Hedging Barrier Options: Current Methods and Alternatives

Dominique Y. Dupont*

EURANDOM - TUE

P.O. Box 513 - 5600 MB Eindhoven

Phone: 040 247 8123

Fax: 040 247 8190

Email: dupont@eurandom.tue.nl

February 2001

Abstract

This paper applies to the static hedge of barrier options a technique, mean-square hedging, designed to minimize the size of the hedging error when perfect replication is not possible. It introduces an extension of this technique which preserves the computational efficiency of mean-square hedging while being consistent with any prior pricing model or with any linear constraint on the hedging residual. This improves on current static hedging methods, which aim at exactly replicating barrier options and rely on strong assumptions on the availability of traded options with certain strikes or maturities, or on the distribution of the underlying asset.

Derivatives markets have become an essential part of the global economic and financial system in the past decade. Growth in trading activity has been particulary fast in over-the-counter (OTC) markets, which have become the dominant force in the industry.[1] One factor in the success of OTC markets has been their ability to offer new, "exotic" derivatives customized to customers' needs. Among the most successful OTC option products are barrier options. The payoff of a barrier option depends on whether the

---

[1] As of December 1999, OTC transactions accounted for around 86 percent ($88 trillion) of the total notional value of derivatives contracts, exchanges for about 14 percent ($14 trillion). (source: IMF International Capital Markets, Sept. 2000.)

price of the underlying asset crosses a given threshold (the barrier) before maturity. The simplest barrier options are "knock-in" options which come into existence when the price of the underlying asset touches the barrier and "knock-out" options which come out of existence in that case. For example, an up-and-out call has the same payoff as a regular, "plain vanilla," call if the price of the underlying asset remains below the barrier over the life of the option but becomes worthless as soon as the price of the underlying asset crosses the barrier.

Barrier options fulfill different economic needs and have been widely used in foreign-exchange and fixed-income derivatives markets since the mid 1990s. (Steinherr (1998) and Taleb (1996) discuss issues linked with using barrier options.) Barrier options reduce the cost of modifying one's exposure to risk because they are cheaper than their plain-vanilla counterparts. They help traders who place directional bets enhance their leverage and investors who accept to keep some residual risk on their books reduce their hedging costs. More generally, barrier options allow market participants to tailor their trading strategies to their specific market views. Traders who believe in a market upswing of limited amplitude prefer to buy up-and-out calls rather than more expensive plain vanilla calls. In contrast, fund managers who use options to insure only against a limited downswing prefer down-and-out puts. Using barrier options, investors can easily translate chartist views into their trading strategies. (Clarke (1998) presents detailed applications of these instruments to technical trading in foreign-exchange markets.) Options with more complex barrier features are also traded. Naturally, traders who use barrier options to express their directional views on the market do not want to hedge their positions; they use such options to gain leveraged exposure to risk. In contrast, their counterparties may not be directional players with opposite views but rather dealers who provide liquidity by running barrier option books and hedge their positions.

Barrier options are difficult to hedge because they combine features of plain vanilla options and of a bet on whether the underlying asset price hits the barrier. Merely adding a barrier provision to a regular option significantly impacts the sensitivity of the option value to changes in the price or the volatility of the underlying asset: The "Greeks" of barrier options behave very differently from those of plain vanilla options (Figure 1).[2] The discontinuity in the payoff of barrier options complicates hedging, especially those that are in the money when they come out of existence (such as up-and-out calls and down-and-out puts). Delta hedging these options is particularly unpractical. For example, when the price nears the barrier and the option is about to expire, the Delta (and the Gamma) of an up-and-out call takes large negative values because the option payoff turns into a spike in this region. Vega also turns negative when the price of the underlying asset is close to the barrier because a volatility pickup near the barrier increases the like-

---

[2]The Greeks are the sensitivities of the option price to changes in the parameters and are used to assess risk. They include Delta (the first-order derivative of the option value with respect to the price of the underlying asset), Gamma (its second-order derivative with respect to the same parameter), and Vega (the derivative of the option value with respect to the volatility).

lihood of the price passing through it. In contrast, the Delta, Gamma, and Vega of regular options are always positive and "well behaved" functions.

Despite the difficulties associated with hedging barrier options, banks write large amounts of those instruments to accomodate customer demand and typically prefer to hedge their positions. Moreover, when devising hedging strategies, banks have to face the limitations of real markets, like transaction costs, discrete trading, and lack of liquidity. One way of addressing these concerns is to limit the frequency of trading and, in the limit, to solely consider static hedging. To compensate the associated loss of flexibility, one may prefer to hedge barrier options with regular options instead of with the underlying asset and a riskless bond because regular options are more closely related to barrier options. Derman, Ergener, and Kani (1995) (thereafter DEK) and Carr, Ellis, and Gupta (1998) (thereafter CEG) model static hedging of barrier options with regular options on the same underlying asset. They are able to achieve perfect replication of the barrier option but, to achieve this result, they need strong assumptions on the availibility of regular options with certain strikes or maturities, or on the distribution of the underlying asset. Their assumptions are not all satisfied in actual markets and DEK and CEG do not present a structured approach on how to adapt their techniques in this case.

The paper presents an alternative methodology. Instead of aiming at perfectly replicating a barrier option, we choose to approximate it. The metric used to gauge the fit of the hedge is the mean of the square of the hedging residual (the difference between the payoff of the claim and that of the hedging portfolio), leading to the "mean-square hedging" method introduced in Duffie and Richardson (1991) and Schweizer (1992). We also extend the mean-square hedging methodology to make the hedging portfolio consistent with linear constraints. For example, we can incorporate the requirement that the value of the portfolio coincide with the value that a given pricing model affects to the barrier option.

Financial economists have developed other hedging methods in incomplete markets. The "super-hedging" strategy of El Karoui and Quenez (1995) satisfies investors who want to avoid any shortfall risk (the risk that the value of the hedging portfolio be below that of the asset it is designed to hedge). For investors willing to hedge only partially, "quantile hedging" achieves the lowest replication cost for a given probability of shortfall and fits well the popular value-at-risk (VaR) approach (Föllmer and Leukert (1999)). However, Vorst (2000) indicates that portfolio optimization under VaR constraint can lead to unattractive risk profiles.

An important difference between the present paper and the approaches cited above, or the current literature on mean-square hedging, is that this paper aims at providing a framework that investors can use to implement mean-square hedging to a case of practical importance: barrier options. To do so, it uses only discrete-time techniques and can be adapted to any pricing model. In contrast, current research focuses on the theoretical issues of introducing mean-square or quantile hedging in continuous-time models where jumps or stochastic volatility render markets incomplete.

The remainder of the paper is organized as follows. Section 1 briefly

3

reviews hedging principles and summarizes the contributions of DEK and CEG. Section 2 introduces the theory of mean-square hedging and section 3 applies this technique to hedging an up-and-out call in an environment similar to the one used in DEK. Section 4 presents our own extension to mean-square hedging and section 5 applies it to hedging an up-and-out call under some constraints on the residual.

1   Hedging Barrier Options

Financial modelers typically rely on the assumption that markets are complete and compute the prices of derivative assets via hedging arguments. In contrast, practitioners of finance face the limitations of real markets and the impossibility of perfect hedging.

1.1   Hedging and Pricing in Complete and Incomplete Markets

Pricing derivatives securities typically uses a hedging argument based on the assumption that markets are complete or dynamically complete, that is, that any claim can be replicated by trading marketable assets because there are enough instruments or because trading can take place sufficiently often. Under these conditions, the no-arbitrage price of any derivatives product is the cost of the replicating portfolio and securities can be valued independently of investors' preferences. Black and Scholes (1973) derive the price of plain vanilla options and Merton (1973) that of barrier options (and other derivatives) by showing how continuous trading in the underlying stock and in a riskless bond replicates the option's payoff in every state of the world.

In contrast, when claims cannot be perfectly replicated by trading marketable assets, any hedging strategy leaves some residual risk and investors' attitudes toward risk affect the pricing of the claims. For example, Hull and White (1987) introduce stochastic volatility in a model otherwise similar to Black-Scholes. As a result, markets are incomplete and to price a call option on the stock, they assume that investors demand no compensation for the risk generated by the randomness of the volatility of the stock returns. Introducing transaction costs or non-continuous trading into the basic Black-Scholes framework also renders markets incomplete.

In practice, market frictions are unavoidable and hedging barrier options (or other assets) dynamically would be prohibitively costly; practitioners rebalance their hedging portfolios less frequently than possible as a way to save on transaction costs, adopting instead static hedging strategies. Pricing models, however, rarely take market frictions into account but are instead sophisticated enhancements of the Black-Scholes approach. As a result, pricing and hedging, which are two sides of the same coin in complete-market models, are rather disconnected in practice.

In the following, we consider the problem of hedging an up-and-out call. Static hedging strategies have been introduced in DEK and CEG. In this context, a "static strategy" involves the constitution of a hedging portfolio in the current period together with the liquidation of this portfolio when the

4

barrier is attained and is essentially a constrained dynamic strategy. "Pure static strategies" would require the hedge to be liquididated only at the expiration of the option even if the barrier has been hit before. Models of static hedging allowing liquidation neglect the liquidation risk, that is, the risk that the prices at which one can liquidate the hedging portfolio differ from those anticipated at the creation of the hedge. Pure static strategies have no liquidation risk but cannot replicate the essentially dynamic properties of barrier options. In contrast, if the trader can unwind the hedge at model-based prices, strategies allowing liquidation more closely replicate the barrier option. In the following, replication is deemed "perfect" if the hedging error is zero conditional on assets being traded at model-implied prices.

Liquidation risk is due to the failure of models to forecast market prices with certainty. Hence, static hedging strategies with liquidation contain more market risk than pure hedging strategies. However, the later are not exempt from such risk. For example, they typically rely on assumptions on the probability distribution of the underlying asset. Pure static strategies are hardly used in practice. For example, few traders would plan not to unwind the hedge when the underlying asset price crosses the barrier and the knock-out option goes out of existence. However, pure static strategies are easier to visualize than strategies allowing liquidation and will be briefly used in the paper as didactic tools.

The static hedging strategies presented in DEK and CEG achieve perfect replication of the barrier option payoff by relying on the availability of traded options with arbitrary strikes or maturities. DEK acknowledge that the hedge is imperfect when only a limited array of maturities is available.[3] CEG can replicate the barrier option with few regular options but make strong assumptions on the probability distribution of the underlying asset. They point to the relaxation of such assumptions and the derivation of an approximate hedge as the main potential extension of their research.

The current paper paper is an extension of both these papers and is meant to be "user friendly." To hedge a given barrier option, the investor chooses the instruments he wishes to include in the hedging portfolio; the probability distribution of the risk factor driving the prices of the hedging instruments and of the barrier option (in the paper, this factor is the underlying asset price); a pricing model to evaluate all assets; and, possibly, some linear constraints on the hedging residual. In return, the procedure outlined in the paper gives him the optimal hedging portfolio. Our approach yields the same results as DEK or CEG when their assumptions are imposed, but extends their results to cases where their assumptions are violated.

Such an extension is necessary because of the real practical limitations of DEK's and CEG's methods. For example, they cannot be used to hedge an up-and-out call with a barrier above the strikes of traded regular options

---

[3]In that case, the hedging portfolio based on their method exactly replicates the value of the barrier option only at some future time and market levels, which are determined by the the maturities and the strikes of regular options available in the market. However, market participants may be more interested in an average measure of fit rather than into perfectly replicating the barrier option for some times and price levels.

whereas this type of barrier option is actually offered to customers.[4]

We now examine in more details DEK's and CEG's procedures. To simplify, we assume that the underlying asset is a stock and that its price has zero drift.

1.2   Static Hedging with Instruments of Different Maturities: The Derman, Ergener, and Kani Model

Taking as given the law of motion of the underlying asset price, DEK aim at replicating the dynamics of the barrier option by taking positions in regular options with a rich enough array of maturities. They achieve perfect replication of the barrier option, for a chosen range of times and stock prices, provided the stock price is allowed to move only at discrete intervals and that one can trade regular options maturing at those times and with appropriate strike prices. A drawback of this method is that, when choosing a finer price grid in the tree to price the barrier option, the hedger must take positions in an increasing number of regular options with intermediate maturities to accomodate the corresponding increase in the number of subperiods.

The hedging strategy chosen by DEK is to replicate the value of the barrier option on its boundary, that is, at expiry and at the barrier. Since the value of any asset is determined by its payoff on the boundary, replicating the value of the barrier option on its boundary also guarantees the replication of the barrier option at every point within the boundary. Before giving a detailed example of DEK's method applied to an up-and-out call, let's summarize their approach in four steps. First, hedge the up-and-out call at expiry with two regular options: one with the same strike as the barrier option to replicate its payoff below the barrier and another to cancel out the payoff of the regular call at the barrier. Second, compute the value of the hedging portfolio the preceding period. Third, set to zero the value of the hedging portfolio at the barrier that period by taking a position in a regular option with intermediate maturity. Fourth, iterate the previous two steps until the current period.

We now present DEK's method in some detail because we will use the same framework to present our alternative method. The DEK method relies on trading options at model-implied prices before expiry; it does not provide a pure static hedging strategy and is subject to liquidation risk. Table 1 illustrates DEK's procedure. Like them, we use a binomial tree to model the price process of the underlying asset: each period, the asset price increases or decreases by 10 units with probability 1/2 (Tree 1). The riskless interest rate is assumed to be zero. The goal is to hedge a 6-year up-and-out call with strike 60 and barrier 120 by trading three regular options: a 6-year call with strike 60, a 6-year call with strike 100, and a 3-year call with strike 120. Using the fact that the value of an asset equals its expected value the following period and proceeding backwards from expiry, we recursively compute the prices of the regular options at every node. The result is displayed in Trees 2

---

[4]Clarke (1998) points out that, when up-and-out calls and down-and-in puts are used by directional traders to gain exposure to movements in the underlying asset price, the barrier is set far away from the current spot price.

6

and 3 for the long-term options, those for the short-term option are omitted. Trees 4 and 5 present the hedging portfolio as it is progressively built. Tree 6 shows the price of the barrier option at each node. In each tree, the shaded region represents the boundary of the barrier option.

Looking at Trees 2 and 6 reveals that the regular and the barrier options have identical payoffs at expiry for stock prices below the barrier but not for those at (or above) it: the regular call pays out 60 when the stock price is 120 while the barrier option is worth zero. Using a standard 6-year call struck at 100, we can offset the regular option's payoff at expiry on the barrier without affecting its payoff below the barrier. Tree 4 shows the value of a the portfolio constituted by adequate long and short positions in the two regular options. The portfolio has the proper payoff on the boundary except in year 2. We correct this using a regular call option struck at 120 and maturing in year 3. Tree 5 shows that the new portfolio matches the value of the barrier option on the barrier but not above it. Liquidating the hedging portfolio at market prices when the underlying asset price hits the boundary makes the value of the hedging strategy coincide with that of the barrier option, transforming Tree 5 in Tree 6.

The binomial trees above are only meant as an illustration of DEK's method. Their technique is flexible enough to be applied on more complex trees, for example, trees calibrated to match the market prices of liquid options, using for example the method in Derman and Kani (1994).

1.3 Static Hedging with Instruments of Identical Maturities: The Carr, Ellis, and Gupta Method.

DEK match the barrier option's dynamics with regular options of intermediate maturities. However, such options may not be available in the market. CEG offer an alternative: by making strong enough assumptions on the distribution of the stock price, they are able to replicate the barrier option with regular options that all mature at the same time as the barrier option. Assuming that the stock price has zero drift and that the implied volatility smile is symmetric, CEG derive a symmetry relation between regular puts and calls:

$$C(K)\,K^{-1/2} = P(H)\,H^{-1/2} \ \ \text{if} \ \ (K\,H)^{1/2} = S, \tag{1}$$

where $C(K)$ is the value of a call struck at K, $P(H)$ the value of a put struck at H, and S is the price of the asset.[5] This relation, which is an extension of Bates (1988) is valid at any time before and including expiry.

Using the symmetry relation, CEG are able to perfectly hedge any European barrier option by building a portfolio of plain vanilla options whose value is zero when the price of the underlying asset crosses the barrier and

---

[5] The implied volatility is the volatility that equates the Black-Scholes formula with the actual price of the option. Implied volatilities form a "smile" when they increase at strikes further away from the forward price. The horizontal axis in implied volatility graphs is often the logarithm of the ratio of the strike prices relative to the forward price of the underlying asset and the symmetry of the smile in CEG has to be understood is this context.

has the same payoff as the barrier option when the barrier has not been hit before expiry. For example, one can hedge a down-and-out call with strike K and barrier H (with $H < K$) by buying a standard call struck at K and selling $K\,H^{-1}$ puts struck at $H^2\,K^{-1}$. The puts offset the value of the standard call when the stock price equals H because of the put-call symmetry. When the underlying asset price remains above H until maturity, the puts expire worthless (since $H^2\,K^{-1} < H$) while the standard call and the barrier call have identical payoffs. Hedging other barrier options is more complex but uses the same logic.

There are some limitations on the applicability of CEG's method. First, options with strikes at levels prescribed by the theory at not always available in the market. For example, in the preceding example, even if the barrier option's strike (K) and barrier (H) correspond to strikes of traded plain vanilla options, it is less certain that $H^2\,K^{-1}$ is a traded strike. More generally, the barrier must belong to the range of marketable options' strikes for the put-call symmetry to be applied. This is not always the case. Second, and more importantly, implied volatilities in actual markets often exhibit non-symmetric smiles (Bates (1991 and 1997)).

2   Alternative method: Mean-square Hedging

DEK and CEG rely on strong assumptions on the availibility of standard options with given strikes or maturities, or on the distribution of the stock price. Furthermore, they offer no clear way to adapt their methods to cases where their assumptions are violated.

An alternative approach is to approximate as best as possible the random payoff of the barrier option with a given set of instruments. Since the option may be imperfectly replicated, one has to choose a metric to measure how close the hedge approximates the option. One possibility is to use the mean of the square of the hedging error. Minimizing this quantity yields "mean-square hedging."

Remark 1 *Mean-square hedging perfectly replicates the barrier option when the assumptions in DEK or CEG are satisfied.*

The advantage of this method compared to DEK and CEG is that it provides guidance to the hedger when their assumptions are not satisfied, for example, when one of the standard options needed to perfectly replicate the claim is not available in the market.

To simplify, we only consider the hedging error at maturity. More generally, regulations may dictate whether one decides to minimize the discrepancy between the value of the hedge and that of the asset only at maturity or over the life of the hedge. New American and international regulations (FAS 133 of the U.S. Financial Accounting Standards Board and IAS 39 of the International Accounting Standards Committee) restrict the use of "hedge accounting," that is, the possibility of deferring losses and gains on the hedging strategy until the risk is realized (e.g. the maturity of the barrier option). As a consequence, companies have to include more often in their

reported income the difference between the value of the hedge and that of the asset they want to hedge. In such cases, one should take into account the hedging residual at each reporting time.

2.1 Principle: Minimizing the Mean of the Hedging Error Squared

The following introduces in some generality the principles behind mean-square hedging. Naturally, mean-square hedging can be used to hedge any type of claim, not just barrier options.

Assume that the underlying asset price can take n values and that there are no short-sale constraints, taxes, or other market frictions. We model assets as random variables taking values in R^n. Let E be the set of all such random variables; it is an n-dimensional linear space and is called the "payoff space." Assume that one can trade k assets $z_1, \ldots, z_k$ with $k \leq n$. Write z the k-dimensional vector composed of those assets and $L(z)$ the set of all the linear combinations of elements of z; it is called the space of marketable assets. Assume that the z_i's are linearly independent, in the sense that no component of z can be replicated by a linear combination of the others. The vector z is a basis of $L(z)$ and this space is k-dimensional. When $n = k$, markets are complete because any claim can be duplicated by a linear combination of the z_i's (or other marketable assets). When $n > k$, markets are incomplete. We focus on this case and assume $n > k$.

Call x the claim we want to hedge, β the asset allocation in the hedging portfolio; the portfolio payoff is $z'\beta$, where the "prime" superscript denotes transposition. We choose β to minimize the mean-square of the residual against a given probability measure. The objective is hence to

$$\min_{\beta} E[(x - z'\beta)^2] \tag{2}$$

and the optimal β is

$$\beta = E[zz']^{-1} E[zx]. \tag{3}$$

Equation (3) defines the estimator of the coefficient on z using ordinary least squares. This property makes mean-square hedging very intuitive. Moreover, when the riskless bond is included in the marketable assets, the hedging residual has mean zero and the "R Squared" can be used to gauge the quality of the optimal hedge since it measures the goodness of fit in a regression. Finally, the optimal hedge of some claim $x \in E$ can be geometrically interpreted as the orthogonal projection of x on $L(z)$, written $proj(x|L(z))$.

Remark 2 *Because of the linearity of mean-square hedging,*

1. *Hedging a portfolio of assets is equivalent to hedging each asset in the portfolio individually,*

2. *One can mean-square hedge any asset by combining the hedging portfolios of n artificial assets called contingent claims.*

The first property is of interest to banks because they typically hedge their option books as whole portfolios. The second property uses the linearity of mean-square hedging to save on computer time by computing $E[zz']^{-1}$ only once. The "state-contingent" claims are artificial assets which are not traded but form a practical basis in which to express all assets of interest. Each of the n contingent claim pays \$1 in a given state and \$0 in all the others. Call x_i the state-contingent claim with non-zero payoff in state i and let $X = (x_i)_{i=1}^n$. This random vector can be hedged in a mean-square sense using the $k \times n$ matrix $\gamma = E[zz']^{-1} E[zX']$. Any asset $x \in E$ can be written as $x = X'\alpha$. The mean-square hedging portfolio of x is defined by $z'\beta$ with $\beta = \gamma\alpha$.

Mean-square hedging can also easily accommodate position limits. First, run the hedging program and list the assets for which the prescribed holdings are higher than the limits. Then, rerun the hedging program constraining the holdings of these assets to match the position limits. For example, let x be the asset to be hedged, assume that the position limits are binding for the first q assets, and write $z' = (z_1', z_2')$ and $\beta' = (\beta_1', \beta_2')$ where z_1 is the q-dimensional vector of those assets and β_1 is their holdings in the hedging portfolio. Fix β_1 equal to the position limits and write $\tilde{x} = x - z_1'\beta_1$. Run the hedging program on $\tilde{x}$ to get the optimal β_2.

Preferring static hedging over dynamic hedging because of transaction costs and then letting mean-square hedging determine the optimal static strategy might be seen as contradictory because mean-square hedging ignores transaction costs. On the other hand, optimal strategies that take into account market frictions like the bid-ask spread appear much harder to compute. Applying mean-square hedging to a static hedging strategy can be viewed as a compromise: one acknowledges the existence of transaction costs by choosing to hedge statically and obtains some computable answer by using mean-square techniques.

2.2 Pure Static Strategies and Strategies Allowing Liquidation

Mean-square hedging can accomodate the case where the hedging portfolio is held until maturity and that where it is liquidated at model-implied prices as soon as the underlying asset price crosses the barrier, with the proceeds of the liquidation being rolled over until maturity. To compute the hedging portfolio in that case, we substitute the liquidation values of the marketable assets to their original last-period payoffs for all paths in the tree for which the price of the underlying asset crosses the barrier. This substitution does not affect the prices of marketable assets before the barrier is hit.[6]

Alternatively, one can choose pure static strategies. Using only pure static strategies and instruments with the same maturity as the barrier option greatly simplifies the hedging problem because one need only hedge the conditional expectation of the option payoff conditioned on the underlying

---

[6]This follows from the fact that the price of any asset is the mean of its discounted payoff and from the law of iterated expectations. The riskless interest rate is assumed to be zero.

asset price at expiry (called the "conditional payoff").[7] Conditioning on the stock price at maturity removes the path dependency of the option payoff from the analysis and reduces the dimension of the hedging problem: instead of considering the number of possible paths between the initial period and expiry, we only have to consider the number of nodes at expiry. Accordingly, the number of states is reduced from 2^n to $n+1$ for binomial trees and from 3^n to $2\,n+1$ for trinomial trees.

3 Application

We apply mean-square hedging using binomial and trinomial trees. Strike and barrier levels can be assumed to correspond to nodes in the tree because an option with a strike (or a barrier) between two nodes has the same payoff as the option with the strike (or the barrier) adequately placed at one of the nodes.

We choose the framework used to present DEK's method but we constrain the strategies to be purely static or allow liquidation but exclude from the marketable assets one of the three regular options used to hedge the barrier option. We also formally introduce the riskless bond as a traded asset. It was not part of the hedging portfolio before, because the barrier option could be perfectly replicated using only the regular options. It is used to roll forward the proceeds of the liquidation until the maturity of the barrier option.

When hedging is imperfect, hedging errors typically are path dependent and cannot be represented on a tree. We use the conditional expectation of the value of the hedging portfolio conditioned on the price of the stock (called the "conditional value") to visualize the fit of the mean-square hedging portfolio. We plot this conditional value at maturity against the possible values of the stock; we also display the conditional values at every node.

The upper panel in Figure 2 shows the conditional values of the three regular options potentially available to hedge the barrier option (two 6-year calls with strikes 60 and 100 and one 3-year call with strike 120) conditional on the last-period stock price ("the conditional payoff"). Conditional value and intrinsic value coincide for the options expiring at the last period but not for the short-maturity option. Every last-period stock price is the end-node of several paths (except the highest and the lowest price to each of which only one path converges). The higher the last-period stock price, the more likely it is that the short-term option expires in the money along the paths leading to it. Consequently, the growth in the conditional value of the short-term option accelerates as the last-period stock price increases.

If regular options are available with strikes equal to all possible stock price levels between (and including) the strike and the barrier of the barrier option, it is possible to perfectly hedge the conditional value of the barrier option. Doing so would not zero out the hedging error on the barrier option

---

[7]This is because $proj(x|L(z)) = proj(E[x|S]|L(z))$ when the random vector z is a function of S, the stock price at maturity.

itself but would lead to the best purely static strategy.[8] Most importantly, the regular option struck at the barrier allows one to set to zero the value of the hedging portfolio for stock prices above the barrier. In contrast, when only two regular options are available, the conditional values of the hedging portfolio and of the barrier option necessarily differ, even if the riskless bond is added to the marketable assets. The middle panel of Figure 2 shows the conditional value of the barrier option (the dashed line) and that of the optimal hedging portfolio using only the long-term options (the solid line). Including the short-term option does not significantly change the quality of the hedge. When no regular options with strikes at or above the barrier are available, the conditional hedging error grows linearly with the asset value when the latter increases above the barrier because, for asset prices above the barrier, the value of the barrier option is zero whereas the payoff of the hedging portfolio is linear in the stock price. However, the likelihood of large hedging errors is small and decreases with the magnitude of the losses, reflecting the thining tails of the probability distribution of the stock price (the lower panel).

Table 2 displays the conditional value of the barrier option (Tree 1) and of the optimal hedging portfolio based on all three options and the riskless bond when only purely static strategies are allowed (Tree 2). Comparing Trees 1 and 2 reveals significant expected discrepancies that arise early on and grow with time. The fit is particularly bad above the barrier.

Liquidation provides added flexibility in the hedging strategy and results in a significant improvement in the quality of the hedge. As shown in the upper panel of Figure 3, liquidating the regular options at model-based prices when the stock price hits the barrier creates more kinks in their conditional values, reflecting their increased usefulness as hedging tools. Although hedging the conditional value of the barrier option would not yield the correct hedge for the barrier option itself, we compare the conditional value of the hedging portfolio to that of the barrier option as a way of visualizing the quality of the hedge.

As shown in Tree 3 in Table 2, liquidation allows the conditional value of the hedging portfolio to be very close to that of the barrier option at every node in the tree even when the 3-year call option is not included in the portfolio. The middle panel of Figure 3 points to a good fit in the tails of the conditional value of the barrier option. The results above suggest that the 6-year calls, which allow the replication of the barrier option on or below the barrier at maturity, are the major hedging instruments. This is confirmed by including the short-term option but excluding the 6-year call with strike 100. As shown in Tree 4, including the short-term option in the hedging portfolio results in a better conditional fit up until the second year, but the exclusion of the longer-term option creates large conditional

---

[8]The hedging procedure can be implemented graphically: first, use the regular call with the same strike as the barrier option to match the conditional value of the barrier option between the strike and the stock price immediately above it. Second, take a position in a regular option struck at that price to equate the conditional values of the hedging portfolio and of the barrier option between that price and the next. Iterate until the barrier is reached.

deviations for periods near maturity. The lower panel of Figure 3 reveals a bad conditional fit in the lower tail.

Table 3 shows the asset allocation in the hedging portfolio in the different scenarios used above. When liquidation is allowed, excluding the short-term option has a limited impact on the amounts of the long-term options held in the hedging portfolio (column 2) compared to when the hedger is allowed to trade all three options (column 1). The "R Squared" suggests a near-perfect hedge (neglecting liquidation risk). In contrast, excluding the long-term option with the highest strike or allowing all three options but excluding liquidation greatly impacts the composition of the hedging portfolio and dramatically lowers the goodness of fit (colums 4 and 5). The riskless bond plays a noticeably larger role when the fit is poor. In the limit, if the values of the hedging instruments were uncorrelated with that of the barrier option, the best hedging strategy would be to invest the value of the barrier option in the riskless bond and nothing in the other hedging instruments.

3.1 The Relation Between Hedging and Pricing

The value of the mean-square hedging portfolios obtained above coincides with the price of the barrier option derived from the tree. Choosing the risk-neutral probability distribution implied by the tree to compute mean-square hedges guarantees this desirable result. Other choices may create discrepancies between pricing and hedging the barrier option.

Call π the pricing functional on $L(z)$, that is, the linear function which to any marketable asset associates its price. This function is uniquely defined by the exogenously given prices of the marketable assets and can be extended to E by defining the price of any asset x in E as the value of the best matching portfolio obtained through mean-square hedging. In mathematical terms, calling $\hat{\pi}$ the extension of π to E formed in this maner, for all $x \in E$,

$$
\begin{aligned}
\hat{\pi}(x) &= \pi(\,proj(x|L(z))\,) \\
&= \pi(\beta'z) \\
&= \beta'\pi(z),
\end{aligned}
\tag{4}
$$

with $\beta = E[zz']^{-1}E[zx]$. Naturally, $\hat{\pi}$ depends on the probability distribution used to compute mean-square hedges.

Conversely, any exogenously given pricing functional $\tilde{\pi}$ on E defines a probability distribution on E. Write $q_i = \tilde{\pi}(x_i)$ where x_i is the state-contingent claim with non-zero payoff in state i. The q_i's are positive because $\tilde{\pi}$ is assumed to admit no arbitrage opportunities and they sum up to 1 because the riskless bond, which pays \$1 in every state and is worth \$1 (the riskless interest rate is assumed to be zero) can be replicated by holding all the contingent claims. Hence, the q_i's define a probability distribution, Q, on E. It is the risk-neutral implied by $\tilde{\pi}$ because, for all $x \in E$, $\tilde{\pi}(x) = E^Q[x]$ where E^Q denotes the expectation operator against Q.

Theorem 1
When Q is used to compute the hedges, the pricing functional derived from mean-square hedging coincides with $\tilde{\pi}$.

13

Proof: For all $x \in E$.

$$\begin{aligned}
\hat{\pi}(x) &= \pi(\, proj(x|L(z))\,) \\
&= \tilde{\pi}(\, proj(x|L(z))\,) \\
&= E^Q[\, proj(x|L(z))\,] \\
&= E^Q[\, x\,] \\
&= \tilde{\pi}(x)
\end{aligned} \qquad (5)$$

The first line comes from the definition of $\hat{\pi}$; the second from the fact that all pricing functionals on E coincide with π on $L(z)$; the third line stems from the definition of Q; the fourth line uses the fact that $E^Q[x - proj(x|L(z)] = 0$ provided the riskless bond is a marketable asset; this is similar to including a constant term as an explanatory variable in a regression to insure that the residual has mean zero.

4   Hedging Consistently with Prices of Non-marketable Assets

The risk-neutral measure is merely a convenient pricing tool and combines properties of both the real-world probability distribution and of the compensation that investors demand for taking on additional risk. An investor might be more interested in minimizing its hedging uncertainty using the real-world measure than using the risk-neutral measure. However, if the investor uses mean-square hedging with the real-world measure, the value of the hedging portfolio could differ from the value of the barrier option implied by pricing models. Moreover, mean-square hedging might affect negative prices to non-marketable claims with positive payoffs in all states because of its linearity. The key is that investors may be interested in adopting mean-square hedging because of its simplicity but may not be willing to replace current pricing tools.

The following introduces a method that preserves the computational simplicity of mean-square hedging and makes the value of the hedging portfolio consistent with the (exogenously given) price of the security it is supposed to hedge. This is done by requiring the price of the hedging error to be zero. The method presented below is flexible enough to incorporate any other linear constraint on the hedging residual. We present the method in some detail below before introducing some applications.

4.1   Principle: a Geometric Approach

This section introduces the new method in an intuitive, geometric fashion. Theorem 2 in the next section generalizes the approach and presents the results in a more rigorous way. Proofs are relegated to the appendix. As before, let E be the n-dimensional payoff space, z the k-dimensional vector of marketable assets, $L(z)$ the linear space spanned by z, π a given pricing functional on E. The objective is, for any given x, to build a portfolio of marketable assets that best replicates asset x and has price $\pi(x)$. We are

14

hence interested into writing any $x \in E$ as

$$x = z'\beta + \varepsilon \tag{6}$$

with $\pi(\varepsilon) = 0$. For convenience, write $\ker \pi = \{u \in E, \pi(u) = 0\}$ and $F^\perp$ the set of all elements of E orthogonal to elements of F, itself a linear subspace of E.

Since the hedging residual generated by mean-square hedging is orthogonal to $L(z)$, a natural extension of this method that incorporates the new constraint $\pi(\varepsilon) = 0$ is to impose that, within $\ker \pi$, ε be orthogonal to $L(z)$. Figure 4 shows mean-square hedging and its extension in a 3-dimensional space. Assets are represented as vectors and random variables that are orthogonal with respect to the probability measure are shown as orthogonal vectors. The space of marketable assets, $L(z)$, is assumed to be a plane. The upper panel represents mean-square hedging as the orthogonal projection on $L(z)$. The middle panel introduces the set of payoffs with zero price, $\ker \pi$, to which the hedging residual must belong. Since E is a 3-dimensional space and $\pi(y) = 0$ imposes one constraint on y, $\ker \pi$ is a plane. Its intersection with $L(z)$ is a line because the no-arbitrage condition precludes all marketable assets from having zero price, so that $\ker \pi$ and $L(z)$ cannot coincide. The hedging error must belong to $\ker \pi$ and, within that set, be orthogonal to $L(z)$, that is, be orthogonal to $L(z) \cap \ker \pi$. We conclude that the hedging error belongs to $\ker \pi \cap [L(z) \cap \ker \pi]^\perp$. The lower panel shows the result of the new hedging method: an asset x is hedged using the decomposition $x = z'\beta + \varepsilon$ with $\varepsilon \in \ker \pi \cap [L(z) \cap \ker \pi]^\perp$. This decomposition is unique, linear in x, and defines a projection on $L(z)$, call it $\tilde{P}[\,|L(z)]$, characterized by its image set $(\mathrm{Im}(\tilde{P}[\,|L(z)]))$ and its null set $(\ker(\tilde{P}[\,|L(z)]))$.

To summarize, one hedging strategy consistent with the pricing functional π is characterized by the linear projection $\tilde{P}[\,|L(z)]$ defined by:

$$\begin{cases} \mathrm{Im}(\tilde{P}[\,|L(z)]) & = & L(z) \\ \ker(\tilde{P}[\,|L(z)]) & = & \ker \pi \cap [\ker \pi \cap L(z)]^\perp \end{cases} \tag{7}$$

This result is generalized to hedging consistently with multiple linear constraints and the associated projection is characterized further in theorem 2 in the next section. If π is the pricing functional obtained from mean-square hedging, $\tilde{P}[\,|L(z)]$ is the orthogonal projection on $L(z)$.[9] If π is not the mean-square hedging pricing functional, $\tilde{P}[\,|L(z)]$ is a non-orthogonal linear projection on $L(z)$ with direction $\ker \pi \cap [\ker \pi \cap L(z)]^\perp$.

4.2 Hedging Consistently with p Linear Constraints

The mean of the hedging residual obtained from the hedging strategy above may fail to be zero. We can easily incorporate this new constraint or impose p linear constraints on the hedging residual. These constraints can be expressed as $\Psi(\varepsilon) = 0$, where Ψ is a p-dimensional linear function with $p < k$.

---

[9]If π is the mean-square pricing functional, $\ker \pi = L(z)^\perp$ so that $\ker(\tilde{P}[\,|L(z)]) = L(z)^\perp \cap [L(z)^\perp \cap L(z)]^\perp = L(z)^\perp$.

The linear functional Ψ is defined on E, we can also apply it on vectors of elements of E by using the following convention: if x is a $r \times 1$ vector of random variables in E, $\Psi(x)$ is a $r \times p$ matrix. Consequently, if x_1 and x_2 are $r_1 \times 1$ and $r_2 \times 1$ vectors and A is a $r_1 \times r_2$ matrix such that $r_1 = Ar_2$, then, $\Psi(r_1) = A\Psi(r_2)$.

Assumption 1

1. Ψ is such that $E = L(z) + \ker(\Psi)$, that is, for all $x \in E$, there exist $v \in L(z)$ and $\varepsilon \in E$ such that $x = v + \varepsilon$ and $\Psi(\varepsilon) = 0$,

2. $\Psi(z)$ is full rank.

The first assumption guarantees that it is possible to hedge claims in E using marketable assets such that the hedging residual satisfies the constraints encapsulated in Ψ. We assume that such constraints can be satisfied; the objective is to find the best way to do so. The second assumption imposes that the rank of $\Psi(z)$ is p and means that none of the p restrictions implied by Ψ are redundant. We assume without loss of generality that the last p components of z are linearly independent.

Using the same geometric principle as above, we define the hedging strategy consistent with the constraints.

Theorem 2

1. One hedging strategy consistent with the constraint $\Psi(\varepsilon) = 0$ is characterized by the linear projection $\hat{P}[\,|L(z)]$ defined by:

$$
\begin{cases}
\mathrm{Im}(\hat{P}[\,|L(z)]) &=& L(z) \\
\ker(\hat{P}[\,|L(z)]) &=& \ker(\Psi) \cap [\ker(\Psi) \cap L(z)]^{\perp}
\end{cases}
\tag{8}
$$

2. There exists some $u \in E$ such that

$$
\hat{P}[x|L(z)] = z'\beta \quad \text{with} \quad \beta = E[uz']^{-1}E[ux]
\tag{9}
$$

3. The projection $\hat{P}[\,|L(z)]$ coincides with the constrained least squares regression of x on z.

Proof: See appendix.

When the hedging residual has mean zero, the "R squared" can also be used to measure the fit of the hedging portfolio with the barrier option. The advantage of our method compared to a direct use of constrained least squares lies in its computational efficiency, a crucial property to practitioners. In contrast, off-the-shelf formulae for constrained least squares regression coefficients appear much more complex. Another advantage is its intuitive geometrical interpretation.

Remark 3 *Because of the linearity of mean-square hedging under constraints,*

1. *Hedging a portfolio of assets under linear constraints on the hedging residual is equivalent to hedging each asset in the portfolio individually,*

16

2. *One can hedge any asset under such constraints by combining the hedging portfolios of the contingent claims.*

Relevant linear constraints on ε include

$$\begin{cases} \pi(\varepsilon) & = & 0 \\ E[\varepsilon] & = & 0 \\ E[\varepsilon|S \in G] & = & 0 \\ E[\varepsilon|S \in G] & = & a\, E[x|S \in G] \end{cases} \tag{10}$$

where S is the price of the underlying asset at maturity, G is an interval on the real line, and a is a scalar. The last two constraints above become risk-management tools when they are seen as saturated versions of economically more meaningful constraints like:

$$\begin{array}{rcl} E[\varepsilon|S \in G] & \leq & 0 \\ E[\varepsilon|S \in G]/E[x|S \in G] & \leq & a \end{array} \tag{11}$$

Constraints like

$$\begin{array}{rcl} P(\varepsilon > \alpha) & = & a_1 \\ E[\varepsilon\,|\varepsilon > \alpha] & = & a_2 \end{array} \tag{12}$$

cannot be accommodated in the present framework because they are not linear. (The parameters a_1 and a_2 above are the "value-at-risk" and the "mean-excess loss" on the hedging strategy at a given "significance level" α, assuming the trader loses money if the hedging residual is positive). However, if the realizations of the hedging error are large when $S \in G$, then constraining $E[\varepsilon|S \in G]$ indirectly imposes some tail condition on the hedging error. Intuitively, hedging a barrier option or other non-linear claims with standard options should generate larger errors when the price of the underlying asset takes values outside the range of marketable strikes or far from the barrier because the hedger has fewer "degrees of freedom" in those areas.

Linear constraints insure the linearity of the hedging operator $\hat{P}[\;|L(z)]$ and the computational efficiency of the hegding method. In contrast, constraints like $\Psi(\varepsilon) = a$ are not linear when $a \neq 0$ even if Ψ is.

5 Application

We return to the example illustrating DEK's method. The objective is to hedge the barrier option excluding the short-term call from the hedging portfolio, which insures that the hedge is imperfect. We use the risk-neutral probability distribution implied by the tree and impose some restrictions on the hedging error. First, we constrain the hedging residual to be zero when the stock price at maturity reaches its maximum; second, we impose the additional constraint that the mean of the residual, and therefore its price, be zero.

Figure 5 shows the conditional payoffs at maturity of the barrier option (the dashed line), of the hedging portfolio when no constraint is imposed

(the light grey line), when the hedging residual is constrained to be zero at the maximum of the stock price (the upper panel) and under the additional constraint that the mean of the hedging residual be zero (the lower panel). When no constraint is imposed, the hedging error arising from mean-square hedging has mean zero because the hedging portfolio includes the riskless bond. The first constraint shifts up the conditional value of the optimal hedging portfolio for high levels of the stock price but has little effect for lower stock prices (where the conditional fit is very good). This leads to a non-zero mean for the hedging error (and a lower price for the hedging portfolio). Imposing a zero-mean condition on the hedging residual shifts down the conditional value of the hedging portfolio for lower stock prices, which offsets the upward shift when the stock price is high. Comparing columns 2 and 3 in Table 3 reveals that imposing the additional constraint that the hedging error be zero when the stock price reaches its maximum at maturity and has mean zero also has a moderate impact on the portfolio allocation and the goodness of fit. This is partly because, in this example, the hedging residual is already very small when no constraint is imposed. In trees used in practice, the underlying asset price may take values farther above the barrier with positive probabilities and imposing a tail constraint on the hedging residual may more significantly impact asset allocation.

To provide another illustration of our extension to mean-square hedging, we assume now that the returns on the stock follow a geometric Brownian motion with zero drift. We choose a trinomial tree to evaluate the assets because such tree offers more flexibility in the position of the nodes with respect to the barrier than binomial trees (Cheuk and Vorst (1996)). Since the distribution of the stock satisfies CEG's assumptions, any barrier option can be replicated if regular options with strikes called for by CEG are traded.

We require the hedging portfolio to be consistent with the pricing model implied by the tree, but choose a different probability distribution to compute hedges. The new distribution affects the same probability to every path in the tree. This increases the likelihood of tail events relative to the risk-neutral distribution as shown in the top panel of Figure 6. We then compare the hedging portfolios based on the risk-neutral distribution to those obtained using the alternative probability distribution under the constraint that hedging portfolios be consistent with the pricing functional implied by the tree and generate mean-zero hedging errors. In each case, both purely static strategies and strategies allowing liquidation are considered. The asset to hedge is an at-the-money up-and-out call option with barrier above the strikes of all marketable options so that hedging is necessarily imperfect.

The lower panels of Figure 6 shows the conditional payoffs of the barrier option and of the hedging portfolios, conditioned on the underlying stock price at expiry assuming admissible hedging strategies are purely static (the middle panel) and allowing the liquidation of the hedging portfolios at model-based prices when the barrier is reached (the lower panel). When liquidation is precluded, optimal hedging portfolio strongly depend on the probability distribution used. Two determining factors in the shape of the hedging portfolio are the probability mass between the barrier (H) and the highest possible asset value smaller than the barrier (H^-) and the probability mass

to the right of H. Between H^- and H, the conditional payoff of the barrier option drops from its maximum to zero. In our example, the highest strike available to the hedger is H^- and, consequently, the conditional payoff of the hedging portfolio is linear in the asset price for prices above H^-. The hedger has to decide whether to fit better the drop in the value of the barrier option between H^- and H or the flat portion of its value above H.

The probability distribution based on the equiprobability of paths affects more mass to asset values above H^- than the lognormal distribution. However, the increase in the probabilities is relatively greater for stock values between H and H^- than for those strictly above H.[10] Consequently, the hedger is more weary of deviations between the conditional value of the barrier option and that of the hedging portfolio for values of the underlying asset between H^- and H than for those above H. Therefore, hedging errors are larger in the tails when the hedging strategy assumes that each path is equiprobable than using the lognormal distribution, even though assuming equiprobability of paths puts less mass in the center of the distribution and more mass in its periphery.

When hedging strategies allow liquidation, the choice of the probability distribution makes little difference in the resulting hedging portfolios. Intuitively, allowing liquidation significantly improves the fit between the payoff of the barrier option and that of the hedging portfolio so that hedging errors are small and the probability distribution used has less impact on the hedging strategy. In the limit, if one was able to perfectly mimic the payoff of the barrier option, the probability distribution would be irrelevant.

6 Conclusion

Barrier options have become widely traded securities in the past recent years. Market professionals can easily price these instruments using existing evaluation tools but find them hard to hedge. This illustrates the dichotomy that market incompleteness introduces between pricing and hedging. Moreover, to save on transaction costs, an unavoidable feature of real markets, practitioners prefer to adjust their hedging portfolios only infrequently. In response, researchers have devised "static strategies" where the hedging portfolio is liquidated if the barrier is reached (in the case of a knock-out option) but is otherwise not modified before maturity. However, current research in static hedging still aims at perfectly replicating the barrier option. To achieve this goal, one has to make strong assumptions on the availability of regular options with certain strikes or maturities, or on the distribution of the underlying asset. This greatly impairs the implementability of the proposed hedging procedures in real-world situations (where these assumptions are not valid).

This paper proposes an alternative approach. It aims at replicating "as best as possible" the barrier option while taking into account more features

---

[10]There are only two possible stock prices strictly above H. They are located at the end of the possible price range and at the last kink of the two probability distributions. The difference between the two probability distributions is fairly small at those two points.

of the real world than current research in the field. The tool of choice is
"mean-square hedging" and is extended to incorporate linear constraints on
the hedging residual. One advantage of the method is that the value of the
hedging portfolio can be made to coincide with the price attributed to the
barrier option by any given pricing model. Another advantage is that it can
incorporate constraints on the tail of the hedging residual, which contributes
to managing risk.

Appendix

Proof of Theorem 2

1. Let's show that

$$\begin{cases} L(z) + \ker(\Psi) \cap [\ker(\Psi) \cap L(z)]^{\perp} & = & E, \\ L(z) \cap \ker(\Psi) \cap [\ker(\Psi) \cap L(z)]^{\perp} & = & \{0\}, \end{cases} \tag{13}$$

where 0 is the null vector in E. (In the following, 0 refers to null vectors of various dimensions.) If equation (13) above is true, we can define a projection, $\hat{P}[\,|L(z)]$, with

$$\begin{cases} \text{Im}(\hat{P}[\,|L(z)]) & = & L(z) \\ \ker(\hat{P}[\,|L(z)]) & = & \ker(\Psi) \cap [\ker(\Psi) \cap L(z)]^{\perp} \end{cases} \tag{14}$$

The second line in equation (13) is obviously true. Given assumption 1, to obtain the first line, it is sufficient to show that $\ker(\Psi)$ can be decomposed as:

$$\ker(\Psi) = \ker(\Psi) \cap L(z) + \ker(\Psi) \cap [\ker(\Psi) \cap L(z)]^{\perp} \tag{15}$$

To show this, let $\varepsilon \in \ker(\Psi)$ and project ε on $\ker(\Psi) \cap L(z)$.

$$\varepsilon = P[\varepsilon|\ker(\Psi) \cap L(z)] + \eta = v + \eta, \tag{16}$$

where $P[\,|\ker(\Psi) \cap L(z)]$ is the orthogonal projection. By definition, $\eta \perp \ker(\Psi) \cap L(z)$. Moreover, $\Psi(\eta) = 0$ because Ψ is linear and ε and v are elements of $\ker \Psi$, Hence

$$\eta \in \ker(\Psi) \cap [\ker(\Psi) \cap L(z)]^{\perp} \tag{17}$$

and equation (15) follows.

2. Let show that there exists some $u \in E$ such that

$$\hat{P}[x|L(z)] = z'\beta \ \text{ with } \ \beta = E[uz']^{-1}E[ux]. \tag{18}$$

Any $x \in E$ can be written as $x = z'\beta + \varepsilon$ where $\varepsilon \in \ker(\hat{P}[\,|L(z)])$. By definition, for any basis in $\ker(\hat{P}[\,|L(z)])^{\perp}$, $E[u\varepsilon] = 0$. Hence,

$$E[uz']\beta = E[ux] \tag{19}$$

To prove that the $k \times k$ matrix $E[uz']$ is invertible, let α be any k-dimensional real vector such that

$$\alpha'E[uz'] = 0, \tag{20}$$

and show that $\alpha = 0$. Writing $v = \alpha'u$, equation (20) is equivalent to $E[vz'] = 0$, that is, $v \in L(z)^{\perp}$. Hence, $v \in \text{Im}(\hat{P}[\,|L(z)])^{\perp}$ and $v \in \ker(\hat{P}[\,|L(z)])^{\perp}$. Since, for all two subspaces W_1 and W_2 of a finite-dimensional space, $W_1^{\perp} \cap W_2^{\perp} = (W_1 + W_2)^{\perp}$, $v \in (\text{Im}(\hat{P}[\,|L(z)] +$

$\ker(\hat{P}[\,|L(z)])^{\perp}$, or equivalently, $v \in E^{\perp}$, that is, $v = 0$. Now, since u is a linearly independent vector, $\alpha' u = 0$ implies $\alpha = 0$. Hence, the matrix $E[uz']$ is invertible and equation (19) is equivalent to

$$\beta = E[uz']^{-1} E[ux] \tag{21}$$

The coefficient β is unique even though u is not. If u_1 and u_2 are two bases in $\ker(\hat{P}[\,|L(z)])^{\perp}$, there exists a k-dimensional invertible matrix T such that $u_1 = T\,u_2$. This matrix drops out of the computation of β.

3. Let's show that $\hat{P}[\,|L(z)]$ defined in (14) coincides with the constrained least squares projection. First, we characterize the constrained least squares projection. Then, we prove it is the same as $\hat{P}[\,|L(z)]$.

   (a) Let's apply constrained least squares to find β solving

   $$x = z'\beta + \varepsilon \ \text{ s.t. } \ \Psi(\varepsilon) = 0. \tag{22}$$

   Since Ψ is a p-dimensional linear functional on E, there exists a p-dimensional random vector $m \in E$ so that, for any vector y of random variables in E.

   $$\Psi(y) = E[ym'] \tag{23}$$

   The constraint on the regression residual can be written:

   $$E[xm'] = \beta'E[zm'] \tag{24}$$

   where $E[zm'] = \Psi(z)$ is a $k \times p$ full-rank matrix (assumption 1). The Lagrangean for the constrained optimization problem is

   $$L = \beta'E[zz']\beta - 2\beta'E[zx] + x^2 + 2\{E[xm'] - \beta'E[zm']\}\lambda \tag{25}$$

   where 2λ is the $p \times 1$ vector of Lagrangean multipliers. The first-order condition with respect to β implies that

   $$\beta = E[zz']^{-1} \{E[zx] + E[zm']\lambda\} \tag{26}$$

   Substituting the expression above for β in equation (24), we obtain

   $$\lambda = \{E[mz']E[zz']^{-1}E[zm']\}^{-1}\{E[mx] - E[mz']E[zz']^{-1}E[zx]\} \tag{27}$$

   The matrix $E[mz']E[zz']^{-1}E[zm']$ is invertible because $E[zm']$ is full rank. The coefficients λ and β are linear in x. Consequently, constrained least squares defines a linear projection on $L(z)$, call it $\bar{P}[\,|L(z)]$.

(b) Let's show that $\bar{P}[\ |L(z)] = \hat{P}[\ |L(z)]$. Since $\text{Im}(\hat{P}[\ |L(z)]) = \text{Im}(\bar{P}[\ |L(z)])$, the two projections are equal if and only if

$$\ker(\hat{P}[\ |L(z)]) = \ker(\bar{P}[\ |L(z)]). \tag{28}$$

Moreover, since these two linear spaces have the same (finite) dimension, it suffices to show that one is included in the other. Let's show that $\ker(\bar{P}[\ |L(z)]) \subset \ker(\hat{P}[\ |L(z)])$. Let β and ε be the regression coefficient and the residual in the constrained least square regression. Let's show that $\varepsilon \perp \ker(\Psi) \cap L(z)$. A random variable $y \in E$ is an element of $\ker(\Psi) \cap L(z)$ if and only if $y = z'\theta$ for a k-dimensional real vector θ and $\Psi(y) = 0$, that is, if and only if

$$\begin{cases} y & = & \theta'z \\ \theta'E[zm'] & = & 0 \end{cases} \tag{29}$$

Showing $\varepsilon \perp y$ is equivalent to showing $E[yx] = E[yz'\beta]$. This is done below.

$$\begin{aligned} E[yz'\beta] & = & E[yz']\beta \\ & = & \theta'E[zz']\beta \\ & = & \theta'\{E[zx] + E[z'm]\lambda\} \\ & = & E[(\theta'z)x] + \theta'E[z'm]\lambda \\ & = & E[yx] \end{aligned} \tag{30}$$

Using equation (30) and $\Psi(\varepsilon) = 0$, we get that $\varepsilon \in \ker(\hat{P}[\ |L(z)])$. Consequently, $\ker(\bar{P}[\ |L(z)]) \subset \ker(\hat{P}[\ |L(z)])$ and, hence, $\hat{P}[\ |L(z)] = \bar{P}[\ |L(z)]$.

Computing u in $\beta = E[uz']^{-1}E[ux]$.

Let $y \in \ker(\Psi) \cap L(z)$ and θ such that $y = \theta'z$. Then,

$$\Psi(z)'\theta = 0. \tag{31}$$

This equality imposes p linear constraints on θ and allows us to write p of its components as linear combinations of the others. More precisely, decompose $\Psi(z)$ and θ as

$$\Psi(z) = \begin{pmatrix} \Psi(z_1) \\ \Psi(z_2) \end{pmatrix}, \quad \theta = \begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix}, \tag{32}$$

where $\Psi(z_1)$ is $(k-p) \times p$; $\Psi(z_2)$ is $p \times p$; θ_1 is $(k-p) \times 1$; θ_2 is $p \times 1$.

$$\begin{aligned} \theta'\Psi(z) = 0 & \Leftrightarrow & \theta_1'\Psi(z_1) + \theta_2'\Psi(z_2) = 0 \\ & \Leftrightarrow & \theta_2 = -[\Psi(z_2)^{-1}]'\Psi(z_1)'\theta_1 \end{aligned} \tag{33}$$

(Note that Ψ is defined on column vectors only so that we must work with $\Psi(z)$, not $\Psi(z')$.) Equation (33) implies that

$$y \in \ker(\Psi) \cap L(z) \Leftrightarrow y = z'\theta \text{ with } \theta = A\,b \tag{34}$$

23

where b is a $(k - p) \times 1$ real vector and

$$A = \begin{pmatrix} I_{k-p} \\ -[\Psi(z_2)^{-1}]'\Psi(z_1)' \end{pmatrix} \tag{35}$$

where I_{k-p} is the $(k - p)$-dimensional identity matrix. Consequently, any $y \in \ker(\Psi) \cap L(z)$ can be written as $y = b'(A'z)$ with A the $k \times (k-p)$ matrix defined above and b a $(k - p)$-dimensional vector.

$$\begin{aligned} \varepsilon \in [\ker(\Psi) \cap L(z)]^{\perp} &\Leftrightarrow& b'E[\varepsilon\,(A'z)] = 0 \ \text{ all } b \in R^{k-p} \\ &\Leftrightarrow& E[\varepsilon\,(A'z)] = 0 \end{aligned} \tag{36}$$

Besides, $\varepsilon \in \ker \Psi$ if and only if $\Psi(\varepsilon) = E[\varepsilon m'] = 0$. We conclude that

$$\varepsilon \in \ker(\Psi) \cap [\ker(\Psi) \cap L(z)]^{\perp} \Leftrightarrow E[u\varepsilon] = 0 \ \text{ with } \ u = \begin{pmatrix} m \\ A'z \end{pmatrix} \tag{37}$$

which implies that $\hat{P}[x|L(z)] = z'\beta$ with $\beta = E[uz']^{-1}\,E[ux]$.

References

Bates, David, 1988, The crash premium: Option pricing under asymmetric processes, with applications to options on the Deutschemark futures, Working Paper, the University of Pennsylvania.

Bates, David, 1991, The crash of 87: Was it expected? The evidence from options markets, *Journal of Finance* 46, 1009-1044.

Bates, David, 1997, The skew premium: option pricing under asymmetric processes, in *Advances in Futures and Options Research*, vol 9, 51-82.

Black, Fischer, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *The Journal of Political Economy*, 81, 637-659.

Carr, Peter, Katrina Ellis, and Vishal Gupta , 1998, Static hedging of exotic options, *Journal of Finance*, 53, 1165-1190.

Cheuk, Terry, and Ton Vorst, 1996, Complex Barrier Options, *Journal of Derivatives*, 4(1), 8-22.

Clarke, David, 1998, Global foreign exchange: introduction to barrier options, unpublished manuscript (http://www.hrsltd.demon.co.uk/marketin.htm).

Derman, Emanuel, Deniz Ergener, and Iraj Kani, 1995, Static Options Replication, *Journal of Derivatives*, 2 (4), 78-95.

Derman, Emanuel, and Iraj Kani, 1994, Riding on a smile, *Risk*, 7 (2), 32-39.

Duffie, Darrell and Henry Richardson, 1991, Mean-variance hedging in continuous time, *Annals of Applied Probability*, 1, 1-15.

El Karoui, Nicole and M. Quenez, 1995, Dynamic programming and pricing of contingent claims in an incomplete market, *SIAM Journal of Control Optimization*, 33 (1), 29-66.

Hans Föllmer and Peter Leukert, 1999, Quantile hedging, *Finance and Stochastics*, 3 (2), 251-273.

Hull, John, and Alan White, 1987, The pricing of options on assets with stochastic volatilities, *Journal of Finance*, 42, 281-299.

International Monetary Fund (2000), *International capital markets*, IMF, Washington, DC.

Merton, Robert, 1973, Theory of rational option pricing, *Bell Journal of Economics and Management Science*, 4, 141-83.

Schweizer, Martin, 1992, Mean-square hedging for general claims, *Annals of Applied Probability*, 2, 171-179.

Steinherr, Alfred, 1998, *Derivatives, the wild beast of Finance* (Wiley).

Taleb, Nassim, 1996, *Dynamic hedging, managing vanilla and exotic options* (Wiley series in financial engineering).

Vorst, Ton, 2000, Optimal portfolios under a value-at-risk constraint, working paper, Erasmus University Rotterdam.

Time

Tree 1 — Stock

0	1	2	3	4	5	6
						160
					150	
				140		140
			130		130	
		120		120		120
	110		110		110	
100		100		100		100
	90		90		90	
		80		80		80
			70		70	
				60		60
					50	
						40

Tree 2 — 6-year call with strike 60

0	1	2	3	4	5	6
						100
					90	
				80		80
			70		70	
		60		60		60
	50		50		50	
40.3		40		40		40
	30.6		30		30	
		21.25		20		20
			12.5		10	
				5		0
					0	
						0

Tree 3 — 6-year call with strike 100

0	1	2	3	4	5	6
						60
					50	
				40		40
			30		30	
		21.3		20		20
	14.4		12.5		10	
9.4		7.5		5		0
	4.4		2.5		0	
		1.25		0		0
			0		0	
				0		0
					0	
						0

Time

Tree 4 — 6-year call with strike 60 / 6-year call with strike 100

0	1	2	3	4	5	6
						-80
					-60	
				-40		-40
			-20		-20	
		-3.7		0		0
	6.9		12.5		20	
12.2		17.5		25		40
	17.5		22.5		30	
		17.5		20		20
			12.5		10	
				5		0
					0	
						0

Tree 5 — 6-year calls with strike 60 and 100 / 3-year call with strike 120

0	1	2	3	4	5	6
						-80
					-60	
				-40		-40
			-12.5		-20	
		0		0		0
	8.7		12.5		20	
13.1		17.5		25		40
	17.5		22.5		30	
		17.5		20		20
			12.5		10	
				5		0
					0	
						0

Tree 6 — Up-and-out call with strike 60 and barrier 120

0	1	2	3	4	5	6
						0
					0	
				0		0
			0		0	
		0		0		0
	8.75		12.5		20	
13.1		17.5		25		40
	17.5		22.5		30	
		17.5		20		20
			12.5		10	
				5		0
					0	
						0

Table 1: **The Derman, Ergener, and Kani model.** The trees represent the values at each node of some asset portfolios. The components of the portfolio figure at the top-left corner of each tree.

Tree 1
Up-and-out call with strike 60 and barrier 120

Time 0	1	2	3	4	5	6
						0
					0	
				0		0
			0		0	
		0		0		0
	8.7		8.3		10	
13.1		17.5		20.8		28
	17.5		22.5		27	
		17.5		20		18.7
			12.5		10	
				5		0
					0	
						0

Tree 3
6-year calls with strikes 60 and 100

Time 0	1	2	3	4	5	6
						-1.6
					-1.6	
				-1.6		-0.4
			-1.6		-0.1	
		-1.6		0.2		0.6
	8.3		8.6		10.4	
13.1		18.2		20.9		27.8
	17.9		22.7		26.8	
		17.6		20		18.6
			12.5		10	
				5.1		0.1
					0.1	
						0.1

Tree 2. Pure static strategies
6-year calls with strikes 60 and 100
3-year call with strike 120

Time 0	1	2	3	4	5	6
						-22.9
					-15.2	
				-7.6		-6.3
			0		1.5	
		7.3		9.5		9.6
	11.5		14.5		17.5	
13.1		15.7		19		25.3
	14.8		16.9		20.1	
		13.8		14.8		14.8
			10.8		9.5	
				6.8		4.2
					4.2	
						4.2

Tree 4
6-year call with strike 60
3-year call with strike 120

Time 0	1	2	3	4	5	6
						0
					0	
				0		5.7
			0		6.9	
		0		8.5		10.2
	8.7		11.5		12	
13.1		17.4		14.5		13.9
	17.5		17.6		15.8	
		17.7		17.7		16.6
			17.9		17.9	
				18		18.1
					18.1	
						18.1

Table 2: **Conditional values conditioned on the stock price at time t.** The trees represent the conditional values at each node of the hedging portfolios. The components of the portfolio figure at the top-left corner of each tree. The asset allocation is shown in Table 3. All numbers are rounded to the first decimal. Hedging strategies allow liquidation unless stated otherwise.

Assets	Asset allocation				
		Including Liquidation			Excluding Liquidation
	(1)	(2)	(3)	(4)	(5)
6-year call with strike 60	1.00	0.99	0.98	-0.02	0.53
6-year call with strike 100	-3.00	-2.88	-2.74	n.a.	-1.29
3-year call with strike 120	0.75	n.a.	n.a.	-3.41	-0.25
Riskless bond	0.00	0.01	-0.01	0.18	0.04
R Squared (percent)	100	99.5	98.8	21.7	43.2

Table 3: **Hedging an up-and-out call with strike 60 and barrier 120.** Columns (1) to (5) show the asset allocations in the hedging portfolio and a measure of fit (the "R squared") in five different scenarios: Allowing liquidation and trading the riskless bond and the three options used in Derman, Ergener, and Kani (1995) (column 1); excluding the short-term option from the marketable assets (column 2); imposing the additional constraint that the hedging residual has zero mean and be zero when the stock price reaches its maximum at maturity (column 3); imposing no such constraint but excluding the long-term option with the highest strike from the marketable assets (column 4); or excluding liquidation while trading in all three options and the riskless bond (column 5).
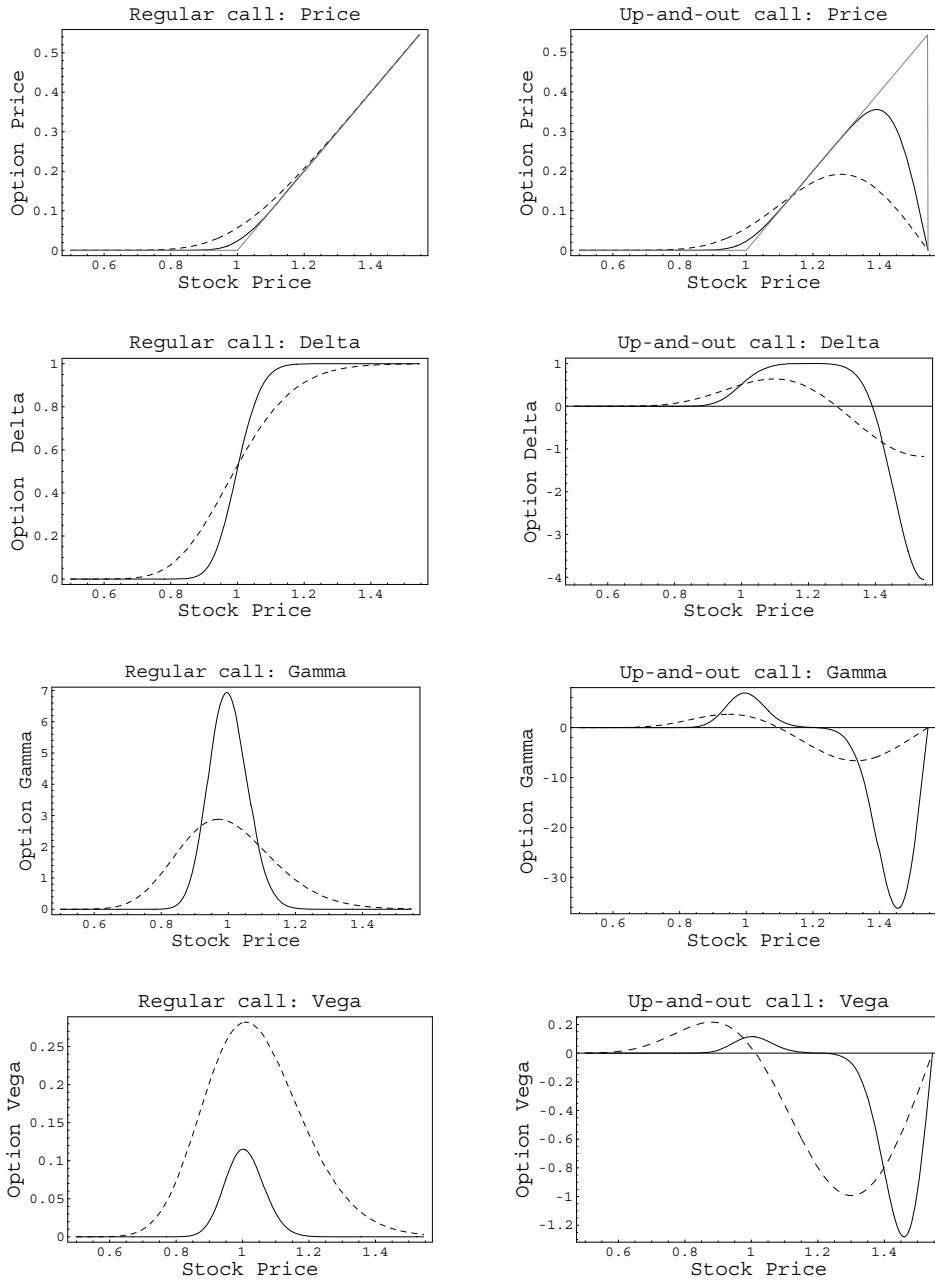
Figure 1: **The value and the Greeks of a regular call and of an up-and-out call.**
The strike price of both calls is 1.0, the barrier is 1.55, maturities are 6 months (dashed line)
and 1 month (solid line); the underlying asset is assumed to follow a geometric Brownian
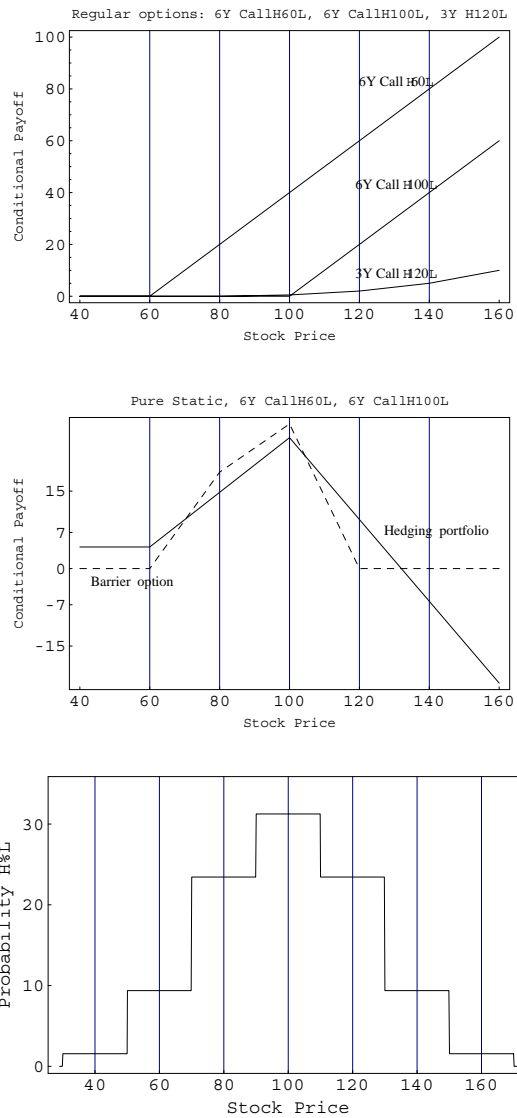motion with volatility $\sigma = 0.2$; the interest rate is assumed to be zero.

Figure 2: **Pure static hedging strategies.** Conditional values at maturity of the regular options potentially available for hedging (upper panel), of the barrier option, and of the hedging portfolio when the short-term option is excluded (middle panel). Probability distribution of the stock price (lower panel).
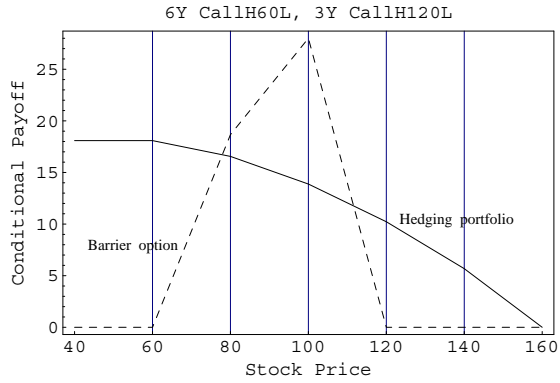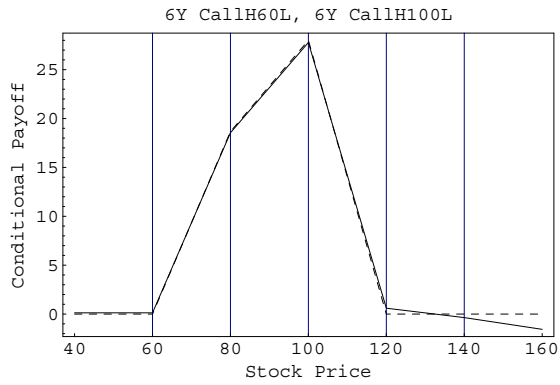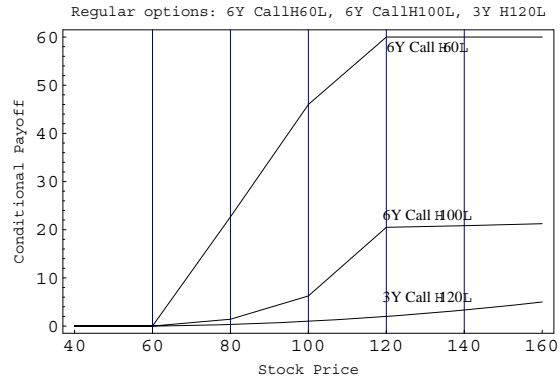
Figure 3: **Static hedging strategies allowing liquidation.** Conditional values at maturity of the regular options potentially available for hedging (upper panel), of the barrier option, of the hedging portfolio when the short-term option is excluded (middle panel), and when the long-term option with the highest strike is excluded (lower panel).
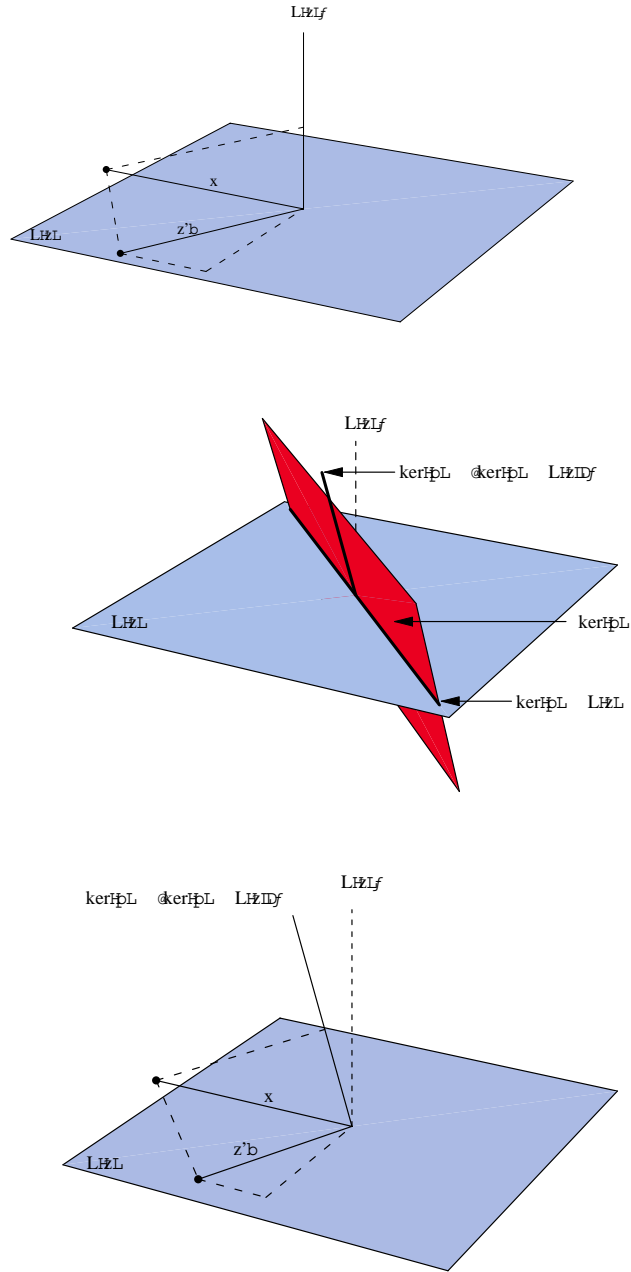
Figure 4: **Geometric approach to hedging.** Mean-square hedging (upper panel) and hedging consistently with a given pricing functional (middle and lower panels).
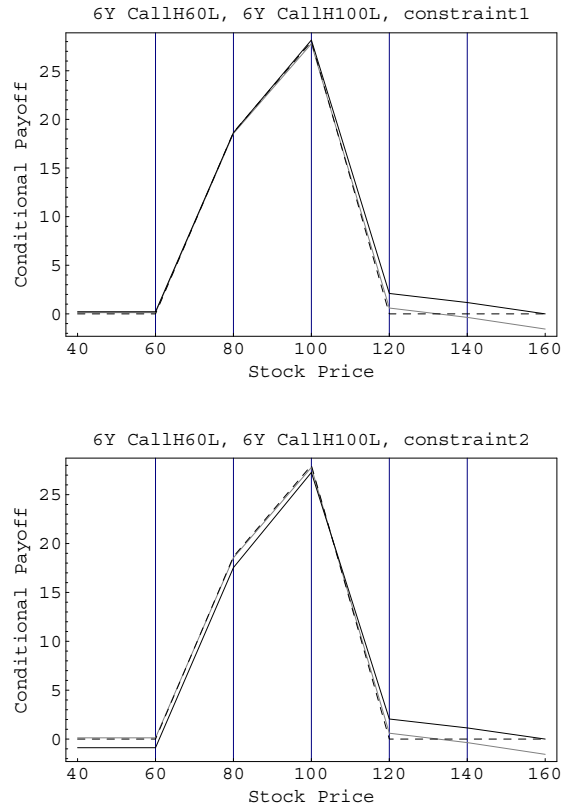
Figure 5: **Constrained static hedging strategies allowing liquidation.** Conditional values at maturity of the barrier option, of the hedging portfolio with the short-term option excluded and under the constraint that the hedging residual be zero at the maximum price level (upper panel), and under the additional constraint that the mean of the hedging residual be zero (lower panel). In both graphs, the light-gray curve represents the conditional value of the hedging portfolio when no constraint is imposed.
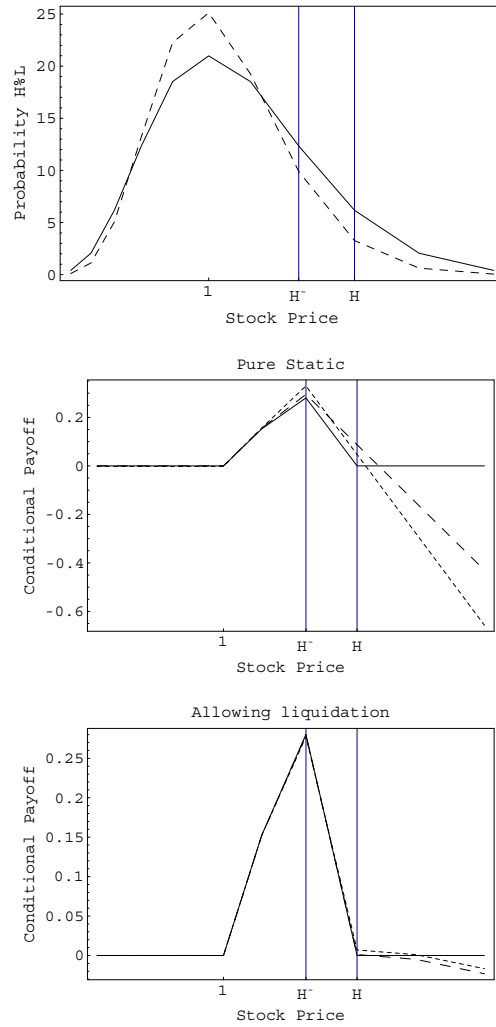
Figure 6: **Hedging consistently with a pricing functional using the risk-neutral probability distribution and another probability distribution.** The upper panel shows the risk-neutral, lognormal, probability distribution (dashed line) and the one based on the equiprobability of paths (solid line). The conditional portftolio payoffs following pure static strategies are displayed in the middle panel; those based on strategies allowing liquidation in the lower panel. In both cases, the long-dashed curves represent the conditional payoffs when mean-square hedging uses the risk-neutral (lognormal) probability distribution; the short-dashed curves show the conditional payoffs when hedging is based on the equiprobability of paths; H denotes the barrier and H^- the highest possible stock price lower than the barrier. H^- is the strike of a regular option used in hedging; H is not.