

Auto-Static for the People: Risk-Minimizing Hedges of Barrier Options

Johannes Sivéń* Rolf Poulsen†

February 18, 2008

Abstract

We present a method for computing risk-minimizing static hedge strategies. The method is straightforward, yet flexible with respect to the type of contingent claim being hedged, the choice of risk-measure, and the underlying asset dynamics. Experimental investigations for barrier options show that in a stochastic volatility model with jumps the resulting hedges outperform previous suggestions in the literature. We also illustrate that the risk-minimizing static hedges work in an infinite intensity Levy-driven model, and that the performance of the hedges is robust to model risk.

Key words: Risk-minimization, static hedge, barrier option, Bates model, NIG model, model risk.

AMS subject classification: 91B28, 91B30

JEL classification: G13, C61

1 Introduction

In this paper we present a risk-minimizing framework for static hedging under general asset dynamics. The idea is simply this: Let P be the time- T (T could be random) payoff of an exotic option we want to hedge, and let $H = (H_1, \dots, H_M)$ be the time- T

*Centre for Mathematical Sciences, Lund University, Box 118, SE-22100 Lund, Sweden. E-mail: jvs@maths.lth.se

†Centre for Finance, University of Gothenburg, Box 640, SE-40530 Gothenburg, Sweden. E-mail: rolf.poulsen@economics.gu.se

value of some hedging instruments; think of these as plain vanilla options. Given a loss function u that specifies our risk-measure (for instance $u(x) = x^2$ or $u(x) = x^+$), choose the hedge weights c that minimize

$$E[u(P - H \cdot c)] \tag{1}$$

subject to constraints on the portfolio weights and the cost of the portfolio; such constraints are linear by nature. Use the sample average method: Generate N independent samples $(P^{(n)}, H^{(n)})$, approximate $E[u(P - H \cdot c)]$ by

$$\frac{1}{N} \sum_{n=1}^N u(P^{(n)} - H^{(n)} \cdot c),$$

and minimize this expression numerically subject to the constraints; when u is convex this is an easy numerical problem. Besides the flexibility regarding the choice of risk-measure and the conceptual and numerical simplicity illustrated by the 10-line description above, this way of constructing hedge portfolios has several other advantages. First, real markets are not complete, so exotic options cannot be perfectly hedged; statically or otherwise. Contrary to previous constructions of static hedges, the risk-minimizing approach gives an ex-ante measure of the residual risk of the hedge. It even provides a full (cost, risk)-profile; an efficient frontier. Second, the proposed method is truly general with respect to the dynamics of the underlying. There are no hidden assumptions (such as Markovianity, continuity, zero-correlation, ...) that cause it to break down or yield only trivial results.

The main contribution of this paper lies in the application of the risk-minimizing approach to static hedging of barrier options. In the 90's¹ Peter Carr and Emanuel Derman (and various co-authors) demonstrated that in the Black-Scholes model there are simple ways to construct portfolios of plain vanilla options that perfectly match the payoff from barrier options without the need for dynamic adjustments; one sets up the portfolio at initiation, monitors when the barrier option expires or is knocked out, and then unwinds the portfolio of vanilla options. Since then, such static hedges have been studied extensively in the literature. Most works take outset in the results and techniques from the basic Black-Scholes framework, which are then tweaked and extended. In this

¹Early papers are Carr, Ellis & Gupta (1998) and Derman, Ergener & Kani (1995); a recent survey is Poulsen (2006). Carr used symmetry to prove equivalence between barrier options and certain simple claims, which are then replicated by different-strike calls and puts. Derman described a numerical method that uses options with different expiries, and is very easy to understand.

paper we take a different view by starting from general principles and then demonstrating their practical applicability. Some sources do take a minimization or regression approach to static hedging; examples are Dupont (2001), Allen & Padovani (2002), and Pellizzari (2005). However, compared to the techniques and investigations in this paper, those studies are limited by considering some combination of (i) only the quadratic criteria, (ii) effectively unconstrained portfolios, and (iii) only the Black-Scholes framework.

In our experiments we look at four different risk measures: quadratic ($u(x) = x^2$), positive part ($u(x) = x^+$), value-at-risk and expected shortfall. The two latter cases do not immediately present themselves on the form (1), but given the samples $(P^{(n)}, H^{(n)})$ we can nevertheless compute the risk-minimizing hedge strategies by solving minimization problems — a convex problem in the case of expected shortfall and a non-convex for value-at-risk. We investigate the performance of the static hedge strategies on up-and-out call-options using European call-options as hedging instruments. Discontinuity (the corresponding plain vanilla call is in-the-money when the barrier option knocks out) makes this a hard hedging problem, but also the most relevant test for practical purposes: by far the largest trading volume in barrier options is in such “reverse” or “live-out” options. The method we propose easily reproduces any static hedge portfolio based on Black-Scholes assumptions, and we show that in a more complex model — the Bates model with stochastic volatility and jumps — it gives hedge portfolios with superior performance to those recently suggested in the literature. This experiment also highlights the importance of being able to handle the cost/risk trade-off, and that no universally optimal static hedge strategy exists; it depends on which risk-measure is used.

We end the paper by looking at model risk aspects through two *gedankenexperimente*. These demonstrate that (i) it is beneficial (but not crucial) to take into account the joint dynamics of state variables and plain vanilla options (“both \mathbb{P} and \mathbb{Q} matter”), (ii) static hedging through risk-minimization is also feasible in infinite intensity Levy-models where previously suggested techniques do not work, (iii) the performance of the risk-minimizing static hedges is robust to model risk in the sense that a Bates-optimal hedge performs well in the Normal Inverse Gaussian model, and vice versa, provided that the models agree on vanilla option prices.

2 Risk-minimizing static hedges

Consider a formal setting like this: Let T be some finite final time-point, look at a filtered probability space, and let $(S_t)_{t \in [0, T]}$ be an adapted process modelling a traded asset, *the stock* for easy reference. Let P denote the (one and only) time- T payoff of a contingent claim, an exotic, illiquid, or path dependent option on the stock. Furthermore, let $H = (H_1, \dots, H_M)$ denote the time- T value of M liquid hedging instruments: typically plain vanilla options, but more generally any contingent claims. A *static hedge strategy* is some $c = (c_1, \dots, c_M) \in \mathbb{R}^M$, where c_i represents the position in the i th hedging instrument ($c_i > 0$ is a long position), so the time- T value of the corresponding hedge portfolio is $H \cdot c = \sum_{m=1}^M c_m H_m$.

2.1 Definition of a risk-minimizing hedge

Look at a function $u : \mathbb{R} \mapsto \mathbb{R}$, and let D denote a compact set in \mathbb{R}^M . We define a *risk-minimizing hedge strategy* corresponding to u and D as a solution to the stochastic programming problem

$$\min_{c \in D} E[u(P - H \cdot c)]. \quad (2)$$

The definition of the hedging strategy is quite natural: $P - H \cdot c$ is the loss on a short position in the exotic contract, and short is the typical direction for an exotic option hedger; buyers most often use the contracts for speculative or insurance purposes, and thus have little interest in (or ability to) hedge their option position.

The requirement that the set D of admissible strategies is a compact is very natural from a practical point of view: hedgers cannot take arbitrarily large positions. Furthermore, note that a restriction on the price of the hedge portfolio gives a simple linear restriction.

An approximation of the minimizer of $g(c) = E[u(P - H \cdot c)]$ can be computed by the so-called sample average approximation method. To this end N independent samples $(P^{(n)}, H^{(n)})$ of the pair (P, H) are generated, and $g(c)$ is approximated by

$$\widehat{g}_N(c) = \frac{1}{N} \sum_{n=1}^N u(P^{(n)} - H^{(n)} \cdot c). \quad (3)$$

In cases where u is convex and D is defined by linear constraints, this can be minimized with little difficulty with standard numerical methods. Most off-the-shelf software packages have efficient solution algorithms. Not only is the sample average approximation

method intuitively appealing, it can be shown, see Shapiro (2008) and the references therein, that under mild conditions the minimizer of (3) converges to the minimizer of (2) as $N \rightarrow \infty$. In fact, refined asymptotic normality and large deviations results are available.

2.2 Scaling and non-linearity of the hedges

For a moment let us assume that u has the form

$$u(x) = \beta_1(x^-)^\gamma + \beta_2(x^+)^\gamma, \quad (4)$$

where $\beta_1, \beta_2 \geq 0$ and $\gamma \geq 1$ are constants. This means that u has the positive homogeneity property $u(\alpha x) = \alpha^\gamma u(x)$, $\alpha > 0$. With the convention $\alpha D = \{\hat{c} \in \mathbb{R}^M; \hat{c} = \alpha c \text{ for some } c \in D\}$, we see by direct inspection that $c^* \in D$ minimizes $E[u(P - H \cdot c)]$ over D if and only if αc^* minimizes $E[u(\alpha P - H \cdot c)]$ over αD . So assuming that budget and weight restrictions change appropriately, we can compute the risk minimizing hedge of a single contract and then scale the hedge weights by $\alpha > 0$ to obtain the risk minimizing hedge for α contracts. This is very important from a practical point of view, since this is exactly what traders do. Functions of the form $u(x) = (x^+)^\gamma$ are homogeneous and do not punish gains; they are one-sided. The parameter γ is related to the risk-aversion of the hedger: the higher γ , the more averse is he to large losses. However, any other choice than $u(x) = x^2$ gives hedges that are not linear: The risk-minimizing hedge of $P_1 + P_2$ is not the sum of the risk-minimizing hedges of P_1 and P_2 . This is troubling from a practical point of view, since it implies that a trader should take the whole bank's assets into account when hedging the risks on his own book.

2.3 Risk-measures used in the experiments

A variety of risk-measures have been suggested; Artzner, Delbaen, Eber & Heath (1999) present an axiomatic approach and a recent work with descriptions, discussions and suggestions is Cherny (2006). In our experiments we focus on four that are commonly used and/or have tractable properties: the quadratic ($u(x) = x^2$), the positive part ($u(x) = x^+$), value-at-risk, and expected shortfall.

Quadratic hedging has been extensively studied, in dynamical settings, see Schweizer (2001) and the references therein, as well as in the context of static hedging, see Dupont (2001). Minimizing $E[(P - H \cdot c)^2]$ is clearly a natural thing to do, although from a

financial point of view it is strange to punish gains and losses symmetrically. Without portfolio restrictions this is a standard least squares regression.

With $u(x) = x^+$ we have a one-sided risk-measure — only losses are punished. Minimizing expected loss with such a convex loss function is advocated in for instance Föllmer & Schied (2002). The associated optimization problem is clearly convex, and can in fact be reformulated as a fully linear programming problem at the cost of including extra variables.

The value-at-risk at level α of the loss $P - H \cdot c$ is the upper α -quantile of the loss distribution,

$$\text{V@R}_\alpha = \inf\{z \in \mathbb{R}; \mathbb{P}(P - H \cdot c \geq z) \leq \alpha\},$$

where typically $\alpha = 0.01$ or 0.05 . This risk-measure has some well-known shortcomings,² but it is widely used. One reason for this is probably its repeated suggestion in the Basel capital directives. Another “pro” argument is that quantiles are more robust than moments. In a risk estimation context Cont, Deguest & Scandolo (2007) have recently demonstrated that robustness and coherency are conflicting objectives. V@R does not fit directly in the $E[u(\cdot)]$ -formulation, but we can easily compute the discrete analogue from the order statistics of the samples. The associated optimization problem is non-convex.

Expected shortfall (also known as tail-V@R or conditional V@R) is the mean of the loss beyond value-at-risk:

$$\text{ES}_\alpha = E[P - H \cdot c | P - H \cdot c \geq \text{V@R}_\alpha].$$

As shown by Tasche (2002), expected shortfall is a coherent³ risk measure in the sense of Artzner et al. (1999). This approach does not immediately fit the $E[u(\cdot)]$ -formulation. However, Rockafeller & Uryasev (2000, Theorem 1) show that with some sleight-of-hand (extra variables and constraints) minimization of expected shortfall can be formulated as a linear programming problem.

3 Hedge Performance in the Bates Model

In this section we conduct experiments with the risk-minimizing static hedge techniques and document their superior performance to what has previously been suggested in the

²V@R fails coherency on subadditivity. It may be the case that $\text{V@R}(X + Y) > \text{V@R}(X) + \text{V@R}(Y)$; Tasche (2002) gives an example involving Pareto-distributions.

³A small adjustment is needed in the definition above to maintain coherency for distributions with point mass at V@R_α .

literature.

As target option we choose a daily-monitored, zero-rebate up-and-out call-option; that is it pays

$$P = (S_T - K)^+ \mathbf{1}_{\max_{t \in \{dt, 2dt, \dots, T\}} S_t < B},$$

with $dt = 1/252$. This is a so-called reverse barrier option; the corresponding plain vanilla call is in-the-money when the barrier option knocks out. This creates a discontinuity that makes the hedging difficult. We set the initial stock price to $S_0 = 100$, and for the target option we choose expiry $T = 1$, strike $K = 110$ and barrier $B = 130$. For hedging we use the time- T value of vanilla call options with expiry less than or equal to T that are liquidated if the barrier is crossed; this is the usual way to statically hedge barrier options. To be precise, let $\tau = \min\{\text{first barrier crossing}, T\} = \min\{\min\{t \in \{dt, 2dt, \dots, T\}; S_t \geq B\}, T\}$. Any cash flows prior to T are held in the bank account with constant interest rate $r = 0.02$. If the k th vanilla call has strike $K^{(k)}$ and expiry $T^{(k)} \leq T$, then

$$H_k = \begin{cases} e^{(T-\tau)r} \times (\text{time-}\tau \text{ market price of } k\text{th call option}) & \text{if } \tau < T^{(k)}, \\ e^{(T-T^{(k)})r} (S_{T^{(k)}} - K^{(k)})^+ & \text{otherwise.} \end{cases}$$

Following Nalholm & Poulsen (2006) we restrict ourselves to hedge instruments that are most 35% out of the money,⁴ we use vanilla calls with strikes 110, 130, 131, \dots , 135.

As data-generating process for our experiments we use the model of Bates (1996) which combines stochastic volatility and jumps. Under the real-world probability measure \mathbb{P} the dynamics are

$$\begin{aligned} dS_t &= \mu S_t dt + \sqrt{V_t} S_t dW_t^S + J_t dN_t, \\ dV_t &= \kappa (\theta - V_t) dt + \sigma \sqrt{V_t} dW_t^V, \\ dW_t^S dW_t^V &= \rho dt, \\ dN_t &\sim \text{Po}(\lambda dt), \\ \ln(1 + J_t) &\sim N\left(\log(1 + \alpha) - \frac{\gamma^2}{2}, \gamma^2\right). \end{aligned}$$

⁴To have any place in a static hedge, an instrument must have a strike that equals the strike of the barrier option or lies at or beyond the barrier, otherwise it will make a time- T contribution in the case where the barrier option is not knocked out. If lower-strike calls are included as potential hedge instruments, their weights are put to — or very close to — zero through the optimization.

This is an incomplete model, so there is a multitude of equivalent martingale measures. We assume the (pricing) martingale measure \mathbb{Q} preserves the parametric structure of the model. As values for the \mathbb{P} - and \mathbb{Q} -parameters we use those reported in the comprehensive study in Eraker (2004); they are reproduced in the left columns of Table 1 (the parameters without superscripts are the same under \mathbb{P} and \mathbb{Q}). The values of the risk-adjusted parameters in Table 1 reflect the stylized facts that (i) there is a positive equity risk-premium ($\mu^{\mathbb{P}} > \mu^{\mathbb{Q}}$), (ii) implied at-the-money volatility is higher than typical historical volatility ($\theta^{\mathbb{Q}} > \theta^{\mathbb{P}}$), and (iii) there is a “fear of (downward) jumps” ($\alpha^{\mathbb{Q}} < \alpha^{\mathbb{P}} < 0$). Note how the real-world and the risk-adjusted parameters are intertwined in the determination of the static hedge portfolio: Paths of the state variables, S and V , are simulated with \mathbb{P} -parameters and the \mathbb{Q} -parameters enter when the hedge instruments are valued (remember, the hedging portfolio is liquidated if the barrier is crossed).

The characteristic function of the log-returns in the Bates model is known in closed form (see Bates (1996)), so prices of plain vanilla options can be computed efficiently with Fourier inversion techniques.⁵

We will use the notation $P(0) = e^{-rT}E^{\mathbb{Q}}[P]$ and $H(0) = e^{-rT}E^{\mathbb{Q}}[H]$, i.e. $P(0)$ and $H(0)$ are time-0 values of the barrier option and the hedge instruments.

In the experiments we report in the following we only use hedge instruments with expiry T . Doing this is justified by the result from the following small experiment⁶ (see Table 2): Using vanilla calls with expiries $T/4, \dots, T$ and the budget restriction $H(0) \cdot c \leq P(0)$, we calculate the risk-minimizing hedges corresponding to the three one-sided⁷ risk-measures from section 2.3 — $E[(\cdot)^+]$, $V@R_{\alpha}$ and ES_{α} , with $\alpha = 0.05$. For each risk-measure, we then repeatedly compute the risk-minimizing hedge using only the expiry- T calls, while gradually relaxing the budget restriction until the performance of the only-expiry- T -hedge equals the performance of the hedge that includes earlier expiries. The difference in price between these two hedges can be seen as the cost of

⁵The link between Fourier inversion and option pricing is investigated in Carr & Madan (1999); a text-book presentation is given in Cont & Tankov (2003, Chapter 11). We do not know in advance which (time to maturity, spot)-combinations we need to price (this is determined by when and how the barrier is crossed), so we use brute-force quadrature integration.

⁶Both Nalholm & Poulsen (2006) and Giese & Maruhn (2007) also find the expiry- T component to be the most important one; in the later comparison we also restrict their hedges in that way.

⁷We cannot do this for the quadratic risk-measure, since the performance of this hedge does not increase when we relax the budget restriction.

Statistical parameter	Value	Design parameter	Set to
$\kappa^{\mathbb{P}}$	4.788	S_0	100
$\kappa^{\mathbb{Q}}$	2.772	V_0	$\theta^{\mathbb{P}}$
$\theta^{\mathbb{P}}$	0.205^2	r	0.02
$\theta^{\mathbb{Q}}$	0.269^2	K	110
σ	0.512	B	130
ρ	-0.586	T	1
$\alpha^{\mathbb{P}}$	-0.004	$K^{(1)}, \dots, K^{(7)}$	110, 130, 131, \dots , 135
$\alpha^{\mathbb{Q}}$	-0.020	$T^{(1)}, \dots, T^{(7)}$	T
γ	0.066	N	10,000
λ	0.504		
$\mu^{\mathbb{P}}$	0.066		
$\mu^{\mathbb{Q}}$	$r - \lambda\alpha^{\mathbb{Q}}$		

Table 1: Default setting for simulation experiments. The two columns to the left give parameters for the Bates model; both real-world (\mathbb{P}) and risk-adjusted (\mathbb{Q}). These have been taken from Eraker (2004) and converted to non-percentage and annualized terms.

Risk-measure	$E[(\cdot)^+]$	$V@R_{0.05}$	$ES_{0.05}$
Cost of excluding early maturities	2%	1%	2%

Table 2: Cost of excluding prior-to- T expiries in the risk-minimizing hedge portfolios, relative to the time-0 price of the barrier option.

excluding earlier expiries in the hedge portfolio. The costs for all three risk-measures are reported in Table 2, and as we see they are small: less than 2% of the time-0 price of the barrier option.

We investigate the quality of the risk-minimizing hedges for the four risk-measures from Section 2.3 — the quadratic, the positive part, $V@R_\alpha$ and ES_α , with $\alpha = 0.05$. The risk-minimizing hedge strategies are compared to two recent suggestions in the literature.

1. Nalholm & Poulsen (2006) combine Carr’s “adjusted payoff” idea with pointwise matching, regularization by singular value decomposition and structural knowledge of the Bates model.
2. Giese & Maruhn (2007) use super-replication, which in our notation amounts to




Figure 1: Payoffs for the static hedge portfolios. Top: risk-minimizing hedges corresponding to $u(x) = x^2$ (dotted), $u(x) = x^+$ (solid) and $V@R_{0.05}$ (dashed). Bottom: the $ES_{0.05}$ -minimizing hedge (dotted) and the hedges from Nalholm & Poulsen (solid) and Giese & Maruhn (dashed).

solving

$$\min H(0) \cdot c \text{ s.t. } H^{(n)} \cdot c \geq P^{(n)} \text{ for all } n.$$

Note that for both these hedges, the hedger does not — contrary to the suggested risk-minimizing hedges — directly control the cost of the portfolio.

Having computed the strategy from Nalholm & Poulsen (2006) we see that the hedge weights are bounded by 9.9, so to make the comparison fair we use this as a constraint on the magnitude of the weights in the risk-minimizing hedges and the hedge from Giese & Maruhn (2007).

We compute the risk-minimizing hedges using 10,000 simulated paths which in our experience gives very reliable results. That takes about four minutes on a laptop. Almost all that time is spent on valuation of the hedge instruments; the path generation takes about ten seconds, while the actual optimization is almost instant.

Because the hedge portfolios only contain contracts of a single expiry, T , they can conveniently be represented by their payoff functions. This is done Figure 1. We see how the super-replicating portfolio of Giese & Maruhn (2007) tries to mimic the truncated call-payoff $(x - K)^+ \mathbf{1}_{x < B}$, which must be the cheapest truly super-replication portfolio

Hedge technique	Hedge error statistics							
	price	m	sd	$E[(\cdot)^+]$	$V@R_{0.05}$	$ES_{0.05}$	$\mathbb{P}(\text{loss} > 0)$	$\max(\text{loss})$
Risk-minimization								
$E[(\cdot)^2]$	99.6%	4.3%	66%	13%	66%	99%	0.34	830 %
$E[(\cdot)^+]$	100%	-1.8%	73%	7.5%	49%	90%	0.14	800 %
$V@R_{0.05}$	100%	0.4%	69%	9.0%	34%	76%	0.38	800 %
$ES_{0.05}$	100%	-0.3%	69%	9.5%	38%	72%	0.37	770 %
Payoff matching as Nalholm & Poulsen	119%	-14%	90%	4.3%	25%	76%	0.09	870%
Super-replication as Giese & Maruhn	164%	-70%	157%	0.00%	0.00%	0.01%	0.00	100%

Table 3: Hedge error behavior in the Bates model for an expiry-1, strike-110, barrier-130 up-and-out call-option, for which $P(0) = e^{-rT} E^{\mathbb{Q}}[P] = 1.33$. All numbers are given as percentages of $P(0)$.

in the case of unbounded jumps, see also Kraft (2007). But other than that, differences are hard to make out with the naked eye. Therefore Table 3 reports descriptive statistics for all the hedge strategies.⁸ In the table

- “price” is the initial price of the (estimated) optimal static hedge, $H(0) \cdot c$,
- “m” is the average loss, $E^{\mathbb{P}}[P - H \cdot c]$,
- “sd” is the standard deviation of the loss, $\sqrt{\text{Var}[P - H \cdot c]}/P(0)$,
- “ $\mathbb{P}(\text{loss} > 0)$ ” is the probability that the hedge loses money, $\mathbb{P}(\{P - H \cdot c > 0\})$,
- “max(loss)” is the largest observed loss, $\max\{P^{(n)} - H^{(n)} \cdot c\}$.

The first thing to note from Table 3 is that pay-off discontinuity makes the up-and-out call difficult to hedge; for the risk-minimizing strategies the standard deviation of the loss is around 70% of the time-0 value of the barrier option. For comparison: Daily delta-hedging of the up-and-out call gives a standard deviation of about 150%, while the standard deviation of the loss is around 5% of the option value for delta-hedging of a plain vanilla call. Second, there are visible — though not vast — differences between the performance of optimal strategies depending on the type of risk-measure. Or briefly put: We can tell what is being minimized. These differences are more pronounced for value-at-risk and expected shortfall, that put more focus on the tail of the loss distribution. The table also gives performance statistics for the hedges suggested by Nalholm & Poulsen

⁸The sample average method has an inherent in-sample bias (positive, performance-wise), so to evaluate hedges out-of-sample, we simulated 10,000 new paths and used the proposed optimal strategies.

and Giese & Maruhn. A direct and fair comparison based on the table is not possible because of the differences in initial prices.

However, as suggested in Krokmal, Palmquist & Uryasev (2002), that can be remedied by varying the budget restriction and each time calculating the corresponding risk-minimizing hedges. In this way we get an efficient frontier in cost/risk-space; in fact we get one for each risk-measure. These are shown in Figure 2. The four panels correspond to different risk measures, and each panel has four curves corresponding to the four (optimal) strategies. First, this figure illustrates what we mean when we say that $E[(\cdot)^2]$ is a financially strange risk-measure to use. It does not capture the cost/risk trade-off; you cannot reduce risk by increasing the budget. That is possible when a one-sided risk-measure is used. We see that optimal strategies w.r.t. the three one-sided measures behave fairly similarly, especially for high budgets. For a bank selling exotic options, this the realistic case; if the option cannot be sold at some mark-up, then it is not sold at all. The efficient frontiers enable us to compare the risk-minimizing strategies to the two previous suggestions, and we see the improvements. The Nalholm & Poulsen strategy clearly does not lie on the efficient frontier for any risk-measure; similar performance can be achieved at approximately 10—20% lower cost (the typical horizontal distance to frontier). For the one-sided risk-measures the Giese & Maruhn hedge is very close to the efficient frontier, but at its extreme right. We find it quite likely that a trader or a bank would like to have the possibility to take a little more risk in return for a significantly lowered cost.

4 Model risk

In this section we investigate how robust risk-minimizing hedges are with respect to “model risk.” We ask (i) what happens if the hedger ignores the difference between the measures \mathbb{P} and \mathbb{Q} , and (ii) how does the Bates-optimal hedge perform in an infinite activity Levy-model, and vice versa?

4.1 \mathbb{P} vs. \mathbb{Q}

The computation of the risk-minimizing hedge depends on both real-world and risk-adjusted parameters. But what happens if we ignore this difference? In our experiment we consider two not-too-hypothetical hedgers: The Quant and The Econometrician.

Risk-measure	The Quant		The Econometrician		true risk
	conjectured risk	actual risk	conjectured risk	actual risk	
$\sqrt{E[(\cdot)^2]}$	69%	66%	118%	110%	66%
$E[(\cdot)^+]$	8.2%	8.0%	5.9%	8.2%	7.6%
$V@R_{0.05}$	35%	36%	34%	48%	34%
$ES_{0.05}$	80%	72%	68%	73%	72%

Table 4: Conjectured and actual performances of risk-minimizing hedge strategies for the two slightly ignorant hedgers, relative to the time-0 price $P(0) = 1.33$ of the barrier option.

- **The Quant** has obtained a correct view of the \mathbb{Q} -parameters from Table 1 by calibrating the Bates model to market prices of vanilla options, and takes as $V(0)$ the squared implied volatility of a one-month at-the-money call-option. Then he takes $\mathbb{P} = \mathbb{Q}$.
- **The Econometrician** has estimated the \mathbb{P} -parameters from Table 1 correctly, say from a long time series of frequent stock price observations. He then takes all \mathbb{Q} -parameters equal to the \mathbb{P} -parameters, but changes the drift to $\mu^{\mathbb{Q}} = r - \lambda\alpha^{\mathbb{P}}$ in order to make \mathbb{Q} a martingale measure.

Both hedgers come in four different versions corresponding to the different risk-measures. All hedgers run simulations and determine their risk-minimizing hedges with the budget restriction $H(0) \cdot c \leq P(0)$. They run additional simulations (again using their own views of the two measures) to get out-of-sample estimates of the performance of their optimal hedges; this we call the “conjectured risk”. Then we estimate the “actual risk” of the strategies by running simulations as in Section 3 with the full \mathbb{P} - \mathbb{Q} parameter combination. The results are reported in Table 4; the “true risk” in the last column is the performance of the benchmark portfolios from Table 3 evaluated on this table’s out-of-sample paths.

We see that irrespective of the choice of risk measure the Quant’s portfolio suffers very little compared to the truly optimal one; at most risk increases by a factor of 1.05. The Econometrician’s performance deteriorates somewhat more. That is understandable; when options are used as hedge instruments, it is important to take their pricing into account — this is done when the \mathbb{Q} -parameters are calibrated to the market. What is

more surprising is that the difference between actual and true risk for the Econometrician varies a lot across risk measures. Last, we note that there is no general pattern in the relation between “conjectured risk” and “true risk”; neither hedger is systematically optimistic or pessimistic. This can be further confirmed by running similar experiments on down-and-out puts.

4.2 Bates vs. NIG

As pointed out by several authors, see for instance Detlefsen & Härdle (2007) and their references, models may produce similar prices of plain vanilla options and yet give very different prices for e.g. barrier options. We consider the Normal Inverse Gaussian (NIG) model which is a pure jump, infinite activity Levy process,⁹ and test the robustness of the risk-minimizing hedge by investigating how well the Bates-optimal hedge performs in the NIG model and vice versa.

In the NIG model, the stock price is given by $S_t = S_0 e^{X_t}$ where $(X_t)_{t \geq 0}$ is a so-called NIG process. This means that the Levy-Khinchine representation of the characteristic function of X_t is $E[e^{isX_t}] = e^{t\psi(is)}$, where

$$\psi(z) = mz + d(\sqrt{a^2 - b^2} - \sqrt{a^2 - (b+z)^2}).$$

With $m = r - d(\sqrt{a^2 - b^2} - \sqrt{a^2 - (b+1)^2})$ the discounted stock price is a martingale, so the model has three (risk-adjusted) parameters. It is possible to simulate X_t efficiently and without bias on a grid (see Glasserman (2004, Section 3.5)) and since the characteristic function is available in closed form, vanilla option prices can be computed with Fourier inversion techniques.

We calibrate the NIG model to the Bates model’s prices for the hedge instruments across expiries.¹⁰ The resulting parameters are $(a, b, d) = (15.0, -8.9, 0.5)$ and from Figure 3 we see that there is a quite good fit over the 3-month to 1-year spectrum of expiries. Note that the NIG model achieves this fit with only three parameters; the Bates model has seven. To avoid the difference between the two models under the real-world measure we set $\mathbb{P} = \mathbb{Q}$ in what follows.

As before, we want to hedge an up-and-out call option with expiry $T = 1$, strike $K = 110$ and barrier $B = 130$. The hedging instruments are expiry- T vanilla calls with

⁹Financial modelling with NIG processes was first suggested in Barndorff-Nielsen (1998). A recent survey/investigation of infinite activity Levy models is Carr, Geman, Madan & Yor (2007)

¹⁰Only expiry- T options are used for the initial hedge portfolio composition, but a poor fit of the NIG model at shorter expiries will cause problems at liquidation.

Data-generating model	Risk-measure							
	$\sqrt{E[(\cdot)^2]}$		$E[(\cdot)^+]$		V@R _{0.05}		ES _{0.05}	
	Hedge model		Hedge model		Hedge model		Hedge model	
	Bates	NIG	Bates	NIG	Bates	NIG	Bates	NIG
Bates	54%	59%	7.3%	9.4%	33%	38%	67%	86%
NIG	72%	64%	6.4%	6.2%	21%	13%	120%	100%

Table 5: Risk-measures for Bates- and NIG-optimal hedge errors relative to the time-0 price of the maturity-1 , strike-110, barrier-130 up-and-out call-option.

strikes 110, 130, 131, ..., 135. For each risk-measure we compute the risk-minimizing hedges in both the Bates and the NIG model with the budget restriction $H(0) \cdot c \leq P_{Bates}(0) = 1.33$. We evaluate the quality of all the different hedges out-of-sample in the Bates model, and report the corresponding risk-measures relative to $P_{Bates}(0)$ in the first row of Table 5. Next, we compute the Bates- and NIG-optimal hedges for all risk-measures, but this time with the budget restriction $H(0) \cdot c \leq P_{NIG}(0) = 1.43$ (the true barrier value in the NIG model). We evaluate the quality of the hedges in the NIG-model and report the risk-measures relative to $P_{NIG}(0)$ in the second row of Table 5.

The first thing we see from Table 5 is that the numbers in the two rows are of the same order of magnitude, which means static hedging is also feasible in the NIG model. Reassuringly, there is a detectable — but again, not tremendous — benefit from using the right model to calculate the risk minimizing strategies. Using the wrong model’s risk-minimizing strategy increases risk by a factor of 1.05 to 1.7, with the tail-focusing measures value-at-risk and expected shortfall displaying greater sensitivity. If ones does not know the true data-generating process, Cont (2006) suggests to the hedge model that minimizes the worst-case error. But in that respect the table gives no clear advice; different models give the highest worst-case error for different risk-measures.

5 Conclusion

We have described a new risk-minimizing framework for static hedging of exotic contingent claims. A main strength of the method is that it is easy to explain, yet quite general: We can build hedges for any contingent claim in any model, as long as we can

simulate trajectories of the underlying and payoffs of the hedge instruments. Another feature of the risk-minimizing hedge is the possibility for the hedger to choose his risk measure; we studied various popular choices.

The performance of the method was evaluated for a discretely observed up-and-out call option on an underlying following the Bates model (stochastic volatility and jumps), and we found that it compared favorably to methods previously suggested in the literature.

Moreover, the method also works in the NIG model: this is a contribution in itself since previous methods could not handle infinite activity Levy models. From simulation experiments we also concluded that the risk-minimizing hedges are reasonably robust to model risk: Bates-optimal hedges work in the NIG model, and vice versa, if the two models price vanilla options similarly.

The analysis, especially the model risk cases in Section 4, also illustrates that there is no “universally optimal” risk-minimizing hedge; optimality and behaviour depends what risk-measure you use. And on how much you are willing to pay to reduce risk. This makes it all the more important to use a method that can handle the cost/risk-tradeoff.

Besides the “obvious” suggestion for future research: “test empirically”, let us mention three intriguing aspects that we are currently working on. First: Proportional frictions (transaction costs and bid/ask-spreads) are straightforward to handle optimization-wise; treat buying and selling separately and linearity and convexity is retained. How does this affect the construction and performance of static hedges; does it “level the playing field” compared to dynamic hedging with the underlying? Second: How large is the effect of non-linearity of the hedges? Could it be that hedging contracts one by one with some $u(x) \neq x^2$ (and thus sub-optimally) is preferable to hedging (optimally) with $u(x) = x^2$? Third: Can we — as suggested by Cont (2006) — set up portfolios that are designed to be robust to model and parameter uncertainty? Or more specifically, can we minimize the worst-case error across possible, non-nested models and plausible ranges of parameter-values?

References

Allen, S. & Padovani, O. (2002), ‘Risk Management Using Quasi-static Hedging’, *Economic Notes* **31**, 277–336.

- Artzner, P., Delbaen, F., Eber, J.-M. & Heath, D. (1999), ‘Coherent Measures of Risk’, *Mathematical Finance* **9**, 203–228.
- Barndorff-Nielsen, O. E. (1998), ‘Processes of normal inverse Gaussian type’, *Finance and Stochastics* **2**, 41–68.
- Bates, D. (1996), ‘Jumps and Stochastic Volatility: the Exchange Rate Process Implicit in Deutchemark Options’, *Review of Financial Studies* **9**, 69–107.
- Carr, P., Ellis, K. & Gupta, V. (1998), ‘Static Hedging of Exotic Options’, *Journal of Finance* **53**, 1165–1190.
- Carr, P., Geman, H., Madan, D. & Yor, M. (2007), ‘Self-decomposability and option pricing’, *Mathematical Finance* **17**, 31–57.
- Carr, P. & Madan, D. (1999), ‘Option Pricing and the Fast Fourier Transform’, *Journal of Computational Finance* **2**, 61–73.
- Cherny, A. S. (2006), ‘Weighted V@R and its properties’, *Finance and Stochastics* **10**, 367–393.
- Cont, R. (2006), ‘Model uncertainty and its impact on the pricing of derivative instruments’, *Mathematical Finance* **16**, 519–542.
- Cont, R., Deguest, R. & Scandolo, G. (2007), Robustness and sensitivity analysis of risk measurement procedures. Working paper. Centre for Financial Engineering. Columbia University.
- Cont, R. & Tankov, P. (2003), *Financial Modelling with Jump Processes*, Chapman & Hall.
- Derman, E., Ergener, D. & Kani, I. (1995), ‘Static Options Replication’, *Journal of Derivatives* **2**, 78–95.
- Detlefsen, K. & Härdle, W. K. (2007), ‘Calibration Risk for Exotic Options’, *Journal of Derivatives* **14**, 47–63.
- Dupont, D. (2001), Hedging Barrier Options: Current Methods and Alternatives. Economics Series from Institute for Advanced Studies, no. 103. Available at <http://econpapers.repec.org/paper/ihsihesp/103.htm>.

- Eraker, B. (2004), ‘Do stock prices and volatility jump? Reconciling evidence from spot and option prices’, *Journal of Finance* **59**, 1367–1403.
- Föllmer, H. & Schied, A. (2002), ‘Convex Measures of Risk and Trading Constraints’, *Finance and Stochastics* **6**, 429–447.
- Giese, A. & Maruhn, J. (2007), ‘Cost-optimal static super-replication of barrier options: an optimization approach’, *Journal of Computational Finance* **10**, 71–97.
- Glasserman, P. (2004), *Monte Carlo Methods in Financial Engineering*, Springer-Verlag.
- Kraft, H. (2007), ‘Pitfalls in static superhedging of barrier options’, *Finance Research Letters* **4**, 2–9.
- Krokhmal, P., Palmquist, J. & Uryasev, S. (2002), ‘Portfolio Optimization with Conditional Value-at-Risk Objective and Constraints’, *Journal of Risk* **4**, 43–68.
- Nalholm, M. & Poulsen, R. (2006), ‘Static hedging of barrier options under general asset dynamics: Unification and application’, *Journal of Derivatives* **13**, 46–60.
- Pellizzari, P. (2005), ‘Static Hedging of Multivariate Derivatives by Simulation’, *European Journal of Operational Research* **166**, 507–519.
- Poulsen, R. (2006), ‘Barrier Options and Their Static Hedges: Simple Derivations and Extensions’, *Quantitative Finance* **6**, 327–335.
- Rockafeller, R. T. & Uryasev, S. (2000), ‘Optimization of conditional value-at-risk’, *Journal of Risk* **2**, 21–41.
- Schweizer, M. (2001), A guided tour through quadratic hedging approaches, in E. Jouini, J. Cvitanic & M. Musiela, eds, ‘Option Pricing, Interest Rates and Risk Management’, Cambridge University Press, pp. 538–574.
- Shapiro, A. (2008), ‘Stochastic programming by Monte Carlo simulation methods’, *Mathematical Programming, Series B* **112**, 183–220.
- Tasche, D. (2002), ‘Expected Shortfall and Beyond’, *Journal of Banking and Finance* **26**, 1519–1533.

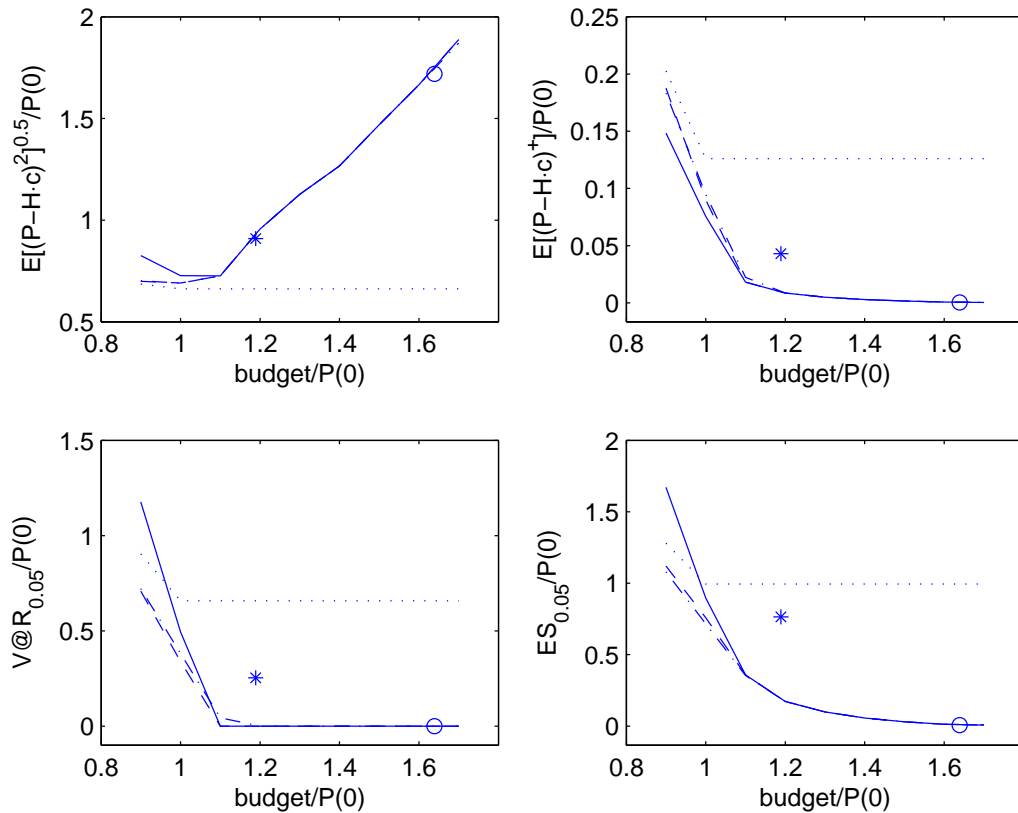


Figure 2: The efficient frontiers for the risk-minimizing hedges in different cost/risk-spaces. The risk-measures corresponding to the curves are: the quadratic (dotted lines), the positive part (solid lines), $V@R_{0.05}$ (dashed lines) and $ES_{0.05}$ (dash-dotted lines). The $*$'s show the performance of the Nalholm & Poulsen portfolio and the \circ 's show the Giese & Maruhn portfolio.

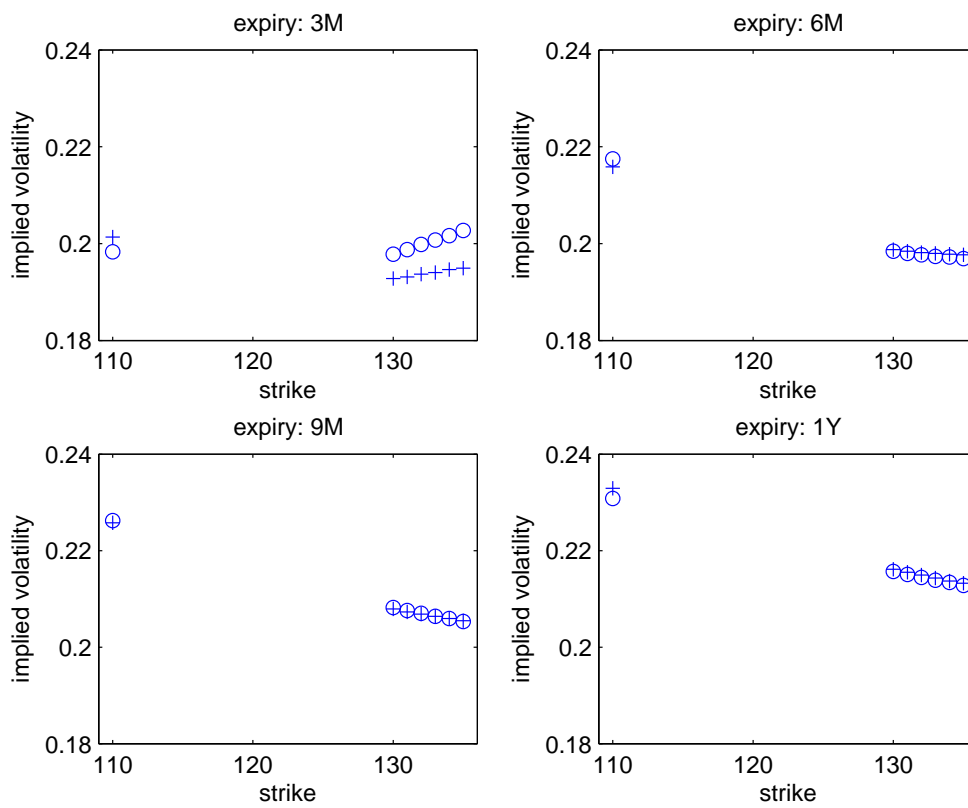


Figure 3: Black-Scholes implied volatilities for call-options across expiries and strikes for the Bates (\circ) and NIG ($+$) models. The points in the graphs correspond to the options used in the risk-minimizing static hedges.