

FINANCIAL OPTIONS RESEARCH CENTRE

University of Warwick

Rational Bounds on the Prices of Exotic Options

**Anthony Neuberger
and
Stewart Hodges**

November 1998

*Financial Options Research Centre
Warwick Business School
University of Warwick
Coventry
CV4 7AL
Phone: (01203) 524118*

FORC Preprint: 98/94

Rational Bounds on the Prices of Exotic Options

Anthony Neuberger and Stewart Hodges

London Business School and University of Warwick

August 1998

corresponding author: Anthony Neuberger
London Business School
Regent's Park
London NW1 4SA
United Kingdom
tel: +44-171-262-5050 x 3360
fax: +44-171-724-3317
email: aneuberger@lbs.ac.uk

Rational Bounds on the Prices of Exotic Options

Abstract

In this paper we provide a technique for pricing exotics relative to the instruments used for hedging them, while making minimal assumptions about price processes. The issue we address is this: given the prices of a set of hedging assets (such as a stock and a set of traded European options on that stock), what restrictions can be placed on the price of an exotic option? The question has a natural formulation as a linear program. We show how to get price bounds. We also show the corresponding robust strategies which enforce those bounds. The strategies allow agents to write exotics and put a floor on their losses in all states of the world. For some common exotics such as a digital barrier option and a lookback, we provide simple characterizations for the bounds and the hedging strategies. For others, notably the forward straddle, we show numerical solutions.

Rational Bounds on the Prices of Exotic Options

Introduction

Suppose that the price both of an asset and of a set of European options on that asset are known; what restrictions does that impose on the price of some exotic option? In the spirit of Merton (1973), we are searching for “rational bounds” on the price of the exotic option, that is bounds which are enforced by arbitrage. Since the bounds are enforced by arbitrage, no assumptions are made about the price dynamics of the securities. All we assume is that securities can be traded at appropriate dates, and there are no market frictions.

The question is interesting for several reasons. First, arbitrage bounds are closely relating to hedging. Given a rational bound, there is a hedging strategy which enforces it. The hedging strategy may not be very *precise*; the profits from selling the exotic and hedging it may be rather volatile. But the hedge will be *robust*; the strategy puts a floor on the maximum loss. Contrast this with a typical model-based hedging strategy. If the model is correct, if asset price volatilities conform with the model’s predictions, if the instruments used for dynamic hedging can be bought and sold at model prices, then the hedge will lock in precisely the difference between the price of the exotic and its theoretical value. But if the model is false, the hedging strategy may lead to virtually unlimited losses.

A second reason relates to the reliability of option pricing models. In general, given the market prices of a set of European call options, it is possible to construct a pricing model which is consistent with them; see for example Derman, Ergener and Kani (1995). Such models are widely used to price exotics. The fact that the model correctly prices options which are traded may lead users to place great reliance on the model’s estimate of the fair price of other securities. Yet if the rational bounds on the price of the exotic option are very wide, this should engender a

sense of caution in the use of the model. Conversely, if the bounds are narrow (suggesting that the exotic is in some sense close to the space spanned by the set of reference assets) then one would use the model prices, and its hedge ratios, with much greater confidence.

A third reason for being interested in rational bounds is that they draw attention to the best and worst case scenarios for a particular option, and thus cast light on the sorts of beliefs or fears one might have which justify the purchase or sale of an exotic option rather than a conventional option. For example, an up-and-out call appears to be quite similar to a conventional call less a second call with a strike close to the barrier. By looking at the asset price processes which justify the bounds one can see the kind of processes which make the exotic particularly valuable relative to the European options.

We show how rational bounds can be found by searching for the least costly strategy which dominates the claim using the techniques of linear programming (LP). The use of LP has two advantages. First, it provides a well understood technique for getting a numerical solution in specific cases. Second, the LP approach provides a useful way of thinking about the problem. The saddle point theorem of LP suggests powerful techniques for finding an analytic solution by moving between the primal and the dual problem. In this case the primal problem is the search for the cheapest super-replicating strategy. The dual problem is also of economic interest; it is the search for the martingale process, consistent with the prices of the European options, which maximizes the value of the exotic. We will make much use of the duality in the present paper.

Related Literature

Carr, Ellis and Gupta (1998) address a rather similar question; they show how a variety of exotic options - mainly barrier options - can be hedged by holding portfolios of European options. Their analysis is not a rational bounds analysis. Crucially, they assume something they name put-call symmetry. In effect, they assume that at all times the expected future volatility of the forward price of the asset is symmetrical about the current forward price. This ensures that at any future date, two puts or calls with the same maturity and whose strikes have a geometric mean equal to the forward rate will trade on the same implied Black-Scholes volatility. This assumption, for

which there is no obvious theoretical justification, is important because it fixes the rate of exchange between two options with different strikes at some point in the future when the forward price reaches a barrier. It enables the authors to arrive at simple pricing results, enforced by hedging strategies which involve very limited trading. Any hedging error will then be related to future deviations in put-call symmetry.

Another approach to the issue which falls between model building and rational bounds is that of Cochrane and Saá-Requejo (1996), Hodges (1997) and Bernardo and Ledoit (1997) where prices are such as to exclude not only arbitrages but also excessively good deals (where a good deal is defined to be one where the Sharpe ratio of return to risk in the first two cited papers, or the ratio of gain to loss in the last is too high). This approach has obvious attractions. If rational bounds are very wide because of some extreme scenarios which are most unlikely to occur, it is useful to get tighter bounds which hold with high probability. However, the approach does require the user to specify a benchmark model, and this is itself quite subjective.

Ahn, Muni and Swindle (1997) show how to find hedging strategies which are robust in the sense that they maximize expected utility in the worst case outcome from a restricted set of volatility scenarios; our approach can be seen as an extreme form of robustness for an agent who is very risk averse and where volatility is unrestricted. The work of Christofedes and Salkin (1996) also bears some relationship to our analysis. They employ an approach based on a graph theory algorithm. Our martingale probabilities have similar “conservation of flow” properties, but other aspects of the formulation in their paper are rather different from ours.

Hobson (1997) finds rational bounds on the price of a lookback option. Our results on this are similar to his. However the focus of the two papers is different; his primary interest is in finding rational bounds for the price of a specific exotic option while we are interested in a general approach to the pricing of exotics, and also in the hedging strategies which support rational bounds.

The use of LP formulations to find super-replicating strategies is not novel (see for example Edirisinghe, Naik and Uppal, 1993). The contribution of this paper is twofold. First, we deal with

super-replication under very weak assumptions. Rather than make assumptions concerning particular stochastic processes we focus on the implications of having market prices for a range of other instruments including conventional call options. Second, within this framework we are able to derive useful valuation bounds for a number of exotic options, and suggest numerical schemes for others.

Structure of the Paper

In the first section, we give a simple numerical illustration of our approach. Then in Section 2 we show how the general problem can be formulated as a linear program. Section 3 uses LP to get solutions for a specific exotic option (the forward straddle). Section 4 shows how LP provides the critical insights to characterize the bounds on the prices of some common exotics (a digital barrier option and a lookback option), and identifies the corresponding bounds on the price of these options. Section 5 concludes.

1. A Simple Example

In this section we illustrate our general approach by means of a numerical example. There is a risky asset, which we call a share. It is priced in whole numbers of dollars. Its current price is \$10. Interest rates are zero. The price next month will be in the range \$9 - \$11, and the following month \$8 - \$12. There are European call options with strikes of \$9, \$10 and \$11 maturing in 1 and 2 months. The exotic option to be priced is a forward straddle. This has a pay-off $|S_2 - S_1|$.

The prices of the European call options are as follows:

Maturity	1 month	2 months
Strike (K)		
\$9	\$1	\$17/16
\$10	\$1/4	\$3/8
\$11	0	\$1/16

These prices are chosen for arithmetical convenience. They do not allow arbitrage, but otherwise there is nothing special about them. They are consistent with a world where investors are risk neutral and believe that the asset price follows an arithmetic trinomial process, with probabilities of 25% of moving up or down by \$1 each month, and 50% of staying unchanged. In such a world, the value of the forward straddle is just \$0.50.

Clearly then \$0.50 is a feasible price for the forward option, but if the option traded at some other price, would there be an arbitrage? What would the arbitrage strategy be?

Consider the following strategy:


- buy two 2 month calls and write two 1 month calls at each strike;
- after 1 month, short 1, 3 or 5 forward contracts depending on whether S_1 is \$9, \$10 or \$11.

The cost of the 1 month contracts is \$2.50 and of the 2 month contracts is \$3.00, so the net cost of the strategy is thus \$0.50. It is easy to verify that the strategy dominates the exotic on all possible paths. Since the price of \$0.50 is a feasible valuation and since any higher price would give rise to arbitrage, it must be the least upper bound on the price of the exotic.

But it is not the only feasible price. Consider the following process:

at the end of month 1 the price goes to \$9, \$10 or \$11 with probabilities 0.25, 0.5 and 0.25 respectively;

if the price at the end of month 1 is either \$9 or \$11, the price remains constant over the next month; if it is \$10 it goes to \$8, \$10 or \$12 with probabilities 0.125, 0.75 and 0.125 respectively.



The process is consistent both with risk neutrality (it is a martingale) and with the European call option prices. It leads to a value for the forward option of only \$0.25.

\$0.25 is actually the lower bound on the value of the forward straddle. To see this, consider the strategy of buying a one month straddle, and writing a two month straddle, both with a strike of \$10. The one month straddle, which consists of a put and a call with the same strike, costs \$0.50; the two month straddle gives \$0.75. So establishing the portfolio generates cash of \$0.25. The net liability is dominated by the straddle:

$$|S_2 - S_T| \geq |S_2 - \$10| - |S_T - \$10|$$

So the bounds on the exotic are \$0.25-0.50; outside those bounds, there is an arbitrage, while inside there is none. This example illustrates a number of points which we will demonstrate formally and with greater generality:

- we can find the upper bound on the value of an exotic either by looking for the cheapest super-replicating strategy or by looking for the process which maximizes the value of the exotic. Similarly for the lower bound.
- even for apparently innocuous looking options, the rational bounds may be quite wide; we show in section 3 that the large spread between the upper and lower bound on the forward straddle is real, and not due to the coarse grid we have used.

- the bounding processes hint at the key determinants of the value of the exotic given the prices of European call options. In our example, the process which generates the upper bound has the stock price moving at most one node at a time, while the lower bound process has the stock price either constant or making large jumps. Pushing the analysis, this suggests that, given option prices for times t_1 and t_2 , the value of a forward straddle will be sensitive to the degree of kurtosis over the period (t_1, t_2) - the higher the kurtosis the lower the value of the straddle. This interpretation is confirmed by the analysis in section 3 below.

The example was simple enough that one might hope to find the binding strategies and processes by guesswork. In the more general case a more systematic approach is needed and for this we turn to linear programming.

2. The General Problem

Our objective is to find arbitrage bounds on the prices of exotic options where a variety of instruments, which we call “reference assets”, can be used for hedging. These reference assets can be thought of as European options. In this section we formalize the problem and set it up as an LP. We show how the formulation can be used to get insights into the problem.

We consider the problem of super-replication in a finite-period setting where the underlying asset takes one of a finite number of prices at each date. For simplicity we assume that the interest rate is constant. By using a bank account as numeraire we can, without further loss of generality, also assume that the interest rate is zero.

The dates in our model are $t = 0, 1, \dots, T$. The asset price takes values in the set $\{K_i \mid i = -N, \dots, +N\}$. By convention the initial price, at time 0, is K_0 . A *state* or a *price path* ω is a $T+1$ vector whose $t+1$ 'th element, an integer in the range $[-N, +N]$, indicates the asset price level at time t . The first element of the vector is 0 since the initial price on any path is K_0 . The number of paths is denoted by W , which is equal to $(2N+1)^T$.

A *history* h is a vector of integers of any length $t \leq T$, where the first element is 0 and the remaining elements are integers in the range $[-N, +N]$. A path ω has a history h (written as $\omega \in h$) if the first t elements of ω coincide with h . The number of distinct histories is denoted by H (which is equal to $\{(2N+1)^T - (2N+1)\}/2N$).

The *exotic option* \mathbf{b} is defined by specifying its pay-off on every possible path. By imposing arbitrary order on the paths we can characterize the exotic by having \mathbf{b} to be a real-valued column vector of length W where the w 'th element is the pay-off to the exotic on the w 'th path.

A strategy which replicates the exotic option will consist of an initial portfolio position, and subsequent dynamic trading. We assume that at any time the hedger can hold cash, one period forward contracts in the underlying, and positions in a set of J reference assets. Allowing longer dated forward contracts adds nothing to the strategy space since there is no interest rate or dividend uncertainty. As we will prove, it also turns out to be unnecessary to consider strategies where the agent trades the reference assets. A *strategy* can then be characterized by a column vector \mathbf{x} of length L where $L = 1 + H + J$. \mathbf{x} has the following interpretation:

1. The first element is the amount invested initially in the risk free asset
2. The next H elements specify for each possible history h how large a forward position the agent holds. These positions have no initial cost. But if at some time t the history is h , and if the corresponding element of the strategy vector is x_h , then the agent will go long x_h one period forward contracts at time t
3. The last J elements specify the number of units of each of the reference assets are held in the hedge portfolio.

The problem is to find the minimum cost strategy \mathbf{x} which dominates the exotic. The problem can be readily expressed as an LP. The cost of the strategy is $\mathbf{c}'\mathbf{x}$ where \mathbf{c} is a $L \times 1$ vector whose first element is 1 (because the cost of buying 1 unit of the risk free asset), whose next H elements are zero (there is zero cost to going long a forward contract) and whose last J elements are the prices of the J reference assets.

The strategy must dominate the exotic. To express this formally, we need to specify the pay-off to each component of the strategy on each path. This we do using a $W \times L$ matrix, denoted by A . The first column gives the pay-off to an investment of 1 in the risk free asset, and so is just a column of 1's. The next H columns give the pay-offs to the forward contracts. Consider the column which corresponds to a forward contract taken on when the history is h , the vector h being of length t . If the actual path is ω , the pay-off is:

$$\begin{cases} S_{t+1} - S_t & \text{if } \omega \in h \\ 0 & \text{otherwise} \end{cases}$$

The last J columns then show the pay-offs to each of the J reference assets on each path. The pay-off to the strategy x is then given by the $W \times 1$ vector $A \cdot x$. This must dominate the pay-off to the exotic, b , in all states of the world.

Thus the LP can be written as choosing the column vector x so as to:

$$\begin{aligned} &\text{minimize} && \mathbf{c}' & \cdot & \mathbf{x} \\ &&& 1 \times L & & L \times 1 \\ \text{subject to:} && \mathbf{A} & \cdot & \mathbf{x} & \geq & \mathbf{b} \\ && W \times L & & L \times 1 & & W \times 1 \end{aligned}$$

(the vector x is unrestricted).

It is useful to establish when the LP has a solution.

Proposition 1: *The problem has a finite solution if and only if the asset price and the set of reference assets do not allow arbitrage.*

Proof: We prove two things: that the feasible set is not empty, and that there is a finite minimum to the objective function if and only if there is no arbitrage. Let b_{max} denote the largest element of the vector of pay-offs to the exotic. A strategy vector x whose first element is b_{max} and whose other elements are zero is feasible - an exotic is dominated by a bank account which contains as much as the maximum possible to pay-off to the exotic. So the feasible set is not empty. Note

that this result does not hold true necessarily in the limit as the range of possible stock prices becomes infinite.

Suppose now that no finite solution to the LP exists. Then there must be a feasible strategy which costs less than $-b_{max}$ which dominates the exotic. Consider a new strategy which is identical apart from increasing the investment in the risk free asset by b_{max} . This has strictly negative cost, and the pay-off is non-negative in all states of the world, so it is an arbitrage. The converse is also easily demonstrated.

Thus, provided that the prices of the reference assets do not allow arbitrage, there is always a least cost super-replicating strategy for any exotic. We now make use of the LP Duality Theorem to explore the nature of the problem more deeply. The dual problem is to choose the $W \times 1$ column vector π so as to:

$$\begin{aligned}
 &\text{maximize} && \mathbf{b}' & \cdot & \boldsymbol{\pi} \\
 & && 1 \times W & & W \times 1 \\
 &\text{subject to:} && \mathbf{A}' & \cdot & \boldsymbol{\pi} & = & \mathbf{c} \\
 & && L \times W & & W \times 1 & & L \times 1 \\
 &\text{and also:} && \boldsymbol{\pi} & \geq & \mathbf{0} \\
 & && W \times 1 & & W \times 1
 \end{aligned}$$

The dual problem has a natural interpretation. The π vector can be interpreted as a pricing vector, or as a risk adjusted probability. To see this note first the constraint which requires it to be non-negative. Then consider the equation: $\mathbf{A}' \cdot \boldsymbol{\pi} = \mathbf{c}$. This is a set of $L = I + H + J$ equations, with each equation corresponding to a different component of the hedging strategy. The first equation relates to holding cash, the next H to going long a one period forward contract for each possible history h , and the last J are going long one of the reference assets. Each equation says that the expected pay-off of that element under the distribution π must equal the cost of that element.

So the first equation says that the expected pay-off to a holding of 1 unit of cash should be 1. This means that $\sum_i \pi_i = 1$. The next H equations stipulate that, for any history h of length t :

$$\sum_{\omega \in h} (S_{t+1} - S_t) \pi_{\omega} = 0$$

In other words, the asset price process must be a martingale under the probability measure π . The final J rows specify that the reference assets must be correctly priced under the measure.

We can now state our main result:

Proposition 2: *The following two problems are equivalent:*

1. *find the cheapest strategy which dominates the exotic option, which involves cash, buying and holding a static portfolio of reference assets, and taking forward positions in the underlying asset.*
2. *find the martingale probability measure which correctly prices the underlying and the reference assets, and which maximizes the expected value of the exotic.*

There are two immediate corollaries:

Corollary 1: *it is not possible to tighten the bounds on the value of the exotic by also allowing the agent to trade the reference assets at intermediate times.*

Proof: If the price of the exotic at $t = 0$ does not violate the rational bounds enforced by the restricted set of hedging strategies, then there exists some feasible process π which is consistent with that price. That same process also implies a price for all the reference assets at all nodes which does not allow arbitrage. So the price of the exotic cannot violate rational bounds even if we extend the strategy space to allow future trading in call options.

Corollary 2: *given any martingale probability measure π which correctly values the underlying and the reference assets, and any strategy \mathbf{x} which dominates the exotic, if the expected pay-off to the exotic under measure π (ie $\mathbf{b}' \cdot \pi$) and the cost of the strategy \mathbf{x} (ie $\mathbf{c}' \cdot \mathbf{x}$) are equal, then \mathbf{x} is the solution to the primal, π is the solution to the dual and the bound on the exotic is $\mathbf{c}' \cdot \mathbf{x}$.*

Proof: this follows directly from Proposition 2 using the Saddle Point Theorem of LP.

3. Applying the Technology

We now use the LP to price a forward straddle - the same exotic we used in the initial example. We will show how the technology can be applied, and how the binding processes and the binding strategies come out of the model. We will see that the rational bounds are wide, and will provide some arguments which suggest that this width is a real measure of the uncertainty in valuing the exotic.

3.1 The Bounds on the Forward Straddle

The pay-off to the straddle is $|S_1 - S_2|$ where S_1 and S_2 are the prices of the asset at times t_1 and t_2 . We assume that European options exist with all possible strikes and maturing at times t_1 and t_2 ¹, and they are all priced with an implied (Black-Scholes) monthly volatility of 10% (annual volatility of about 35%).

The option prices determine the risk adjusted probability that $S_1 = K_i$ which we denote by p_{1i} , and similarly p_{2j} denotes the probability that $S_2 = K_j$. The process can be characterized by the matrix $\{\pi_{ij} \mid i, j \in [-N, N]\}$ which is the probability that $S_1 = K_i$ and $S_2 = K_j$. The dual LP problem (finding the process which maximizes the value of the exotic) can be expressed as:

$$\begin{aligned} \text{Maximize} \quad & \sum_{i,j} \pi_{ij} |X_i - X_j| \\ \text{subject to:} \quad & \sum_j \pi_{ij} = p_{1i} \quad \forall i \\ & \sum_i \pi_{ij} = p_{2j} \quad \forall j \\ & \sum_j \pi_{ij} (X_i - X_j) = 0 \quad \forall i \\ & \pi_{ij} \geq 0 \quad \forall i, j \end{aligned}$$

¹ Inclusion of options with other maturities would not affect the problem. The assumption that there is an option for every possible strike could easily be dropped. Its only purpose is to focus the example on the tightness with which European options in general bind exotics. If we assumed a finite number of strikes, the laxness of the bounds would be due in part to the particular strikes chosen.

The first two constraints ensure correct prices for options which expire at times t_1 and t_2 , the third ensures prices follow a martingale and the fourth that probabilities are non-negative.

In implementing the LP we use 40 stock price nodes at time t_1 and 80 at time t_2 . The values at each node are chosen so as to give equal probability to each node at each time. Figure 1 shows the bounds on a three month straddle ($t_2 - t_1 = 0.25$) for various values of t_1 . The ratio between the bounds is wide, exceeding 2:1 unless t_1 is small relative to t_2 ; when $t_1 = 0$, the bounds of course converge.

To get more insight into the bounding process, we look at the transition matrix for a three month straddle starting in one month. For most first period nodes i , the vector π_{ij} has non-zero values for two neighbouring $j > i$ and for two neighbouring $j < i$, but is otherwise empty. This suggests that the process which maximizes the value of the straddle is a binomial process. The top chart of Figure 2 shows the expected up and down prices in the second period given the price in the first period.

The dual gives the dominating strategy. It is shown graphically in the lower chart of Figure 2. To interpret it, suppose the straddle is trading at its theoretical upper bound. Write the straddle. Hedge the position by buying options which mature at time t_2 . The cost of these options is paid for partly by the proceeds from writing the forward straddle and partly by writing a portfolio of options which expire at time t_1 . The pay-offs to the portfolio of options bought and written are the two U-shaped curves illustrated in the chart. At time t_1 the position has to be hedged dynamically by buying or selling forward contracts on the share. The chart illustrates what happens in the particular case when S_1 takes the (arbitrarily chosen) value of 98. The hedger goes long 0.08 of the asset. This is just sufficient to ensure that the liability in the first period together with the liability under the straddle is dominated by the pay-off in the second period. The position breaks even if S_2 is either 83 or 113 (which is where the price jumps to in the limiting process); otherwise the strategy makes money.

The lower bound is rather more complex, and is shown graphically in Figure 3. For large or small values of S_1 the asset price does not move in the second period - $S_2 = S_1$. For intermediate values

the asset price follows a trinomial process, S_2 is S_1 or $u(S_1)$ or $d(S_1)$, where $u(S_1) > S_1 > d(S_1)$, and u and d are decreasing. The process is shown in the upper chart. We have superimposed an “exact” solution on the LP output. The exact solution is obtained from the continuous distributions by solving for the functions u and d which satisfy the above constraints, which match the unconditional distributions at both times, and satisfy the martingale requirement².

We also identify the dominating strategy if the straddle is trading at or below its lower bound of 9. This time the agent buys the straddle, buys a curved pay-off at time 1 and sells one at time 2. The strategy is shown in the lower chart, highlighting in particular the case of $S_1 = 98$. Some interesting points of difference from the previous case emerge. First, the pay-offs at the end of two periods now coincide. That is because the binding process puts positive weight on paths where $S_1 = S_2$; on such paths the arbitrage strategy must break even. That is only possible if the pay-offs exactly offset each other. A second feature of interest is that the pay-off is not strictly U-shaped. The curvature changes sign as can be seen by the fact that the straddle is tangent at the points to which the asset price can jump from 98 - that is 74 and 146.

3.2 The Realism of the Bounds

The rational bounds on the value of the forward straddle are wide. We have seen that the ratio of upper bound to lower bound for a three month straddle starting in one month’s time is almost 2:1, even when there are both one and four month options available with a full range of strike prices to use for hedging. One reason may be the difficulty in hedging some exotics with conventional options, even where the exotic is straightforward and similar to conventional options. An alternative explanation is that by using a criterion as rigorous as dominance, without putting any limitations on processes, we are admitting processes which are pathological and which can safely be ignored for practical purposes.

² The exact and LP solutions match almost exactly except for the extreme values of S_2 . As the grid becomes infinitely fine, the measure of the extreme values falls to zero.

It is hard to see how one can say which explanation is better in general. But it is possible to use the framework we have developed to provide some useful evidence in this specific case. The features of the processes which define the upper and lower bounds help identify the factors to which the value of a forward straddle is sensitive. We can then build models which are consistent with the prices of the reference assets, which are plausible, and which differ substantially on the dimension to which the straddle is sensitive, and see how far their valuations of the straddle diverge.

The two binding processes differ most obviously in the distribution of price changes in the second period, as illustrated in Figure 4. The martingale property and the prices of the options determine both the mean (zero) and the variance. The process which provides the upper bound has virtually a two point distribution with most of the mass at plus or minus one standard deviation. The lower bound process has most of the mass at zero with sufficient probability of large jumps to get the right standard deviation. It suggests that the kurtosis of the price changes in the second period is significant in determining the value of the forward straddle.

Consider then the following set-up: the task is to price and hedge a three month straddle, starting in one month. The spot price is \$100. There are one month and four month options trading. The implied risk neutral density function at each date can be represented as the average of two lognormal distributions (see Mellick and Thomas (1997) for a discussion on the use of combinations of lognormals to model return distributions). The one month options are consistent with a world where the market believes that volatility will be either 10% or 20% with equal probability.³ The four month options are similarly the average of two Black-Scholes price functions, but in this case the two volatilities are 10% and 15%.

There are many models consistent with these prices. Consider the following two which are set out in Table 1: in both models, there are two possible states of the world, A and B. They are

³ The implied volatility of one month options is then a U-shaped function of the strike, being close to 15% and rising to 16.5% for strikes 20% in or out of the money.

equally likely (on a risk-adjusted basis). Once the state is determined, it does not change. Conditional on the state of the world, the stock price follows a geometric Brownian motion with a volatility which takes one value in month 1 and another value in months 2-4. In both models, the volatility in the first month is 20% in state A and 10% in state B. In Model I, the average volatility over the four months is 15% in state A and 10% in state B; in Model II it is the other way round. Both models are internally consistent, and also fit the prices of options. They differ however in their predictions for volatility in months 2-4: in Model I the volatility will be 12.9% or 10%, while in Model II it will be 0% or 16.3%.

Model I, whose distribution of returns over the life of the straddle is the average of two mean zero lognormals with roughly similar variances, has excess kurtosis of only 1.0. By contrast, Model II is the average of a constant and a lognormal with a large variance, and has excess kurtosis of 5.9. The straddle is worth \$15.80 in Model I but only \$11.25 in Model II - a ratio of 1.40:1.

Model I seems quite plausible. It reflects uncertainty in volatility, and it captures the fact that volatility tends, in many markets, to be autocorrelated. Model II seems less natural, involving negative correlations in volatility. But there could well be reasons for this. Suppose for example that there is a 50% possibility of a cash bid being made for the company over the next month. In state A, a bid is made. The price is quite uncertain, so volatility over the next month will be high, at 20%. But if the bid is made, the share value will then be known with certainty, so volatility after the first month will be zero. So in state A volatility will be 20% over 1 month and 10% over 4 months. In state B no bid is made, and volatility is expected to be 10% over one month and 15% over 4 months.

Consider then an institution which holds the forward straddle until t_j . If it values and hedges the straddle using Model I it will conclude that the straddle is worth \$15.80. If it uses the model for hedging, it may believe that the value is locked in. Yet if the world actually conforms to Model II, the straddle is overvalued by more than \$4. A trading loss of that amount is bound to be recognized in due course. For whatever valuation and hedging model is used, it is not possible to

lock in value which is not there. Furthermore, if an incorrect model is used for hedging, this could lead to further error and therefore to substantially higher losses.

This example suggests that the rational bounds analysis may indeed be useful in highlighting the risks in positions which include exotic options, even in cases where models can be calibrated against a liquid conventional options market. It is also helpful in identifying the type of scenarios under which problems may be most severe.

4. Near-Analytic Bounds on Particular Exotic Options

We have seen how linear programming can be used in practice to bound exotics and identify robust hedging strategies. But LP also provides us with a useful framework for getting close to analytic solutions for particular exotics. In this section we establish bounds on two closely related exotic options: a digital barrier option and a lookback option. We show how to reduce the linear programming problem to a simple optimisation problem in one dimension.

We look for a super-replicating strategy which is a plausible candidate for being the cheapest such strategy. It is a feasible solution to the primal problem. We then look for a martingale process which is consistent with the prices of the reference assets, and which minimizes the expected margin between the exotic and the super-replicating strategy. This is a feasible solution to the dual problem. Using Corollary 2, if there are feasible solutions to both primal and dual which give the same values of their respective objective functions then they are optimal solutions to both. In other words, if the expected pay-off to the exotic is equal to the cost of the super-replicating strategy, then we have identified the upper bound on the value of the exotic.

If the feasible solutions to the primal and the dual do not give the same values, then we have not identified the least upper bound. But at least we have bracketed the bound, and the process helps to identify where economies can be made on the cost of the super-replicating strategy. In this way we can move towards the saddle point.

4.1 Bounds on the Digital Barrier Option

An American digital barrier option pays \$1 if the asset price reaches or passes some barrier level B at any time before the expiry date T , and zero otherwise. Suppose the asset price at $t = 0$ is S_0 ($< B$). The set of reference assets is European call options $\{C_i\}$ with maturity T and increasing positive exercise prices $\{X_i\}$ with $i = 1, \dots, M$.

A plausible candidate for an optimal super-replicating strategy is to hold $1/(B-X)$ calls with a strike price of X for some $X < B$. The idea is that if the barrier is breached the agent sells $1/(B-X)$ forward contracts. This locks in at least \$1. In fact if the stock price immediately after the barrier is breached is B^* (B^* would be equal to B if the stock price process is continuous, otherwise it could be larger) then the pay-off to the strategy is:

$$\frac{1}{B-X} \left(\{S_T - X\}^+ + (B^* - S_T) \right) \text{ if the barrier is breached}$$

$$\frac{1}{B-X} \{S_T - X\}^+ \text{ otherwise}$$

The strategy super-replicates the digital. It is therefore a feasible solution to the primal. It replicates the digital exactly on paths where:

- either the barrier is breached, with $B^*=B$, and where $S_T \geq X$;
- or where the barrier is not breached and where $S_T \leq X$.

The issue then is whether we can identify a feasible solution to the dual problem which puts all the probability mass on those paths where the strategy replicates the digital exactly. This is possible, as shown in the Appendix, so we have:

Proposition 3: *The upper bound on the value of a digital option is:*

$$\frac{C_{j^*}}{B - X_{j^*}} \text{ where } j^* = \underset{j}{\text{ArgMin}} \left[\frac{C_j}{B - X_j} \right]$$

The corresponding super-replicating strategy involves buying $1/(B-X_{j^*})$ options with strike X_{j^*} , and selling $1/(B-X_{j^*})$ of the underlying forward the first time that S reaches the barrier.

The lower bound is quite easy to identify:

Proposition 4: *the lower bound on an American digital option is the price of the equivalent European digital option:*

$$\frac{C_n - C_{n+1}}{X_{n+1} - X_n} \quad \text{where } X_n < B \leq X_{n+1}$$

Proof: the European digital is clearly a lower bound since it pays 1 if $S_T \geq X_{n+1} \geq B$ while the American digital pays 1 if the asset price ever reaches B . To show it is the greatest lower bound, consider the following price process: the price remains constant except at a single date, say at T , when it jumps to its final level. Any path which hits the barrier is above the barrier at the end, so in this case the price of the American digital is equal to its European counterpart.

This lower bound is very low because the price jumps straight through the barrier. The bound would be higher if paths are forced to be continuous. Assume in particular that on any path which crosses the barrier, it is possible to rebalance the portfolio at the barrier level. To simplify the notation, it will be helpful to assume that there are European call options available at every strike X with price $C(X)$. To find the new lower bound with continuity we follow our normal procedure of guessing the corresponding strategy and proving it is optimal.

With a complete set of call options we can construct an arbitrary European pay-off $f(S_T)$. We conjecture that the optimal strategy is to buy this pay-off and sell x forward contracts if the barrier is breached. The dominance condition can be expressed as:

$$\begin{aligned} f(S_T) &\leq 0 && \text{for } S_T < B \\ f(S_T) - x(S_T - B) &\leq 1 && \text{for all } S_T \end{aligned}$$

The second expression depends on continuity: if the barrier is breached, we assume the forward contracts are sold at B rather than some higher price. To obtain the highest lower bound for any x it is necessary to set:

$$f(S_T) = \begin{cases} 1 + x(S_T - B) & \text{if } S_T < B - \frac{1}{x} \\ 0 & \text{if } B - \frac{1}{x} \leq S_T \leq B \\ 1 + x(S_T - B) & \text{if } S_T > B \end{cases}$$

(it is easy to verify that it is not optimal to have $x < 1/B$).

Proposition 5: *if prices are continuous, the lower bound on the value of the American digital is found by choosing the x which maximises the value of $f(S_T)$.*

Proof: the proof is in the Appendix. Having identified a strategy which is dominated by the digital, all that is necessary is to show that there exists a process under which the expected value of the digital and of the strategy are the same.

Having identified the bounds on the American digital, we now look at a numerical example to see how tight the bounds are in practice. Suppose that all European options are priced on the same implied monthly volatility of 10%. The maturity of the digital is 1 month. For every set of parameter values there are three bounds: a supremum, an infimum with continuity, and an unrestricted infimum. The first two entail the search for a strike price which leads to a maximum or minimum cost. This requires a one dimensional optimization which is straightforward. The third can be computed analytically. We also calculate the ‘‘Black-Scholes price’’ which assumes that the stock price follows a diffusion with constant volatility; this can be expressed analytically (see Nelken (1996), page 223). The results are in Figure 5. The top chart shows the four prices as a function of the level of the barrier for two different maturities. To give further insight into the width of the bounds, the lower chart expresses the prices in terms of implied volatility. That is we take each price and ask what level of volatility would give a Black-Scholes price of the same magnitude.

The bounds, measured in terms of implied volatility, tighten as the option goes further out of the money. The bounds are not very tight. If for example we take a point where the Black-Scholes value is 0.5 - so there is a 50% chance of the barrier being hit - the width of the bounds whether measured in terms of price or implied volatility is around half the level of the Black-Scholes price.

4.2 Lookback Options and Other Exotics

The analysis of the digital can readily be extended to the lookback option. A lookback pays $\text{Max}\{S_t \mid t \in [0, T]\}$. Assume that the price space is discrete (prices $K_N, \dots, K_0 = S_0, \dots, K_N$) and that all option strike prices are included in the set of possible stock prices.

A lookback option is equivalent to a portfolio of American digital options, with $(K_n - K_{n-1})$ digitals with a barrier set at K_n for all $n > 0$. One way of super-replicating the lookback is to sum the super-replicating strategies for each of the individual digitals. Thus for each n there is a $j^*(n)$ such that the strategy involves buying $(K_n - K_{n-1})/(K_n - K_{j^*(n)})$ calls with a strike of $K_{j^*(n)}$ for all $n > 0$. Similarly, one candidate for the highest cost strategy dominated by the lookback is the sum of the optimal dominated strategies for the individual digitals.

As shown in the Appendix, there are processes under which these feasible strategies have the same price as the lookback option, and hence we get:

Proposition 6: *The three bounds (the upper bound, the unrestricted lower bound and the lower bound assuming continuity) on the price of a lookback option are the sum of the corresponding bounds on the constituent American digitals.*

We have looked at digitals and lookbacks because they have particularly simple bounds. We can readily apply the same technology to closely related exotics like barrier options (for example up-and-out puts and calls), and again we have been able to identify upper and lower bounds. We have also found bounds for other simple exotics like options on the sum of two assets, or claims which pay the greater of two assets when the set of reference assets is the set of European options

on each asset separately. In all these cases the solution is near-analytic in the sense that the identification of the optimal strategy involves a one dimensional optimization problem.

It should be recognized that the solutions are quite sensitive to the specification of the problem. For example, in the case of digital options we have assumed that the barrier is flat. In conjunction with the zero interest rate assumption this means we are assuming a barrier with a constant level in present value terms. Similarly we are assuming that the present value of the pay-off if the barrier is breached is independent of the time at which it is breached. If these assumptions are relaxed, the problem becomes more complicated. We have not been able to identify a near-analytic solution for this general problem. Nevertheless for many cases of interest we believe that there are features which can be exploited to reduce the dimensionality of the LP and so make it possible to get accurate numbers at reasonable computing cost.

5. Conclusions

Exotic options are often hard to price and hard to hedge. Traditional methods generally make strong assumptions about the stochastic process in order to come up with an exact price; practitioners frequently add in a generous margin fixed by some rule of thumb. We have shown how to find rational bounds on exotic option prices given the prices at which suitable hedging assets are currently traded.

In some cases these bounds will be useful as a guide to pricing. However the bounds may often be so wide that most trading will take place well within the bounds; but even then the bounds serve a number of useful functions. First they show the potential for gain or loss. However sophisticated the pricing and hedging strategy, an agent who sells an exotic for less than the rational upper bound cannot exclude the possibility of making a loss. Second, the bounds are enforced by a robust hedging strategy. The strategy is robust in two senses. It is robust in that it super-replicates however the prices of the underlying and of its derivative securities develop in future. It is also robust in that it does not require dynamic trading of the call options. Any call

options used for hedging are bought at the outset and held to maturity. The only dynamic trading is in the underlying asset itself.

A strategy which puts a floor on potential loss, independent of price process, may be much more attractive than one which minimises variance of profit provided that prices conform to some theoretical model. The fact that the robust strategy does not involve dynamic trading in options may also be advantageous. For whereas hedgers may be prepared to rely on the existence of a deep and liquid market in the underlying asset at future times, they may be much more reluctant to depend on always being able to trade options at tight spreads and in sufficient size.


The processes which support the bounds are also worthy of interest. They indicate the scenarios under which the exotic is more or less valuable. We saw for example that a short forward straddle is effectively a long position in forward kurtosis. Knowing this, traders can make informed decisions about whether they actually want to take a position in these exotics, and what level of reward is required to compensate them for the risks.

This paper has illustrated what can be done. We have shown explicit solutions for particular exotic options - the digital and the lookback. We have shown how the problem can be formulated as a linear program for the case of a generalized exotic with a generalized set of reference assets and we have provided numerical solutions for the case of a forward straddle for which we cannot provide a full explicit solution.

There are a number of areas in which the work can be taken forward. It can be extended to find analytic (or near-analytic) bounds for other exotics. If the technique is to be used for hedging portfolios of exotic options, it is essential that we understand how to use the LP approach in an efficient manner. In this context the fact that the conclusions from the very crude analysis of the forward straddle in section 1 were borne out by the full LP in section 3 was heartening.

APPENDIX

Proof of Proposition 3: We need to show that there is a feasible process which gives zero measure to paths where the strategy strictly dominates the exotic. The obvious candidate is a process as shown below:



with the jumps occurring say at times $T/2$ and T .

Complete the set of reference assets by adding two options at the lower end, and two at the upper end with strike prices $X_{-1} < X_0 = 0$, and $\text{Max}[B, X_M] < X_{M+1} < X_{M+2}$. These extra options are assumed to trade at their intrinsic value (so their prices are $S_0 - X_{-1}$, S_0 , 0 and 0 respectively⁴. We will not actually wish to trade these extra options, but their existence helps simplify the algebra.

Let π_j denote the probability of ending up at node X_j . Then to fit the option prices:

$$C_j = \sum_{k=j}^M \pi_k (X_k - X_j)$$

A feasible solution for π (non-negative, adding to 1) is guaranteed by the set of option prices themselves not allowing arbitrage. The solution is in fact unique and given by:

⁴ To avoid arbitrage $X_{M+1} \geq (C_{M-1} X_M - C_M X_{M-1}) / (C_{M-1} - C_M)$

$$\pi_j = \frac{(X_j - X_{j-1})C_{j+1} - (X_{j+1} - X_{j-1})C_j - (X_{j+1} - X_j)C_{j-1}}{(X_j - X_{j-1})(X_{j+1} - X_j)}$$

(for $j = 0$ to $N+1$)

For the process to be a martingale, we need to have $X = X_{j^*}$ where:

$$j^* \text{ is the maximum } j \text{ such that } B \sum_{k=j}^M \pi_k \geq \sum_{k=j}^M X_k \pi_k$$

This expression can be rearranged in terms of option prices rather than probabilities to give:

$$j^* \text{ is the maximum } j \text{ such that } \frac{C_{j-1}}{B - X_{j-1}} \geq \frac{C_j}{B - X_j}$$

In other words, this is equivalent to fixing the strike of the options so as to minimise the cost⁵, $C_j / (B - X_j)$, of the super-replicating strategy. To complete the specification of the process, we need to specify the probability of the path (S_0, B, X_{j^*}) and also the value of Y . Calling the former α , we must have:

$$B\pi_B = X_{j^*}\alpha + \sum_{k=j^*+1}^M X_k \pi_k$$

$$\text{where } \pi_B \equiv \left(\alpha + \sum_{k=j^*+1}^M \pi_k \right), \text{ the probability of hitting the barrier}$$

while Y is determined by:

$$Y(1 - \pi_B) + B\pi_B = S_0$$

⁵ It is indeed a global minimum; the second derivative of the minimand is positive over the whole range.

Proof of Proposition 5: first, we establish the condition on x to maximize the value of the dominated strategy. If the risk adjusted probability that $S_T=S$ is $\pi(S)$, then the price of the strategy is:

$$f(x) = \int_0^{B-1/x} + \int_B^{\infty} \{1 + x(S - B)\} \pi(S) dS$$


This is maximized when $x=x^*$ where:

$$f'(x^*) = \int_0^{B-1/x^*} + \int_B^{\infty} (S - B) \pi(S) dS = 0$$

This condition can be interpreted as meaning that:

$$E[S_T | S_T \in [0, B - 1/x^*] \cup [B, \infty)] = B$$

So C is dominated for any x and $f(x^*)$ is the maximum price of C ; it is a lower bound on the value of the American digital. To show it is the highest lower bound we need to identify the corresponding process under which the pay-off to the strategy is always equal to the pay-off to the digital. C is *strictly* dominated either if the barrier is hit and $S_T \in (B-1/x^*, B)$, or if the barrier is not hit and $S_T \in (0, B-1/x^*)$. Now consider the following process. The stock price is constant except at times $T/2$ and T ; at $T/2$ the price is either B or else declines to some value Y ; if it hits the barrier, then at time T the price jumps to a terminal value X_j ; where either $X_j \geq B$ or $X_j \leq X$ for some X ; if it fails to hit the barrier, the price jumps to a terminal value X_j ; where $B > X_j \geq X$. The process looks like:




The process has to be martingale and fit the prices of European options; this imposes a restriction on X . We have seen that this restriction is equivalent to saying that $X = B - I/x^*$.

Proof of Proposition 6: If the super-replicating strategy does indeed give the least upper bound, this means that there must be a price process whereby with probability 1 the replication is exact. From the earlier discussion on the digital, it follows that the only possible paths have the following characteristics: if they attain S_n then they must end up at or above $X_{j^*(n)}$, and conversely if they fail to attain S_n then they must end up at or below $X_{j^*(n)}$. The terminal value is thus a function of the maximum value.

There is such a process, it is a martingale, and it is consistent with the prices of the reference assets. Partition the interval $[0, T]$ using arbitrary increasing times $0 < t_1, < \dots < T$ to ensure there are as many times as asset prices above S_0 . At time t_1 the asset price either rises to S_1 or falls to some X in the range $[0, X_{j^*(1)}]$. If the price falls, it stays at that level until time T . Otherwise it then either rises again to S_2 or it falls to some X in the range $[X_{j^*(1)}, X_{j^*(2)}]$. And so the process continues. When the stock price reaches X_{M+1} it moves no more until time T . The probability that the stock price will reach or surpass S_n is $\pi(n)$.

So the price goes up one tick each period unless it falls. If it falls it will fall to a price which is an option strike price. It will then stay constant. The higher the price reached before the fall, the higher the terminal value. Depending on the distribution, it is possible for paths to end up at the same node having fallen from different nodes, or end up at different nodes having fallen from the same node. Graphically the process looks like this:



We have shown that the lookback is equivalent to a portfolio of barrier options, that it can be super-replicated by replicating the individual barrier options, and that there is a feasible process under which the expected pay-off to the exotic is equal to the cost of the super-replicating strategy. We have therefore proved that the cost of the super-replicating strategy is the rational bound on the price of the lookback.

For the lower bound on the lookback when the asset price can jump, consider a process where the asset price at some time t jumps to its terminal value and stays there. The dominated portfolio, which pays $\text{Max} \{S_0, S_T\}$, has the same pay-off as the digital on all possible paths.

If the asset price is assumed not to jump then we need to show the existence of a process where the price each interval either rises to hit a new bound or falls to a level $X(B)$ where B is the bound reached so far. It looks very similar to the upper bound process above except that $X(\cdot)$ is a decreasing function.

Table 1: Two Models which are consistent with the same set of Call Prices

State:	Model I		Model II	
	A	B	A	B
Volatility:				
Month 1	20.0%	10.0%	20.0%	10.0%
Month 1-4	15.0%	10.0%	10.0%	15.0%
Months 2-4	12.9%	10.0%	0%	16.3%
Straddle Price	\$15.80		\$11.25	

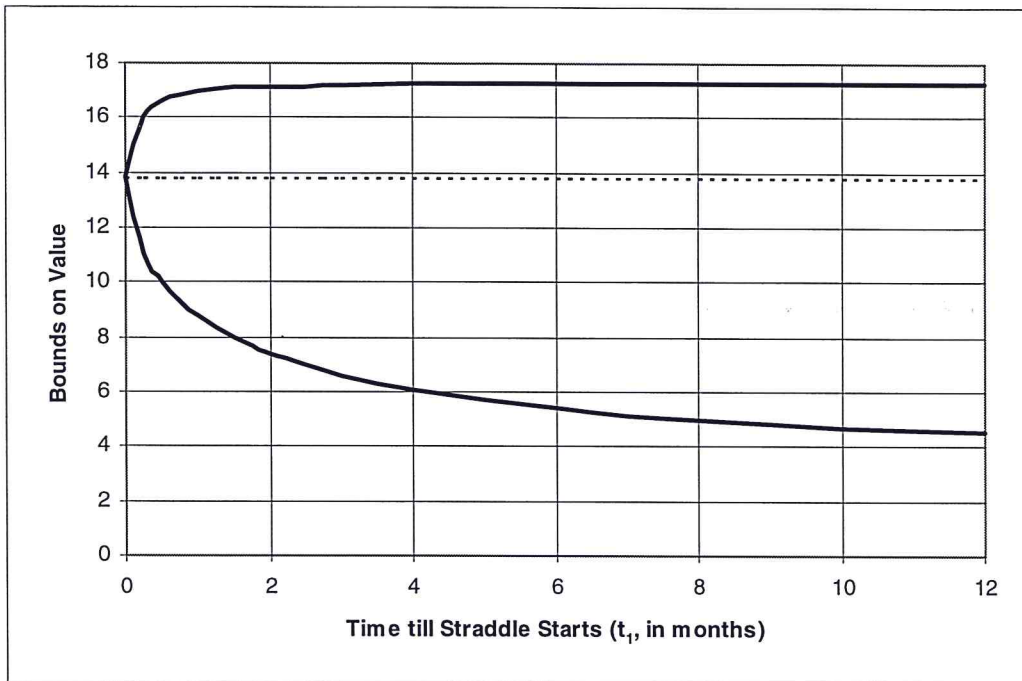


Figure 1: Bounds on the Value of a Forward Straddle

The Chart shows the rational bounds on a forward straddle which pays $|S_1 - S_2|$, where S_1 is the price at time t_1 and S_2 at time t_2 . The current stock price is 100. The duration the straddle, $t_2 - t_1$, is 3 months. There exist call options which mature at times t_1 and t_2 for every possible strike. They all trade on a Black-Scholes implied volatility of 10%. The dotted line shows the Black-Scholes value of the forward straddle.

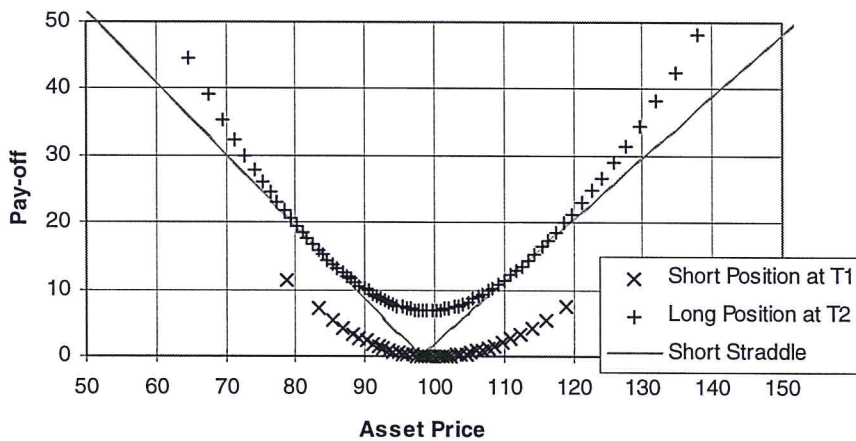
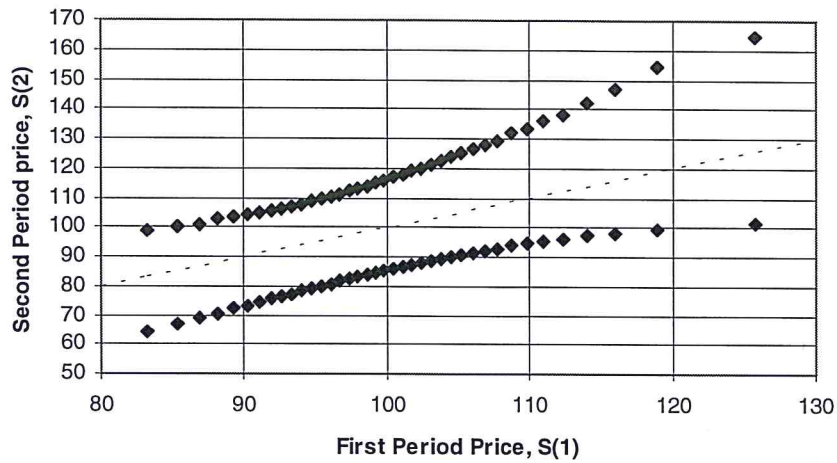


Figure 2: The Upper Bound on a Forward Straddle

The top chart shows the process which is consistent with the upper bound on a forward straddle. It is a binomial process and the diamonds show the two possible prices at t_2 as a function of price at t_1 . The dashed line is a 45° line through the origin and is for reference. The lower chart shows the arbitrage strategy which dominates the straddle. It involves taking a short position at t_1 shown by the x's, and a long position at t_2 shown by the +'s. At t_1 , an appropriate forward position is taken out. The V-shaped line shows the negative pay-off at t_2 of the short straddle, the short position at t_1 , and the forward contract, on the assumption that S_1 is 98. The parameters are: a three month straddle, starting in 1 month, initial stock price 100, and all options trading on an implied Black-Scholes volatility of 10% monthly.




Figure 3: The Lower Bound on a Forward Straddle

The top chart shows the process which is consistent with the lower bound on a forward straddle. For extreme values of S_1 , $S_2 = S_1$; for intermediate values the stock price follows a trinomial process. The chart shows the values of S_2 as a function of S_1 . The diamonds show the actual LP solution while the line shows the limit case with an arbitrarily fine grid. The lower chart shows the arbitrage strategy which is dominated by the straddle. It involves taking a long position at t_1 shown by the x's, and a short position at t_2 shown by the +'s. At t_1 , an appropriate forward position is taken out. The V-shaped line shows the positive payoff at t_2 of the long straddle, the long position at t_1 , and the forward contract on the assumption that S_1 is 98. The parameters are as in Figure 2.

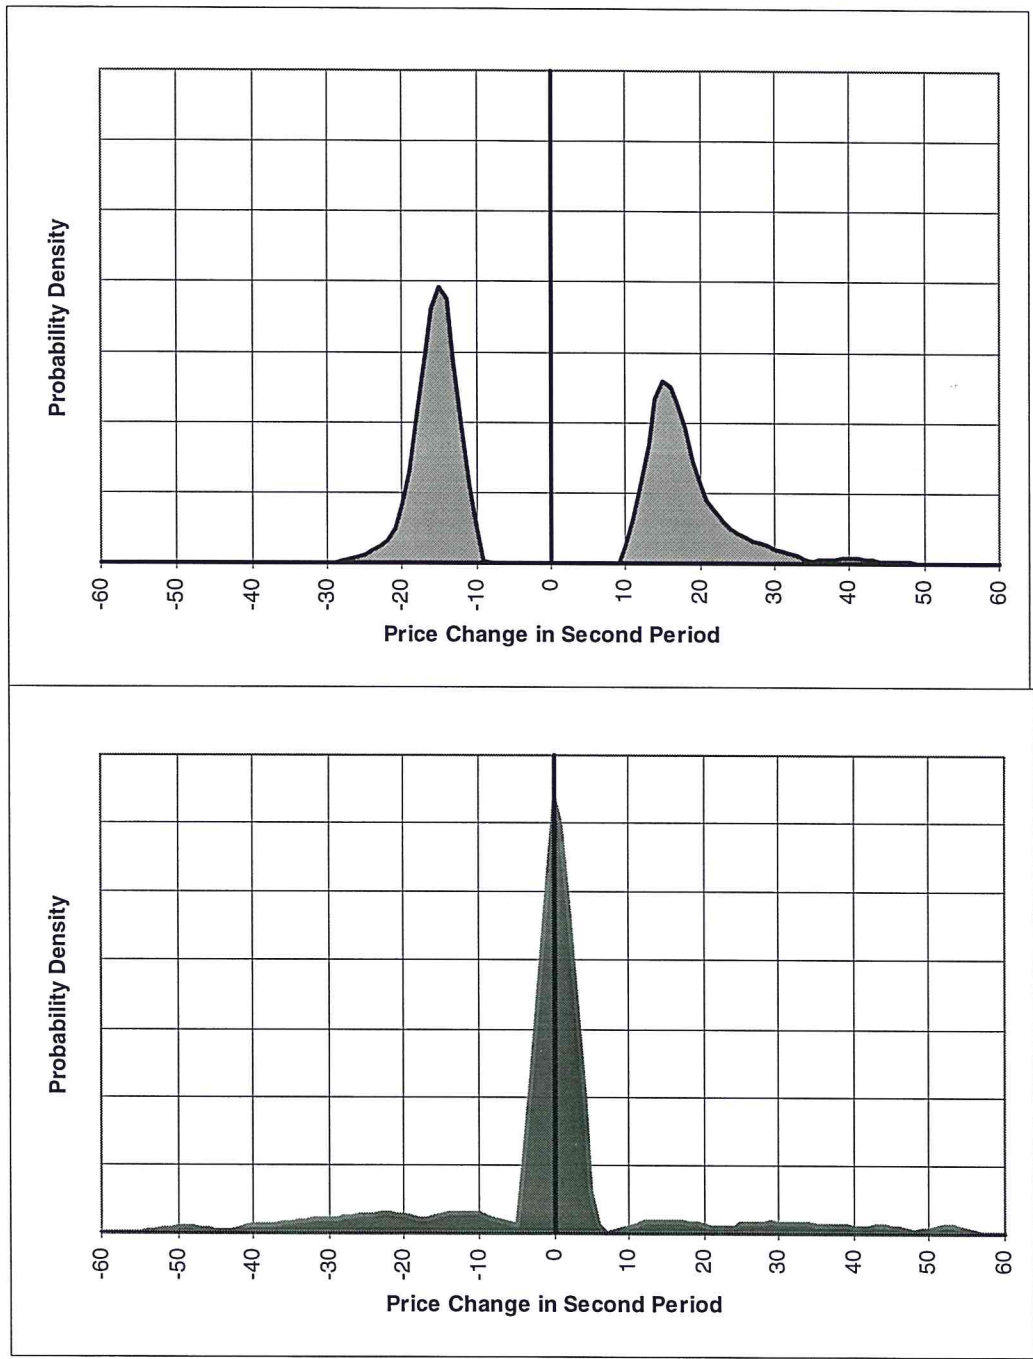


Figure 4: Distribution of Price Change in Bounding Processes

The two charts show the distribution of the price change over the life of straddle, $S_2 - S_1$, under the two processes which define the no arbitrage bounds on the value of the straddle. The top chart shows the upper bound and the bottom shows the lower bound. The two distributions have the same mean and variance. The parameters are as in Figures 2 and 3.




Figure 5: Bounds on the price of an American Digital Option

The top chart shows the bounds on a 1 month American digital option as a function of the barrier level. The initial price is 100, and the option pays 1 if and only if the price breaches the barrier during the month. There exist 1 month options for every strike, and they all trade on a Black-Scholes implied monthly volatility of 10%. The dotted line is the Black-Scholes value. The two lower bounds differ according to whether it is assumed that the price process is continuous or not. The lower chart carries the same information, but prices are expressed as Black-Scholes implied volatilities.

References

- Ahn, H, A. Muni and G. Swindle, 1997, "Misspecified Asset Price Models and Robust Hedging Strategies", *Applied Mathematical Finance*, 4, 21-36.
- Berger, E, 1996, "Barrier Options", in I Nelken (ed.), *The Handbook of Exotic Options*, Irwin, Burr Ridge, IL, 213-243.
- Bernardo, A and O. Ledoit, 1997, "Gain, Loss and Asset Pricing", Working Paper, UCLA (May 1997)
- Carr, P, K Ellis, and V Gupta, 1998, "Static Hedging of Exotic Options", *Journal of Finance*, 53:3, 1165-1190.
- Christofides, A and G R Salkin, 1996, "A Graph-Theoretic Model for Option Pricing in Imperfect Markets", Working Paper, Imperial College, June 1996.
- Cochrane and Saá-Requejo, 1996, "Beyond Arbitrage: 'Good Deal' Asset Price Bounds in Incomplete Markets", Working Paper 5489, National Bureau of Economic Research, March 1996.
- Dantzig, G B, and P Wolfe, 1960, "Decomposition Principle for Linear Programs", *Operations Research*, 8, 101-111.
- Dembo, R, 1994, "Hedging in Markets that Gap", in R A Klein and J Lederman, (Eds.), *Handbook of Derivatives and Synthetics*, Probus Publishing, 1994.
- Derman, E, D Ergener and I Kani, 1995, "Static Options Replication", *Journal of Derivatives*, 2, 78-95.
- Hobson, D G, 1997, "Robust Hedging of the Lookback Option", *Finance and Stochastics*, forthcoming.

Edirisinghe, C, V Naik and R Uppal, 1993, "Optimal Replication of Options With Transactions Costs and Trading Restrictions", *Journal of Financial and Quantitative Analysis*, 28, 117-138.

Gilmore, P C and R E Gomory, 1961, "A Linear Programming Approach to the Cutting Stock Problem", *Operations Research*, 9, 849-859.

Hodges, S D, 1997, "A Generalization of the Sharpe Ratio with Applications to Valuation Bounds and Risk Measures in Incomplete Markets", Working Paper, University of Warwick, June 1997.

Maros, I and G Mitra, 1996, "Simplex Algorithms", Chapter 1 in Beasley, J. E (Ed), *Advances in Linear and Integer Programming*, Clarendon Press, Oxford, 1996.

Mellick, W R and C P Thomas, 1997, "Recovering an Asset's Implied PDF from Option Prices: An Application to Crude Oil during the Gulf Crisis", *Journal of Financial and Quantitative Analysis*, 31:1, 91-115.

Merton, R C, 1973, "The Theory of Rational Option Pricing", *Bell Journal of Economics*, 4, 141-183.