	3.1	6.5	17.0
Constant	3.2	3.6	18.2

	Options Turnover (#'s of contracts)	
	D-Gamma	**D-Kappa**
Implied	25.1	38.3
Constant	20.5	39.6

Figure 1

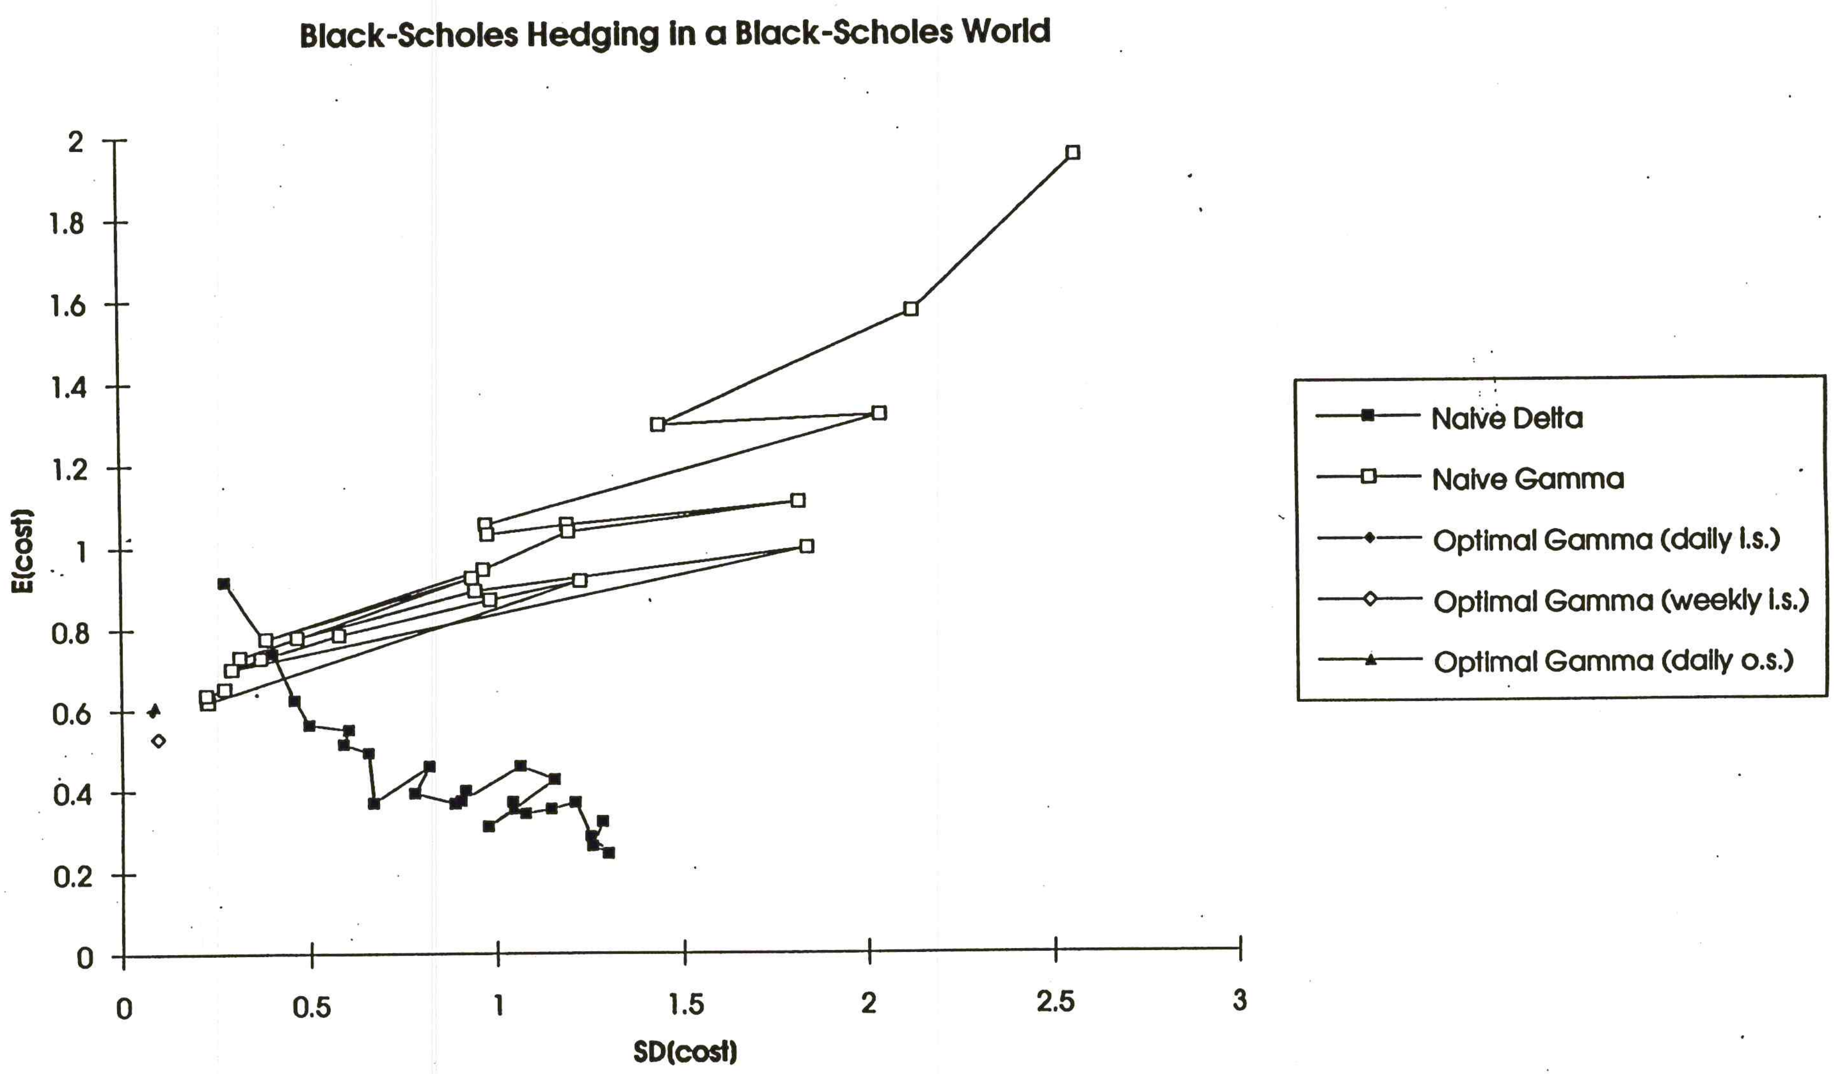
Black-Scholes Hedging in a Black-Scholes World
E(cost)
2
1.8
1.6
1.4
1.2
1
0.8
0.6
0.4
0.2
0
0.5
1
1.5
2
2.5
3
SD(cost)
Naive Delta
Naive Gamma
Optimal Gamma (daily i.s.)
Optimal Gamma (weekly i.s.)
Optimal Gamma (daily o.s.)

Figure 2

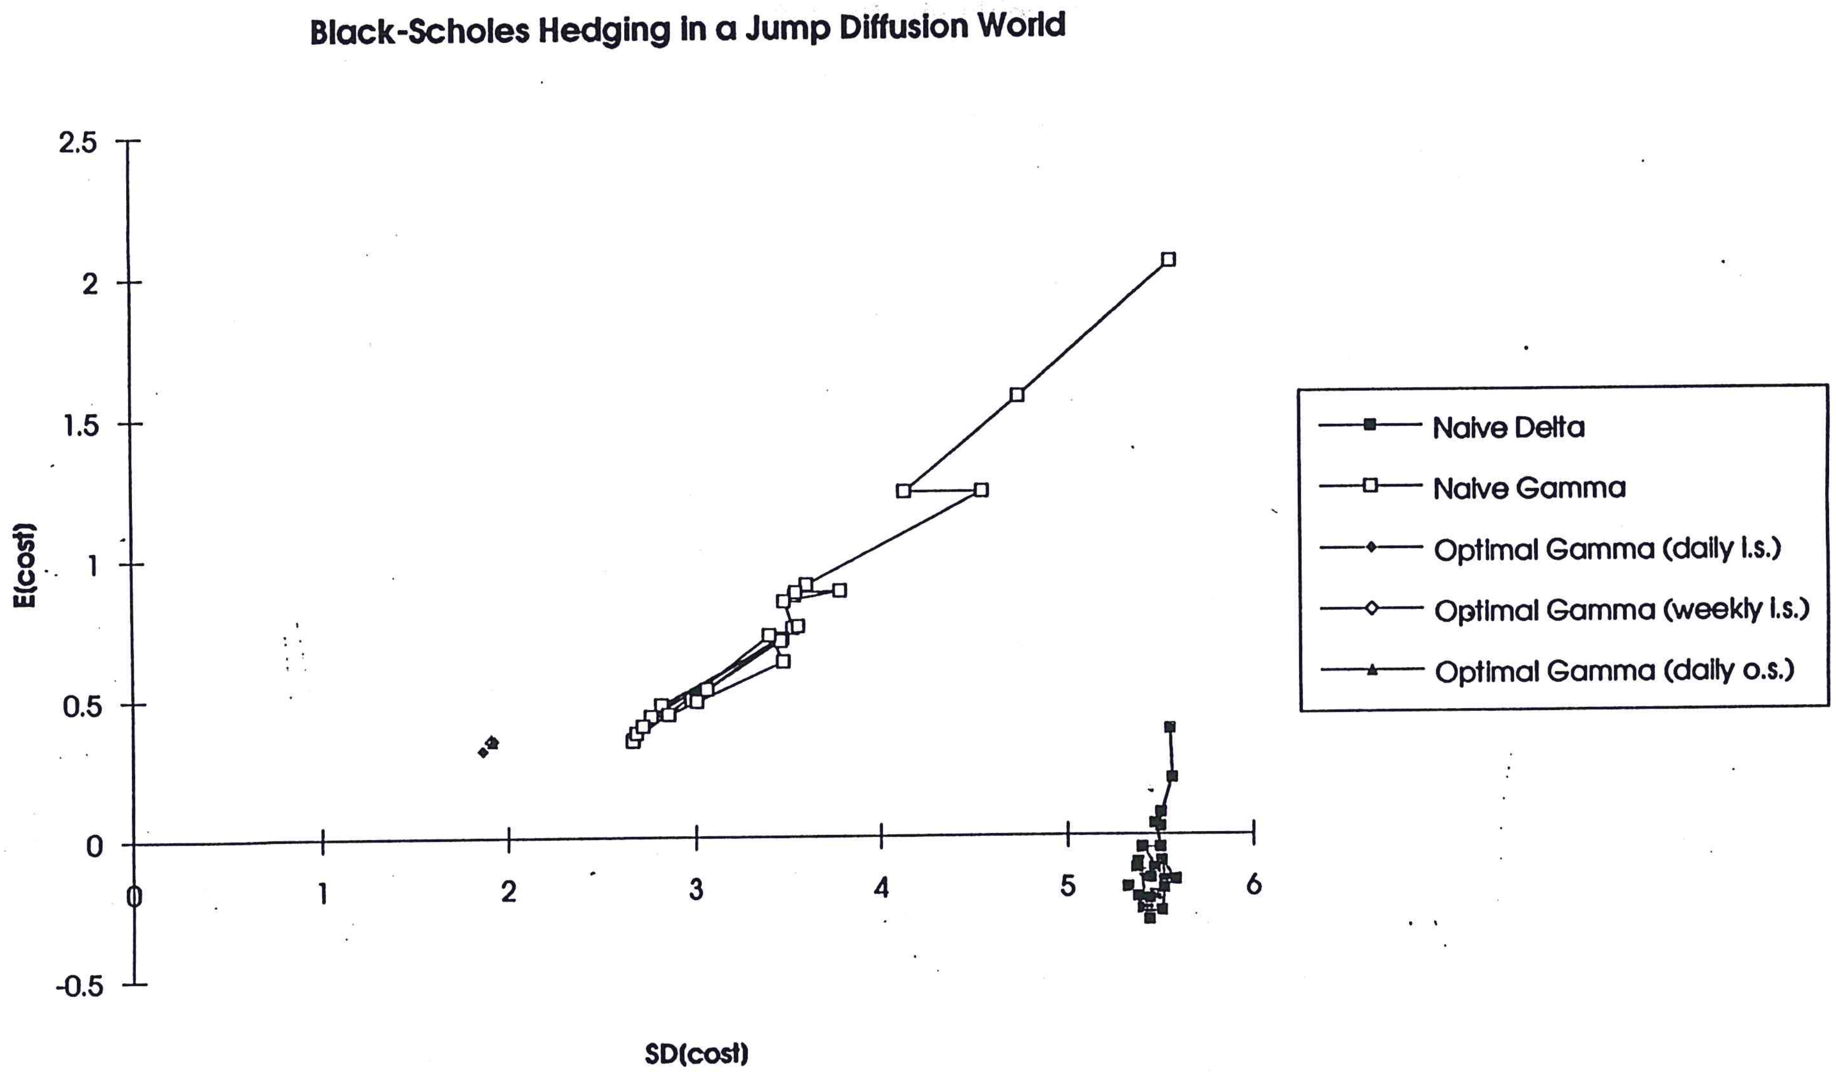

Black-Scholes Hedging in a Jump Diffusion World
2.5
2
1.5
1
0.5
0
-0.5
E(cost)
0
1
2
3
4
5
6
SD(cost)
Naive Delta
Naive Gamma
Optimal Gamma (daily l.s.)
Optimal Gamma (weekly l.s.)
Optimal Gamma (daily o.s.)

Figure 3

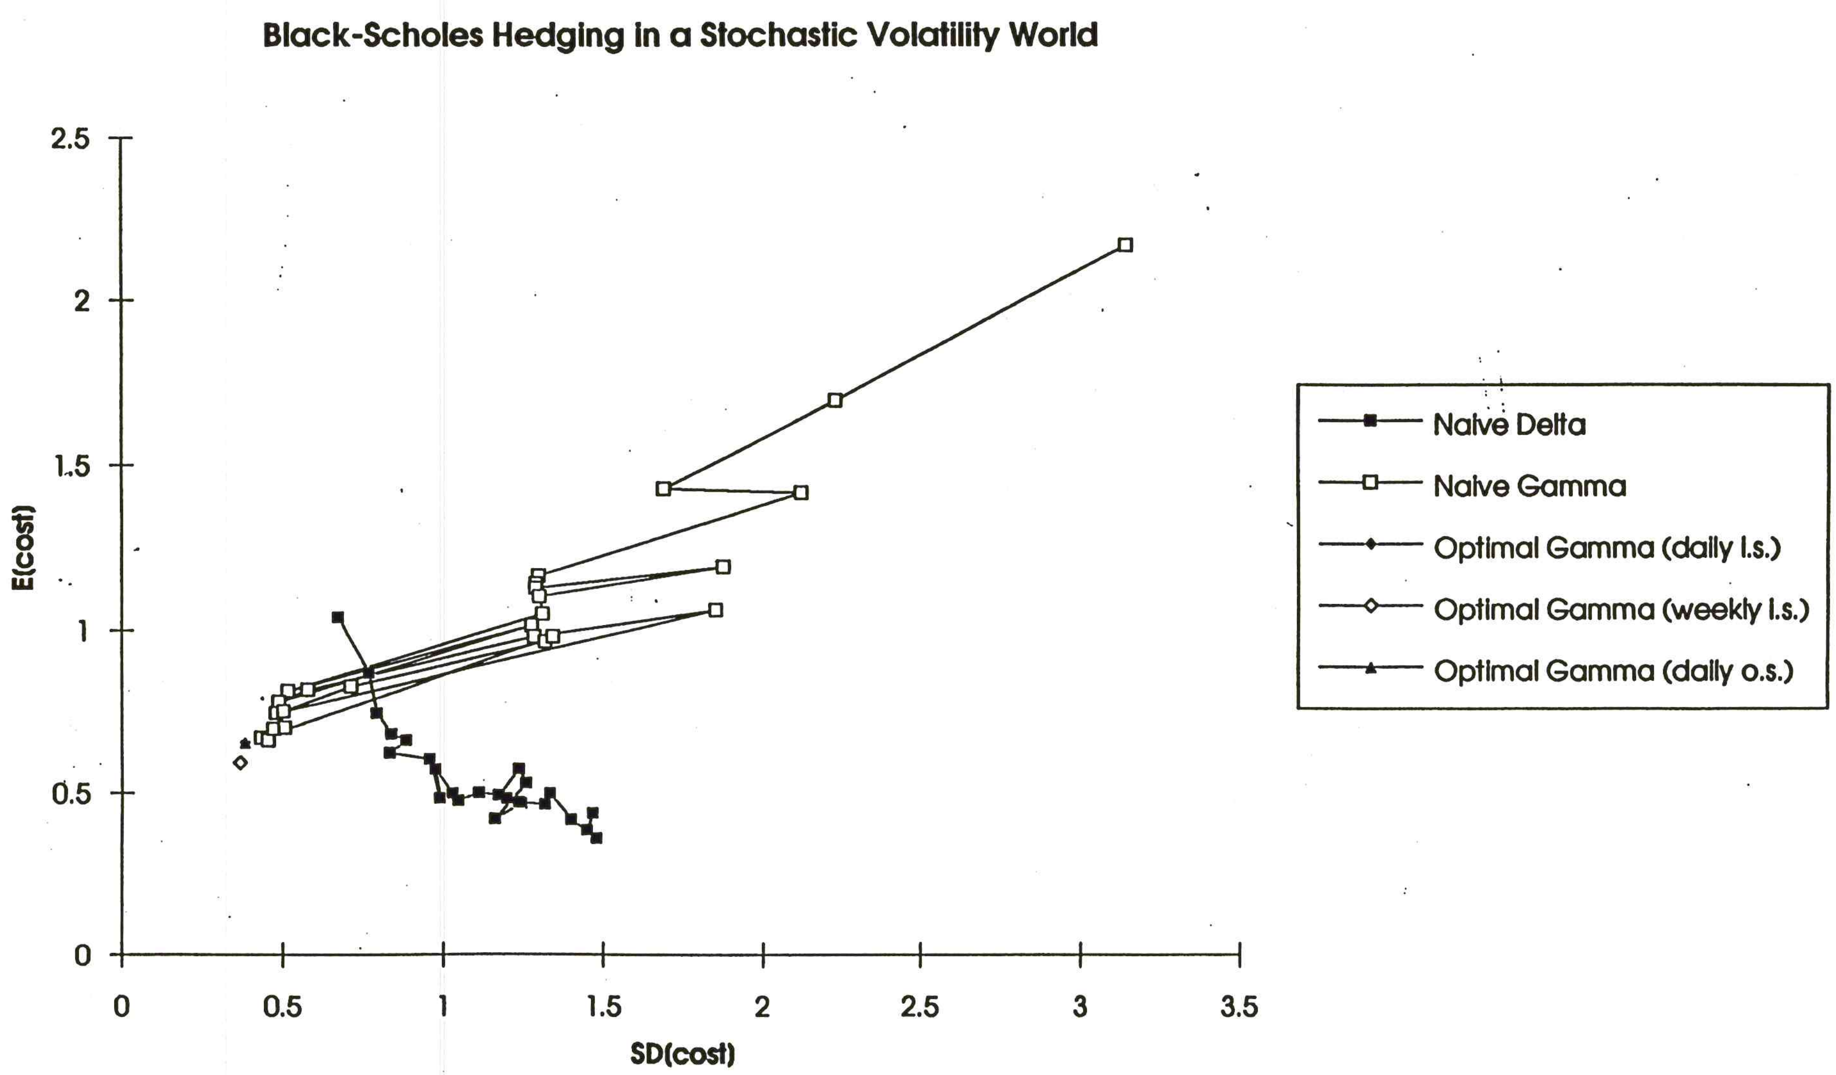

Black-Scholes Hedging in a Stochastic Volatility World
E(cost)
SD(cost)
2.5
2
1.5
1
0.5
0
0
0.5
1
1.5
2
2.5
3
3.5
Naive Delta
Naive Gamma
Optimal Gamma (daily l.s.)
Optimal Gamma (weekly l.s.)
Optimal Gamma (daily o.s.)

Figure 4

Expected Utility of Optimal Gamma Strategy in Black Scholes World

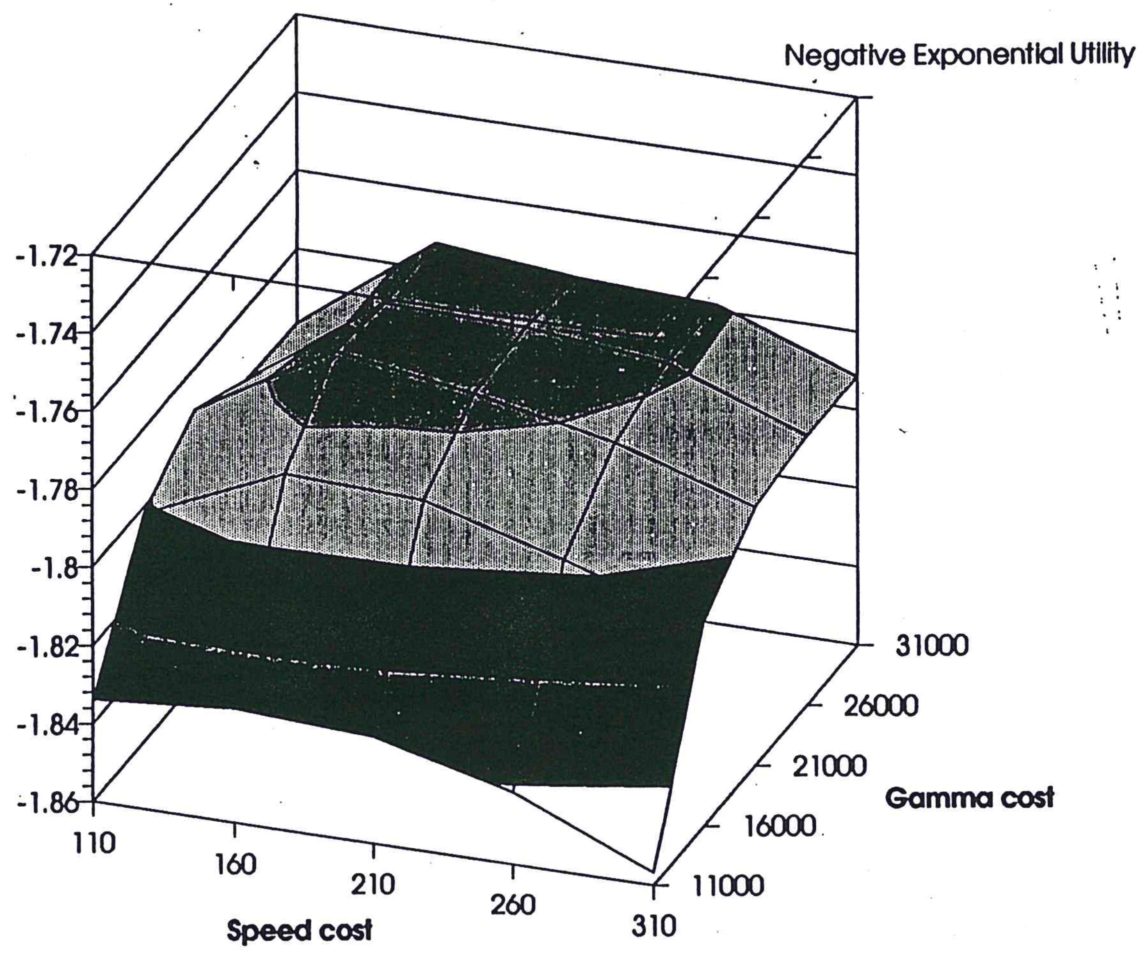

Figure 5

Optimal Gamma Strategy in Jump Diffusion World