MANAGING OPTIONS RISK FOR EXOTIC OPTIONS

An exotic derivative is one for which no liquid market exists. As a general rule, the only liquid options are European-exercise calls and puts, including interest rate caps, floors, and European swaptions. Since exotic derivatives can rarely be marked-to-market based on publicly available prices at which they can be bought and sold, they usually have to be marked based on some hedging strategy which involves combinations of dynamic hedging using forwards and European-exercise options.

The general approach we will be following is one of building up techniques for marking-to-market and calculating risk on less liquid instruments by designing hedging strategies involving more liquid instruments. A less liquid instrument is marked based on the cost of executing the hedging strategy, and risk is then calculated by estimating the degree of uncertainty around this cost. We will do this in successive stages - at each stage, instruments for which we developed marks at previous stages will be considered available for use in creating hedging strategies to mark other instruments. However, we must be careful to realize that any risks due to uncertainty of how to mark an instrument will be inherited by instruments which it is then used to mark. So we will often be faced (once again) with tradeoffs between basis risk (employing a hedging strategy which uses more liquid instruments but which leaves a lot of room for uncertainty) and liquidity risk (employ a hedging strategy which does not introduce much uncertainty by using less liquid instruments).

As one example of how the stages will work: USD/JPY FX

- We will use more liquid forwards to create a hedge for less liquid forwards on USD/JPY FX and estimate the resulting basis risk.
- We will use both less liquid and more liquid forwards on USD/JPY FX along with more liquid European-exercise options on USD/JPY FX to create a hedge for less liquid European-exercise options on USD/JPY FX.
- We will use both less liquid and more liquid European-exercise options on USD/JPY FX to create a hedge for barrier options on USD/JPY FX.
- We will use barrier options on USD/JPY FX to create a hedge for lookback options on USD/JPY FX.

The following table, taken from RISK Magazine, May, 2000, shows the principal forms of exotic products and how widely they are used in different markets. Those which are classified as correlation-dependent, as well as virtually all interest rate exotics, have payoffs based on more than one underlying security. We will address the risk management of these exotics in another section. In this section we will focus on exotics whose payoff is based on just one underlying security, but many of the principles we establish here are applicable to multi-security exotics as well.

Intensity of use of option structures in various market

	Interest rate options		Forex options		Equity options		Commodity options	
	OTC	Exchanges	OTC	Exchanges	OTC	Exchanges	OTC	Exchanges
First-generation options								
European-style	A	A	A	A	A	O	A	A
American-style	A		A		R	A	A	A
Bermuda-style	A		A		R		O	
Second-generation options								
Path-dependent options								
Average price (rate)	A		A		A		A	A
Barrier options	A		A		A		A	
Capped	O		O		A		O	
Lookback	R		R		O		R	
Ladder	O		O		A		O	
Ratchet	O		O		A		O	
Shout	R		R		R		O	
Correlation-dependent options								
Rainbow options	R		O		O		O	
Quanto options	A		A		A		A	
Basket options	R		A		A		A	R
Time-dependent options								
Chooser options	R		R		R		O	
Forward start options	R		R		A		A	
Cliquet options	R		R		A		O	
Single payout options								
Binary options	A		A		A		A	
Contingent premium options	A		A		R		A	

A = actively used; O = occasionally used; R = rarely used; blank = not used

Here’s another table of exotic option types, showing which ones can be hedged using static combinations of vanilla options and which require dynamic hedging with vanilla options. In addition to the options listed in the RISK Magazine chart, we have also added a few others which have a reasonable amount of use. Dynamic hedging is divided up into the coverage of 3 primary risks: shape of the volatility surface, liquidity risk, and correlation risk. Management of the first two risk types will be covered in this section; management Of correlation risk will be covered in the section on multi-factor exotics.

		Dynamic Hedge Needed		
	Static Hedge possible	Shape of volatility surface	Liquidity Risk	Correlation Risk
First-generation options				
European-style				
American-style		✓		✓
Bermuda-style		✓		✓
Second-generation options				
Path-dependent options				
Average price (rate)		✓		small
Barrier options		✓	Reverse barriers	
Lookback	by barriers			
Ladder	by barriers			
Ratchet	by barriers			
Shout		✓		
Correlation-dependent options				
Rainbow options				✓
Quanto options				✓
Basket options				✓
Time-dependent options				
Chooser options	price risk only	✓		
Forward start options		Strikeless vo		
Cliquet options		Strikeless vol		
Compound options	price risk only	✓		
Volatility swaps		Strikeless vol		
Variance swaps	✓			
Single payout options				
Binary options			✓	
Contingent premium options	✓			
Power options	✓			

It is reasonably straight-forward to see that the payoff of any derivative whose payments are a function solely of the price of one underlying security at a single future time can be replicated as closely as one likes with the payoff of a combination of vanilla forwards, calls, and puts, since any (smooth) function can be approximated as closely as you wish, by a piecewise-linear function, and vanilla options at different strikes can produce any desired piecewise linear payoff. Hence, this portfolio of vanillas can be held as a nearly perfect hedge of the exotic derivative (see Carr & Madden , "Towards a Theory of Volatility Trading", available on Peter Carr's website). For a concrete example of how this can be used in practice, let's consider gold-in-gold options.

A plain vanilla gold option would be, for example, an option to exchange 1,000 ounces of gold for $300,000 in one year. Another way of viewing it is that the buyer has an option to receive 1,000 x (G - 300), where G is the gold price in one year. If G = 330, the buyer can exchange $300,000 for 1,000 ounces of gold, under the option, and then sell the 1,000 ounces of gold for $330,000, a net profit of $330,000 - $300,000 = $30,000.
A gold-in-gold option calls for this price increase to be received in gold. So if the price of gold goes up 10% from $300 to $330, the option buyer receives 10% x 1,000 ounces = 100 ounces of gold. Since gold is worth $330 in ounces, this is a profit of 100 x $330 = $33,000, a payoff of 11% rather than 10%.

In general, the payoff is x% * (1 + x%) = x%+x^2%.

This can be hedged with 101% of a vanilla option struck at par plus 2% of a call option struck at each 1% up in price, e.g.., to hedge a 1,000 ounces gold-in-gold contract with a strike of $300, a 1,001 ounce call with a strike of $300, a 2 ounce call at $303, a 2 ounce call at $306, a 2 ounce call at $309, and so on. The payoff on these calls, if gold rises by x%, is

$$101\% * x\% + 2\% * \sum_{I=1}^{x-1} I \;=\; 101\% * x + 2\% * \frac{x^2 - x}{2} \;=\; x\% + x^2\%$$

In practice, there is some degree of underhedging based on the spacing between strikes used in the hedge and that there will be some highest strike above which you won't buy any hedges. This leads to basis risk, which could be calculated fairly easily by an assumed probability distribution of final gold prices multiplied by the amount of mis-hedge at each price level. This basis risk can be reduced as low as one likes, in principle, by going to a sufficiently fine and broad range of strikes, but then we will be encountering liquidity risk in finding liquid prices for this large a set of vanilla calls. The basis risk can also be reduced by dynamic hedging in gold forwards the residual payments of the gold-in-gold option less the hedge package.

Even if this is not selected as a desirable hedge from a trading viewpoint, it still makes sense as a way to represent the trade from a risk management viewpoint, for the following reasons:

1) It allows realistic marking-to-market based on liquid, public prices. Alternative marking procedures would utilize an analytic pricing model for the gold-in-gold option, which is easily derivable using the PDE, but a level of volatility needs to be assumed and there is no straightforward procedure for deriving this volatility from observed market volatilities

of vanilla options at different strikes. The hedge package method recommended will converge to this analytic solution as you increase the number of vanilla option hedges used, provided all vanilla options are priced at a flat volatility (it is recommended that this comparison be made as a check on the accuracy of the implementation of the hedging package method). But the hedging package method has the flexibility to price the gold-in-gold option based on any observed volatility curve (in fact you don't use the volatility curve but the directly observed vanilla option prices, so the pricing is not dependent on any option model).

2) The hedge package method gives an easier calculation of remaining risk than the analytic method, which requires Monte Carlo simulation of dynamic hedging.

3) The hedge package method gives an easy means of integrating gold-in-gold options into standard risk reports, such as vega exposure by strike and maturity.

4) It is sometimes raised as an objection to the hedge package method that it requires transactions in vanilla options at strikes in which there is no available liquidity. But any trading desk sophisticated enough to deal in exotic options should be sophisticated enough to have a system for hedging and pricing desired positions in vanilla options at illiquid strikes with vanilla options at liquid strikes (methods we discussed in our section on Vanilla Options Risk).

These points are ones which hold generally for the replication of exotic derivatives with vanilla options. By representing the exotic derivative as closely as possible with a hedge package of vanilla options, you can minimize the remaining basis risk which needs to be managed using techniques specific to the exotic derivative and maximize the amount of risk which can be combined with and managed as part of the vanilla options book, utilizing established risk management tools such as the spot-vol matrix. Placing as much risk as possible within a single context also increases the chances that risks from one position may offset risks in another position – it is only the net risks which need to be managed.

The Excel workbook BASKETHEDGE calculates a basket of vanilla options which can replicate any non-standard payoff which is solely dependent on the price of one underlying security at a single future time. Different worksheets in this workbook illustrate the application of this technique to a variety of exotics. In addition to the gold-in-gold option, these are:

- A contract which makes payments based on the logarithm of the final price. This log contract is a particularly important product, since it has been shown to provide a static hedge for variance swaps (see Demeterfli, Derman, Kanal, and Zou, "A Guide to Variance Swaps", RISK, June, 1999).

- The convexity risk on a conststant maturity swap

- A compound option. In this case only the exposure to the price of the underlying security is being hedged, not the exposure to changes in the implied volatility of the option on which an option is being written. Hedging the implied volatility requires more general techniques of the type which we will discuss later in this section, when we treat barrier options.

Digital Options

European binary options come very close to fulfilling the previously stated condition of having a payout which is a function of the price of an asset at one definite time. Therefore it can be treated by the methodology just stated, using a basket of vanilla options to hedge it and using this hedge package to mark-to-market, including skew impact, calculate remaining risk, and incorporate into standard risk reports. Details are in Taleb's Chapter 17, in which he demonstrates that a European binary call can be replicated as closely as one wishes by the proper ratio of call spread (a vanilla call just below the barrier offset by a vanilla just above the barrier). So why do we need a special treatment? Because the discontinuous nature of the payment at the strike leads either to unrealistically large hedge positions in vanilla calls (liquidity risk, since market prices would be impacted by an attempt to transact so many calls) or significant hedge slippage (basis risk) between the binary option and its hedge.

For example, let's say a customer approaches a trading desk wanting to buy a 1-year binary call which will pay $10MM if the S&P index is above the current 1-year forward level at the end of 1 year. Let us start by assuming that all vanilla calls are priced at a 20% flat implied volatility. The straight analytical formula for the value of the binary is

$10MM x N($d_2$), where $d_2 = (\ln(price/strike) - \frac{1}{2}\sigma^2 t) / \sigma\sqrt{t}$

$= (\ln(1) - \frac{1}{2}20\%^2)/20\%$ = 10% N(-.1) = .46017, giving a price of the binary of $4,601,700.

Replicating the binary option using a vanilla call spread, the exact choice of vanilla calls to be used makes virtually no difference to the price (as long as we assume a flat implied volatility), but does make significant difference to the mix between liquidity risk and basis risk. For example:

- Buy a vanilla call on $100 billion at a strike of 99.995% of the forward level at a price of 7.9678802% for $100 billion x 7.9678802% = $7,967,880,200 and sell a vanilla call on $100 billion at a strike of 100.005% of the forward level at a price of 7.9632785% for $7,963,278,500, for a net cost of $7,967,880,200 - $7,963,278,500 = $4,601,700.
- Buy a vanilla call on $2 billion at a strike of 99.75% of the forward level at a price of 8.0812430% for $161,624,900 and sell a vanilla call on $2 billion at a strike of 100.25% of the forward level at a price of 7.8511554% for $157,023,100, for a net cost of $4,601,800.
- Buy a vanilla call on $500MM at a strike of 99% of the forward level at a price of 8.4357198% for $42,178,600 and sell a vanilla call on $500MM at a strike of 101% of the forward level at a price of 7.5152765% for $37,576,400, for a net cost of $4,602,200.

Note the inverse relationship between the width of the call spread (.01%, .50% and 2%, respectively) and the size of the legs of the call spread ($100 billion, $2 billion, and $500MM respectively).

The first combination offers the smallest basis risk — it will exactly replicate the binary option as long as the S&P index at the end of one year is outside the range 99.995% - 100.005%, that is, as long as the S&P index does not finish within about one-half point of its current forward level. But liquidity risk is heavy — purchases and sales in the size of $100 billion would be certain to move market prices if they could be accomplished at all. (Even if the trading desk does not expect to actually buy this call spread, its use in representing the risk profile of the trade will lead to illiquid dynamic hedging requirements because they are too large). At the other end of the spectrum, the third combination is of a size which could possibly be transacted without major market movement, but basis risk is now much larger. Exact replication of the binary option only takes place in a range outside 99% - 101% of the current forward, so there are about 100 points of market movement on either side of the current forward level in which replication would be inexact. And replication could be very inexact — if the index ended at 100.1% of the forward, for example, the customer would be owed $10MM but the vanilla call at 99% would only pay $500MM x 1.1% = $5.5MM, a net loss of $4.5MM.

Of course, the basis risk can be dynamically hedged with purchases and sales of S&P futures. But the large payment discontinuity of the binary can lead to unmanageable hedging situations. For example, suppose you are close to expiration and the S&P is one point below the forward level. If there is no further movement, you will make about $4.995MM on the vanilla call and owe nothing on the binary, but an uptick of just two points will lead to a loss of about $5MM. Should you put on a delta hedge of a size which will make $5MM for a two point uptick? The problem is that a position of this size will cost you $10MM for a four point downtick, and you do not gain anything from option payouts to offset this loss. While in theory, in a world of complete liquidity and no transaction costs, you could put on this hedge only at the exact moment you approach the binary strike and take it off as soon as you move away from that strike, in practice such strategies are wholly implausible. Actual experience of trading desks caught needing to delta hedge a sizable binary position which happens to be near the strike as expiration approaches is excruciatingly painful. Traders have their choice of gambles, but they must decide on a large bet in one direction or another.

In light of this, risk managers will always seek to place some sort of controls on binary positions. These controls, which may be complementary, come in the form of both limits and reserves. Limits are placed on either the size of the loss which can occur for a certain size price move, or the maximum delta position which can be required for a hedge, or the maximum gamma, the change in delta, which can be required for a given price move. Delta and gamma limits are based on the anticipated liquidity and transaction costs of the underlying market in which hedging is being done. Limits on loss size are designed to allow traders to take a purely insurance approach to binaries, hoping to come out ahead on the long run. This requires that no one binary be too large. Such an approach needs to be combined with eliminating binaries close to a strike and expiration from delta and gamma reports, so that delta hedging is not attempted. It also requires decisions about how binaries should be combined for limit purposes.

To operate like insurance, binaries need to be widely scattered as to maturity date and strike level, and limits need to bucket strikes and maturities in a manner which forces this scattering. However, bucketing should only combine binaries in one direction (bought or sold) — it is dangerous to allow netting of one binary with another except when date and strike (and any other contract terms, such as exact definition of the index) exactly match.
A valuation and reserve policy should also be consistent with the insurance approach to binaries — P&L should be recognized only the extent it can come close to being locked in. Gains which have great uncertainty attached to them should only be recognized when realized. There are several methods for accomplishing this. I will provide a detailed example of one which I consider particularly elegant in its balancing of liquidity and basis risks, its maximal use of static hedge information, and good fit with dynamic hedging risk reporting. In this approach, every binary has assigned to it an "internal" representation which is designed to be as close as possible to the binary in its payouts while still being capable of liquid hedging, and which is designed to be "conservative" relative to the binary in that the internal track will always produce a lower P&L for the firm than the binary. All risk reports for the firm are based on the internal representation, not the true representation of the binary. The MTM difference between the true and internal representation, which by design must always be a positive value to the firm, is booked to a reserve account Since the reserve is always positive, this policy sometimes results in the firm recognizing "windfall" profits but never windfall losses.

Let's see how this policy would work in the case we have been considering. A call spread will be selected as the internal representation of the binary by choosing the smallest spread with which results in a position size which is considered to be small enough to be liquid, either by representing a real possibility for purchase in the market or by being representable in the firm's risk reports by delta positions which can be achieved with reasonable liquidity. But rather than choosing a call spread which straddles the binary, and which therefore has payouts greater than the binary in some scenarios, we choose a call spread which is on one side of the binary and therefore always has payouts greater than the binary. If 2% is the width of call spread we select as the smallest consistent with a liquid position, then we will use as an internal representation a call spread consisting of a sale of $500MM at a strike of 98% and a purchase of $500MM at a strike of 100% (notice that the internal representation has the opposite sign from the hedge which would extinguish it). The resulting MTM would be $500MM x8.9259724% - $500MM x 7.9655791 = $44,629,900 - $39,827,900 = $4,802,000. This is the MTM of the internal representation. The actual binary continues to be MTM at $4,601,700 — the difference of $200,300 is placed into a reserve. If the actual sale price of the binary to a customer is $5MM, then only $200M of the profit from the difference between the price and MTM goes into immediate P&L recognition, the other $200M goes into a reserve against anticipated liquidity costs of managing the digital risk.

What happens to this reserve? There are several possibilities:

- The firm might decide to actually buy the static overhedge which costs $4,802,000. The internal hedge reports of the firm will now show no net position between the internal representation of the binary and the actual call spread hedge. If the S&P index ends up below 98% or above 100%, there will be no difference between the eventual payout under the binary and the payin due to the call spread, and the reserve will end up at 0. If the S&P index ends up between 98% and 100%, the call spread will have a payin while there is no payout on the binary. For example, if the S&P index ends at 99%, the call spread will pay $5MM, which will be the final value of the reserve. At expiry of the options, this $5MM will be recognized in P&L as a windfall gain.

- The firm might not do any static hedging and just delta hedge based on the internal representation of the static overhedge. Since the static overhedge was selected to be of a size which allows liquid delta hedging, the results in this case should be close to the results in the case that the static overhedge is actually purchased, but with some relatively small variance. As an example, suppose that we are very close to expiry and the S&P index forward is at 99%. Based on the internal representation of the call spread overhedge, the appropriate delta will be a full $500MM long in the S&P index forward and roughly $5MM in dynamic hedging profits should already have been realized but held in reserve. If the index ends at 99%, the $5MM in dynamic hedging profits will be taken from the reserve and recognized in P&L as a windfall gain. If the index ends just above 100%, the $5MM in dynamic hedging profits realized to date plus the $5MM gain from the 1% increase on the $500MM long in the S&P index will be exactly enough to pay the $10MM owed on the binary. Note that keeping the $5MM in dynamic hedging profits realized to date in reserve is necessary to avoid having to reverse a previously recognized gain in order to pay off on the binary.
- Other combinations are possible, such as static hedges which are not overhedges, but all produce similar results.

This technique of representing a binary internally as a static overhedge is sometimes objected to by front office personnel as trading off a very probable gain in order to achieve security. In this view, the $400M which was originally realized on the transaction was real P&L and $200M was sacrificed in order to achieve security in the very small minority of cases in which the index finishes very close to the strike. The idea that $200M has been thrown away is, in fact, an “optical” illusion caused by focusing only on those cases in which the index finishes outside the 99% to 101% range. The trade still has a $400M expected value — it just consists of a sure $200M in the vast majority of cases in which the index finishes outside 99% - 101% and a set of windfall profits up to $10MM when the index finishes within this range. The front office view would be correct if there were some means, such as dynamic hedging, of being “almost” sure of achieving this $400M result in all cases, but it was exactly the lack of such means, the fact that dynamic hedging to try to come close to achieving $400M in all cases results in some cases with disastrous losses, which cased us to seek an alternative approach. This reserve methodology can be seen to be consistent with moving the front office away from viewing these trades as normal derivatives trades which can be approached in an isolated manner and towards viewing them as necessarily being part of a widely diversified portfolio of binaries. In this context, over a long enough time period, the sum of occasional windfall gains can become a steady source of income. If limits can insure a wide enough diversification, then reserves may not be necessary.

So far in the example we have assumed a lack of volatility skew. In the presence of skew, the binary will price quite differently. Let's see the impact of using a 20.25% implied volatility for a strike of 99% and a 20% volatility for a strike of 101%. The cost of the 99% vanilla call is now 8.534331%, resulting in a net cost of $5,095,274. Just as with the other non-standard payment cases previously discussed, the reduction to a package of vanilla options let's us pick up the impact of volatility skew. We can see that binary options are highly sensitive to skew.

Taleb, on p. 286, says "the best replication for a digital is a wide risk reversal (that would include any protection against skew). There will be a trade-off between transaction costs and optimal hedges. The trader needs to shrink the difference between the strikes as time progresses until expiration, at a gradual pace. As such an optimal approach consumes transaction costs, there is a need for infrequent hedging." Using a call spread (also known as a "risk reversal") which is wide reduces the size of the vanilla options which are needed, reducing transaction costs and liquidity concerns, and also capturing the volatility skew more accurately, since a wide spread could utilize more liquid strikes. As we have seen, the width of the spread should not materially impact the total hedge cost.

In many cases, the underlying price will finish nowhere near the strike and no further transactions are needed. But in those cases where the underlying is threatening to finish close to the strike, basis risk will get too large and the trader will need to roll from the original call spread into a tighter call spread, incurring transactions costs both due to the need to purchase and sell options and because the size of the options transactions is growing as the spread narrows. This potential transactions cost needs to be factored into the valuation of binary options. Following Taleb, on p. 286, "when the bet option is away from expiration, the real risks are the skew. As it nears expiration, the risks transfer to the pin. In practice, the skew is hedgable, the pin is not.

Barrier Options

So far we've dealt strictly with exotic options whose payment is based on the price of an asset at a single time period, i.e., European-style options. Next we want to look at how an option which is based on asset prices at many time periods can be handled. Barrier options are a good choice because they illustrate dependence on the entire volatility surface, both in terms of time and strike level,because they have a large range of variants, because they are overwhelmingly the most traded exotic options among foreign exchange options and are also used with equities, commodities, and interest rates, and because they can be used as building blocks in forming static hedges for other exotic options, such as lookback options, roll-down options, and ratchet options.

A barrier option is one whose payoff is equal to that of a standard call or put but which only pays off under the condition that some price level (called the barrier) has been breached (or not) at some time period prior to the time the call or put payoff is determined. Options which only pay if a barrier has been breached are called knock-in ("down and in" if the barrier is below the asset's price at the time the option in written, "up and in" otherwise). Options which only pay if a barrier has not been breached are called knock-out (either "down and out" or "up and out"). Variations include double barrier options which either knock out if either a down and out or an up and out condition has been reached or which knock in if either a down and in or up and in condition has been reached. Another variation is a delayed start barrier, where the barrier condition can be activated only during a specified time period which begins after the option start date.

Barrier options may or may not have significant digital features. Table 19.1 in Taleb and the surrounding text offer a good discussion of this point. When significant digital features are involved, the same type of limit, hedge adjustment, and reserving issues we discussed for European digitals are needed. But the most important feature of barrier options, shared by all barrier options, is that they are dependent on large parts of the implied volatility surface, which is a sharp contrast with European options. For example, to price a one-year European call at a strike of 100, you only need to know how the market is pricing probability distributions for the price one year from now, with almost all the sensitivity being to the part for the probability distribution close to 100. By contrast, a one-year up-and-out call with a strike at 100 and barrier at 110 is sensitive to probability distributions of price at all times up to one year (since you need to know how likely it is that the knock-out condition will be met) and is very sensitive to the price distribution close to 110 as well as close to 100.

While good analytic models based on PDEs have been developed for both barrier and double barrier options (see Hull, p. 462-464 for the single barrier equations), these models have the drawback that they need to assume a single level of volatility and there are no good rules for translating a volatility surface observed for European options into a single volatility to be used for the barrier options. In fact, cases can be shown where no single volatility assumption can be utilized with the PDE approach to give a reasonable price for the barrier option. We will illustrate this point with the following example:

Asset price (S)	**100.00**
Strike price (X)	**100.00**
Barrier (H)	**120.00**
Cash rebate (K)	**0.00**
Time to maturity (T)	**0.25**
Risk-free rate (r)	**0.00%**
Cost of carry (b)	**0.00%**
Volatility (s)	**20.00%**

Value of Up-and-Out Call with skewed volatility of 20% at 100, 18% at 120

Using Local Volatility Model	3.090
Using Carr Static Hedge Model	3.095

Value of Up-and-Out Call Using Standard PDE Formula

Volatility	Value of Up-and-out Call
1.00%	**0.1995**
2.00%	**0.3989**
3.00%	**0.5984**
4.00%	**0.7979**
5.00%	**0.9973**
6.00%	**1.1968**
7.00%	**1.3962**
8.00%	**1.5956**
9.00%	**1.7942**

10.00%	1.9897
11.00%	2.1772
12.00%	2.3499
13.00%	2.5008
14.00%	2.6242
15.00%	2.7166
16.00%	2.7771
17.00%	2.8070
18.00%	2.8087
19.00%	2.7858
20.00%	2.7421
21.00%	2.6816
22.00%	2.6080
23.00%	2.5245
24.00%	2.4340
25.00%	2.3390
26.00%	2.2415
27.00%	2.1432
28.00%	2.0455
29.00%	1.9492
30.00%	1.8552

Note that the analytic result, which utilizes the formula shown in Hull, has option values that first increase as the volatility level rises, since rising volatility causes the call value to increase. At higher volatility levels the option values decrease as the volatility level rises, since rising volatility increases the probability of knock-out. Since the barrier level starts far away from the current price, it is only at high volatilities that the impact of rising volatility on probability of knock-out dominates the impact of rising volatility on the value of the call.

Methods for utilizing the full volatility surface, which we shall discuss shortly, would agree with those analytical results for flat volatility surfaces. But if we assume a skewed volatility surface, with implied volatility of 20% for a European call struck at 100 and of 18% for a European call struck at 120, approaches which utilize the full volatility surface (either the Derman-Kani trinomial tree approach or the Carr static hedging approach) would price the barrier option at 3.10, which is 10% higher than the 2.81 maximum value the barrier option reaches at any volatility level using the analytic approach. The reason for this is that the lower volatility as you approach the barrier decreases the chance of penetrating the barrier without simultaneously lowering the value of the call.

The following approaches to pricing and hedging barrier options using the full volatility surface are possible:

A. <u>Dynamic Hedging Models</u>

These models price barrier options (or any other exotic option whose payoff is a function of a single underlying asset) based on the cost of dynamically hedging the exotic with a portfolio of the underlying asset and vanilla European options. This is analogous to the Black-Scholes model pricing of vanilla European options based on the cost of dynamically hedging with the underlying asset. These models utilize the full set of current prices of vanilla European options, and so make use of the full volatility surface, along with a theory of how these vanilla option prices can evolve with time. If you utilize an actual dynamic hedging strategy consistent with the model, you will be successful in replicating the model's price for the exotic to the extent that the model's theory about the evolution of the vanilla options prices is correct and to the extent that transactions costs are reasonably insignificant.

There are two principal types of dynamic hedging models used for exotics:

1) Local volatility models which assume that volatility is a known and unvarying function of time and underlying price level. These models are natural extensions of the Black-Scholes model which assumes that volatility is known and unvarying but also assumes it is the same at all times and underlying price levels. Based on the assumption of the local volatility model, you can derive a definite price at any future time and underlying price level of any vanilla or exotic option. The cost of the dynamic hedge will therefore differ from the originally derived price only to the extent that future volatilities prove to follow a varying function of time and underlying price level (or that transaction costs are significant).

2) Stochastic volatility models which assume that volatilities will vary over time based on some assumed model. The cost of the dynamic hedge will differ from the derived price to the extent that the process of actual volatility variation differs from that assumed by the model (or that transaction costs are significant).

B. <u>Static Hedging Models</u>

These models price barrier options based on the cost of a replication strategy which calls for an almost unvarying hedge portfolio (at least of the vanilla options; it would be possible to use a dynamic hedge of the underlying, though the particular static hedging models we will discuss only utilize vanilla options in the hedge portfolio). These models utilize nearly static hedge portfolios both as a way to reduce transaction costs (which can be considerably higher when buying and selling vanilla options than when buying and selling the underlying) and as a way to reduce dependence on assumptions about the evolution of volatility. Three approaches to static hedging of barriers can be distinguished:

1) The approach of Derman, Ergener, and Kani, which is broadly applicable to all exotic options whose payoff is a function of a single underlying asset, but which has considerable exposure to being wrong about future volatility levels.

2) The approach of Peter Carr and his colleagues which is more specifically tailored to barrier options, utilizing an analysis of the Black-Scholes formula to form a hedge portfolio which is immune to changes in overall volatility level. However, the Carr approach is still vulnerable to changes in the volatility skew. It is easier to implement than the Derman-Ergener-Kani approach in cases of single barriers in the absence of drift

(i.e., forward equal to spot), but harder to implement for double barriers and in cases where forward does not equal spot.

3) Approaches which utilize optimal fitting give solutions close to those provided by the Carr approach for single barriers in the absence of drift, but are more flexible in handling drift and double barriers, and are less vulnerable to changes in volatility skew.

The following table compares these different approaches to pricing and hedging barriers on a number of criteria:

		Barrier Options managed using			
Risk exposure to:	Vanilla options book managed using portfolio approach	Dynamic hedge	Derman-Engener-Kani static hedge	Carr static hedge	Optimization static hedge
Price jump	Low	None	None	None	None
Parallel vol shift	Low	None	High	None	None
Vol time shift	None	None	High	None	Low
Vol smile shift	None	None	High	None	Low
Vol skew shift	None	None	High	High	Low
Dynamic hedging costs	Low and moderately easy to estimate using Monte Carlo	Hard to estimate	None	None	None
Transaction costs	Low and moderately easy to estimate using Monte Carlo	High and hard to estimate	Low and easy to estimate	Low and easy to estimate	Low and easy to estimate
Flexibility to extend to other cases		Extends easily to all structures	Extends moderately easily to all structures	Can extend with greatly added complication to barriers with drift, double barriers, no obvious extension to other cases	Extends moderately easily to all structures

A relatively straightforward implementation for a local volatility model is the trinomial tree approach of Derman and Kani (see Derman and Kani, “Riding on a Smile”, RISK, Feb. 94), which builds the unique trinomial tree for modeling the price diffusion of the underlying and which meets the following two criteria:

a) Volatility is a known and unvarying function of time and underlying price level

b) The tree correctly prices **all** European calls and puts on the underlying at different strike levels and time to expiry.

Dynamic hedging utilizes the full volatility surface in pricing barrier options and can be readily utilized for representing the barrier option in risk reports through its vanilla option hedges, and can easily be applied to any derivative based on a single underlying. Its drawback is its vulnerability to incorrect assumptions about volatility evolution and to transaction costs of buying and selling vanilla options. Vanilla options typically have far lower liquidity and far higher transaction costs than non-option underlyings such as spot and forward contracts. It is also quite difficult to calculate the potential risks of incorrect assumptions and probable transaction costs relative to similar calculations for the dynamic hedging of vanilla options. When performing a Monte Carlo simulation of the dynamic hedging of a vanilla option, the underlying hedge at each time step along each random path can be easily calculated using the Black-Scholes delta formula. By contrast, a Monte Carlo simulation of the dynamic hedging of a barrier requires the recalibration of the entire local or stochastic volatility model, along with its sensitivities to changes in vanilla option prices, at each time step along each random path in order to calculate the required change in hedge.

Both the Derman-Ergener-Kani approach and the Carr approach are based on the idea of finding a basket of vanilla options which statically replicate the differences between the barrier option and a closely related vanilla option. To facilitate discussion, we will confine ourselves to the case of a knock-out call, since a knock-in call can be handled as a vanilla call less a knock-out call, and puts can be priced from calls using put-call parity. The idea is to purchase a vanilla call with the same strike and expiration date as the knock-in being sold and then reduce the cost of creating the knock-in by selling a basket of vanilla options (this basket may have purchases as well as sales in the Derman-Ergener-Kani approach, but either way the net initial cash flow on the basket is positive to the barrier option seller). The basket of vanilla options must be constructed so that:

a) It has no payoff if the barrier is never hit. In this case the payout on the barrier option, which has not been knocked-out, is exactly offset by the payin from the vanilla call which was purchased, so there is nothing left over to make payments on the basket.

b) Its value when the barrier is hit is an exact offset to the value of the vanilla call. When the barrier is hit, you know you will not need to make any payments on the barrier option, so you can afford to now sell the vanilla call you purchased. You do not want to later be vulnerable to payouts on the basket of vanilla options you sold, so you must purchase this basket. In order for cash flows to be zero, the basket purchase price must equal the vanilla call sale price.

You can guarantee the first condition by only using calls struck at the barrier in the case of a barrier higher than the current price and by only using puts struck at the barrier in the case of a barrier lower than the current price. If the barrier is never hit, then certainly you won't be above the up barrier at expiration, so you won't owe anything on a call, and you certainly won't be below the down barrier at expiration, so you won't owe anything on a put.

Where the Derman-Ergener-Kani and Carr approaches differ is in how they attempt to assure that the option package will be equal in value to the vanilla call at the time the barrier is hit. Both take advantage of knowing that at the time you are reversing your position in these vanilla options the underlying must be at the barrier. Note that both approaches will have the advantage that the only possible transaction costs involve a single event, the reversing of all positions when the barrier is hit. Therefore, the possible size of transaction costs is a straightforward calculation.

The Derman-Ergener-Kani approach (see Derman, Ergener. And Kani, "Forever Hedged", RISK, Sep. 94) uses a package of vanilla options which expire at different times. The algorithm works backwards, starting at a time close to the expiration of the barrier option. If the barrier is hit at this time, the only vanilla options still outstanding will be the vanilla call and the very last option to expire in the package. Since both the underlying price is known (namely, the barrier) and the time to expiry is known, the only remaining factor in determining the values of the vanilla options is the implied volatility, which can be derived from a local or stochastic volatility model (if from a stochastic volatility model, it will be based on expected values over the probability distribution). Thus the Derman-Ergener-Kani approach can be viewed as the static hedging analogue of the dynamic hedging approaches we have been considering. Once the prices of the vanilla options at the time the barrier is hit are calculated, you can easily determine the amount of the option that is part of the basket which needs to be sold in order to exactly offset the sale of the vanilla call with the purchase of the option in the basket. You then work backwards time period by time period, calculating the values of all vanilla options if the barrier is hit at this time period and calculating what volume of the new option in the basket is needed to set the price of the entire basket equal to the price of the vanilla call. At each stage, you only need to consider unexpired options, so you only need to consider options for which you have already computed the volumes held. A detailed example illustrating this technique can be found in Hull, section 18.8.

The following points about the Derman-Ergener-Kani approach should be noted:

(1) If the barrier is hit in between two time periods for which vanilla options have been included in the package, the results are approximated by the nearest prior time period. The inaccuracy of this approximation can be reduced as closely as you wish by increasing the number of time periods used.
(2) The approach can easily accommodate all sorts of complexities, such as the existence of drift (dividend rate unequal to risk free rate), time-varying barrier levels (such as forward-starting barriers) and double barriers, since a separate computation is made for each time the barrier could potentially be hit.

(3) Since the approach relies on the results of a local or stochastic volatility model to forecast future volatility surface levels and shapes, it is vulnerable to the same issue as when these models are used for dynamic hedging — the hedge will only work to the extent that the assumptions underlying the model prove to be true. This is illustrated in the following table, which shows the potential mismatch in unwind cost at a period close to expiry based on differences between model assumed volatilities and actual volatilities at the time the barrier is hit.

Asset price (S)	**100.00**
Strike price (X)	**100.00**
Barrier (H)	**95.00**
Time to maturity (T)	**0.25**
Risk-free rate (r)	**0.00%**
Cost of carry (b)	**0.00%**
Volatility (s)	**20.00%**

Value of Purchased Down-and-Out Call Using Standard PDE Formula
Hedged Using Vanilla Options According to Derman-Ergener-Kani Method Based on 20% Volatility

VOL	DO Call	Hedge	Net
10.00%	**1.9576**	1.8606	0.0970
11.00%	**2.1274**	2.0117	0.1157
12.00%	**2.2862**	2.1573	0.1289
13.00%	**2.4337**	2.2981	0.1356
14.00%	**2.5703**	2.4348	0.1355
15.00%	**2.6964**	2.5680	0.1284
16.00%	**2.8128**	2.6981	0.1147
17.00%	**2.9201**	2.8256	0.0945
18.00%	**3.0192**	2.9508	0.0684
19.00%	**3.1108**	3.0740	0.0368
20.00%	**3.1955**	3.1955	0.0000
21.00%	**3.2740**	3.3151	-0.0411
22.00%	**3.3469**	3.4335	-0.0866
23.00%	**3.4146**	3.5507	-0.1361
24.00%	**3.4778**	3.6668	-0.1890
25.00%	**3.5367**	3.7819	-0.2452
26.00%	**3.5917**	3.8960	-0.3043
27.00%	**3.6433**	4.0094	-0.3661
28.00%	**3.6916**	4.1220	-0.4304
29.00%	**3.7371**	4.2340	-0.4969
30.00%	**3.7798**	4.3453	-0.5655

Note that the Derman-Ergener-Kani approach is vulnerable to model errors both as to level of volatility surface and skew of volatility surface and skew of volatility surface.

The Carr approach (see Carr, Ellis and Gupta,"Static Hedging of Exotic Options", Journal of Finance, 1998, vol. LIII(3), p. 1165 – 1190) avoids this dependence on projecting future volatility surfaces and is much simpler to implement, but at a price — it cannot handle volatility skews (though it can handle volatility smiles) and its simplicity depends on an absence of drift (dividend rate equals risk-free rate).

The Carr approach achieves a degree of model independence by working directly with the Black-Scholes equations and determining hedge package which will work providing only that volatility is flat. In these circumstances, one can calculate exactly a single vanilla put which will be selling at the same price as the vanilla call in the case that a down barrier is hit. It is based on the principle of put-call symmetry, which states that for all strike pairs, K_1 and K_2, such that $\sqrt{K_1 K_2}$ =forward price.

$$\sqrt{K_2}Call(K_1) = \sqrt{K_1}Put(K_2)$$

This is a direct consequence of the Black-Scholes formulas. Since there is no drift, the forward price is equal to the spot price, which is the barrier level, H. Since the call is struck at K, we can find a reflection strike, R, such that $\sqrt{KR}$ = H and by put-call symmetry,

$$\sqrt{R}Call(K) = \sqrt{K}Put(R)$$

Since $\sqrt{KR} = H$, $R = H^2/K$, $\sqrt{R} = H/\sqrt{K}$ so you need to purchase $\frac{\sqrt{K}}{\sqrt{R}} = \frac{K}{H}$ puts struck at H^2/K

For an up barrier, one must separately hedge the intrinsic value and the time value of the vanilla call at the time the barrier is hit. The intrinsic value can nearly be perfectly offset by selling binary options which pay 2 x I, the intrinsic value. Any time the barrier is hit, there will be nearly a 50-50 chance that the binary will finish in the money, so its value is close to 50% x 2 x I = I. In fact, standard lognormal pricing of a binary will result in assuming slightly less than a 50% chance of finishing above the barrier so we need to supplement the binary with I of a plain vanilla call struck at the barrier. The exact value of the binary when the barrier is hit is $2 \times I \times N(-\frac{\sigma\sqrt{\tau}}{2})$ and the value of the vanilla call struck at the barrier, and hence exactly at-the-money when the barrier is hit is $I \times (N(\frac{\sigma\sqrt{\tau}}{2}) - N(-\frac{\sigma\sqrt{\tau}}{2}))$

$$= I \times (1 - 2N(-\frac{\sigma\sqrt{\tau}}{2}))$$

The sum of these two terms is then exactly I.

The Carr approach has several advantages:

- It shows that there is at least a plausible way of pricing the barrier based on options with tenor equal to the final tenor of the barrier, indicating that this is probably where most of the barrier's risk exposure is coming from.
- By finding an equivalent package of options we already know how to represent in standard risk reports, it allows the barrier to be accommodated in standard risk reports. This is true independent of whether we actually choose to perform this hedge.
- Since we know how to price all the pieces of the hedge package using the volatility skew, this provides us with pricing for the barrier which accommodates skew.
- Having a large binary component of the hedge is an excellent means of highlighting and isolating the "pin" risk contained in this barrier which dies in-the-money. Techniques we have already developed for managing pin risk on binaries can now easily be brought into play. For example, we could establish a reserve against the pin risk of the binary. This approach is quite independent of whether the trading desk actually sells a binary as apart of the hedge — the risk of the binary is present in any case.
- Because the Carr approach uses a small number of options in the hedge package, it is very well suited for intuitive understanding of how changes in the shape of the volatility surface impact barrier prices.
- Even if you choose to hedge and price using a dynamic hedging approach, the Carr methodology can still be useful in identifying cases which are relatively insensitive to all the assumptions which need to be made in choosing between competing dynamic hedging models. Since the Carr methodology constructs its static hedge using vanilla options, any model which is calibrated to market prices has identical prices for these vanilla options and so must have identical prices for the barrier option.
- Neither the presence of volatility smiles, nor uncertainty as to future volatility smiles, impacts the Carr approach. Since it deals with options which are symmetrically placed relative to the at-the-money strike, all smile effects cancel out.

The simplicity of the Carr approach is lost in the presence of drift of for double barriers (see the appendix to Carr & Chou "Breaking Barriers", on Carr's website, for a method of using a large number of vanilla options to create a volatility independent static hedge of barrier options in the presence of drift; see Carr, Ellis, and Gupta, "Static Hedging of Exotic Options" for a method of using a large number of vanilla options to create a volatility level-independent static hedge of double barrier options). The Carr approach cannot handle changes in drift or any volatility shape other than a symmetric smile.

A more general approach to static hedging, which can handle all drift and volatility shape conditions, is optimization, in which a set of vanilla options is chosen which fits as closely as possible, the unwind of the barrier option at different possible times, drifts, and volatility levels and shapes which may prevail when the barrier is hit. The optimization approach is discussed in Dembo's "Hedging in Markets that Gap", Handbook of Derivatives and

Synthetics. It is often the case that no perfect static hedge can be found, but in these cases the optimization produces information on the distribution of possible hedge errors which can be useful input to determining a reasonable reserve. A similar approach can be taken to many different types of exotic structures.

The spreadsheet STATICOPT illustrates how optimization can be used to find a static hedge for a barrier option. If the possible conditions when the barrier is hit are restricted to zero drift and volatility smile but no skew, then the EXCEL Solver will find a set of vanilla options which almost exactly matches the barrier unwind for all volatility levels and times to expiry. Of course, this is not a surprise, since we know from the Carr approach that a perfect static hedge is possible under these circumstances. When different non-zero drift and volatility skew conditions are allowed, the match of the barrier unwind is no longer as exact.

Combinations of Exotic Options

Some exotic options can be constructed as static combinations of other exotics. If we already know how to use vanilla options to hedge the exotic options which comprise the static combination, then the static combination automatically extends to a vanilla hedge for the other exotics.

A simple example is a contingent premium option which entails no initial payment by the option buyer, who only pays at option termination under the circumstances that the option finishes in the money. This type of option is popular with some clients both because of the deferral of cash payment and because the client will not need to pay for an option which turns out to be "useless," although it should be noted that an option which finishes just slightly in the money will still require a net payment by the option buyer, since the payment due from the option seller will be less than the option's cost. It is easy to see that a contingent premium option is just a standard vanilla option plus a forward to defer payment of the option premium plus a binary option to offset the option premium due in the event the price finishes below (in the case of a call) the strike of the vanilla option.

As one further example, let's see how to construct a static hedge for a type of exotic option called a lookback option. In this case, we will construct the static hedge using a package barrier options.

Lookback calls come in two varieties: (1) those which pay the difference between the maximum price which an asset achieves during a selected period and the closing price, and (2) those which pay the difference between the maximum price which an asset achieves during a selected period and a fixed strike. Symbolically, the lookback either pays)1) S_{max} - S_T of (2) max $[0,S_{max} - K]$. We can exactly reproduce the payoffs of a lookback of the first type by buying a lookback of the second type with a strike equal to the current price of the asset, selling the asset forward to time T and buying a forward delivery of S_0 dollars at time T. Since S_{max} is certainly $\geq S_0$, max $[0,S_{max} - S_0] = S_{max} - S_0$, the total payoff if this combination at time T is $\max[0,S_{max} - S_0] - S_T + S_0 = (S_{max} - S_0) - S_T + S_0 = S_{max} - S_T$. So if we can create the second type of lookback option by static hedging, we can create the first type of lookback option by static hedging as well.

Lookback options have a closely related product called ratchet options which pay max[0, S_{max} - K] rounded down by a specified increment. For example, if K = 100 and S_{max} = 117.3, the lookback call of the second type would pay 17.3, a ratchet call with increments of 1 would pay 17, a ratchet call with increments of 5 would pay 15, and a ratchet call with increments of 10 would pay 10. Since a lookback call can be approximated as closely as we want by a ratchet call with a small enough increment, it is sufficient to show how to statically hedge a ratchet call.

In section 3.3 of "Static Hedging of Exotic Options" by Carr, Ellis, and Gupta, they show how to create a static hedge for a ratchet call using barrier calls. Using the notation of KI (strike, trigger) for a knock-in call, a ratchet call with strike of K and increment of I can be reproduced by the following sum of knock-in calls:

$$(\sum_{J=1}^{N} KI(0, K + J * I) - KI(I, K + J * I))$$

There is no protection above K+N*I, so N must be chosen large enough that this is not a serious risk. For any given increment K+J*I, if $S_{max} \le$ K+J*I, the knock-ins don't occur and the payouts on both KI(0,K+J*I) and KI(I,K+J*I) are 0. If S_{max} >K+J*I, the knock-ins do occur, and the net payouts of KI(0,K+J*I) -KI(I,K+J*I) = I exactly match the payouts due on the ratchet.