

FRU

Finance Research Unit

Static Replication and Model Risk: Razor's Edge or Trader's Hedge?

Morten Nalholm and Rolf Poulsen

No. 2005/02

**Finance Research Unit
Institute of Economics
University of Copenhagen
<http://www.econ.ku.dk/FRU>**

Static Replication and Model Risk: Razor's Edge or Trader's Hedge?*

Morten Nalholm[†]

Rolf Poulsen[‡]

Abstract

We investigate how sensitive a variety of dynamic and static hedge strategies for barrier options are to model risk. We find that using plain vanilla options to hedge barrier options offers considerable improvements over usual Δ -hedges. Further, we show that the hedge portfolios involving options are relatively more sensitive to model risk, *the Devil is in the detail*, but that the degree of misspecification sensitivity is quite robust across commonly used models.

1 Introduction

Models may produce similar plain vanilla option prices, yet give markedly different prices of exotic options. This is documented for instance in Hirta, Courtadon & Madan (2002). Focusing on barrier options, we investigate the natural follow-up question: How does this affect hedge portfolios? Qualitatively, if barrier options are derivatives of plain vanilla options, then hedge portfolios are second derivatives, and you could fear that a further order of accuracy was lost. The question is made all the more interesting because there are several alternatives to traditional dynamic Δ -hedging. We look at static hedges that are portfolios involving plain vanilla options constructed in such cunning ways that no dynamic trading is necessary to achieve replication of the barrier option. (It is not obvious that this is possible at all. The last decade's literature shows that it is.) It has been argued, and the sheer presence of the word "static" suggests it, that the use of options as hedge instruments reduces sensitivity to model risk. However, if you inspect the derivations of the static hedge portfolios, you'll see that they hinge very strongly on *something*, typically a Black/Scholes model assumption; the razor's edge of the title.

*File: NalholmStatic/nalholmpoulsen_1603.tex, version March 17, 2005. We thank Søren Willemann, Søren Lolle, and Peter Ritchken

[†]Department of Finance, Copenhagen Business School. Email: mn.fi@cbs.dk

[‡]Department of Applied Mathematics and Statistics, Institute for Mathematical Sciences, Universitetsparken 5, University of Copenhagen, DK-2100, Denmark. Email: rolf@math.ku.dk

We carry out a simulation study of the performance of hedge strategies, static and dynamic, Black/Scholes-based or in some way market data adjusted (what the trader of the title would do), against different classes of non-Black/Scholes models. We find that static hedges are relatively more model risk sensitive than Δ -hedges, that they may very well outperform these anyway, that simple adjustments to static may help considerably (and that they may not) and that the model risk sensitivity is fairly robust across commonly used models producing similar prices of plain vanilla options.

The rest of the paper is organized as follows. In Section 2 we briefly summarize the literature on pricing and statically hedging barrier options. Section 3 describes the experimental design and the Black/Scholes benchmark case. Section 4 is a simulation study of hedge performance against the (smile-generating) models constant elasticity of variance, Heston's stochastic volatility, Merton's jump diffusion, and the infinite jump intensity model Variance Gamma. Section 5 concludes. (In case you are wondering: If we had a definitive answer to the question in the paper's title, we would have put it in there.)

2 Barrier Options and Static Hedges

First, consider a model where the interest rate is constant (r) and the stock-price dynamics under the (bank-account denominated) equivalent martingale measure \mathbb{Q} are

$$dS(t) = (r - d)S(t)dt + \sigma(S(t), t)S(t)dW^{\mathbb{Q}}, \quad (1)$$

where d is a constant dividend yield. We always refer to this asset as “the stock”, but by putting $r = d$ we can model futures prices, and if we use $d = \text{“the foreign short rate”}$, then we have an arbitrage-free exchange rate model. This semi-flexible form of volatility preserves completeness of the model. It includes the Black-Scholes (B/S) model ($\sigma(x, t) = \sigma$), the constant elasticity of variance (CEV) model ($\sigma(x, t) = \sigma x^{\alpha-1}$) as well as the local volatility models originally suggested in Dupire (1994) where the functional form of σ is inferred directly from a double continuum over strikes and expiry-dates of observed option prices.

The price of a knock-out option (strictly speaking: provided it is still alive) is of the form $\pi(t) = F(S(t), t)$ where the function F solves the partial differential equation (PDE)¹

$$F_t + \frac{1}{2}x^2\sigma^2(x, t)F_{xx} + (r - d)xF_x = rF \text{ on the } \textit{alive region}$$

$$F(B, t) = R \text{ for } t < T \text{ (rebate at the barrier)}$$

$$F(x, T) = g(x) \text{ (pay-off at expiry),}$$

¹The short way to prove this to say “obviously, we must have price = rebate at the barrier, and away from it, we can use the standard Black/Scholes hedge argument”, see Wilmott (1998) for instance. A more rigorous/probabilistic approach starts with the martingale formulation in Equation (2), and then studies its connection to boundary value problems, see Øksendal (1995, Chapter 9).




Figure 1: The BVPs for down-and-out (left) and up-and-out call options

and the option can be replicated by a Δ -hedging strategy that holds

$$\Delta(t) = F_x(S(t), t)$$

units of stock time t and is kept self-financing via the bank-account.

In more complicated models (with stochastic volatility or jumps) arbitrage-free knock-out claim prices can still be found from the martingale relation

$$\pi(t) = \mathbf{E}_t^{\mathbb{Q}}(\exp(-r(\tau - t))\text{Pay-off}(\tau)), \quad (2)$$

where τ is the minimum of expiry and the first hitting time to the barrier. Analytical representations of the price become more complicated such as PDEs with multi-dimensional space variables or non-local terms.

We will look at two cases that capture the essence of pricing, hedging and model risk problems for barrier options: The down-and-out call and the up-and-out call, both with zero rebates. Closed-form expressions for B/S model prices of these two barrier options² were given already in Merton (1973), so they are as old as the B/S formula itself. In Figure 1 the boundary value problems (BVPs) are illustrated. Note that the problem for the up-

²It is fairly easy to handle any other piecewise linear pay-off function, any combination of spot, strike and barrier, as well as knock-in versions. With some further sleight of hand rebates, look-backs and double barriers can be handled.

and-out call has a discontinuity at the point (B, T) , and that spills over into unpleasant (or outright nasty, when it comes to "greeks") behavior in the general vicinity of the barrier.

2.1 Static Hedging with Calendar-Spreads

This section describes the workings of the static hedging technique first suggested in Derman, Ergener & Kani (1995). When reading the section it is important to keep in mind the BVP formulation of the barrier pricing problem, or more specifically Figure 1.

We want to construct a portfolio from plain vanilla options such that its value at expiry and along the barrier is equal to that of the barrier option, which is trivially known exactly there. This can be done with (different-expiry, strike=barrier) plain vanilla options. Consider the down-and-out call³ and suppose we have access to strike- B puts with expiry dates $t_1, \dots, t_n = T$. Let $\text{Put}(x, s|y; t)$ denote the time- s price of a strike- y expiry- t put when the stock price at time s is x (and similarly for calls). We assume this function is known; it could be the B/S-formula. Consider the following portfolio:

- 1 strike- K expiry- T call (referred to as the underlying option).
- α_n strike- B expiry- T puts (referred to as auxiliary options), where

$$\alpha_n * \text{Put}(B, t_{n-1}|B; T) + \text{Call}(B, t_{n-1}|K; T) = 0 \Rightarrow \alpha_n = -\frac{\text{Call}(B, t_{n-1}|K; T)}{\text{Put}(B, t_{n-1}|B; T)}.$$

If $S(t_{n-1}) = B$, this portfolio has t_{n-1} -value 0. If we hold the portfolio to T , and the stock-price ends above the barrier, then it pays off exactly as the call.

- α_{n-1} strike- B expiry- t_{n-1} puts, where

$$\alpha_{n-1} * \text{Put}(B, t_{n-2}|B; t_{n-1}) + \alpha_n * \text{Put}(B, t_{n-2}|B; T) + \text{Call}(B, t_{n-2}|K; T) = 0$$

If $S(t_{n-2}) = B$, this portfolio has t_{n-2} -value 0. If we hold the portfolio to t_{n-1} then its value if $S_{t_{n-1}} = B$ is also 0 because the α_{n-1} expiry- t_{n-1} strike- B puts are worthless. And if we hold it to T , then it pays off exactly as the call if $S(T) \geq B$.

- And so on: Buy α_i strike- B expiry- t_i puts, where

$$\alpha_i * \text{Put}(B, t_{i-1}|B; t_i) + \sum_{j=i+1}^n \alpha_j * \text{Put}(B, t_{i-1}|B; t_j) + \text{Call}(B, t_{i-1}|K; T) = 0$$

The time-0 price of this portfolio is $\text{Call}(S(0), 0|K; T) + \sum_{i=1}^n \alpha_i \text{Put}(S(0), 0|B; t_i)$. Suppose we buy it and liquidate it either the first time the stock price crosses the barrier B or, if

³For an up-and-out call simply substitute "call" for "put" for the auxiliary options

that doesn't happen, at time T . Should the stock-price stay above B , we get the regular call payoff. The only way a continuous stock-price process can cross the barrier is by actually hitting it, in which case we receive a payoff that is close to 0 when the spacing between the put expiry-dates is small. In other words, we are close to having replicated the down-and-out call.

In the algorithm above we can use auxiliary options with any strike beyond the barrier; the barrier as strike-level was only considered for ease of exposition. By further noting that nothing prevents us from using several strikes for each expiry-date extensions are obtained. The value of the portfolio may be forced to be 0 between the t_i -expiry-dates. If strikes are more easily available than expiry-dates this is useful. In a stochastic volatility models, as suggested and studied in Fink (2003), we don't know what the volatility is when the barrier is hit, so we may want to make the replicating portfolio's value 0 for several volatility-levels. Yet another approach is to make the portfolio's value 0 at different stock-price levels. In models with continuous sample-paths this would serve little purpose, but in for instance jump diffusion models it could be very useful. This, as well as other generalizations, is suggested in Andersen, Andreasen & Eliezer (2002).

2.2 Adjusted Pay-offs and Static Hedging with Strike-Spreads

In a series of papers Peter Carr has developed an elegant technique for constructing static hedges; see Carr & Chou (1997a), Carr, Ellis & Gupta (1998).

The key is the existence of an expiry- T simple claim⁴ that is *equivalent* to the barrier option in the sense that their values coincide at expiry and along the barrier, and thus in the entire alive region. This simple claim's pay-off function can be written down directly in terms of characteristics of the barrier option. The pay-off function g in the BVP formulation must be adjusted by a suitably reflected version of itself. What is "suitable" depends on strike, barrier and model parameters. In the Black/Scholes model the adjusted payoffs for the down-and-out and up-and-out calls (ie. $g(x) = (x - K)^+$) look like this:

$$h^{\text{DO}}(x) = \begin{cases} g(x) & \text{if } x > B \\ -(\frac{x}{B})^p g(\frac{B^2}{x}) & \text{if } x \leq B \end{cases} \quad \text{and} \quad h^{\text{UO}}(x) = \begin{cases} -(\frac{x}{B})^p g(\frac{B^2}{x}) & \text{if } x > B \\ g(x) & \text{if } x \leq B \end{cases},$$

where we have put $p = 1 - 2(r - d)/\sigma^2$.

The final step towards static replication is then to match the adjusted pay-off function by a portfolio of plain vanilla puts or calls with different strikes. For a general p , the h -functions

⁴A simple claim is a contract that pays off at time T in a way that depends only on $S(T)$. In other words it is characterized by a pay-off *function* $g : \mathbb{R} \mapsto \mathbb{R}$.




Figure 2: Adjusted pay-off functions for down-and-out call (left) and up-and-out call (right). The dotted lines show the approximations achieved using, respectively, 3 and 11 auxiliary options, as in the simulation experiments.

are not piecewise linear,⁵ for up-and-out call it isn't even continuous as seen in Figure 2, so the perfect static hedge portfolio involves a continuum of options, which is of course a problem in practice, but one we largely ignore, or rather learn to live with.

More advanced option structures (partial barrier, double barrier, lookback) are treated in Carr & Chou (1997b). Less is known about generalizations that go beyond the B/S model. Andreasen (2001) and Carr & Lee (2003) give results for stochastic volatility models, but they appear practical only in the case 0-drift case, i.e. when $r = d$.

⁵We call the region where this happens the adjustment region. For the down-and-out call it is $]0; B^2/K]$, for the up-and-out call it is $]B; B^2/K]$.

Quantity	Symbol	Numerical value
Initial stock-price	$S(0)$	100
Interest rate	r	0.06
Dividend yield	d	0.02
B/S-volatility	σ	0.20
Carr/Chou- p	$p = 1 - 2(r - d)/\sigma^2$	-1
Strike of underlying call	B	110
Expiry of underlying call	T	1
Down-and-out barrier	B^{UO}	90
Up-and-out barrier	B^{DO}	140
Δ -hedging time step	Δt	2/252
No. simulated paths	N	10,000
CEV model		
Variance elasticity	α	0.5
CEV-volatility	σ_{CEV}	2.0472
Heston model		
Mean reversion of variance	κ	1.301
Long term variance level	θ_{SV}	0.044 (=0.2097 ²)
Volatility of variance	η	0.105
Correlation (stock, variance)	ρ	-0.608
Merton model		
Jump intensity	λ	1.158
Mean jump size	γ	-0.135
Jump size variance	δ	4.7×10^{-6}
Volatility of diffusion part	σ_{JD}	0.148
Variance Gamma model		
Q-drift of diffusion part	θ_{VG}	-0.514
Q-volatility of diffusion part	σ_{VG}	0.174
Q-variance rate of gamma part	ν	0.048

Table 1: Parameter values. The parameters of the Heston, the Merton and the Variance Gamma models are determined by calibration to the 1-year volatility skew from the CEV model as shown in Figure 3.

3 Experimental Design and the Benchmark Case

The rest of the paper compares variants of hedging strategies by simulation experiments, the overall design of which is like this:

- Specify the true model. Besides the benchmark B/S model, we look at CEV, Heston's stochastic volatility, Merton's jump diffusion and a Variance Gamma model. Table 1 summarizes parameter symbols, values and interpretations.
- Suggest a variety of possible hedging strategies. Which and how many varies from model to model.
- Simulate stock-price paths from the true model. These paths run from time 0 to time τ , which is either expiry or the first time the barrier is hit; whichever comes first. Except for a brief – but we feel informative – digression at the end of this section, all simulations are done under the equivalent martingale measure \mathbb{Q} .
- Since the focus of the paper is model risk, the expiry-dates or strikes of auxiliary options used for the static hedges are chosen uniformly over the life of the barrier option or the adjustment area of modified pay-off function.
- The dynamic strategies are adjusted along the stock-price path. At each trading date we adjust our stock-position such that we hold exactly the number of stocks that the continuous-time Δ -hedge strategy prescribes. The strategy is kept self-financing by trading in the bank-account (ie. borrowing or lending).
- At time τ all strategies are liquidated and we record the discounted hedge error or the Profit/Loss (P/L), by which we mean

$$\epsilon_i^j = \exp(-r\tau_i) (\text{Value of hedge portfolio}_j(\tau_i) - \text{Value of barrier option}(\tau_i)),$$

for the j 'th hedge strategy and the i 'th stock-price path. To calculate this value for hedges involving options we need an efficient way of calculating prices of plain vanilla puts and calls in the true model; a closed-form solution being preferred of course. For many strategies we don't actually need to be able to price the barrier option in the true (and presumably complicated) model.

- Repeat this for many stock-price paths; say N of them. Record discounted hedge errors and estimate their distributional characteristics.

BLACK-SCHOLES MODEL						
Hedge method	Barrier option type					
	down-and-out call			up-and-out call		
	cost	mean	std. dev.	cost	mean	std. dev.
Pure Δ	4.8523	0.1%	10.1%	2.277	-0.7%	94.9%
Mixed Δ	4.8523	-0.1%	2.0%	2.277	-0.7%	93.5%
STR-static	4.7666	-1.8%	1.9%	2.662	16.4%	52.3%
CAL-static	4.7856	-1.4%	1.4%	2.481	9.0%	62.8%

Table 2: Hedging barrier options with a correct model. All moments are given relative to the true barrier option prices. These are equal to the costs of the Δ -hedging strategies. Both static hedges use (1+3) options in the down and out case and (1+11) options in the up and out case.

3.1 The Black/Scholes Model

The rest of this section deals with the B/S model; our specific choice of parameters is shown in Table 1. We compare four hedging strategies for the two barrier options. It should surprise no-one that we look at standard Δ -hedging, the CAL-hedge (the calendar-spreads), and the STR-hedge (the strike-spreads). But we include a fourth strategy called “mixed Δ -hedging”. Here the idea is to buy the plain vanilla option underlying the barrier option (always a strike-110 call), and then Δ -hedge the residual, ie. “barrier - plain vanilla”. Since a large part of the hedge error for discrete Δ -hedging is due simply to the “kink” in the plain vanilla call’s pay-off function, we expect this to be an improvement over standard Δ -hedging, and feel that this is the strategy static hedges should rightfully be compared to.

The down-and-out call

The second to fourth columns in Table 2 show descriptive statistics for hedge error in the case where we’re trying to hedge a down-and-out call. Under the equivalent martingale measure \mathbb{Q} the discounted value process for *any* self-financing trading strategy – continuously as well as discretely adjusted – is martingale.⁶ Therefore the difference between the cost of the hedging strategy and true price of the barrier option equals the mean of the discounted hedge error and can’t really be used as an indication of hedge quality, it merely reflects the difference between the true barrier price and our initial outlay; but it’s a convenient “check-sum”. So when we talk about “accuracy” of the hedge in the following we mean

⁶It doesn’t matter that τ is stochastic; it’s a stopping time bounded by T , so we can use the optional stopping theorem.

standard deviation. This choice is parsimonious, but not the only possible one; for the sake of brevity we have left out risk measures such as short-fall probability or semi-variance that could be relevant, in particular in light of the drift-dependence discussion at the end of this section, and the fact that in practice the hedger would typically be short the barrier option.

We first notice a large accuracy improvement as we go from pure to mixed Δ -hedging; the hedge error standard deviation drops from about 10% of the barrier option price to around 2%. This confirms that much of the inaccuracy in discrete hedging stems not from the barrier, but from the “ordinary” kink in the pay-off function at the strike. We also see that – for either type of static hedge – it takes only 3 options (in addition to the plain-vanilla (T, K) -call) to achieve a hedge that is comparable (slightly better) than what is achieved by a mixed Δ -hedge that is adjusted every second day.

The up-and-out call

Table 2 also shows hedge performance for an up-and-out call in the B/S model. The most notable change from the down-and-out hedging is the much higher standard deviations; these now range from 50% to 100% of the true barrier price, even after raising the number of options in the static portfolios to 11. Its discontinuity makes this option hard to hedge. There is a world of difference between the stock-price finishing just over and just under the barrier (for the down-and-out call, it is o-t-m when it knocks out, and we probably weren’t going to make much money on it anyway). Another way of saying this is that the Δ of the option is unbounded in the (B, T) -region. We see that there is no improvement from mixed Δ -hedging. Looking at pay-off functions, this isn’t surprising. Adding the regular call does nothing to remove the discontinuity and the kink isn’t removed, it’s just shifted (from K to B).

Drift Digression

Table 3 shows hedge errors when the stock-price is simulated with a different drift than its risk-neutral one ($r - d$); you can think of these as different \mathbb{P} -drifts (also known as real-world, statistical or physical drifts). Girsanov’s theorem tells us that with perfect (ie. continuous in time or space) hedging, the \mathbb{P} -drift does not matter. With a discrete set of hedge points/options, this is no longer true. But for plain vanilla options, the drift has very little effect on the standard deviation of the hedge error under any of the equivalent measures in play. For discontinuous barrier options, barrier hits are very critical; that’s where the errors are coming from. The drift strongly affects the probability of the barrier being hit. Thus the standard deviations of discrete hedges are quite dependent on the \mathbb{P} -drift. The less likely a barrier hit is, the better. This is important information for practical hedging: If you are a bank, and you’re selling barrier options, you’d sell them at a price

Barrier option	Hedge type	Drift shift					
		- 0.05		0.0		+ 0.05	
		mean	std. dev.	mean	std. dev.	mean	std. dev.
D/O call	Pure Δ	0.2%	15.3%	0.0%	14.8%	-0.8%	15.9%
	Mixed Δ	-0.2%	3.3%	0.2%	3.1%	-0.2%	2.9%
	STR-static	-2.1%	1.9%	-1.6%	1.9%	-1.4%	1.9%
	CAL-static	-1.4%	1.4%	-1.4%	1.6%	-1.2%	1.4%
U/O call	Pure Δ	-0.4%	89.6%	-4.0%	80.4%	-10.1%	123.0%
	Mixed Δ	-0.4%	95.3%	-0.9%	87.0%	-9.7%	127.4%
	STR-static	7.9%	36.9%	13.2%	45.2%	21.1%	60.6%
	CAL-static	4.4%	39.5%	6.6%	50.5%	-12.7%	81.2%

Table 3: Means and standard deviations of discounted hedge error for stock-prices simulated with drifts different to the risk-neutral one. Specifically, the stock-price is simulated with drift-rate = $r - d$ +drift shift.

of the form “price of theoretically perfect hedge” + “compensation because of residual risk”. The last term depends in some way on the hedge error standard deviation, so with knowledge of the direction of your trade and a good view of what the real-world drift is, you can quote “better prices”. But it is a double-edged sword: If you get the drift wrong, then things get worse. So for the rest of the paper, we simulate with a drift rate of $r - d$; or put differently we simulate under \mathbb{Q} .⁷

4 Misspecification

A direct way to see that option markets do not behave as the Black/Scholes analysis tells us is to look at the implied volatilities across strikes and see that this curve is not flat. In this section we look at four Black/Scholes extensions that are *qualitatively* quite different, but *parametrically* similar in the sense that they all produce implied volatilities that match the skew-curve in Figure 3.

4.1 The CEV Model

In the proportional volatility notation of Equation (1) the CEV model has $\sigma(x, t) = \sigma x^{\alpha-1}$. When α is less than 1 volatility increases when the stock-price falls, which makes sense

⁷In incomplete models there is the further subtlety that the change of measure from \mathbb{P} to \mathbb{Q} changes other parameters than just the drift.




Figure 3: The skew of 1-year implied volatilities. The skew was first generated by the CEV model and all subsequent models (Heston, Merton, Variance Gamma) then calibrated to it. Differences are not visible to the naked eye, and all parameter values are given in Table 1. The \triangle indicates the strike-110, expiry-1 call. The solid triangles show the options used in the STR-hedge hedge of the D/O call and the solid square is the 1-year option strike-at-barrier in the CAL- hedge.

economically when thinking of stocks as options on firm value. The model can explain part of the volatility smile, namely why o-t-m puts have higher implied volatilities than a-t-m puts; see Figure 3. However, it makes o-t-m calls have lower implied volatilities than a-t-m options, so it is sort of a “half-tricky pony”, but in some markets it does quite well empirically. Also, the model is a minimal extension of B/S, so if things break down here, then the red lights should go on. And there are closed-form solutions for plain vanilla options in the CEV model, see Cox (1996).

As for hedge strategies, more choices have to be made.

Dynamic hedge strategies

The naive Δ -hedger just uses the Δ for option from the B/S-model. To make this work a volatility has to be plugged in. We assume he uses the implied volatility of the (K, T) -call; 20% percent just as before in the pure B/S world. We also consider a correct Δ -hedger; he uses Δ calculated correctly in the CEV model. For general α , we have no closed-form solutions for barrier option prices, so we calculate Δ using a numerical PDE-method. We only have to solve a single PDE once (on a fine grid) and then store the results. Then rest

is done by table look-up.⁸

Calendar-spreads

One way to do it is exactly as in the true B/S model with the (K, T) -implied volatility as σ . But for the CEV model we can, almost as easily, do it correctly by applying the true CEV option pricing formula along the boundary when we construct the replication portfolio. Note that this does not require knowledge of barrier option prices in the CEV model.

Strike-spreads

Again, the first idea simply is to use the hedge portfolio from the B/S model with the (K, T) -implied volatility playing the role of σ .

Corrections can be made following this line of reasoning: When we construct the STR-replicating portfolio and sell short options with strikes in the “adjustment area” of the pay-off function, our model tells us that they trade at the B/S price. But the market immediately tells us that they don’t; they trade at some other price given by the volatility smile/skew. In other words, the astute trader sees a (large) discrepancy between the value of (the components of) his hedge portfolio and the value of the barrier option in the B/S model. Two ways to adjust the portfolio of auxiliary options, say w , seem natural.

First, w could be scaled by a factor β such that the price of the STR-portfolio matches the B/S model’s barrier option price, i.e.⁹

$$\beta = \frac{\text{B/S model's barrier option price} - \text{market price of underlying plain vanilla option}}{\sum_{i \sim \text{strikes}} w_i \times \text{market price of auxiliary option } i}. \quad (3)$$

This we call uniform scaling.

But not all options are equally mispriced relative to the Black/Scholes model. Adjustments should reflect that. So suppose the traded amount of auxiliary option i set such that the revenue we get from shorting it is what we would get in the true B/S model, i.e. use the *smile-scaled* portfolio weights

$$\hat{w}_i = \frac{\text{B/S model's price of auxiliary option } i}{\text{market price of auxiliary option } i} \times w_i. \quad (4)$$

Both adjustments are general; we just have to know market prices of the hedge instruments.

⁸Lo, Yuen & Hui (2001) give a series expansion in the “ $\alpha = 1/2$ ”-case that we used to cross-check the numerical PDE solution.

⁹The hedge volatility in our examples is chosen exactly such that the B/S model and market prices of the second term in the numerator in Equation (3) are equal.

CEV MODEL						
Hedge method	Barrier option type					
	down-and-out call			up-and-out call		
	cost	mean	std. dev.	cost	mean	std. dev.
Pure Δ ; B/S-based	4.8523	1.7%	10.3%	2.277	-14.1%	126%
Pure Δ ; CEV-based	4.7700	-0.1%	10.1%	2.651	0.5%	120%
Mixed Δ ; B/S-based	4.8523	1.6%	4.3%	2.277	1.6%	129%
Mixed Δ ; CEV-based	4.7700	0.0%	2.8%	2.651	0.7%	122%
STR-static	3.9405	-17.3%	14.0%	2.942	9.0%	36%
Uniformly scaled STR-static	4.8523	1.7%	6.2%	2.277	-15%	55%
Smile-scaled STR-static	4.5742	-4.1%	0.8%	2.620	-4.0%	26%
CAL-static; B/S-based	4.6837	-1.7%	1.8%	3.184	20.1%	75%
CAL-static; CEV-based	4.7011	-1.4%	1.5%	2.884	8.8%	65%

Table 4: Hedging the down-/up-and-out calls with a possibly misspecified model. All moments are given relative to the true barrier option prices. These are 4.7700 for the down-and-out call, and 2.651 for the up-and-out call. The initial prices of the hedge portfolios are given in the “cost” columns. This information is largely superfluous, but very convenient for debugging and replication purposes. The static hedges use the same (1+3) and (1+11), respectively, options as in Table 2. B/S-based strategies use the implied volatility of a the underlying plain vanilla call as “hedge volatility”.

In Table 4 we report the results of the nine strategies for hedging down-and-out and up-and-out calls in the CEV model. Looking at the down-and-call first, we see that when Δ -hedging with the stock, the gain from using the CEV model is barely visible; the standard deviations of the discounted hedge errors (for short, we simply call this the *error* in the following) are 10.3% and 10.1% of the true price. When options are incorporated in the replicating portfolio the effects of using a misspecified or naively applied model are evident. When going from B/S to CEV the error drops from 4.3% to 2.8% for the mixed- Δ portfolio, and that involves only one option, namely the underlying. For the strike-spread hedges the error reduction is even more dramatic, from 14% (naive) to 6.0% (uniform scaling) to 0.8% (smile-scaled). The effect on calendar-spread hedges is lower (a drop from 1.8% to 1.5%), but still considerable in relative terms. The reason the strike-spread hedges benefit the most from adjustments is that while the calendar-spread hedges use strikes along the barrier, the strike-spread hedges involve options with strikes well beyond the barrier and thus deep o-t-m (see Figure 3) and “most mispriced” by B/S.

Turning to the discontinuous barrier option, ie. the up-and-out call, we first notice that it is still much harder to hedge, no matter if we use static or dynamic hedging. At hedge frequencies that it does not seem possible to go beyond in practice, the error is of the same order of magnitude as the barrier option’s price itself. Again, the effects on dynamic hedges are low (errors range between 120% and 130%; as in the B/S-case mixed- Δ doesn’t help here), although larger than for the down-and-out call. The static hedges clearly outperform the dynamic ones, and in this case the strike-spreads (25% - 50% errors) hold the advantage over calendar-spreads (65% - 75% errors). The effects from adjustments are clearly seen. Calendar-spreads with the true CEV model and smile-scaling (misspecification correction reflecting different degrees of “non-B/S’ness”) causes errors to drop about 10%-points, but the uniform scaling increases the error (from 36% to 55%). There is a large difference between the price of the barrier option in the true CEV model (2.651) and the B/S-price (2.227), which is what the uniformly scaled portfolio’s price matches, and that causes the uniform adjustment to “overshoot”.

4.2 Stochastic Volatility

In virtually all financial markets, price change variances display considerable variation over time. A model that captures this is the Heston (1993) model,

$$\begin{aligned} dS(t) &= (r - d)S(t)dt + \sqrt{v(t)}S(t)dW_1^{\mathbb{Q}} \\ dv(t) &= \kappa(\theta_{SV} - v(t))dt + \eta\sqrt{v(t)}(\rho dW_1^{\mathbb{Q}} + \sqrt{1 - \rho^2}dW_2^{\mathbb{Q}}). \end{aligned} \quad (5)$$

Here volatility, or more precisely local variance, is modeled by a Cox/Ingersoll/Ross-type process, so it is random, mean-reverting, and may be correlated with price movements themselves (something that is strongly supported empirically in stock markets).

Simulation by a discretization scheme is straightforward. We use an Euler scheme with reflecting boundary at 0 for the variance process. Call prices can be expressed in closed form up to the integration of a known function. If the correlation is non-zero we know of no closed-form barrier option formula, but the zero correlation case is treated in Lipton (2001). In this model Δ -hedging is “impossible, but not hard”. Perfect dynamic hedging with the stock and bank-account is no longer possible, but there is no difficulty in creating operational Δ -hedge strategies by plugging in $\sqrt{v(t)}$ or some implied volatility.

So the immediate hedge strategies we consider are the pure and mixed Δ -hedging strategies, the pure, the uniform and smile-scaled STR-hedges and the CAL-hedge all implemented using implied (K, T) -volatility where needed.¹⁰

We consider two further versions of the CAL-hedges. First, we use a hedge volatility equal to the square root of the expected local variance conditional on the underlying being at the

¹⁰Plugging in $\sqrt{v(t)}$ in the dynamic strategies is easily done and matters little.

HESTON MODEL						
Hedge method	Barrier option type					
	down-and-out call			up-and-out call		
	cost	mean	std. dev.	cost	mean	std. dev.
Pure Δ	4.8381	2.0%	15.4%	2.2828	-19.0%	85.9%
Mixed Δ	4.8381	0.9%	4.5%	2.2828	-19.6%	93.2%
STR-static	4.1206	-11.1%	13.3%	2.9955	9.1%	38.0%
Uniformly scaled STR-static	4.8381	1.7%	2.5%	2.2769	-14.9%	58.7%
Smile-scaled STR-static	4.7527	-0.1%	1.3%	2.6664	-1.7%	19.3%
CAL-static	4.7866	0.1%	2.9%	3.0926	11.0%	57.4%
CAL-static with conditional vol.	4.8142	0.6%	3.2%	2.8969	4.3%	49.1%
CAL-static with Fink's extension	4.6166	-4.1%	3.1%	2.8239	2.1%	31.6%

Table 5: Hedging the down-/up-and-out calls with a misspecified model under the Heston model. All moments are given relative to the true barrier option prices. These are 4.8245 for the down-and-out call, and 2.7579 for the up-and-out call. The Fink extension uses 5 strikes beyond and thus hedges 5 volatility levels chosen symmetrically around the implied volatility of the strike-110 call.

barrier. This can be calculated in the spirit of Dupire (1994) or Derman & Kani (1998) from the formula

$$\mathbf{E}^{\mathbb{Q}}(v(t)|S(t) = B) = 2 \frac{\partial_T \text{Call} + (r - d)B \partial_K \text{Call} + r \text{Call}}{B^2 \partial_{KK}^2 \text{Call}}, \quad (6)$$

where we use the notation of Section 2.1 and all the functions on the right hand side must be evaluated at $(S(0), 0|B, t)$.

Second, we employ the extension suggested in Fink (2003), where for each expiry-date, 5 different options with strikes beyond the barrier (we use strike-steps sizes of 5) is used to create a portfolio whose value is 0 for 5 possible possible levels of $\sqrt{v(t)}$ (chosen symmetrically around the (K, T) -volatility implied volatility).

In Table 5 we report the results of the eight hedge strategies. The picture is virtually the same as that from the CEV-analysis in the Table 4: Static hedges outperform Δ -hedges, static hedges are more sensitive the misspecification (too simple corrections may make things worse), and discontinuous options are hard to hedge. Further, we see that all the CAL-hedges behave in the same way for the down-and-out call (the easy case). For up-and-out call it helps the CAL-hedge to use the expected conditional volatility along the barrier (but still only one auxiliary option for each expiry); the error goes from 57% to 49%. It helps even

MERTON JUMP DIFFUSION MODEL						
Hedge method	Barrier option type					
	down-and-out call			up-and-out call		
	cost	mean	std. dev.	cost	mean	std. dev.
Pure Δ	4.8564	2.6%	31.6%	2.2752	-20.2%	132.8%
Mixed Δ	4.8564	-1.5%	6.1%	2.2752	-20.9%	150.0%
STR-static	4.1378	-9.7%	11.2%	2.9942	5.1%	29.7%
Uniformly scaled STR-static	4.8564	0.2%	2.5%	2.2769	-17.1%	66.1%
Smile-Scaled STR-static	4.7705	-1.0%	2.3%	2.6612	-5.6%	31.1%
CAL-static; B/S	4.7777	-1.0%	2.9%	3.0145	5.4%	54.1%

Table 6: Hedging the down-/up-and-out calls with a misspecified model under Merton’s jump-diffusion. All moments are given relative to the true barrier option prices. These are 5.1590 for the down-and-out call, and 2.8685 for the up-and-out call.

more (error of 32%) to use the correction suggested by Fink, where several strikes are used for each auxiliary option expiry-date such that the hedge portfolio’s value is 0 for different $\sqrt{v_t}$ -levels. Anything else would be strange, since this is a “theoretically perfect hedge”, and with the use of extra strikes, it could be said to have an unfair advantage in comparison to the other CAL-hedges.

4.3 Jumps

Merton’s jump diffusion

First, we consider the classical Merton (1976) model where stock returns are hit by the Poisson arrivals of (displaced) lognormal jumps

$$\begin{aligned}
dS(t) &= (r - d - \lambda k)S(t)dt + \sigma_{JD}S(t)dW^{\mathbb{Q}} + (J_t - 1)S(t)dq, \\
\log J_t &\sim N(\gamma, \delta), \\
dq &\sim \text{Poisson}(\lambda), \\
k &= \mathbf{E}^{\mathbb{Q}}[J_t - 1].
\end{aligned}$$

Simulating the process is easy and call prices can be written as (infinite) sums of Black/Scholes prices. Prices for barrier options have to be found numerically, and there are several ways to do this. Since we only need the true price as a benchmark, we just use simulation. Models with random jump sizes are (grossly) incomplete, so perfect Δ -hedging is not possible; qualitatively it is “more impossible” here than in stochastic volatility models. However, the same operational dynamic hedge strategies are considered, that is the pure and mixed

VARIANCE GAMMA MODEL						
Hedge method	Barrier option type					
	down-and-out call			up-and-out call		
	cost	mean	std. dev.	cost	mean	std. dev.
Pure Δ	4.8531	2.3%	29.8%	2.2766	-25.1%	150.5%
Mixed Δ	4.8531	-1.3%	5.1%	2.2766	-25.2%	165.8%
STR-static	4.1296	-9.5%	11.2%	2.9957	4.1%	32.2%
Uniformly scaled STR-static	4.8523	0.3%	2.5%	2.2769	-17.0%	69.5%
Smile-Scaled STR-static	4.7673	-0.9%	2.3%	2.6622	-6.5%	41.6%
CAL-static	4.7752	-1.0%	2.8%	3.0182	4.8%	56.5%

Table 7: Hedging the down-/up-and-out calls with a misspecified model under the Variance Gamma model. All moments are given relative to the true barrier option prices. These are 5.1692 for the down-and-out call, and 2.9430 for the up-and-out call.

Δ strategies. For the static hedges we use the both the pure and the smile-scaled strike-spread hedges and the calendar-spread using implied (K, T) -volatilities in the Black-Scholes formula.

Variance Gamma

Merton's jump diffusion has few but large jumps. There are other ways jumps can occur. The Variance Gamma model from Madan, Carr & Chang (1998) is a pure jump process with many small jumps but infinite arrival rate. With $\gamma(\cdot; \alpha, \beta)$ denoting a Gamma process, ie. the Levy process whose increments over unit intervals follow a $\text{Gamma}(\alpha, \beta)$ -distribution, the basic Variance Gamma model can be written like this:

$$\begin{aligned}
X(t; \theta_{VG}, \sigma_{VG}, \nu) &= \theta_{VG} \gamma(t; 1, \nu) + \sigma_{VG} W^{\mathbb{Q}}(\gamma(t; 1, \nu)), \\
S(t) &= S(0) \exp((r - d + \omega)t + X(t; \theta_{VG}, \sigma_{VG}, \nu)), \\
\omega &= \frac{1}{\nu} \ln \left(1 - \theta_{VG} \nu - \frac{\sigma_{VG}^2 \nu}{2} \right).
\end{aligned}$$

The model can be thought of as a Black/Scholes model that runs after a different, stochastic clock that captures varying trading activity over time. This and related models have recently become popular because to their ability to match both the dynamics of the underlying and option market prices. Simulation is fairly straight-forward, see e.g. Glasserman (2003), and call prices can be expressed with special functions (modified Bessel, degenerate hypergeometric) or found by integration of B/S prices, see e.g. Joshi (2003). We use simulation to find the true prices of the barrier options.

Tables 6 and 7 show the performance of the hedge strategies for jump diffusion and Variance Gamma models. The general picture is much the same as for the CEV and Heston model analyses. That in itself is interesting, but let us just comment on the new effects. First, in both jump models, the hedge quality for adjusted static hedges of the down-and-out call is surprisingly high (errors between 2% and 3%). We are quite far from the B/S-world, and it's not hard to imagine that jumps could cause havoc for barrier options. And sometimes they do, because for the discontinuous barrier option, the corrections to the STR-hedges do not help, but rather tend to make matters worse. Earlier we found that uniform scaling might be too simple, but in this case taking account of different degrees of misspecification through smile-scaling also has an adverse effect (less so in the jump diffusion model; that is probably because jumps are predominantly negative, so knock-out is usually caused by diffusion), which makes the construction of other jump-adjustments a clear topic for further research, but one we will not pursue here.

5 Conclusion

Barrier option hedge portfolios that involve plain vanilla options offer improvements over ordinary Δ -hedging. That's well known. A thorough analysis of model risk does not alter that conclusion. Although barrier options and jump models are a tricky combination, the likeness of Tables 4 to 7 indicates a reassuring robustness across the most commonly used Black/Scholes alternatives. However, these static hedges are relatively more model risk sensitive. That is wise to remember, especially if you are the trader from the title of the paper, because then your competitors probably also do static hedging, and (as traders are confident people) the sensitivity result is really a positive one, because it means that there is an edge for you. And since the academic literature offers little constructive advice beyond Black/Scholes there are wide perspectives for both sides.

References

- Andersen, L., Andreasen, J. & Eliezer, D. (2002), 'Static replication of barrier options: some general results', *Journal of Computational Finance* **5**, 1–25.
- Andreasen, J. (2001), 'Behind the Mirror', *Risk Magazine* **14**(November).
- Carr, P. & Chou, A. (1997a), 'Breaking Barriers', *Risk Magazine* **10**(September).
- Carr, P. & Chou, A. (1997b), Hedging Complex Barrier Options. Working Paper. Available from <http://www.math.nyu.edu/research/carrp/>.

- Carr, P., Ellis, K. & Gupta, V. (1998), ‘Static Hedging of Exotic Options’, *Journal of Finance* **53**, 1165–1190.
- Carr, P. & Lee, R. (2003), Static hedging with Correlated Stochastic Volatility. Working paper, Courant Insitute, NYU.
- Cox, J. C. (1996), ‘The Constant Elasticity of Variance Option Pricing Model’, *Journal of Portfolio Management* pp. 15–17.
- Derman, E., Ergener, D. & Kani, I. (1995), ‘Static Options Replication’, *Journal of Derivatives* **2**, 78–95.
- Derman, E. & Kani, I. (1998), ‘Stochastic Implied Trees: Arbitrage Pricing with Stochastic Term and Strike Structure of Volatility’, *International Journal of Theoretical and Applied Finance* **1**, 61–110.
- Dupire, B. (1994), ‘Pricing with a smile’, *Risk Magazine* **7**(January).
- Fink, J. (2003), ‘An Examination of the Effectiveness of Static Hedging in the Presence of Stochastic Volatility’, *Journal of Futures Markets* **23**, 859–890.
- Glasserman, P. (2003), *Monte Carlo Methods in Financial Engineering*, Springer.
- Heston, S. (1993), ‘A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options’, *Review of Financial Studies* **6**, 327–344.
- Hirsa, A., Courtadon, G. & Madan, D. B. (2002), ‘The Effect of Model Risk on the Valuation of Barrier Options’, *Journal of Risk Finance* **4**(Winter), 47–55.
- Joshi, M. S. (2003), *The Concepts and Practice of Mathematical Finance*, Cambridge.
- Lipton, A. (2001), *Mathematical Methods For Foreign Exchange: A Financial Engineer’s Approach*, World Scientific Publishing Company.
- Lo, C. F., Yuen, P. H. & Hui, C. H. (2001), ‘Pricing Barrier Options with Square Root Process’, *International Journal of Theoretical and Applied Finance* **4**, 805–818.
- Madan, D. B., Carr, P. & Chang, E. C. (1998), ‘The Variance Gamma Process and Option Pricing’, *European Finance Review* **2**(1), 79–105.
- Merton, R. (1973), ‘The Theory of Rational Option Pricing’, *Bell Journal of Economics and Management Science* **4**, 141–183.
- Merton, R. (1976), ‘Option pricing when the underlying stock returns are discontinuous’, *Journal of Financial Economics* **5**, 125–144.

Øksendal, B. (1995), *Stochastic Differential Equations*, 4. edn, Springer-Verlag.

Wilmott, P. (1998), *Derivatives*, Wiley.