Optimal Delta-Hedging Under Transactions Costs

Stewart Hodges*
and
Les Clewlow**

* Director
** Research Fellow

Financial Options Research Centre
University of Warwick

June 1992
Revised: February 1993

Funding for this work was provided by Corporate Members of the Financial Options Research Centre: Bankers Trust, LIFFOE, London Clearing House, London Fox, Midland Montagu, Mitsubishi Finance, Nomura Bank, and UBS Phillips & Drew.

Financial Options Research Centre
Warwick Business School
University of Warwick
Coventry
CV4 7AL
Phone: 0203 523606

FORC Preprint: 93/36

Optimal Delta-Hedging Under Transactions Costs

Stewart D Hodges and Les J Clewlow

* Financial Options Research Centre
Warwick Business School
University of Warwick
Coventry
England

June 1992
Revised: February 1993

Funding for this work was provided by corporate members of the Financial Options Research Centre: Bankers Trust, LIFFE/LTOM, London Clearing House, London Fox, Midland Montagu, Mitsubishi Finance, Nomura Bank, UBS Phillips and Drew.

Optimal Delta-Hedging

Under Transactions Costs

ABSTRACT

The paper examines the problem of delta-hedging under transactions costs, using the stochastic optimal control approach first described by Hodges and Neuberger (1989). Rather than seeking a strategy for exact replication, which is liable to be expensive and may be dominated by other strategies, this approach obtains the optimal hedging strategy to maximise expected utility (or to minimise a loss function defined on the replication error). Under proportional transactions costs this results in policies characterised by control bands within which the hedge delta must be maintained. Only with a fixed cost component would it be appropriate to make large transactions to jump into the interior of the control region.

This method has the advantage over Leland's (1985) approach that it works just as well for hedging mixed portfolios of long and short positions, and also mixed maturity dates. The paper describes the basic approach, and derives a new computational method which substantially reduces the storage required for the calculation. Characteristics of the optimal policies are discussed, and a simulation study is completed to compare the hedging performance of some alternative policies. The strategies we tested were chosen so that we could examine which features of a hedging strategy are most important: hedging to the "correct" delta or hedging only to within a band in order to conserve transactions costs. The simulations show that the optimal control approach is substantially more effective than Leland's method, and that while the target delta and the band around it are both important, surprisingly good hedges can be obtained by hedging using a control region of the right width but based around an incorrect central delta.

Optimal Delta-Hedging Under Transactions Costs

1. Introduction

This paper describes a stochastic optimal control approach for delta-hedging contingent claims. It extends earlier work by Hodges and Neuberger (1989). The paper describes an efficient numerical procedure for this type of hedging, discusses properties of the optimal hedging strategies, and reports on the results of a series of numerical simulations to compare the effectiveness of the approach to some alternative methods.

The construction of hedging strategies which best replicate the outcomes from options (and other contingent claims) in the presence of transactions costs is an important problem. Hedging is central to the theory of option pricing. Arbitrage valuation models, such as that of Black and Scholes (1973), depend on the idea that an option can be perfectly hedged using the underlying asset, so making it possible to create a portfolio which replicates the option exactly. Hedging is also widely used to reduce risk, and the kind of delta-hedging strategies implicit in Black and Scholes are commonly applied, at least approximately, by participants in options markets. Optimal hedging strategies are therefore of direct practical interest. Much of the theory of options assumes that markets are frictionless. This paper considers the impact of transactions costs on delta-hedging and valuation. This is closely related to the valuation issues which arise where the nature of the market dictates that trading is discontinuous, or that the asset process is such that the market is incomplete and contingent claims are not spanned by existing securities.

The first paper to consider the problem of replicating options' payoffs using delta-hedging under transactions costs was Leland (1985). The issue is particularly interesting because under the usual Black-Scholes strategy, implemented as rebalancings at discrete intervals, the expected volume of transactions becomes unbounded as the number of rebalancings is increased. Leland's analysis is set in a continuous-time framework and assumes proportional transactions costs. It describes how by making an adjustment to the variance (which depends on the exogenously specified revision frequency) the Black-Scholes formula can be used to hedge with a zero expected replication error, and with a standard deviation which tends to zero with the length of the rebalancing interval. Neuhaus (1989) contributes some further theoretical insights to this approach. However, this method is in no sense an optimal one.

The method described in this paper follows earlier work by Hodges and Neuberger (1989) and is based on maximising expected utility. Alternatively, we may view it as minimising a loss function defined on the replication error. This approach seems more appropriate since the valuation bounds provided by exact replication may be very wide. Depending on the choice of risk aversion parameter we can obtain either tight or much looser (but also cheaper) hedging. The approach is in a paradigm similar to that of Davis (1988), Davis and Norman (1988) and Taksar, Klass and Assaf (1988) and Dumas and Luciano (1991). These papers describe optimal portfolio policies to maximise expected utility over an infinite horizon. They extend earlier work by Merton (1971) and Constantinides (1986). However, while these papers are concerned with optimal policies, they are not directly concerned with the problems of replicating (or similarly hedging) contingent claims by means of the underlying asset.

Mention should also be made of a number of other recent papers related to ours. The model first proposed by Hodges and Neuberger (1989) has been further studied by Davis and Panas (1991), and Davis, Panas and Zariphopoulou (1992). Dixit (1991) and Dumas (1991) provide useful material on smooth pasting conditions which usually apply to problems such as ours. Figlewski (1987) gives some interesting simulation results. Boyle and Vorst (1992) provide an elegant reworking of Leland's analysis within a binomial framework. This has the following interesting feature. Their variance adjustment differs from Leland's, essentially because the binomial assumption distorts the expected absolute price change in any sub-interval even though it provides the correct variance. Edirisinghe, Naik and Uppal (1991) also provide a binomial replication based approach, and apply a technology based on linear programming. Replication can be a dangerous philosophy. Bensaid, Lesne, Pages and Scheinkman (1991), also in a discrete-time framework, show that it can be cheaper to dominate a contingent claim than to exactly replicate it. Neuberger (1992) shows that in contrast to the diffusion case, under a pure jump process (with fixed jump size), exact replication can provide tight bounds on option values.

Section 2 of the paper describes the general problem of the best replication of a contingent claim (or portfolio of claims) under transactions costs. Exact replication at finite cost, even when possible, is generally too expensive to be desirable. The replication problem must therefore be formulated relative to some loss function (or utility function for marginal wealth changes). We show that this problem is one of stochastic optimal control and can be characterised by a conventional backward recursion (dynamic programming) technique like a binomial procedure. In general it is necessary to use numerical methods of solution as the general problem involves three state variables. By using a suitable utility function (exponential), the number of state variables is reduced to two and we are able to obtain numerical solutions to realistic problems. Section 3 of the paper describes the method we have chosen for numerical implementation of our procedures, and section 4 discusses some of the properties of the optimal hedging strategies we obtain. Our simulation analysis, in section 5, compares the hedging performance of a variety of approaches to hedge a small selection of option exposures. The optimal control based strategies are shown to be considerably better than the alternative strategies described by Leland. Section 6 presents our conclusions, and summarises the advantages of our optimal-control approach.

2. The Model

The structure of the model follows that developed in Hodges and Neuberger (1989). Davis and Panas (1991) and Davis, Panas and Zariphopoulou (1992) have developed some of its theoretical properties. We consider an asset whose price S_t at time t evolves under the diffusion process described by

$$dS_t = \mu(S_t)dt + \sigma(S_t)dz \tag{1}$$

The problem is to replicate the outcomes from a contingent claim whose payoffs at a single future date T are given by $C(S_T)$. The replication is to be accomplished by holding x_t units of the asset plus either borrowing or lending at a constant interest rate r. The holdings in this replicating portfolio are to be actively managed through time, but transactions in the underlying asset involve a transactions cost amounting to $k(v,S)$ where v is the volume of shares transacted (either positive or negative) and S is the (mid) share price. For most purposes we shall specialise this to the case of costs which are a constant proportion of the value transacted:

$$k(v,S) = k|v|S \quad (2)$$

In general it is either impossible or at least undesirable to replicate the contingent claim exactly. For many problems exact replication at finite cost is impossible. For others, while exact replication at finite cost may be possible, it will be too expensive to represent an attractive policy. The replication problem is therefore ill-defined until we have specified a criterion for choosing between alternative replicating strategies. We will assume initially a fairly general expected utility criterion which will later be specialised to a particular function. Thus we assume that an initial amount of money is invested through time (managed between the risky asset and the risk free rate). By the terminal date *T*, after liquidating the asset holding, an amount of cash y_T is available to set against the contingent liability $C(S_T)$. At this date we have an accumulated surplus of

$$w_T = y_T - C(S_T) \quad (3)$$

net of the option value to be replicated. We define a utility function $U(w_T)$, and seek to characterise and calculate replication strategies which maximise the expected value of this utility function. We may also allow the horizon date to be later than the expiry date of the contingent claim. Also, since it is definitely possible for y_T to be negative, we are precluded from using some commonly employed utility functions, such as power or logarithmic utility functions for $U(.)$. We shall assume that $U(w)$ is defined for all real numbers w, that its first two derivatives exist, are continuous, and satisfy the usual properties for a risk averse utility function, ie, that $U_w>0$ and $U_{ww}<0$.

We now describe the structure of the general problem. Using the notation already introduced we define the indirect utility function $J(.)$ as:

$$J(t,S,x,y) = \max E[U(w_T)] \quad (4)$$

as the maximum expected utility possible starting at time t when the asset price is S, with initial holdings of x shares, and an amount y in cash. $E[.]$ is the expectation operator under some suitable probability measure, not necessarily the objective one. We also define μ^* as the rate of drift of S under this measure. The maximum is taken over all feasible transactions policies. At the last date T, it is clear that by definition $J(.)$ is obtained trivially as

$$J(T,S,x,y) = U(w_T) \tag{5}$$

where

$$w_T = xS_T + y_T - C(S_T) \tag{6a}$$

corresponding to no costs at termination (ie. stock settlement permitted), or

$$w_T = xS_T - k(x,S_T) + y_T - C(S_T) \tag{6b}$$

corresponding to cash settlement after transactions costs have been paid.

The indirect utility function *J(.)* is solved recursively backwards through time using the dynamic programming approach of stochastic optimisation. *J(.)* evolves backwards as given by

$$J(t,S,x,y) = \max\{E_{dS}[J(t+dt,S+dS,x^*,y(x^*))]\} \tag{7}$$

where the maximum is taken over the choice of the quantity of shares x^* to hold.

This optimal control problem is characterised by the second order partial differential equation (using subscripts to denote partial derivatives)

$$J_t + \mu^*(S)J_S + \tfrac{1}{2}\sigma^2(S)J_{SS} + yrJ_y = 0 \tag{8a}$$

for interior values of $x \varepsilon X$, subject to the boundary conditions

$$J(T,S,x,y) = U(w_T) \tag{8b}$$

defining J at the terminal date T, and

$$J(t,S,x+u,y-uS-k(u,S)) \leq J(t,S,x,y) \tag{8c}$$

which defines the boundary of the region X on which transactions occur. For the special case of constant proportional transactions costs, equation (8c) simplifies to an inequality relationship between J_x and J_y.

The solution to this problem provides a "reservation selling (or buying) price" for the contingent claim, and also the optimal strategy for hedging it. For our case of proportional transactions costs the optimum strategy is defined by the functions

$$x_-(t,S,y),\, x_+(t,S,y)$$

which define lower and upper bounds for the number of shares to be held: the discovery that $x<x_-$ leads to transactions to re-establish x at the value x_-, and similarly if $x>x_+$ it is re-established to x_+.

The reservation selling (and buying) prices are defined as follows. We define $J^C(t,S,x,y)$ as the expected utility (under an optimal hedging strategy) of assuming the state contingent liability C, (eg. $C(S_T)$ as before). The individual's reservation selling price of C, $V_S(C)$ is defined as the price required to provide the same expected utility as not selling the contingent claim. Thus V_S is defined by the equation

$$J^C(0,S,0,V_S) = J^0(0,S,0,0) \tag{9}$$

where J^0 is defined as J^C, but with no state contingent liability assumed.

Similarly, we can define the buying price V_B as the maximum price worth paying to buy the contingent claim, defined by the equation

$$J^{-C}(0,S,0,-V_B) = J^0(0,S,0,0) \tag{10}$$

In addition to calculating optimal hedging strategies and valuations, we can also calculate recursively as many moments of the distribution of w_T as may be of interest. The moments of w_T about zero simply accumulate as expectations conditional on the state variables involved.

Note that under this general formulation, at each date in the calculation, the indirect utility function (and also any derived moment functions) depend on the three state variables of S, x and y. The computational effort required may be considerable, unless simplifications are found. A general numerical solution would be daunting. We therefore specialise the utility function to the negative exponential

$$U(w_T) = -\exp(-\lambda w_T) \tag{11}$$

and also choose the risk free interest rate r as the risk adjusted rate of drift, μ^*. This reduces the state variables by one and makes the computations relatively straight forward. It also enables us to produce strategies which have the attractively simple properties of not being wealth dependent and not creating risky positions in the absence of any contingent claim to be hedged. The optimal control problem under this specific set of assumptions is developed in Appendix A.

3. Computational Aspects

In this section we give a brief description of some transformations and approximations we use to obtain a robust and quick computational scheme. In their original work, Hodges and Neuberger used a binomial tree for the stock price, with a large vector corresponding to different possible deltas at each node of the tree. At each stage, the value of the indirect utility function $H(.)$ was calculated for each element of each vector by the usual binomial averaging scheme, and these were then modified to reflect the inequality relationships which stem from undertaking transactions when it improves expected utility to do so. Although this scheme works reasonably well, it has a number of disadvantages. The values of expected utility are exponential in terminal wealth, and therefore computationally liable to "blow-up" as either underflow or overflow. It therefore seems better to work with the logarithm of (minus) the utility, which is simply the reservation price.

The indirect utility function H evolves as a normal diffusion. It is a simple matter to work out (using Ito's Lemma) the diffusion for $\ln(-H)$, normalised as the reservation price. The boundary condition on the reservation price is now the obvious linear slope condition that $\partial V/\partial x$ must lie between $-S(1+k)$ and $-S(1-k)$. We still use a conventional binomial scheme for the evolution of S, and it is chosen so that the up and down probabilities are equal. Our numerical scheme is based on the following evolution of reservation values in our binomial tree. The reservation price V_j at a particular node j is related to the corresponding indirect utility function by

$$\exp\{\lambda V_j\} = -H_j$$

H_j diffuses backwards, so V_0 is related to its two successor values V_1, V_2 as

$$\exp\{\lambda V_0\} = \frac{\exp\{\lambda V_1\} + \exp\{\lambda V_2\}}{2}$$
$$= \exp\{\lambda V_1\} \frac{1 + \exp\{\lambda (V_2 - V_1)\}}{2}$$

This gives

$$V_0 = V_1 + \frac{1}{\lambda} \ln\left[\frac{1 + \exp\{\lambda (V_2 - V_1)\}}{2}\right] \tag{18}$$

Finally we note the power series expansion:

$$\ln\left(\tfrac{1}{2} + \tfrac{1}{2} e^x\right) = \tfrac{x}{2} + \tfrac{x^2}{8} - \tfrac{x^4}{192} + \ldots$$

so $$V_0 = \tfrac{1}{2}(V_1 + V_2) + \frac{\lambda}{8}(V_1 - V_2)^2 - \ldots \tag{19}$$

and we can ignore the higher order terms as long as λ is not chosen too large, and our grid size is also sufficiently small.

We can also use this approximation to find the first derivative of V_0, as

$$V_1' = \tfrac{1}{2}\left(V_1' + V_2'\right) + \frac{\lambda}{4}(V_1 - V_2)\left(V_1' - V_2'\right)\ldots \tag{20}$$

in order to take account of the boundary conditions arising from transacting.

The method we adopt is to make a functional approximation to V_j as a function of x at each node. This approximation is based on knowing the values of x and V_j at which $\partial V_j/\partial x$ equals $-S(1+k)$, $-S$, and $-S(1-k)$ respectively. The first and last of these correspond to our control limits x_- and $x+$, and outside these V_j is simply a linear function. From the functional approximations we can use the equations given above to search for the corresponding points at the next node of our tree. We use a combination of Newton-Raphson and bisection to perform this search, as in some parts of the tree the pure Newton-Raphson does not converge. We can thus make the same type of functional approximation on the new node and proceed in this way for the entire tree. This procedure gives a large improvement in speed, and an enormous saving in storage requirements, as we now need only store a handful of parameters at each node of the binomial tree.

4. Properties of Optimal Hedging Strategies

We will now provide some general comments regarding the properties of optimal replication strategies. The problem of finding the optimal policy is one of solving the partial differential equation corresponding to (8). The solution starts from the terminal boundary condition and is also subject to a free boundary condition which corresponds to the position of the control boundary. The solution provides a reservation selling price (or buying price) above which it is advantageous to sell (or buy) the contingent claim and hedge the risk using the prescribed strategy. As we have just seen, the optimal control x is constrained to evolve between control limits which depend on time, and on the asset price. Under the negative exponential utility assumption the amount of cash accumulated into the replicating portfolio is irrelevant. No controlling action is taken until the control parameter x attains one of its limits. We will use some numerical examples to illustrate general features of our solutions.

Figure 1 shows the buy and sell reservation prices computed for a six month call option, with Black-Scholes values for comparison. The asset value is 100p, the exercise prices is 100p, the volatility is 30% and we have used a zero interest rate. Transactions costs are at 2% (each way) and we have chosen $\lambda=1$. We have used these values throughout the numerical work reported in this paper, excepting that for our simulations we have used a transactions cost of 1% instead of 2%.

Figures 2 and 3 plot the control region for fixed time to expiry as a function of the price of the underlying asset. This region evolves through time in a way comparable to the behaviour of the Black-Scholes delta curve for hedging options in the absence of any transactions costs. As S changes, it is necessary to adjust the delta of the hedge only as much as is required to keep within the region defined by the two curves. The shape of the curves depend on the contingent claim to be hedged, on the level of transactions cost and on the degree of risk aversion. They also reflect the variance intuition of Leland (1985). Leland noted that when there is a short gamma exposure (so that an increase in asset price needs to an increased requirement for the underlying) the transactions cost makes it as if the price movement had been even greater. If we are hedging a short call our control region is flattened, corresponding to an increased variance assumption. Note that for out-of-the-money and in-the-money options, the Black-Scholes delta may be outside the optimal control region. In other words, if we inherit an options book which is currently exactly delta-hedged under Black-Scholes and we face delta-hedging costs, it may nevertheless be optimal to move the hedge away from the Black-Scholes value. If we plot Black-Scholes delta curves on our diagram for different values of volatility we find that they cross the curves we have computed: our curves are not simply "Leland" ones for simple constant volatility adjustments. We should think of the volatility adjustment as reflecting the expected cost of future hedging transactions, and this depends on the level of the price itself.

Conversely, for the case where an option has been purchased, the positive gamma means we can sell some of the underlying after a price rise, so here it is as if the variance were smaller, and hence our hedging region slopes at a steeper angle. In this case the region is wider, which reflects the fact that the convex shaped payoff presents much less risk of large losses. For lower levels of cost or higher levels of risk aversion we should expect the width of the no-transactions region to be reduced.

Hedging more complex positions is especially interesting, since with a combination of both long and short positions Leland's method may find it hard to know whether to increase or decrease the variance. The optimal-control approach has no such problems. Conditional on any pre-specified contingent payoffs, which can even occur at differing dates, it works out the best hedge for a given degree of risk aversion. The risk aversion reflects the trade-off to be made between the expected cost of managing the hedge and its variance. We have chosen to hedge the payoffs from a bull spread in order to illustrate hedging a portfolio with mixed positions.

We next turn to the behaviour of the boundary as a function of time. Close to expiry, the shape of the control boundary depends critically on the assumptions about settlement. Where stock settlement is permitted, it is not worth paying transactions' costs immediately prior to expiry just to avoid minor replication errors. In this case the limits x_- and x_+ flare outwards near to the expiry date. However, if settlement must be made in cash, any excess position in the underlying had better be liquidated sooner rather than later. These features are evident from Figures 4-7 which show how the no-transactions regions evolve through time for various fixed values of the underlying.

It is worth noting how the corridors we compute frequently have fairly constant width once we are back from the immediate pre-expiration transient. They are not obviously related to the gamma of the claim, though that is clearly one aspect of their determinants. Overall we are struck by the complexity of these diagrams and we are no longer surprised that we have so far failed to obtain any real analytical results as to their behaviour.

5. Simulation Analysis

In this section we compare the performance of the optimal delta-hedging strategy against other common strategies: Black-Scholes, Leland and a heuristic strategy based on the optimal strategy. For the Black-Scholes and Leland strategies it is necessary to choose a replication interval which reflects the investor's risk aversion. We compute the relevant delta at the start of each interval and adjust the holding in the underlying asset accordingly. The heuristic strategy is designed to highlight the relative importance of the two facets of the optimal strategy: the optimal delta and the width of the no-transactions region around this delta. The heuristic strategy therefore centres the optimal region on the Black-Scholes delta rather than the optimal delta. For the optimal and heuristic strategies we must choose the risk aversion parameter λ, we then wait until the holding in the underlying asset moves outside the no-transactions region at which point we adjust the holding to the nearest boundary of the no-transactions region.

In order to compare the four strategies we require a suitable metric. We simply adopt a mean-variance framework and plot the expected cost relative to the Black-Scholes "fair" value and the standard deviation of that cost. This cost is defined as follows: we sell or buy the contingent claim for the Black-Scholes (no transactions cost) fair value and use the proceeds to replicate the claim under the transactions costs. At maturity we compute the cash value of the portfolio (the value of the underlying asset held less the borrowing and the liability). This value discounted back to the present is the cost of the strategy relative to the no transactions costs case.

Our simulations are based on replication over a year where the minimum revision interval is one day. We simulate replication of selling and buying a European call (with an exercise price of 100) and a bull spread[1] (with a lower exercise price of 100 and an upper exercise price of 110). The price of these options is set, as stated above, at the no transactions costs fair value. For all the simulations the initial underlying asset price is 100, the annualised volatility is 30% and the riskless rate is 0. We set the proportional transactions costs to be 1% of the value of any single trade in the underlying asset.

The simulation proceeds as follows: Firstly we compute the binomial lattice solution to the optimal strategy. We save a table of the no-transactions region boundary values for each binomial lattice value of the underlying asset and each time step (we arrange for the binomial lattice to have the same number of time steps as the simulation). The initial hedge portfolios are then set up and we begin simulation of the underlying price path. After each daily time step we check the optimal and heuristic hedges to see if the holding in the underlying asset is outside the no-transactions region. This is done by interpolating for the region at the current underlying price from the saved table. Note that it is possible for the underlying price to move outside the range in the table. If this occurs we must recompute the binomial lattice solution. However, by arranging for the binomial lattice to have three nodes rather than one at the current time, we almost never have to recompute the solution[2]. If the holding is outside the no-transactions region we rebalance the hedge back to the nearest boundary. At each Black-Scholes/Leland replication interval we rebalance these hedges. Finally, at maturity we compute the cost of each strategy and collect the statistics necessary to compute the mean and variance of the costs. This path simulation is repeated 1000 times.

Figures 8, 9, 10 and 11 summarise the results of the simulations. The curves correspond to rebalancing intervals of one to twelve days for the Black-Scholes and Leland strategies and to λ from 0.2 to 10.0 for the optimal and heuristic strategies.

Consider first the replication of a European call (Figures 8 and 9). For the optimal strategy and the Leland strategy the expected cost is strictly monotonically increasing with decreasing standard deviation of the cost (increasing risk aversion). But for the

---

[1]Note that the Leland strategy is strictly only applicable to globally convex or concave payoff functions. In our implementation of the Leland strategy for a bull spread we use the sign of the Black-Scholes Gamma to determine the direction in which the volatility is adjusted in an attempt to account for this problem.

[2]In the case of a lognormal underlying asset we can of course quantify this probability.

Black-Scholes and heuristic strategies this is not the case. As we rebalance more and more frequently the variance of the cost begins to be dominated by the variation of the underlying asset price. As the risk aversion increases the heuristic strategy tends towards the Black-Scholes strategy. This is because the no-transactions region becomes very narrow and since it is centred on the Black-Scholes delta we obtain the Black-Scholes strategy in the limit. The heuristic strategy helps us to understand the relationship between the Black-Scholes, Leland and optimal strategies. At low risk aversion the no-transactions region is wide and the optimal delta tends towards the Black-Scholes delta. As the risk aversion increases the region becomes narrower and the optimal delta deviates from Black-Scholes.

The curves for the Black-Scholes and Leland strategies exhibit more variability than those for the optimal and heuristic strategies because for low levels of risk aversion they rebalance very infrequently and the times when they do rebalance can be severely sub-optimal. In contrast the optimal (and heuristic) strategies mean rebalancing interval only varies from 4.6 to 1.6 days over the entire range of λ from 0.2 to 10.0.

For the bull spread (Figures 10 and 11) we obtain the same qualitative results but the details are different. The Leland strategy suffers a similar problem to the Black-Scholes and heuristic strategies for high levels of risk aversion in that the standard deviation is not monotonically decreasing with expected cost. In fact the Leland strategy now performs worse than Black-Scholes. This is because (as we noted earlier) the Leland strategy is only strictly applicable for globally concave or convex payoff functions where the volatility is adjusted upwards or downwards respectively.

In summary, it is very important for risk averse investors, to optimise their delta-hedging; the alternative strategies we have examined are poor substitutes.

6. Summary

The paper has examined the problem of delta-hedging under transactions costs, using the stochastic optimal control approach first described by Hodges and Neuberger (1989). Rather than seeking a strategy for exact replication, which is liable to be expensive and may be dominated by other strategies, this approach obtains the optimal hedging strategy to maximise expected utility (or to minimise a loss function defined on the replication error). Under proportional transactions costs this results in policies characterised by control bands within which delta must be maintained. Only with a fixed cost component would it be appropriate to make large transactions to jump into the interior of the control region.

This method has the advantage over Leland's approach in that it works just as well for hedging mixed portfolios of long and short positions, and also mixed maturity dates. The paper describes the basic approach, and derives a new computational method which substantially increases the speed and reduces the storage required for the calculation. Characteristics of the optimal policies are discussed, and a simulation study is completed to compare the hedging characteristics of some alternative policies. The strategies we tested were chosen so that we could examine which features of a hedging strategy are most important: hedging to the "correct" delta or hedging only to within a band in order to conserve transactions costs. The simulations show that the optimal control approach is substantially more effective than Leland's method, and that while the target delta and the band around it are both important, surprisingly good hedges can be obtained (at least for low levels of risk aversion) by hedging using a control region of the right width but based around an incorrect central delta.

References

Aczel M A and J E Broyles, (1986), "Option Pricing Reflecting Transactions Costs", Working Paper, MRP 86/8, Templeton College, Oxford, July 1986.

Bensaid B, J-P Lesne, H Pages and J Scheinkman, (1991), "Derivative Asset Pricing with Transactions Costs", Working Paper 91/1 , Bank of France, Centre de Recherché, March 1991.

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.

Boyle P and T Vorst, (1992), "Option Replication in Discrete Time with Transactions Costs", *Journal of Finance*, 47, pp 271-293.

Constantinides G M, (1986), "Capital Market Equilibrium with Transactions Costs", *Journal of Political Economy* 94, pp 842-862.

Cox J C, Ross S and M Rubinstein, (1979), "Option Pricing: A Simplified Approach", *Journal of Financial Economics* 7, pp 229-263.

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.

Dixit A K, (1991), "A Simplified Treatment of the Theory of Optimal Regulation of Brownian Motion", *Journal of Economic Dynamics and Control*, 15, pp 657-673.

Dumas B, (1991), "Super Contact and Related Optimality Conditions", *Journal of Economic Dynamics and Control*, 15, pp 675-685.

Dumas B and E Luciano, (1992), "An Exact Solution to a Dynamic Portfolio Choice Problem under Transactions Costs", *Journal of Finance*, 46, pp 577-595.

Eastham J F and Hastings K J, (1988), "Optimal Impulse Control of Portfolios", *Mathematics of Operations Research* 13, pp 588-605.

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.

Gilster J and Lee W, (1984), "The Effects of Transactions Costs and Different Borrowing and Lending Rates on the Option Pricing Model: A Note", *Journal of Finance* 39, pp 1215-22.

Garman M B and Ohlson J A, (1981), "Valuation of Risky Assets in Arbitrage Free Economies with Transactions Costs", *Journal of Financial Economics* 9, pp 271-280.

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

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

Magill M J P and G M Constantinides, (1971), "Portfolio Selection with Transactions Costs", *Journal of Economic Theory* 13, pp 245-263.

Merton R C, (1971), "Optimum Consumption and Portfolio Rules in a Continuous Time Model", *Journal of Economic Theory* 3, pp 373-413.

Merton R C, (1973), "Theory of Rational Option Pricing", *Bell Journal* 4, pp 141-183.

Neuberger A, (1992), "Option Replication with Transactions Costs - An Exact Solution for the Pure Jump Process", Working Paper, London Business School.

Neuhaus H, (1989), "Discrete Time Option Hedging", Doctoral dissertation, London Business School.

Shreve S, (1989), "A Control Theorist's View of Asset Pricing", Working Paper, Carnegie Mellon University. (Presented at Workshop on Applied Stochastic Analysis, Imperial College, London, April 1989).

Taksar M, M J Klass and D Assaf, (1988), "A Diffusion Model for Optimal Portfolio Selection in the Presence of Brokerage Fees", *Mathematics of Operations Research* 13, pp 277-294.

Appendix A

We here present the derivation of the optimal control problem under our specific choice of negative exponential utility.
Since from (4) and (11):

$$J(t,S,x,y) = E\left[-\exp\{-\lambda w_T\}\right] \tag{A1}$$

and the management of x through time is independent of y,

$$J(t,S,x,y) = J(t,S,x,0)\exp\left\{-\lambda\, y\, e^{r(T-t)}\right\} \tag{A2}$$

If we define a new indirect utility function

$$H(t,S,x) = J(t,S,x,0) \tag{A3}$$

then we may derive the following new equations and boundary conditions for H, which correspond to our previous equations (8):

$$H_t + r\,S\,H_S + \tfrac{1}{2}\sigma^2(S) H_{SS} = 0 \tag{A4a}$$

$$H(T,S,x) = -\exp\{-\lambda w_T\} \tag{A4b}$$

$$H(t,S,x+u) \geq H(t,S,x)\exp\left\{-\lambda(uS - k(u,S))\, e^{r(T-t)}\right\} \tag{A4c}$$

For the special case of a constant proportional transactions cost, $k(v,S) = k\,|v|\,S$, this last equation translates to

$$H_x = -\lambda\, S(1+k)e^{r(T-t)}H \tag{A4d}$$

for $x \leq x_-$, and

$$H_x = -\lambda\, S(1-k)e^{r(T-t)}H \tag{A4e}$$

for $x \geq x_+$.

Our valuation formulae for selling and buying values V_S and V_B simplify as follows. As before, V_S is defined by the equation

$$J^C(0,S,0,V_S) = J^0(0,S,0,0) \tag{A5}$$

which now can be expressed as

$$H^C(0,S,0)\exp\{-\lambda V_S\, e^{rT}\} = H^0(0,S,0) = -1, \quad \text{so}$$

$$V_S = \frac{1}{\lambda} e^{-rT} \ln\left(-H^C\right) \tag{A6}$$

Similarly, for the buying price V_B, we have

$$V_B = -\frac{1}{\lambda} e^{-rT} \ln\left(-H^{-C}\right) \tag{A7}$$

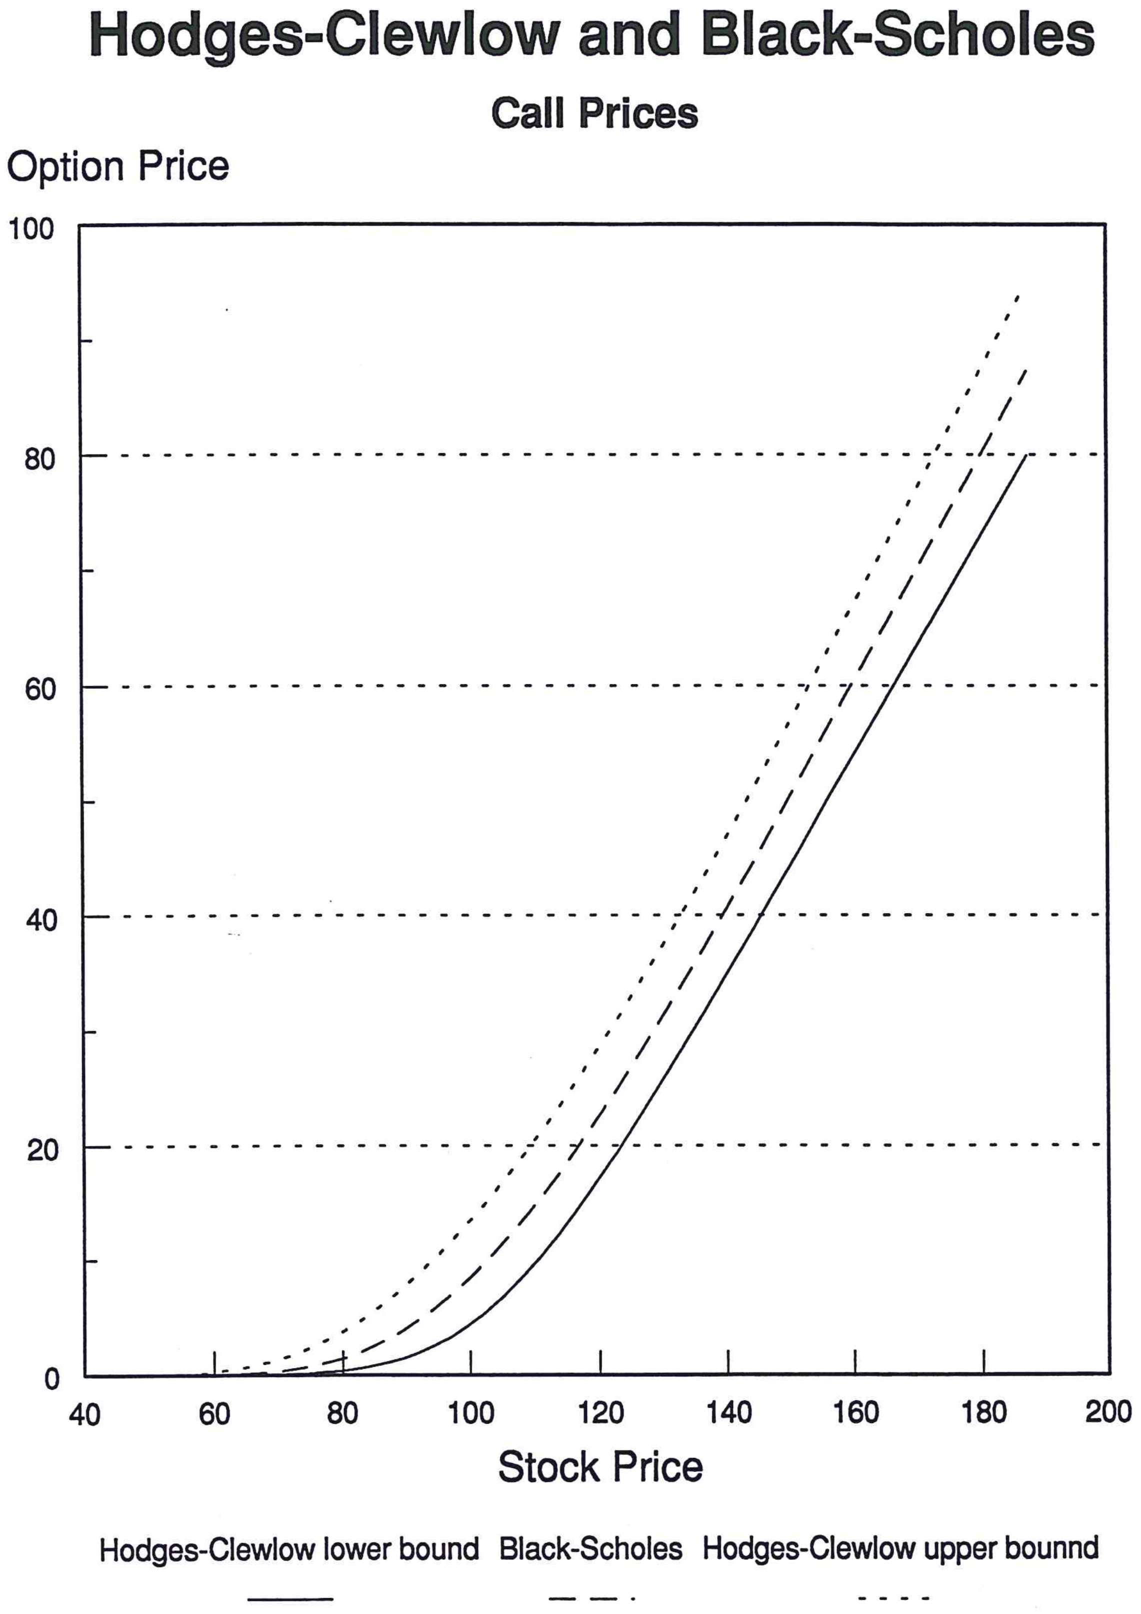
Hodges-Clewlow and Black-Scholes
Call Prices
Option Price
100
80
60
40
20
0
40
60
80
100
120
140
160
180
200
Stock Price
Hodges-Clewlow lower bound
Black-Scholes
Hodges-Clewlow upper bounnd

Figure 1

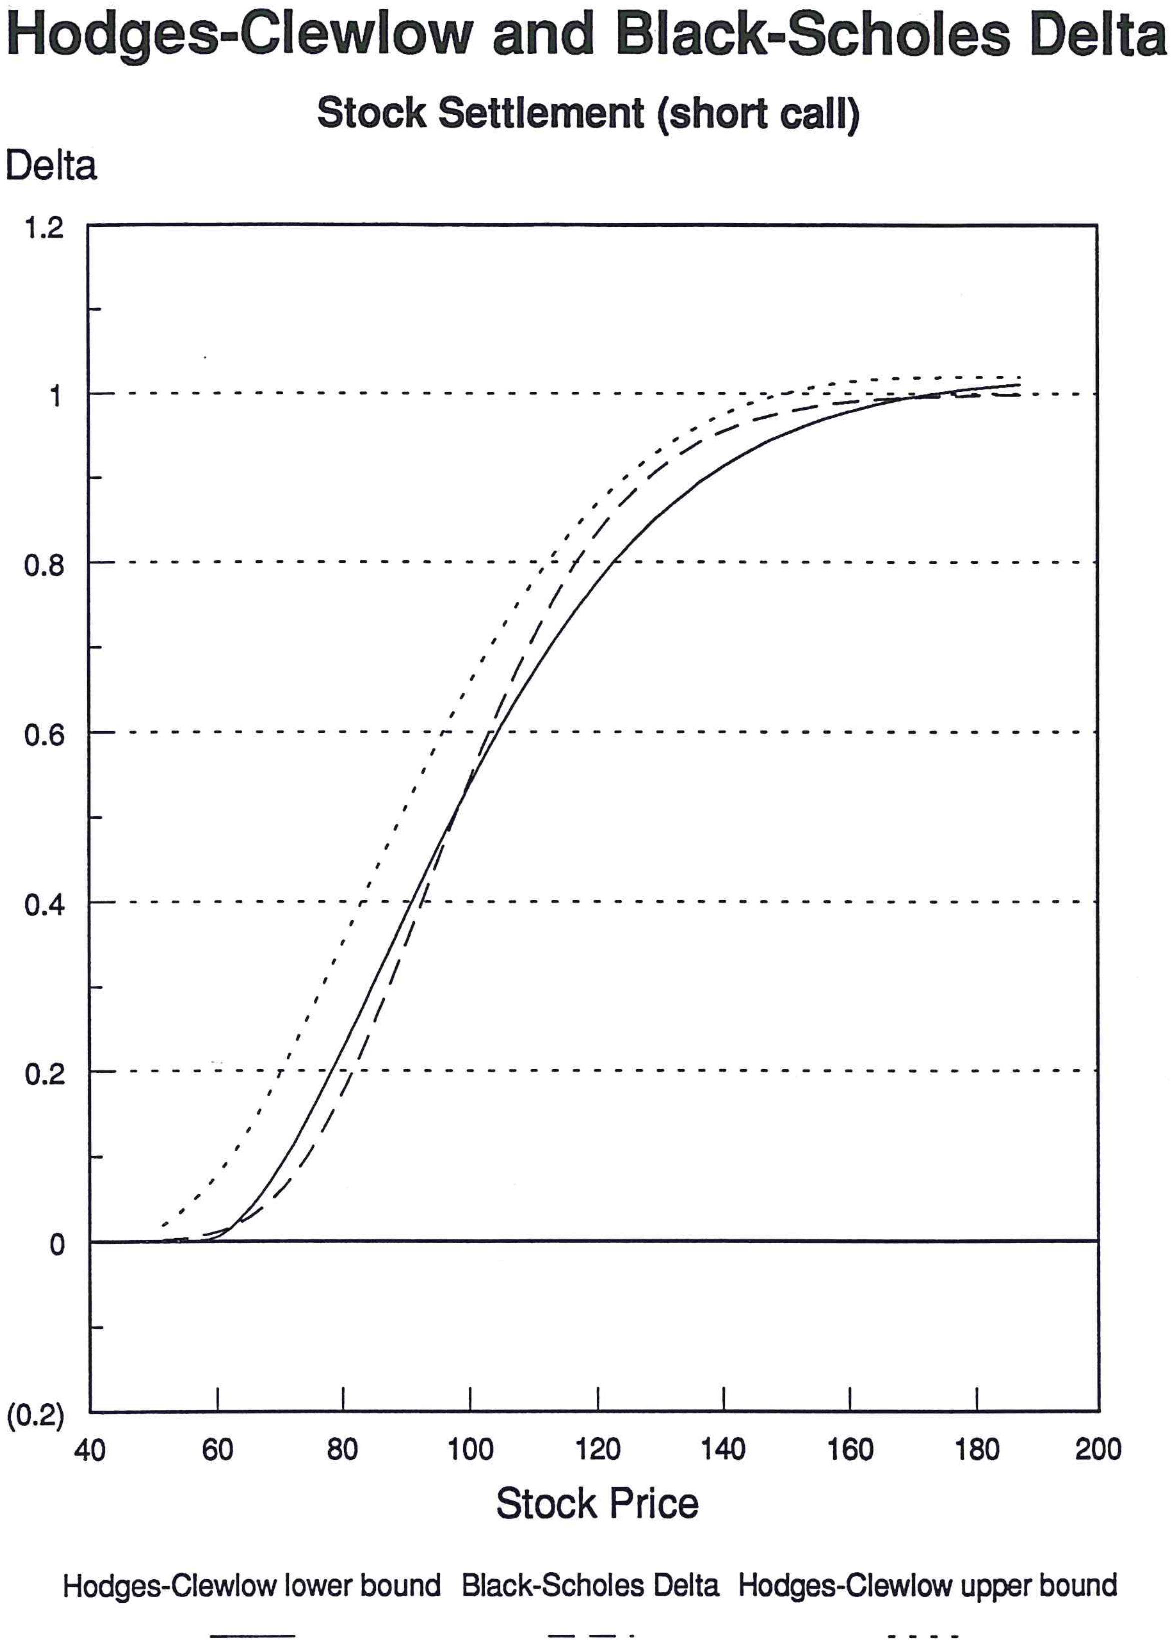
Hodges-Clewlow and Black-Scholes Delta
Stock Settlement (short call)
Delta
1.2
1
0.8
0.6
0.4
0.2
0
(0.2)
40
60
80
100
120
140
160
180
200
Stock Price
Hodges-Clewlow lower bound
Black-Scholes Delta
Hodges-Clewlow upper bound

Figure 2

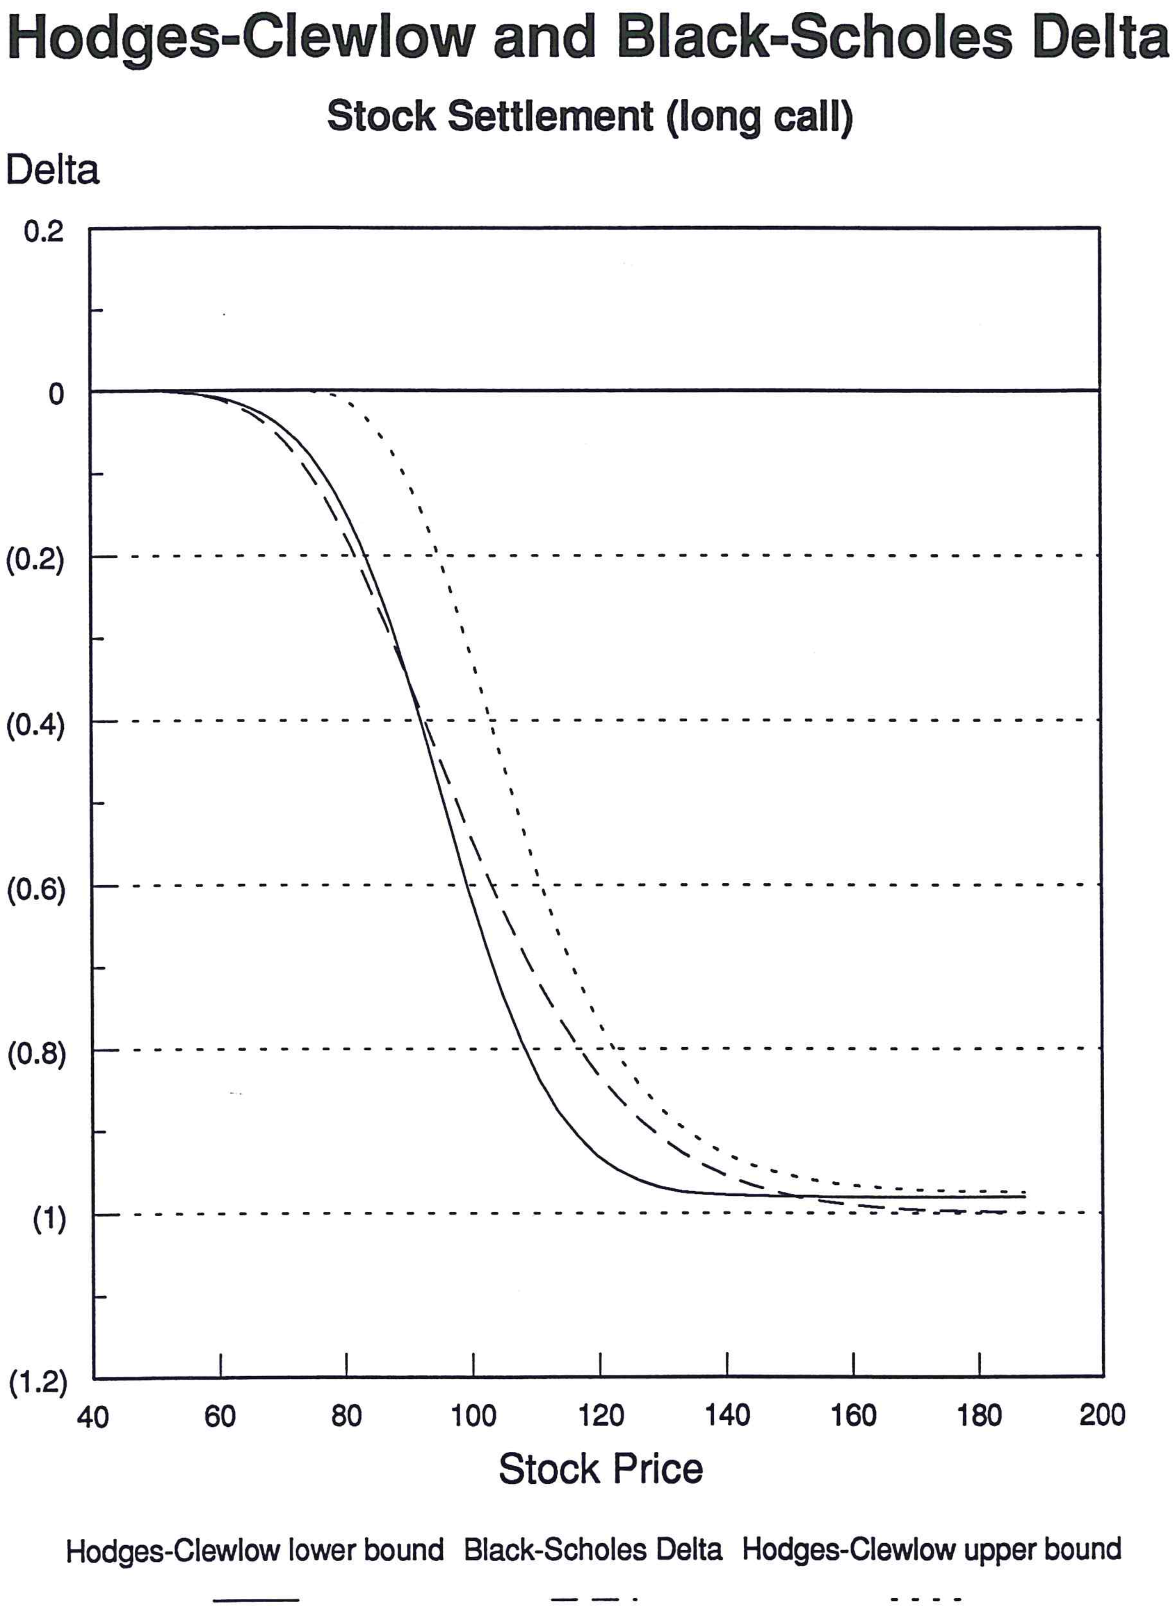
Hodges-Clewlow and Black-Scholes Delta
Stock Settlement (long call)
Delta
0.2
0
(0.2)
(0.4)
(0.6)
(0.8)
(1)
(1.2)
40
60
80
100
120
140
160
180
200
Stock Price
Hodges-Clewlow lower bound
Black-Scholes Delta
Hodges-Clewlow upper bound

Figure 3

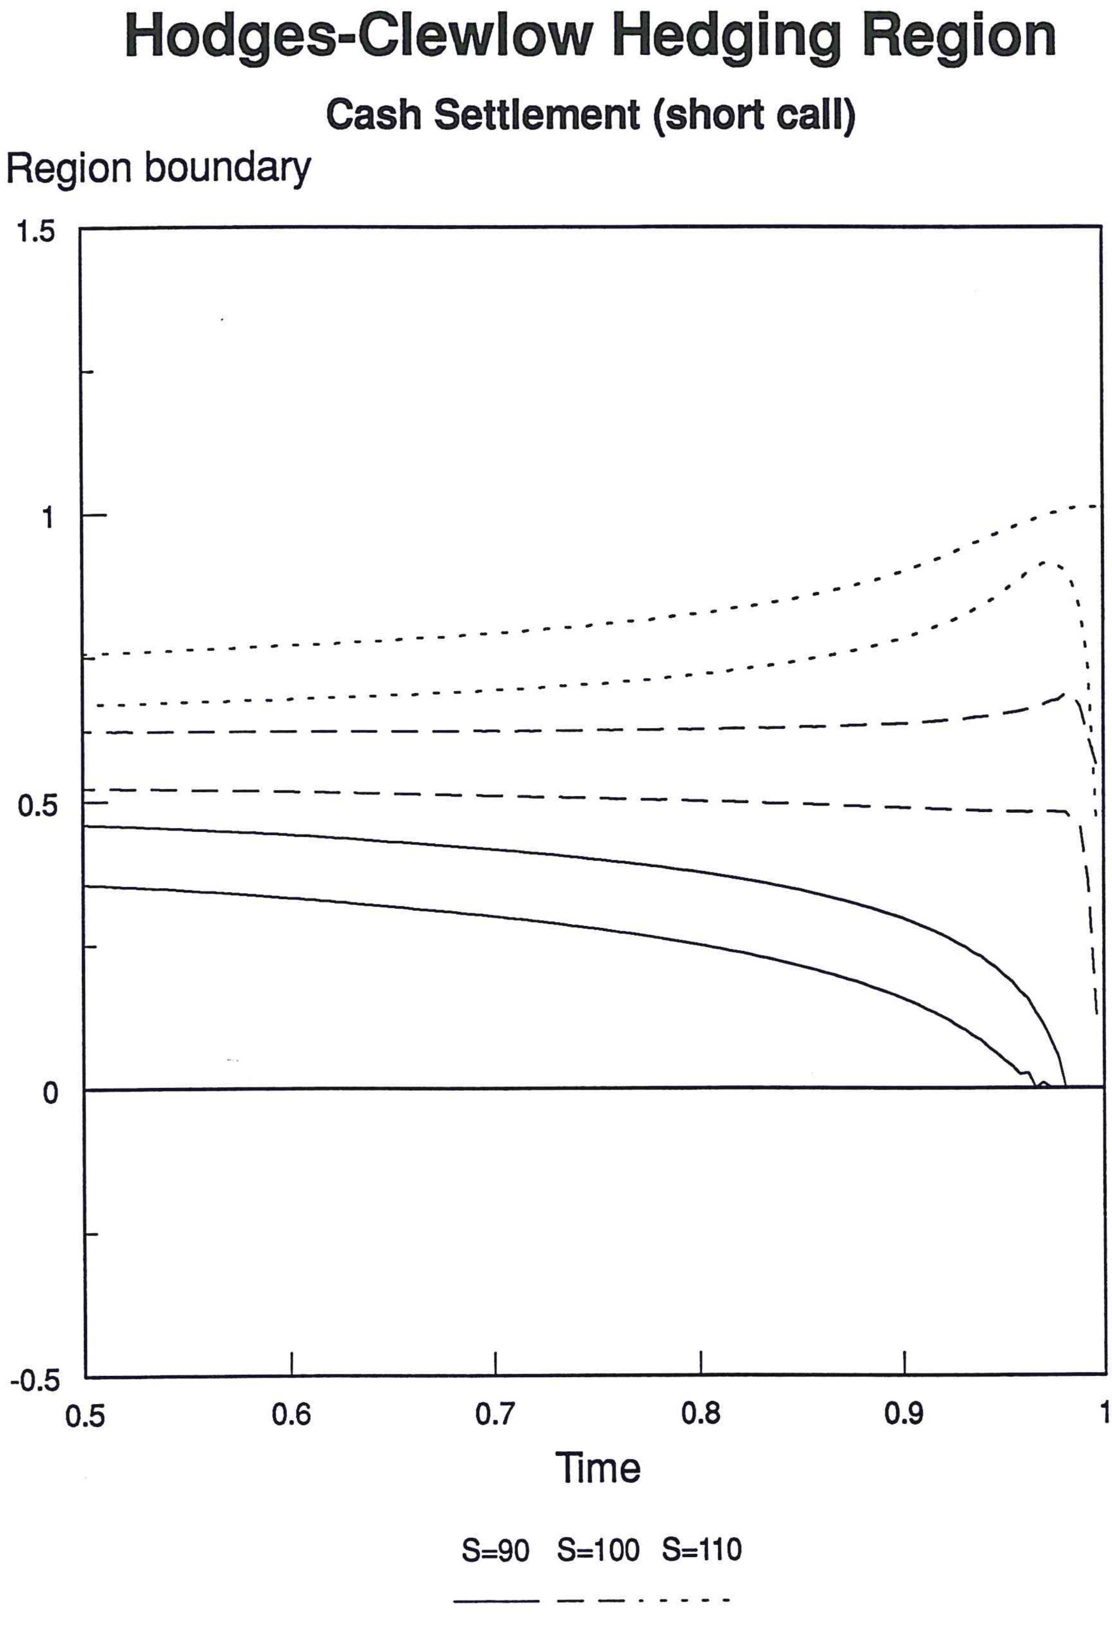

Hodges-Clewlow Hedging Region
Cash Settlement (short call)
Region boundary
1.5
1
0.5
0
-0.5
0.5
0.6
0.7
0.8
0.9
1
Time
S=90 S=100 S=110

Figure 4

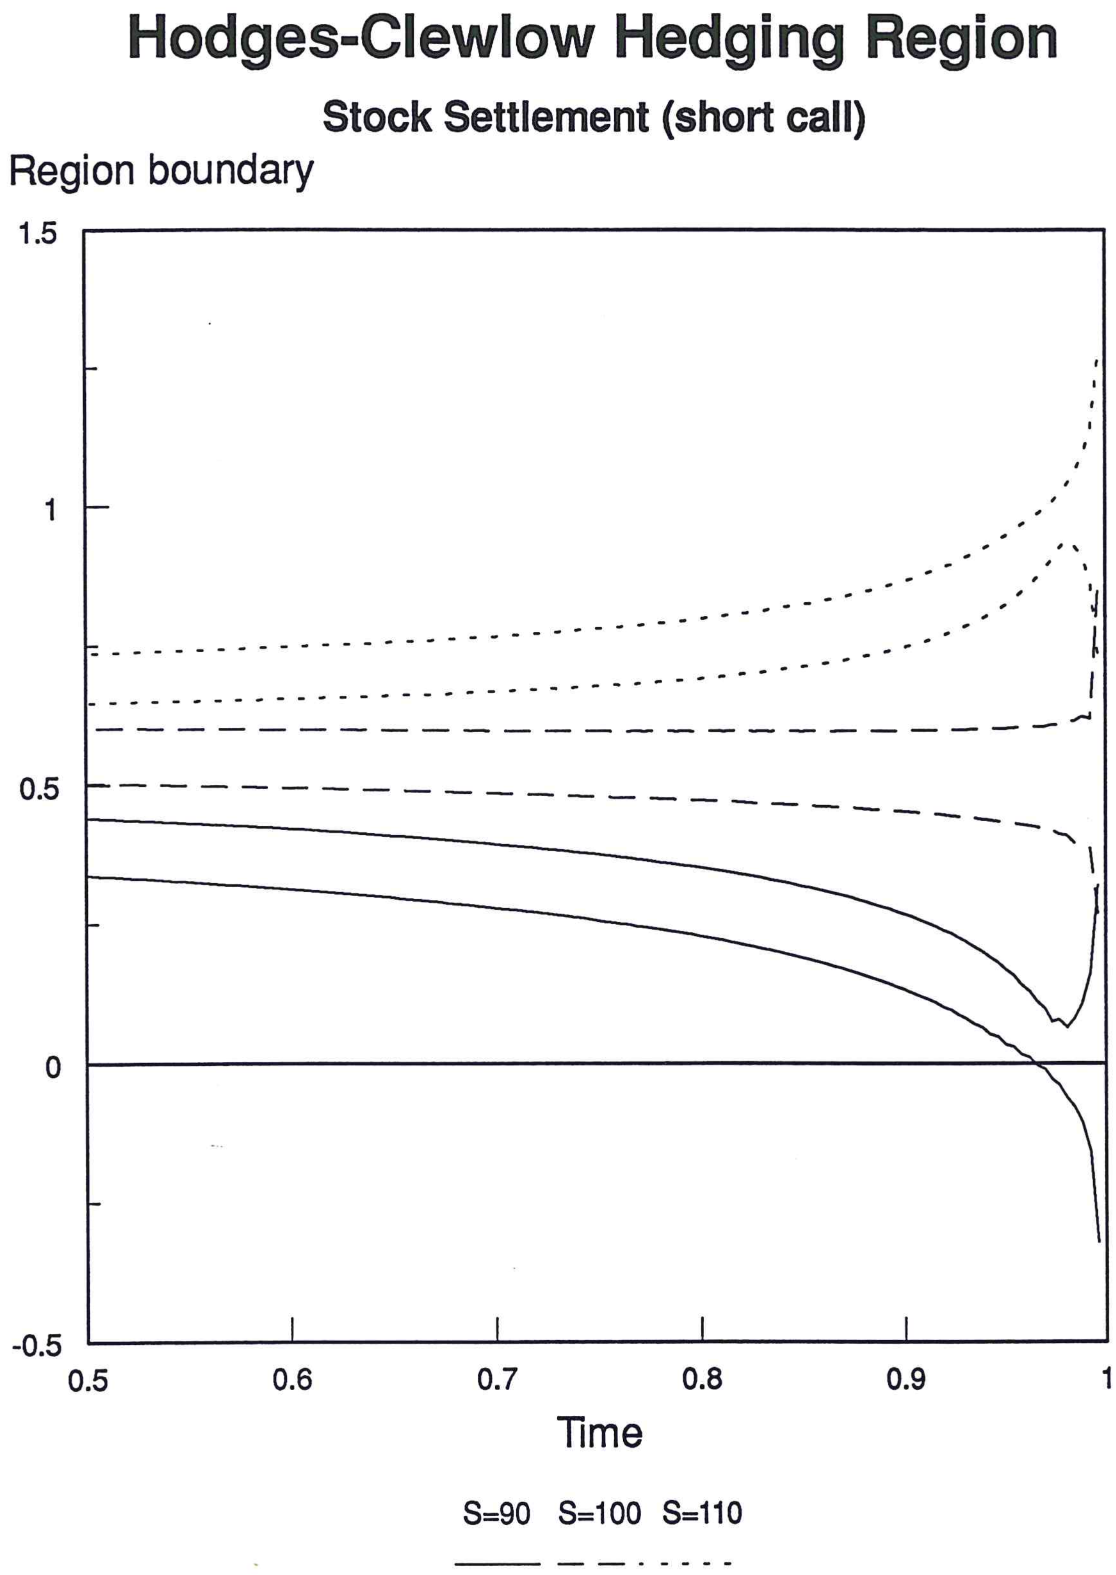
Hodges-Clewlow Hedging Region
Stock Settlement (short call)
Region boundary
1.5
1
0.5
0
-0.5
0.5
0.6
0.7
0.8
0.9
1
Time
S=90 S=100 S=110

Figure 5

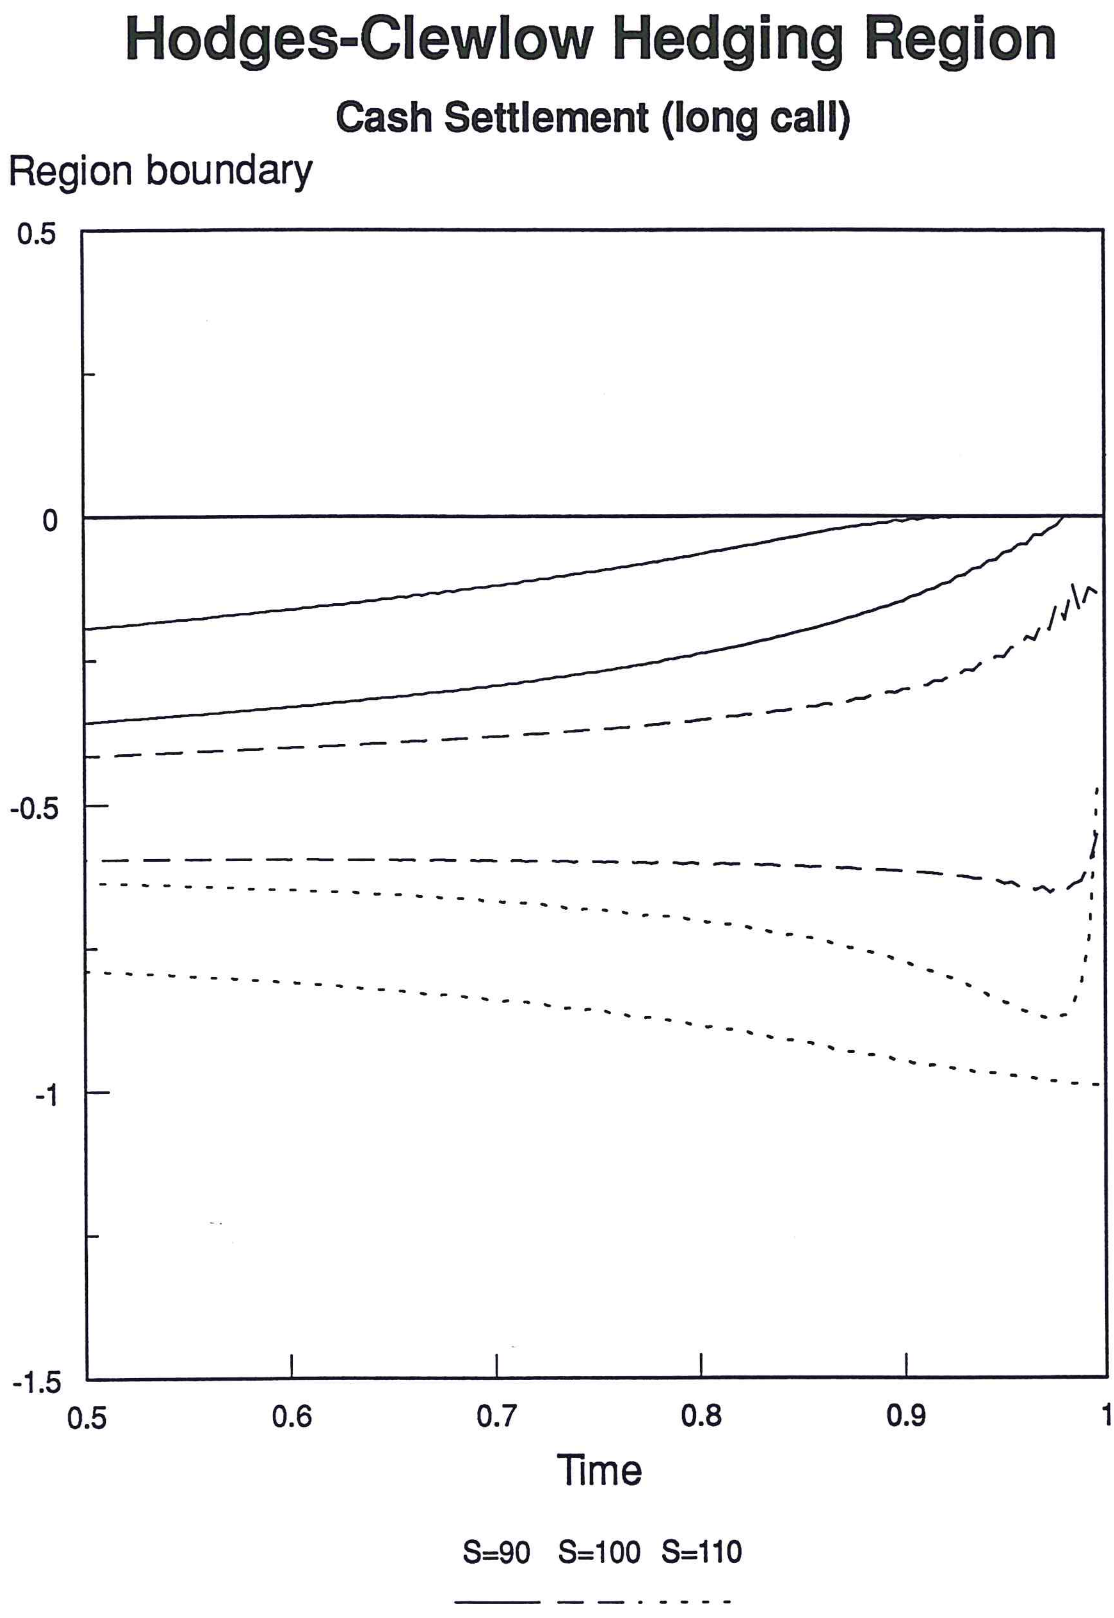
Hodges-Clewlow Hedging Region
Cash Settlement (long call)
Region boundary
0.5
0
-0.5
-1
-1.5
0.5
0.6
0.7
0.8
0.9
1
Time
S=90 S=100 S=110

Figure 6

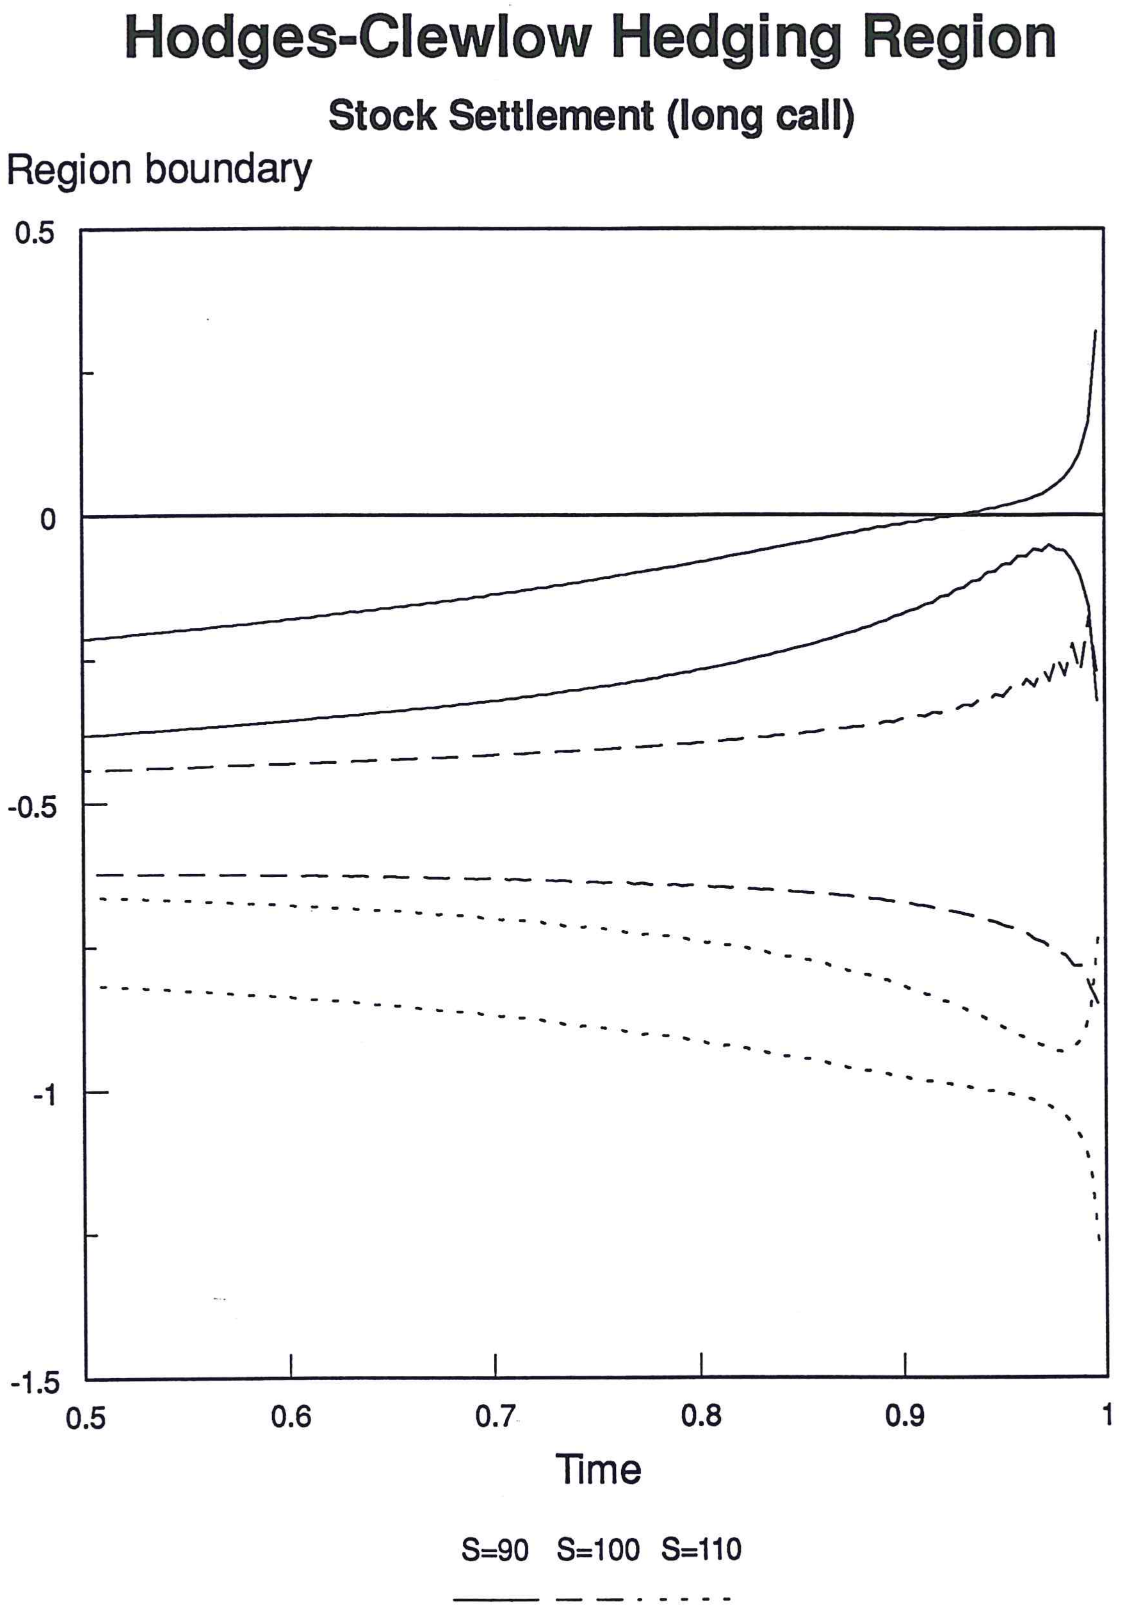
Hodges-Clewlow Hedging Region
Stock Settlement (long call)
Region boundary
0.5
0
-0.5
-1
-1.5
0.5
0.6
0.7
0.8
0.9
1
Time
S=90
S=100
S=110

Figure 7

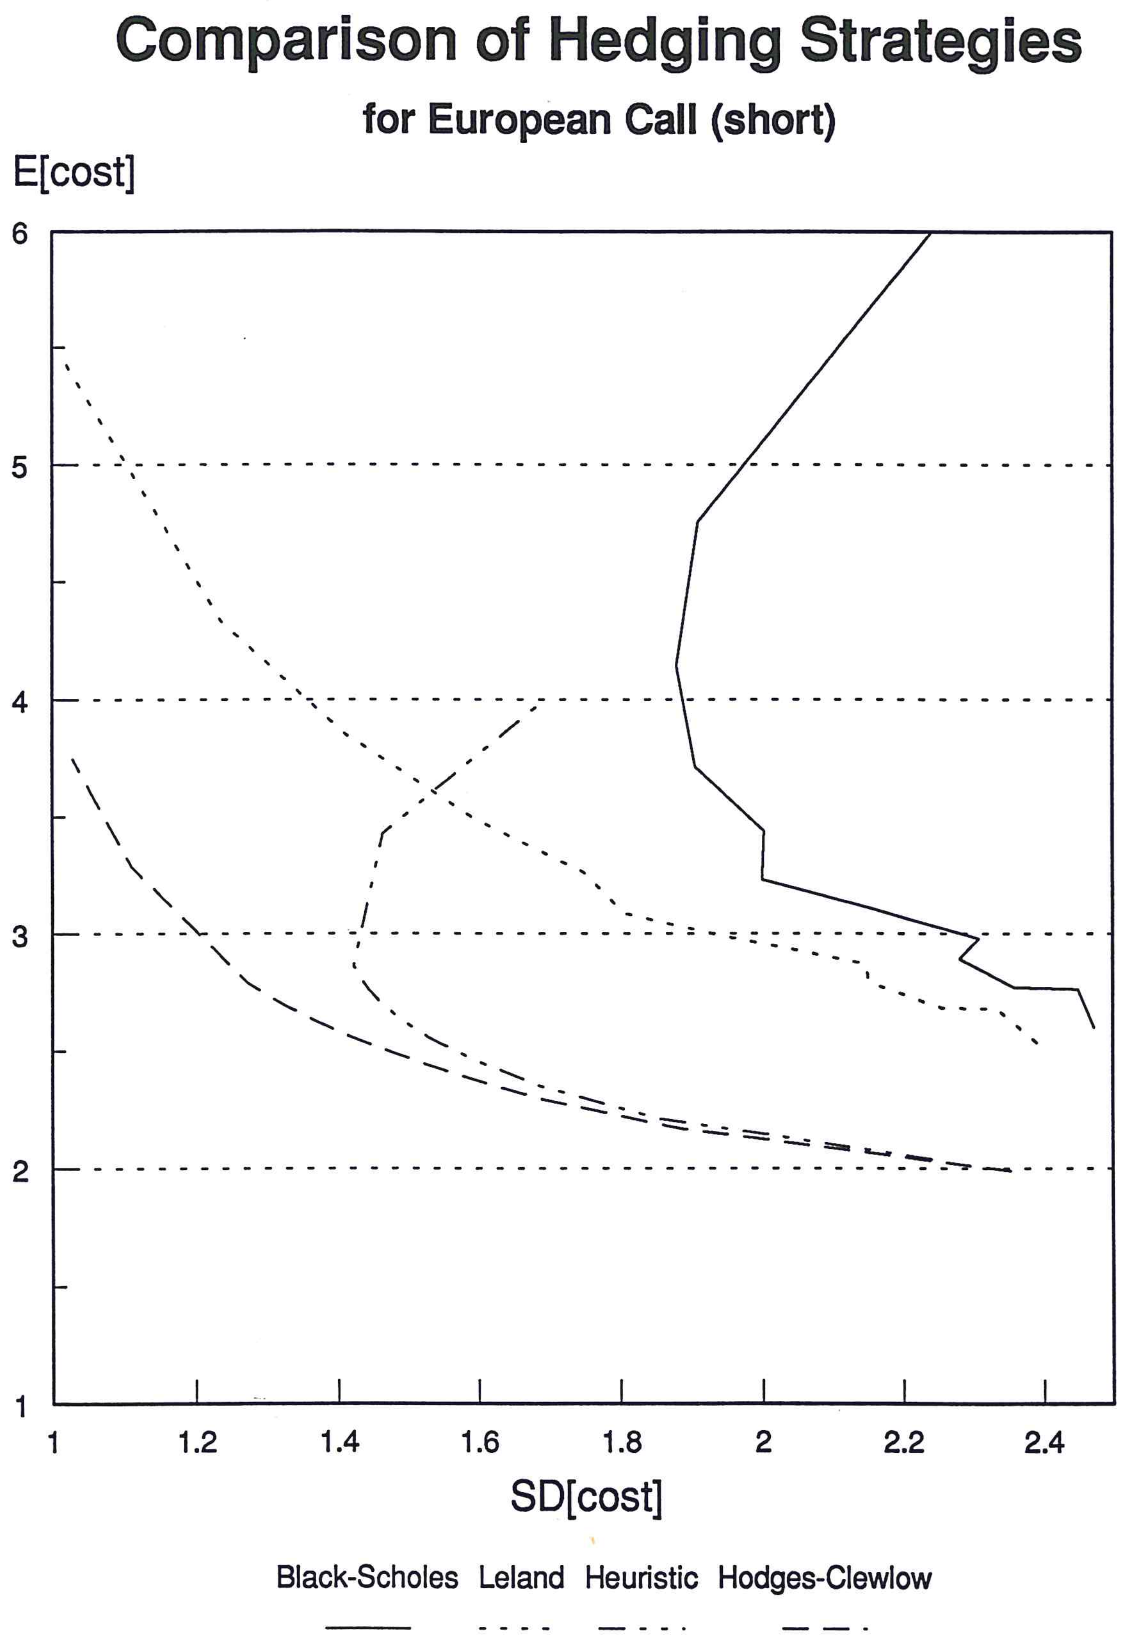

Comparison of Hedging Strategies
for European Call (short)
E[cost]
6
5
4
3
2
1
1
1.2
1.4
1.6
1.8
2
2.2
2.4
SD[cost]
Black-Scholes
Leland
Heuristic
Hodges-Clewlow

Figure 8

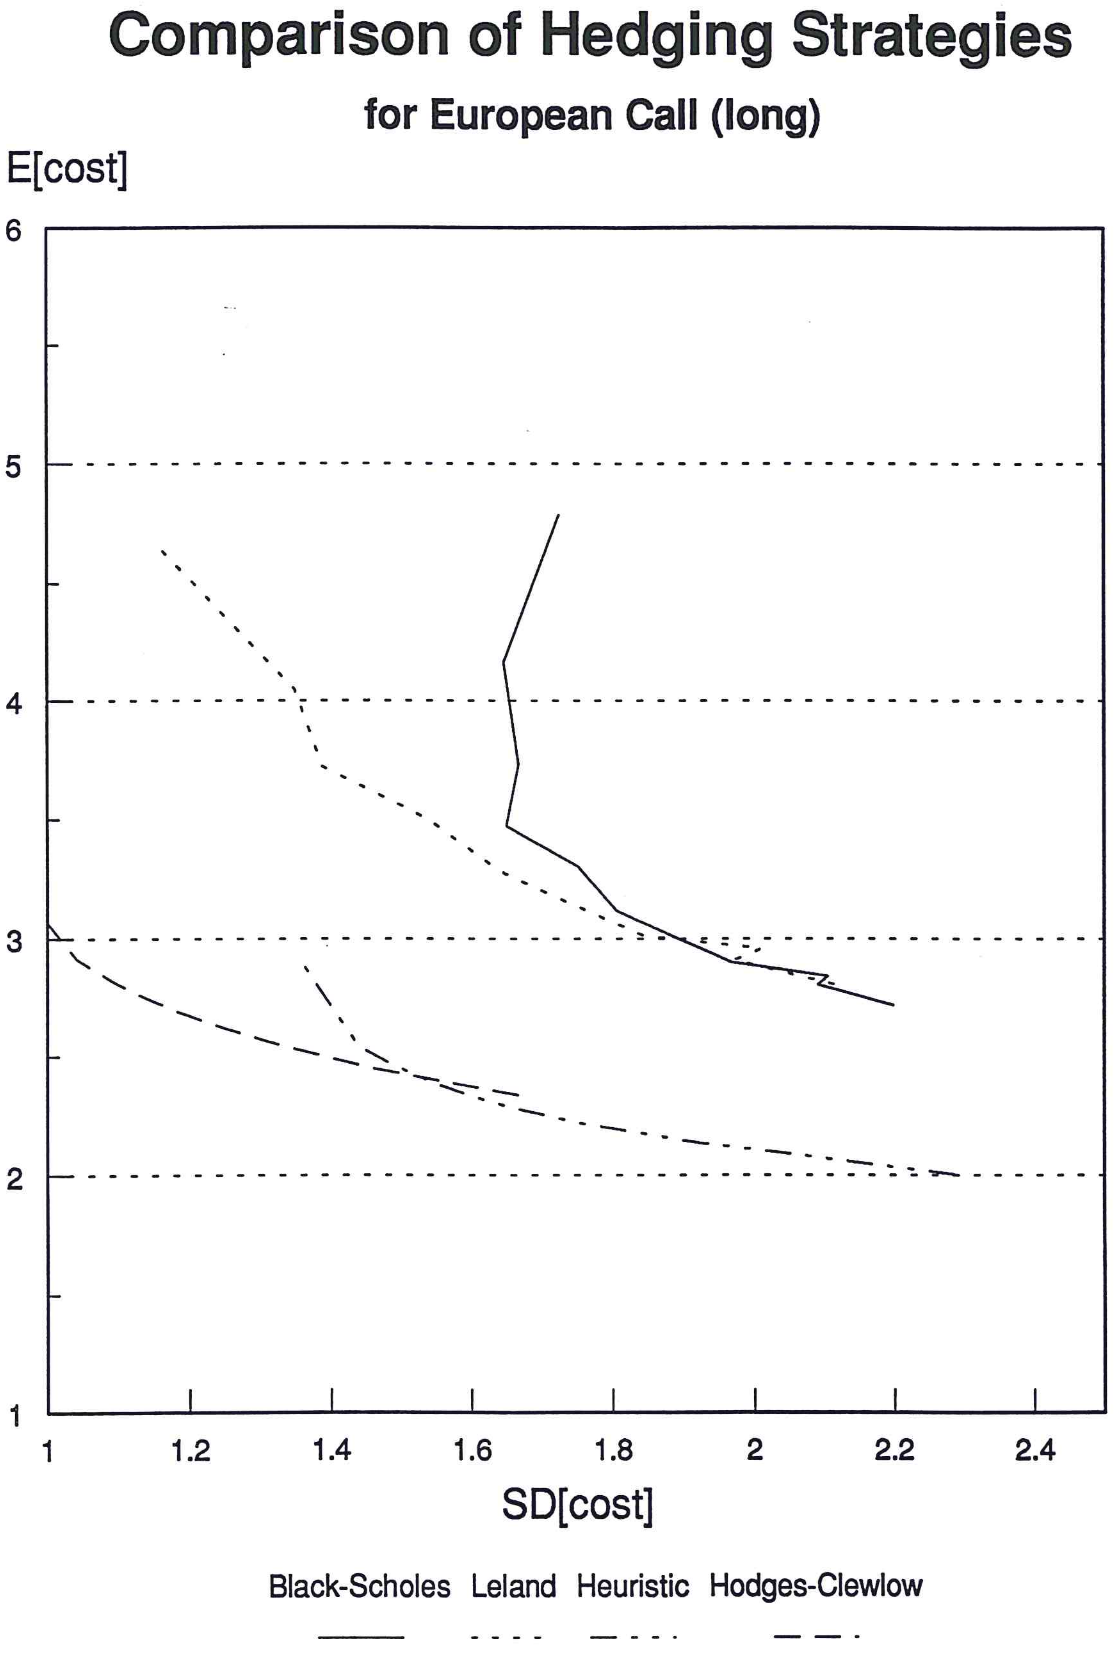
Comparison of Hedging Strategies
for European Call (long)
E[cost]
6
5
4
3
2
1
1
1.2
1.4
1.6
1.8
2
2.2
2.4
SD[cost]
Black-Scholes
Leland
Heuristic
Hodges-Clewlow

Figure 9

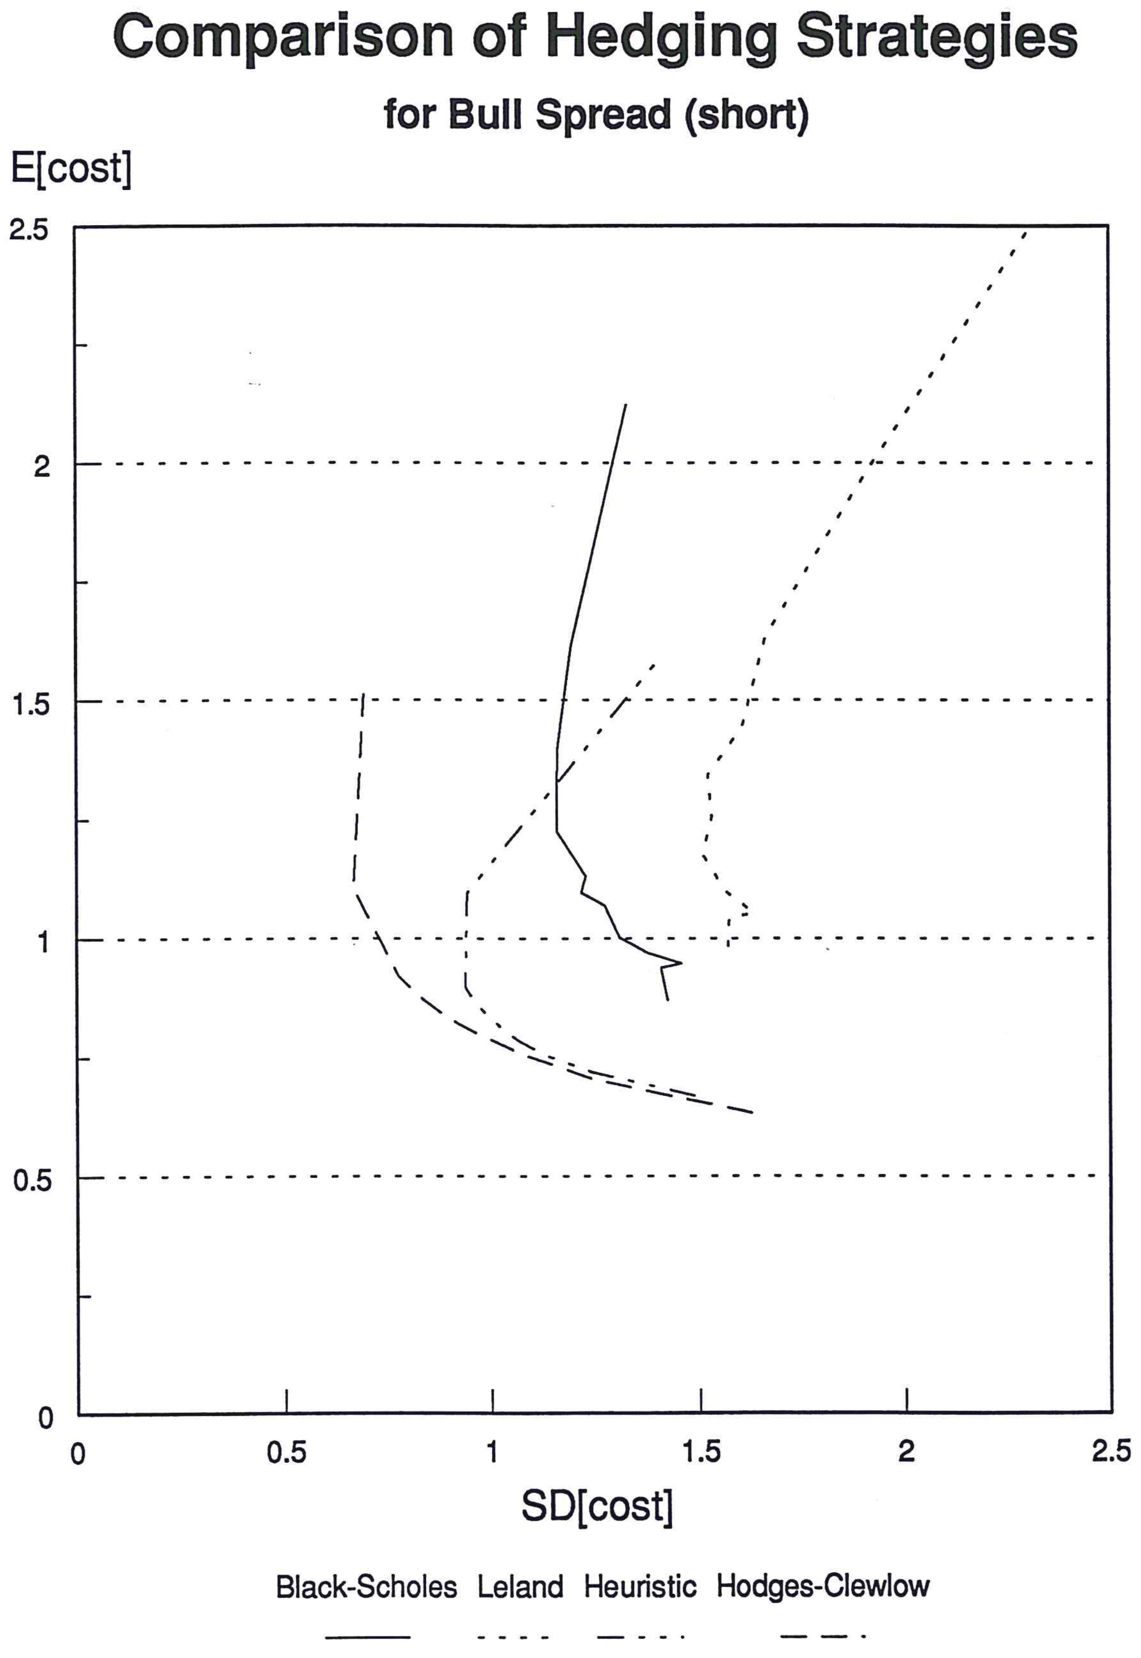
Comparison of Hedging Strategies
for Bull Spread (short)
E[cost]
2.5
2
1.5
1
0.5
0
0
0.5
1
1.5
2
2.5
SD[cost]
Black-Scholes Leland Heuristic Hodges-Clewlow

Figure 10

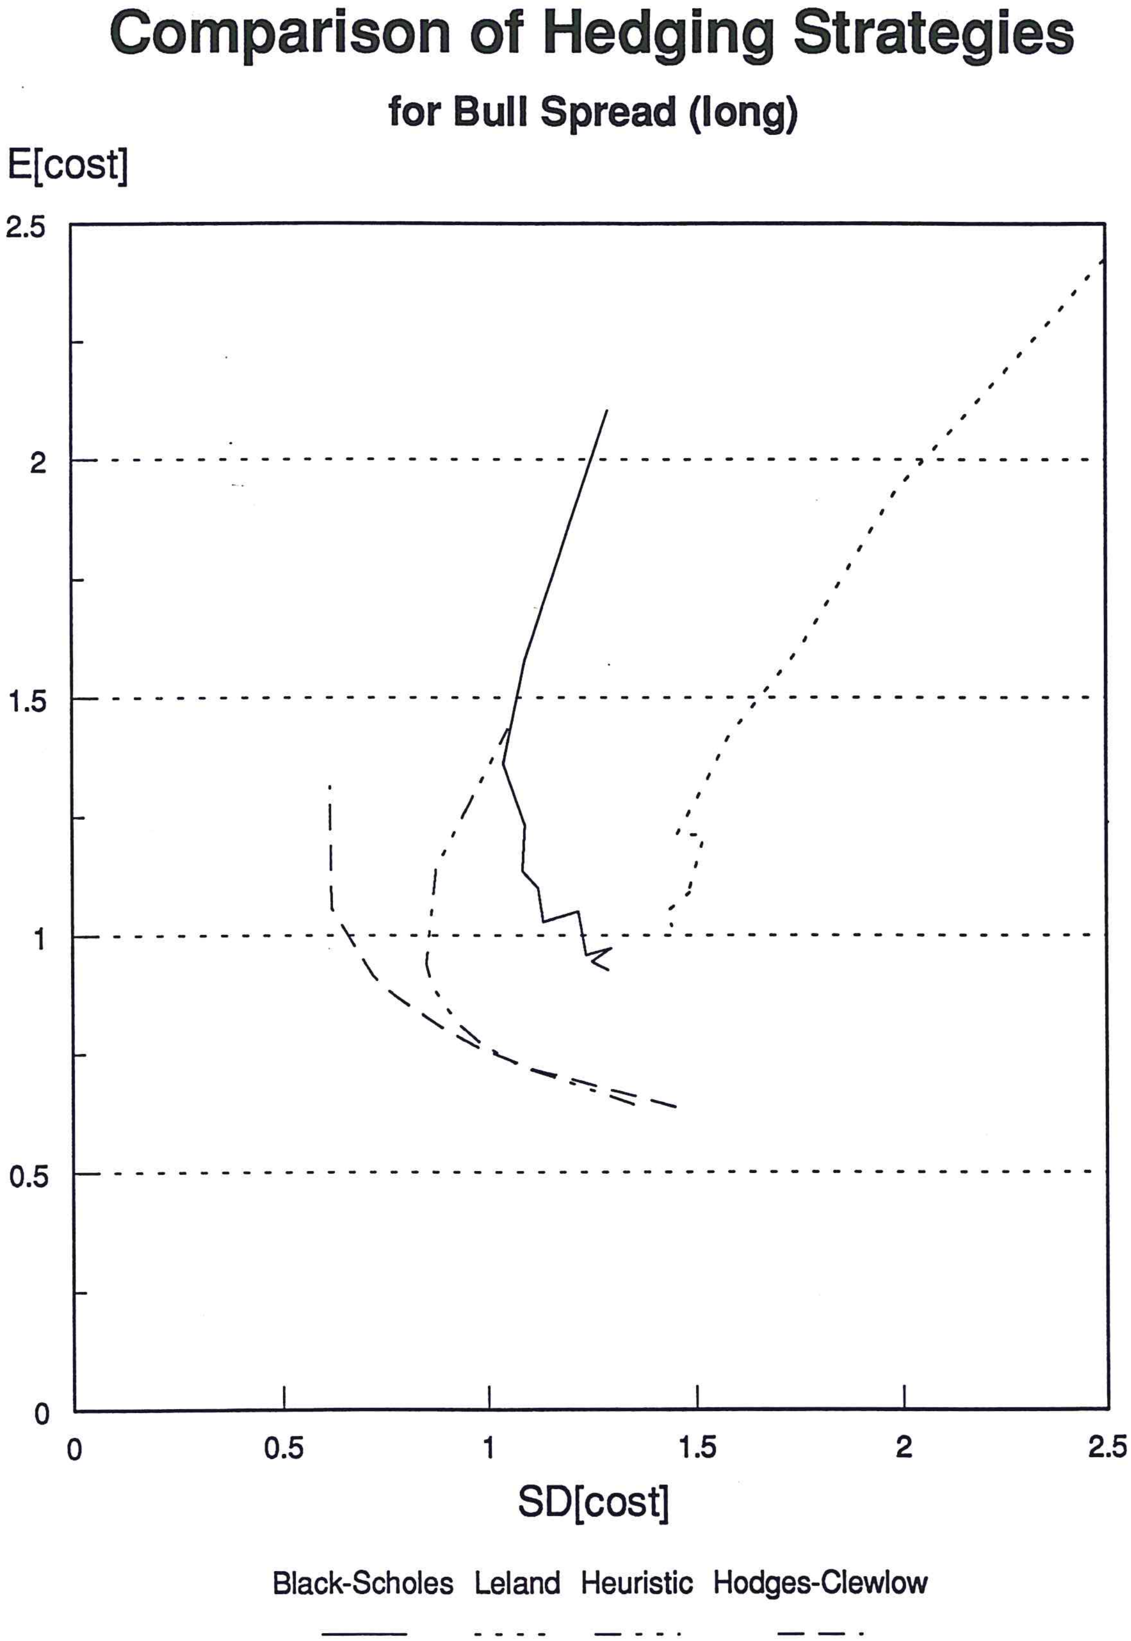

Comparison of Hedging Strategies
for Bull Spread (long)
E[cost]
2.5
2
1.5
1
0.5
0
0
0.5
1
1.5
2
2.5
SD[cost]
Black-Scholes
Leland
Heuristic
Hodges-Clewlow

Figure 11