

Optimal Hedging of Options with Transaction Costs

Valeri. I. Zakamouline

Bodø Graduate School of Business
8049 Bodø, Norway
Tel.: (+47) 75517923; Fax: (+47) 75517268
zakamouliny@yahoo.no

Abstract

One of the most successful approaches to option hedging with transaction costs is the utility based approach, pioneered by Hodges and Neuberger (1989). Judging against the best possible tradeoff between the risk and the costs of a hedging strategy, this approach seems to achieve excellent empirical performance. However, this approach has one major drawback that prevents the broad application of this approach in practice: the lack of a closed-form solution. Since the knowledge of the optimal hedging strategy is of great practical significance, in this paper we suggest and implement two methods for finding the optimal hedging strategy with transaction costs. The first method is the approximation of the utility based hedging strategy which yields a closed-form solution. The second method is the optimization of the parameters of the hedging strategy in some risk-return space. We provide an empirical testing of our hedging strategies against the asymptotic and some other well-known strategies and find that our strategies outperforms all the others.

Key words: option hedging, transaction costs, approximation, optimization, simulations.

JEL classification: G11, G13.

EFM classification: 410, 450

1 Introduction

One of the most successful approaches to option hedging with transaction costs is the utility based approach, pioneered by Hodges and Neuberger (1989). Judging against the best possible tradeoff between the risk and the costs of a hedging strategy, the utility based approach seems to achieve excellent empirical performance (see Mohamed (1994), Clewlow and Hodges (1997), Martellini and Priaulet (2002), and Zakamouline (2005)). However, this approach has one major drawback that prevents the broad application of this approach in practice: the lack of a closed-form solution. Therefore, the solution must be computed numerically. The numerical algorithm is cumbersome to implement and the calculation of the optimal hedging strategy is time consuming.

According to the utility based approach, the qualitative description of the optimal hedging strategy is as follows: do nothing when the hedge ratio lies within a so-called “no transaction region” and re hedge to the nearest boundary of the no transaction region as soon as the hedge ratio moves out of the no transaction region.

Since there are no explicit solutions for the utility based hedging with transaction costs and the numerical methods are computationally hard, for practical applications it is of major importance to use other alternatives. One of such alternatives is to obtain an asymptotic solution. In asymptotic analysis one studies the solution to a problem when some parameters in the problem assume large or small values.

Whalley and Wilmott (1997) were the first to provide an asymptotic analysis of the model of Hodges and Neuberger (1989) assuming that transaction costs are small. Barles and Soner (1998) performed an alternative asymptotic analysis of the same model assuming that both the transaction costs and the hedger’s risk tolerance are small. However, the results of Barles and Soner are quite different from those of Whalley and Wilmott. While Whalley and Wilmott derive only an optimal form of the hedging bandwidth which is centered around the Black-Scholes delta, Barles and Soner show that the optimal hedging strategy has two key elements: a particular form of the hedging bandwidth and a volatility adjustment. The latter means that the middle of the hedging bandwidth does not coincide with the Black-Scholes delta. Since practitioners often use asymptotic solutions in option hedging, the two different results of asymptotic analysis are, at least, confusing.

In this paper we present a detailed study of the exact (numerically calculated) utility based hedging strategy and its comparison with the two asymptotic strategies. We document the fact that the Barles and Soner result is closer to reality than that of Whalley and Wilmott. The comparison of the performance of an asymptotic strategy against the exact strategy reveals another fact that under realistic model parameters an asymptotic

strategy performs noticeably worse than that obtained from the exact numerical solution. The explanation lies in the fact that, when some of the model parameters are neither very small nor very large, an asymptotic solution provides not quite accurate results. In particular, as compared to the exact numerical solution, under realistic parameters the size of the hedging bandwidth and the volatility adjustment obtained from asymptotic analysis are overvalued. What is more important, an asymptotic solution showed to be unable to sustain a correct interrelationship between the size of the hedging bandwidth and the volatility adjustment. The significance of the correct interrelationship could hardly be overemphasized: Our empirical testing of the hedging strategies reveals that either undervaluation or overvaluation of the volatility adjustment (with respect to the size of the hedging bandwidth) results in a drastic deterioration of the performance of a hedging strategy.

Since the knowledge of the optimal hedging strategy is of great practical significance, in this paper we suggest and implement two methods for finding the optimal hedging strategy with transaction costs. The first method is an approximation of the utility based hedging strategy. Under approximation we mean the following: We have a slow and cumbersome way to compute the optimal hedging policy with transaction costs and want to replace it with simple and efficient approximating function(s). To do this, we first specify a flexible functional form of the optimal hedging policy. Then, given a functional form, the parameters are chosen in order to provide the best fit to the exact numerical solution. This second stage is known as “model calibration”.

We implement the approximation of the optimal hedging strategy with transaction costs for a short European call option. Despite the simplicity of our approximating functions and straightforwardness of finding the best fit parameters, our approximation strategy showed to be very effective with the performance close to the exact strategy. We provide an empirical testing of our approximation strategy against the asymptotic and some other well-known strategies within unified mean-variance and mean-VaR frameworks and show that our strategy outperforms all the others.

Despite the impressive results, the approximation method does retain the major drawback we would like to get rid off: To implement the approximation for another type of a European option one needs to start with the numerical calculations of the optimal hedging strategy for a large set of the model parameters. Again, since the numerical algorithm is cumbersome to implement, the approximation methodology is unlikely to be commonly used by the practitioners. To overcome the drawback of the approximation method we suggest an optimization method in a risk-return space. This method have a great advantage over the approximation method in that one does not need to implement the numerical calculations of the optimal hedging strategy. Instead, the optimization method is based on simulations, which are much easier to implement. However, this method heavily relies

on the qualitative knowledge of the form of the optimal strategy which is obtained using the asymptotic and approximation methods.

The reader should be reminded that our paper concerns only with the optimal hedging a single (plain vanilla) European option. To be more specific, our methods are currently limited to hedging a European option whose gamma does not change sign. This limitation is due to the fact that the volatility adjustment varies according to whether the option gamma is positive or negative. The optimal hedging strategy is presented by the example of hedging a short European call option. We conjecture that our risk-return optimization method could be extended for a portfolio of options using the result of Hoggard, Whalley, and Wilmott (1994) which extends the method of Leland (1985).

2 The Option Hedging Problem and First Solutions

We consider a continuous time economy with one risk-free and one risky asset, which pays no dividends. We will refer to the risky asset as the stock, and assume that the price of the stock, S_t , evolves according to a diffusion process given by

$$dS_t = \mu S_t dt + \sigma S_t dW_t,$$

where μ and σ are, respectively, the mean and volatility of the stock return per unit of time, and W_t is a standard Brownian motion. The risk-free asset, commonly referred to as the bond or bank account, pays a constant interest rate of $r \geq 0$. We assume that a purchase or sale of δ shares of the stock incurs transaction costs $\lambda|\delta|S$ proportional to the transaction ($\lambda \geq 0$).

We consider hedging a short European call option with maturity T and strike price K . We denote the value of the option at time t as $V(t, S_t)$. The terminal payoff of the option one wishes to hedge is given by

$$V(T, S_T) = \max\{S_T - K, 0\} = (S_T - K)^+.$$

When a hedger writes a call option, he receives the value of the option $V(t, S_t)$ and sets up a hedging portfolio by buying Δ shares of the stock and putting $V(t, S_t) - \Delta(1 + \lambda)S_t$ in the bank account. As time goes, the writer rebalances the hedging portfolio according to some prescribed rule/strategy.

When a market is friction-free ($\lambda = 0$), Black and Scholes (1973) showed that it is possible to replicate the payoff of an option by constructing a self-financing dynamic trading strategy consisting of the risk-free asset and the stock. As a consequence, the absence of arbitrage dictates that the option price be equal to the cost of setting up the replicating portfolio. According to the Black-Scholes model, the price of a European call option is given by

$$V(t, S_t) = S_t N(d_1) - Ke^{-r(T-t)} N(d_2), \quad (1)$$

where

$$d_1 = \frac{\log\left(\frac{S}{K}\right) + \left(r + \frac{1}{2}\sigma^2\right)(T - t)}{\sigma\sqrt{T - t}}, \quad (2)$$

$$d_2 = d_1 - \sigma\sqrt{T - t},$$

and $N(\cdot)$ is the cumulative probability distribution function of a normal variable with mean 0 and variance 1. The Black-Scholes hedging strategy consists in holding Δ (delta) shares of the stock and some amount in the bank account, where

$$\Delta = \frac{\partial V}{\partial S} = N(d_1). \quad (3)$$

It should be emphasized that the Black-Scholes hedging is a dynamic replication policy where the trading in the underlying stock has to be done continuously. In the presence of transaction costs in capital markets the absence of arbitrage argument is no longer valid, since perfect hedging is impossible. Due to the infinite variation of the geometric Brownian motion, the continuous replication policy mandated by the Black-Scholes model incurs an infinite amount of transaction costs over any trading interval no matter how small it might be. How should one hedge an option in the market with transaction costs?

One of the simplest and most straightforward hedging strategies in the presence of transaction costs is to re hedge in the underlying stock at fixed regular intervals. One would simply implement the delta hedging according to the Black-Scholes strategy, but in discrete time. More formally, the time interval $[t, T]$ is subdivided into n fixed regular intervals δt , such that $\delta t = \frac{T-t}{n}$. The hedging proceeds as follows: at time t the writer of an option receives V and constructs a hedging portfolio by purchasing $\Delta(t) = N(d_1(t, S_t))$ shares of the stock and putting $V - \Delta(t)(1 + \lambda)S_t$ into the bank account. At time $t + \delta t$, an additional number of shares of the stock is bought or sold in order to have the target hedge ratio $\Delta(t + \delta t) = N(d_1(t + \delta t, S_{t+\delta t}))$. At the same time, the bank account is adjusted by

$$\left[\Delta(t + \delta t) - \Delta(t) - |\Delta(t + \delta t) - \Delta(t)|\lambda\right]S_{t+\delta t}.$$

Then the hedging is repeated in the same manner at all subsequent times $t + i\delta t$, $i = 2, 3, \dots$

The choice of the number of hedging intervals, n , is somewhat unclear. Obviously, when n is small, the volume of transaction costs is also small, but the variance of the replication error is large. An increase in n reduces the variance of the replication error at the expense of increasing the volume of transaction costs. Moreover, as $\delta t \rightarrow 0$, the volume of transaction costs approach infinity.

A variety of methods have been suggested to deal with the problem of option pricing and hedging with transaction costs. Leland (1985) was the first

to initiate this stream of research. He adopted the reheding at fixed regular intervals and proposed a modified Black-Scholes strategy that permits the replication of a short option position with finite volume of transaction costs no matter how small the reheding interval is. The hedging strategy is adjusted by using a modified volatility in the Black-Scholes formula as given below (for a short option position)

$$\sigma_m^2 = \sigma^2 \left(1 + \sqrt{\left(\frac{8}{\pi\delta t}\right) \frac{\lambda}{\sigma}} \right). \quad (4)$$

Using the Leland's method one hedges an option with the delta of the modified price calculated in accordance with pricing formula (1), but with adjusted volatility. In other words, the Leland's hedge ratio is given by

$$\Delta = \frac{\partial V(\sigma_m)}{\partial S} = N(d_1(\sigma_m)). \quad (5)$$

It could be easily checked that as $\delta t \rightarrow 0$, the modified volatility becomes unbounded and the Leland's hedge amounts to the super-replicating strategy where a single share of the stock is held to hedge the option.

The intuition behind the Leland's volatility adjustment is not so easy to interpret. For a short option position the Leland's volatility adjustment makes the option delta a flatter function of the underlying asset. Thus, this reduces the sensitivity of the option delta to the underlying asset price. Recall that the sensitivity of the option delta to the underlying asset price is known as option gamma

$$\Gamma = \frac{\partial \Delta}{\partial S} = \frac{\partial^2 V}{\partial S^2} = \frac{N'(d_1)}{S\sigma\sqrt{T-t}}. \quad (6)$$

It seems to be quite intuitive that the Leland's volatility adjustment decreases the option gamma in order to decrease the amount of hedging transaction costs (since we expect to transact more in regions where the option gamma is high). However, for a long option position the Leland's volatility adjustment is counter-intuitive, since it makes the option delta a steeper function of the underlying asset. This follows from the fact that for a long option position the adjusted volatility is given by

$$\sigma_m^2 = \sigma^2 \left(1 - \sqrt{\left(\frac{8}{\pi\delta t}\right) \frac{\lambda}{\sigma}} \right). \quad (7)$$

Note that when $\delta t \rightarrow 0$ the Leland's approach for hedging a long option position is not well-posed as the adjusted volatility becomes negative.

3 The Utility Based Hedging

Since the initiative work of Leland, there has been proposed a variety of alternative approaches to option pricing and hedging with transaction costs. A great deal of them are concerned with the “financial engineering” problem of either replicating or super-replicating the option payoff. These approaches are mainly preference-free models where rehedging occurs at some discrete time intervals whether or not it is optimal in any sense. However, common sense tells us that an “optimal” hedging policy should achieve the best possible tradeoff between the risk and the costs of replication. Recognizing the fact that risk preferences differ among the individuals, the following conclusion becomes obvious: in pricing and hedging options one must consider the hedger’s attitude toward risk.

In modern finance it is customary to describe risk preferences by a utility function. The expected utility theory maintains that individuals behave as if they were maximizing the expectation of some utility function of the possible outcomes. Hodges and Neuberger (1989) pioneered the option pricing and hedging approach based on this theory. The key idea behind the utility based approach is the indifference argument: The so-called “reservation” price of an option is defined as the amount of money that makes the hedger indifferent, in terms of expected utility, between trading in the market with and without writing the option. In many respects such an option price is determined in a similar manner to a certainty equivalent within the expected utility framework, which is a well grounded pricing principle in economics. The difference in the two trading strategies, with and without the option, is interpreted as “hedging” the option.

The starting point for the utility based option pricing and hedging approach is to consider the optimal portfolio selection problem of the hedger who faces transaction costs and maximizes expected utility of his terminal wealth. The hedger has a finite horizon $[t, T]$ and it is assumed¹ that there are no transaction costs at terminal time T . The hedger has the amount x_t in the bank account, and y_t shares of the stock at time t . We define the value function of the hedger with no option liability as

$$J_0(t, x_t, y_t, S_t) = \max E_t[U(x_T + y_T S_T)], \quad (8)$$

where $U(z)$ is the hedger’s utility function. Similarly, the value function of the hedger with option liability is defined by

$$J_w(t, x_t, y_t, S_t) = \max E_t[U(x_T + y_T S_T - (S_T - K)^+)]. \quad (9)$$

Finally, the option price is defined as the compensation p such that

$$J_w(t, x_t + p, y_t, S_t) = J_0(t, x_t, y_t, S_t). \quad (10)$$

¹This assumption is made for simplicity. In practice, there are two types of option settlements: either asset or cash. The type of option settlement affects, to some extent, the option price and hedging strategy.

The solutions to problems (8), (9), and (10) provide the unique option price and, above all, the optimal hedging strategy. Unfortunately, there are no closed-form solutions to all these problems. As a result, the solutions have to be obtained by numerical methods. The existence and uniqueness of the solutions were rigorously proved by Davis, Panas, and Zariphopoulou (1993). For implementations of numerical algorithms, the interested reader can consult Davis and Panas (1994) and Clewlow and Hodges (1997).

It is usually assumed that the hedger has the negative exponential utility function

$$U(z) = -\exp(-\gamma z); \quad \gamma > 0,$$

where γ is a measure of the hedger's (absolute) risk aversion. This choice of the utility function satisfies two very desirable properties: (i) the hedger's strategy does not depend on his holdings in the bank account, (ii) the computational effort needed to solve the problem is low. This particular choice of utility function might seem restrictive. However, as it was conjectured by Davis et al. (1993) and showed in Andersen and Damgaard (1999), an option price is approximately invariant to the specific form of the hedger's utility function, and mainly only the level of absolute risk aversion plays an important role.

The numerical calculations show that when the hedger's risk aversion is rather low, the hedger implements mainly a so-called "static" hedge, which consists in buying Δ shares of the stock at time t and holding them until the option maturity T . When the hedger's risk aversion increases, he starts to rebalance the hedging portfolio in between (t, T) . Recall that in the framework of the utility based hedging approach the option hedging strategy is defined as the difference, $\Delta(\tau) = y_w(\tau) - y_0(\tau)$, between the hedger's optimal trading strategies with and without option liability. When the hedger's risk aversion is moderate, it is impossible to give a concise description of the optimal hedging strategy. However, when the hedger's risk aversion is rather high (or, equivalently, $\mu = r$), then we can assume that $y_0(\tau) \equiv 0$ and the optimal hedging strategy can be conveniently described as $\Delta(\tau) = y_w(\tau)$. It is important now to note that in the rest of the paper we assume that either the hedger's risk aversion is rather high or $\mu = r$.

Now we turn to the analysis of the nature of the optimal hedging policy. The numerical calculations show (see Figure 1) that the optimal hedge ratio Δ is constrained to evolve between two boundaries, Δ_l and Δ_u , such that $\Delta_l < \Delta_u$. As long as the hedge lies within these two boundaries, $\Delta_l \leq \Delta \leq \Delta_u$, no rebalancing of the hedging portfolio takes place. That is why the region between the two boundaries is commonly called the no transaction region. As soon as the hedge ratio goes out of the no transaction region, a rebalancing occurs in order to bring the hedge to the nearest boundary of the no transaction region. In other words, if Δ moves below Δ_l , one should immediately transact to bring it back to Δ_l . Similarly, if Δ moves above

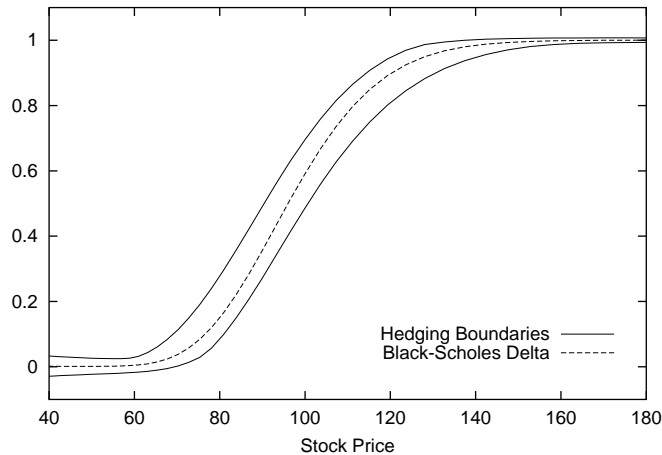


Figure 1: Optimal hedging strategy versus the Black-Scholes delta for the following model parameters: $\gamma = 0.25$, $\lambda = 0.01$, $S_t = K = 100$, $\sigma = 0.25$, $\mu = r = 0.05$, and $T - t = 0.5$.

Δ_u , a rebalancing trade occurs to bring it back to Δ_u .

Despite a sound economical appeal of the utility based option hedging approach, it does have a number of disadvantages: the model is cumbersome to implement and the numerical computations are time consuming. One commonly used simplification of the optimal hedging strategy is known as the *delta tolerance* strategy. This strategy prescribes to re hedge when the hedge ratio moves outside of the prescribed tolerance from the corresponding Black-Scholes delta. More formally, the boundaries of the no transaction region are defined by

$$\Delta = \frac{\partial V}{\partial S} \pm H, \quad (11)$$

where $\frac{\partial V}{\partial S}$ is the Black-Scholes hedge, and H is a given constant tolerance. The intuition behind this strategy is pretty obvious: The parameter H is a proxy for the measure of risk of the hedging portfolio. More risk averse hedgers would choose a low H , while more risk tolerant hedgers will accept a higher value for H .

Since there are no explicit solutions for the utility based hedging model with transaction costs and the numerical methods are computationally hard, for practical applications it is of major importance to use other alternatives. One of such alternatives is to obtain an asymptotic solution. In asymptotic analysis one studies the solution to a problem when some parameters in the problem assume large or small values.

Whalley and Wilmott (1997) were the first to provide an asymptotic analysis of the model of Hodges and Neuberger (1989) assuming that trans-

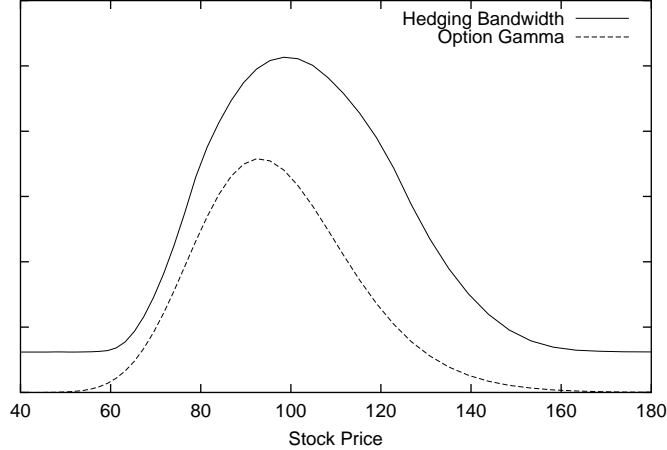


Figure 2: The form of the optimal hedging bandwidth $(\Delta_u - \Delta_l)S$, obtained using the exact numerics, versus the option gamma for $\gamma = 0.1$ and the rest of the model parameters as in Figure 1.

action costs are small. They showed that the boundaries of the no transaction region are given by

$$\Delta = \frac{\partial V}{\partial S} \pm H_{ww} = \frac{\partial V}{\partial S} \pm \left(\frac{3 e^{-r(T-t)} \lambda S \Gamma^2}{2 \gamma} \right)^{\frac{1}{3}}, \quad (12)$$

where, again, $\frac{\partial V}{\partial S}$ is the Black-Scholes hedge.

It is easy to see that the optimal hedging bandwidth is not constant, but depends in a natural way on a number of parameters. As $\lambda \rightarrow 0$, the optimal hedge approaches the Black-Scholes hedge. As γ increases, the hedging bandwidth decreases in order to decrease the risk of the hedging portfolio. The dependence of the hedging bandwidth on the option gamma is also natural, as we expect to re hedge more often in regions with high gamma. Moreover, it agrees quite well with the results of exact numerical computations, see Figure 2.

Barles and Soner (1998) performed an alternative asymptotic analysis of the same model assuming that both the transaction costs and the hedger's risk tolerance are small. They found that the optimal hedging strategy is to keep the hedge ratio inside the no transaction region given by

$$\Delta = \frac{\partial V(\sigma_m)}{\partial S} \pm H_{bs} = \frac{\partial V(\sigma_m)}{\partial S} \pm \frac{1}{\lambda \gamma S} g(\lambda^2 \gamma S^2 \Gamma), \quad (13)$$

where $\frac{\partial V(\sigma_m)}{\partial S}$ is the Black-Scholes hedge with an adjusted volatility given by

$$\sigma_m^2 = \sigma^2(1 + \Sigma_{bs}) = \sigma^2 \left(1 + f \left(e^{r(T-t)} \lambda^2 \gamma S^2 \Gamma \right) \right). \quad (14)$$

The function $f(z)$ is the unique solution of the nonlinear initial value problem

$$\frac{df(z)}{dz} = \frac{f(z) + 1}{2\sqrt{zf(z)} - z}, \quad z \neq 0, \quad f(0) = 0,$$

and the function $g(z)$ is defined by

$$g(z) = \sqrt{zf(z)} - z.$$

For $z > 0$ the function $f(z)$ is a concave increasing function.

The Barles and Soner volatility adjustment works similarly to the Leland's volatility adjustment: for a short option position the adjusted volatility is greater than the original volatility. Note that the Barles and Soner volatility adjustment depends on the option gamma. This agrees quite well with the results of exact numerical computations: Note first that, as it is clearly seen from Figure 1, the middle of the optimal hedging bandwidth does not coincide² with the Black-Scholes hedge. The picture becomes even more obvious if we increase the hedger's risk aversion. We then convert the middle of the hedging bandwidth to a corresponding volatility adjustment³ that is depicted in Figure 3.

Note in addition that the asymptotic analysis of Barles and Soner correctly predicts that volatility adjustment increases when either the level of transaction costs or the hedger's risk aversion increases. The comparative statics for the Barles and Soner optimal hedging bandwidth is similar to that of the Whalley and Wilmott one: the hedging bandwidth increases when either the level of the transaction costs, the hedger's risk tolerance, or the option gamma increases.

Recall that in asymptotic analysis one studies the limiting behavior of the optimal hedging policy as one or several parameters of the problem approach zero. Even though asymptotic analysis can reveal the underlying structure of the solution, under realistic parameters this method provide not quite accurate results. Our comparison of the performance of an asymptotic strategy against the exact strategy reveals that under realistic model parameters an asymptotic strategy performs noticeably worse than that obtained from the exact numerical solution. The explanation lies in the fact that, when some of the model parameters are neither very small nor very large, an asymptotic solution provides not quite accurate results. In particular, as compared to the exact numerical solution, under realistic parameters the size of the hedging bandwidth and the volatility adjustment obtained from asymptotic analysis are overvalued. What is more important, an asymptotic solution showed to be unable to sustain a correct interrelationship between

²This was first observed by Hodges and Neuberger (1989) and emphasized by Clewlow and Hodges (1997) using the results of Monte Carlo simulations.

³The procedure to do it will be described in the subsequent section.

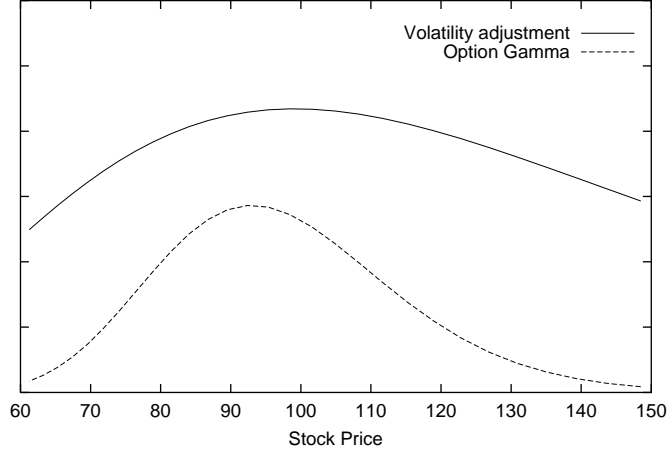


Figure 3: The form of the optimal volatility adjustment $\sigma_m^2 - \sigma^2$, obtained using the exact numerics, versus the option gamma for the same model parameters as in Figure 1.

the size of the hedging bandwidth and the volatility adjustment. The significance of the correct interrelationship could hardly be overemphasized: Our empirical testing of the hedging strategies (see Section 4) reveals that either undervaluation or overvaluation of the volatility adjustment (with respect to the size of the hedging bandwidth) results in a drastic deterioration of the performance of a hedging strategy.

In the conclusion of this section we would like to summarize the stylized facts about the optimal hedging policy and suggest a general specification of the optimal hedging policy. The careful visual inspection of the numerically calculated optimal hedging policy together with the insights from asymptotic analysis advocate for the following general specification of the optimal hedging strategy

$$\Delta = \frac{\partial V(\sigma_m)}{\partial S} \pm (H_w + H_0), \quad (15)$$

where σ_m is the adjusted volatility given by

$$\sigma_m^2 = \sigma^2 (1 + H_\sigma). \quad (16)$$

The term H_w is closely related to H_{ww} in the Whalley and Wilmott hedging strategy (see equation (12)) and to H_{bs} in the Barles and Soner hedging strategy (see equation (13)). The main feature of H_w is that this term depends on the option gamma. Note that the option gamma approaches zero as the option goes farther either out-of-the-money ($S \rightarrow 0$) or in-the-money ($S \rightarrow \infty$). This means that the size of the no transaction region in an asymptotic strategy also approaches zero. On the contrary, the exact

numerics show that, when the option gamma goes to zero, the size of the no transaction region times the stock price approaches a constant value⁴, see Figure 2. It turns out that this constant value is actually the size of the no-transaction region in the model without option liability. To reflect this feature of the optimal hedging policy we introduced the term H_0 . Finally, H_σ is closely related to Σ_{bs} in the Barles and Soner volatility adjustment (see equation (14)) which depends on the option gamma.

4 The Approximation Method

Recall that in asymptotic analysis one studies the solution to a problem when some parameters in the problem assume large or small values. Our empirical testing showed that the performance of an asymptotic strategy is noticeably worse than that of the exact strategy when some of the model parameters are neither very small nor very large. Consequently, for practical applications it makes sense to use alternatives to the asymptotic analysis. One of such alternatives is the approximation method. The general description of the approximation technique we employ can be found in, for example, Judd (1998) Chapter 6. Under approximation we mean the following: We have a rather slow and cumbersome way to compute the optimal hedging policy with transaction costs and want to replace it with simple and efficient approximating function(s). To do this, we first specify a flexible functional form of the optimal hedging policy. Then, given a functional form, the parameters are chosen to provide the best fit to the exact numerical solution. This second stage is known as “model calibration”.

Specifying the right type of functional form of an approximating function is an art rather than science. Our type of specification is based on the detailed graphical inspection of the hedging policy, the comparative statics analysis of the exact numerical solution, and on the results of the asymptotic analysis. Our conclusion is that the following general form is a reasonable specification of an approximating function for the optimal hedging policy

$$\hat{h}(z_1, \dots, z_n) = \alpha \prod_{i=1}^k \varphi_i(z_1, \dots, z_n)^{\beta_i}, \quad (17)$$

where z_1, \dots, z_n are the model parameters, $\varphi_1(\cdot), \dots, \varphi_k(\cdot)$ are appropriately selected basis functions, and $\alpha, \beta_1, \dots, \beta_k$ are parameters to be chosen in order to achieve the best fit. Some of the basis function are fairly simple and have the following form

$$\varphi_i(z_1, \dots, z_n)^{\beta_i} = z_j^{\beta_i}.$$

⁴This feature of the optimal hedging strategy is important when the hedger’s risk aversion is not very high. However, when the hedger’s risk aversion is very high, we can ignore this term.

Observe that after the linearizing log-log transformation of (17) it takes the form

$$\log(\hat{h}(z_1, \dots, z_n)) = \log(\alpha) + \sum_{i=1}^k \beta_i \log(\varphi_i(z_1, \dots, z_n)).$$

We denote the true value of the unknown function at point (z_1^j, \dots, z_n^j) by $h(z_1^j, \dots, z_n^j)$. This value is obtained from exact numerical computations. We measure the goodness of fit using the L^2 norm. This largely amounts to using the techniques of ordinary linear regression after the log-log transformation. That is, we find the parameters $\alpha, \beta_1, \dots, \beta_k$ by solving the problem

$$\min_{\alpha, \beta_1, \dots, \beta_k} \sum_{j=1}^m \left(\log(h(z_1^j, \dots, z_n^j)) - \log(\alpha) - \sum_{i=1}^k \beta_i \log(\varphi_i(z_1^j, \dots, z_n^j)) \right)^2,$$

where m is the number of different data points.

Our study shows (we refer the reader to Zakamouline (2004) for the analysis and the detailed description of the approximation methodology) that the following model parameters z_1, \dots, z_n are to be expected in a reasonable explicit formula for the optimal hedging strategy: the product γS , the quotient $\frac{S}{K}$, the risk-free rate of return r , the volatility σ , the time to maturity $T - t$, and the level of transaction costs λ . There is no good reason to believe that a simple functional specification like (17) can produce a decent approximation for all possible sets of model parameters. However, such a functional form yields quite a nice fit to a true function over some narrow intervals of parameters. We restrict our attention to the following intervals of the *realistic* model parameters: $r \in [0, 0.1]$, $\sigma \in [0.1, 0.4]$, $T - t \in [0, 1.5]$, $\lambda \in [0.001, 0.02]$. The hedger's risk aversion γ is largely unknown. Consequently, for a fixed S we want the interval for γ to be as wide as possible. Note that an "exact" solution computed numerically is subject to some discretization error ϵ . This error does not allow us to carry on the measurements of the volatility adjustment with sufficient precision when γS is very low, and the measurements of the hedging bandwidth when γS is very high. As a result, our model calibration is limited to the interval $\gamma \in [0.05, 15]$ when $S = 100$. Nevertheless, the results of our empirical testing presented in the subsequent section show that our approximation can be "extrapolated" to a wider range for γ .

We assume the following functional form of the approximating function for H_0

$$H_0 = \alpha \sigma^{\beta_1} \lambda^{\beta_2} (\gamma S)^{\beta_3} (T - t)^{\beta_4}. \quad (18)$$

The functional form of the approximating functions for H_w and H_σ is assumed to be as follows

$$H_w = H_\sigma = \alpha \sigma^{\beta_1} \lambda^{\beta_2} (\gamma S)^{\beta_3} (N'(d_1))^{\beta_4} e^{-\beta_5 r (T-t)} (T - t)^{\beta_6}, \quad (19)$$

where d_1 is defined by (2) with the original volatility σ .

The two bandwidths H_0 and H_w are computed in accordance with the following procedure: for some fixed set of parameters $T - t, r, \sigma, \lambda, \gamma S, \frac{S}{K}$ we calculate numerically the upper y_0^u and the lower y_0^l boundaries of the no transaction region without option liability, and the upper y_w^u and the lower y_w^l boundaries of the no transaction region with option liability. Then

$$H_0 = \frac{y_0^u - y_0^l}{2},$$

$$H_w = \frac{y_w^u - y_w^l}{2} - H_0.$$

The volatility adjustment function H_σ is computed as follows: After the following transformation of equation (16)

$$\log \left(\frac{\sigma_m^2}{\sigma^2} - 1 \right) = \log (H_\sigma)$$

it takes a linear form where the unknown parameters α and β_i can be estimated using ordinary regression. The adjusted volatility σ_m is measured in accordance with the following procedure: for some fixed set of parameters $T - t, r, \sigma, \lambda, \gamma S, \frac{S}{K}$ we calculate numerically the upper y_w^u and the lower y_w^l boundaries of the hedging bandwidth with option liability. According to equation (15), the middle of the hedging bandwidth is given by

$$\frac{y_w^u + y_w^l}{2} = \frac{\partial V(\sigma_m)}{\partial S} = N(d_1(\sigma_m)).$$

Consequently, the value of $d_1(\sigma_m)$ can be calculated as

$$d_1(\sigma_m) = N^{-1} \left(\frac{y_w^u + y_w^l}{2} \right),$$

where $N^{-1}(\cdot)$ is the inverse of the cumulative normal distribution function. From the other side, $d_1(\sigma_m)$ is defined in accordance with equation (2). Therefore, to find the unknown σ_m we need to solve the following quadratic equation

$$\frac{1}{2}(T - t)\sigma_m^2 - \sqrt{T - t} d_1(\sigma_m)\sigma_m + \log \left(\frac{S}{K} \right) + r(T - t) = 0.$$

The general idea for estimating the parameters $\alpha, \beta_1, \dots, \beta_k$ is as follows: One first localizes the problem on the bounded space $(Z_{min}^1, Z_{max}^1) \times \dots \times (Z_{min}^n, Z_{max}^n)$ for the model parameters z_1, \dots, z_n . Then one defines a grid on this space. That is, every interval (Z_{min}^i, Z_{max}^i) of the model parameter z_i is divided into M regularly spaced values. For every possible combination of model parameters z_1^j, \dots, z_n^j , the true value of unknown function

$h(z_1^j, \dots, z_n^j)$ is obtained from exact numerical computations. Finally, one runs some algorithm (in our case it is the linear regression) to find the best fit parameters.

The results of our estimation of the best-fit parameters for H_0 , H_w , and H_σ , after rounding off the values of parameters α and β_i and keeping the most significant of them, give the following approximating functions

$$H_0 = \frac{\lambda}{\gamma S \sigma^2 (T - t)}, \quad (20)$$

$$H_w = 1.08 \frac{\lambda^{0.31}}{\sigma^{0.25}} \left(\frac{\Gamma}{\gamma} \right)^{0.5}, \quad (21)$$

$$H_\sigma = 6.85 \frac{\lambda^{0.78}}{\sigma^{0.25}} (\gamma S^2 \Gamma)^{0.15}. \quad (22)$$

Note that for H_w and H_σ we have gathered together some terms to show the dependence on the option gamma.

5 The Simulation Analysis

In this section we present an empirical testing of our approximation strategy against the asymptotic and some other well-known strategies. In order to compare the performance of different hedging strategies we need to choose a suitable unified risk-return framework. The expected replication error seems to be the only sensible candidate for the return measure. As for the risk measure, there are many metrics of risk. In the context of option hedging, the two following metrics are most popular: the variance and the value-at-risk (VaR) of the replication error. Since some practitioners prefer one metric to another, we have chosen to employ both of them. That is, we evaluate the performance of the different hedging strategies within both the mean-variance and mean-VaR frameworks. The reader is reminded that neither of the two frameworks favors a priori any one hedging strategy.

We provide an empirical testing of our approximation strategy against the Black-Scholes, Leland, delta tolerance, Whalley and Wilmott, Barles and Soner strategies, and a so-called *asset tolerance* strategy⁵. The latter strategy prescribes rehedging to the Black-Scholes delta when the percentage change in the value of the underlying asset exceeds the prescribed tolerance. More formally, the series of stopping times is recursively given by

$$\tau_1 = 0, \quad \tau_{i+1} = \inf \left\{ t > \tau_i : \left| \frac{S(t) - S(\tau_i)}{S(\tau_i)} \right| > h \right\}, \quad i = 1, 2, \dots,$$

where h is a given constant percentage. The intuition behind this strategy is similar to that of the delta tolerance strategy.

⁵This is the model of Henrotte (1993).

The different hedging strategies were simulated and the results are presented below. The option was a 1 year short European call with $S_t = K = 100$, the volatility $\sigma = 0.25$, and the drift $\mu = r = 0.05$. The proportional transaction costs were $\lambda = 0.01$. The simulation proceeds as follows: At the beginning, the writer of a European call option receives the Black-Scholes value of the option and sets up a replicating portfolio. The underlying path of the stock is simulated according to

$$S(t + \delta t) = S(t) \exp \left(\left(\mu - \frac{1}{2} \sigma^2 \right) \delta t + \sigma \sqrt{\delta t} \varepsilon \right),$$

where ε is a normally distributed variable with mean 0 and variance 1. At each δt a check is made to see if the option needs to be rehedge. If so, the rebalancing trade is performed and transaction costs are drawn from the bank account. Finally, at expiration, we compute the replication error, that is, the cash value of the replicating portfolio minus the due exercise payment. For the Black-Scholes and Leland strategies, we vary the parameter δt (rehedging interval). For the other strategies, each path consists of 250 equally spaced trading dates over the life of the option. In a utility based strategy we vary the risk aversion parameter γ in $[0.01, 50]$. In both the delta tolerance and the asset tolerance strategies, H and h take values in $[0.01, 0.35]$. For each value of the parameter of a hedging strategy we generate 100,000 paths and compute the following statistics of the replication error: the mean, variance, and VaR at a 95% confidence level (the latter is defined as the loss that is expected to be exceeded with a probability of 5 %). Hence, for a given parameter value, the results of simulations can be represented by a point in a risk-return space. By varying the value of the parameter of a hedging strategy, we span all the possible combinations of risk and return. Thus, the obtained curve for a given strategy can be intuitively interpreted as the tradeoff between the risk and the costs of a hedging strategy.

Figures 4 and 5 summarize the results of simulations. In the both risk-return frameworks our approximation strategy outperforms all the others. We would like to point out the reasons why our approximation strategy is better than the two strategies obtained via asymptotic analysis: Neither of the two asymptotic strategies has the hedging bandwidth H_0 . Consequently, both deep in-the-money and out-of-the-money options, irrespective of the hedger's risk aversion, are hedged with almost zero bandwidth. This is clearly suboptimal and explains why our approximation strategy is better for a hedger with a low risk aversion.

For a hedger with high risk aversion it is very important to hedge with optimal volatility adjustment. The delta tolerance, asset tolerance, and Whalley and Wilmott strategies have no volatility adjustment. Recall that as γ increases, the hedging bandwidth decreases in order to decrease the risk of the hedging portfolio. Without volatility adjustment, any strategy

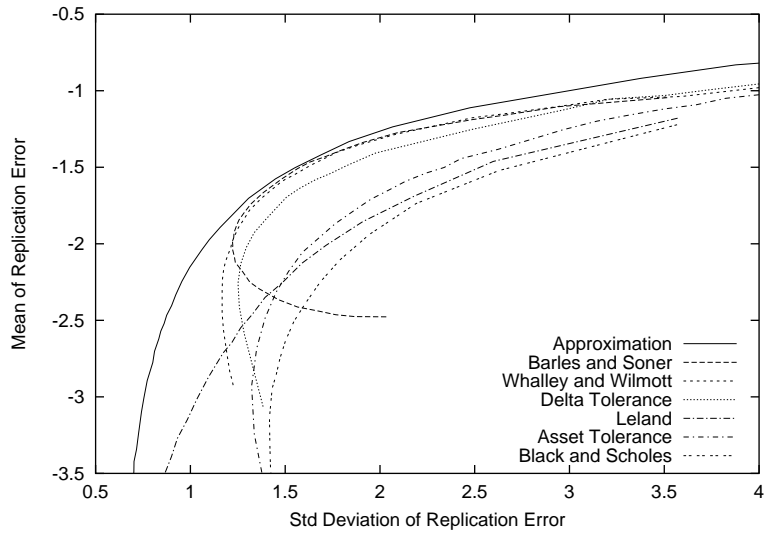


Figure 4: Comparison of hedging strategies in the mean-variance framework.

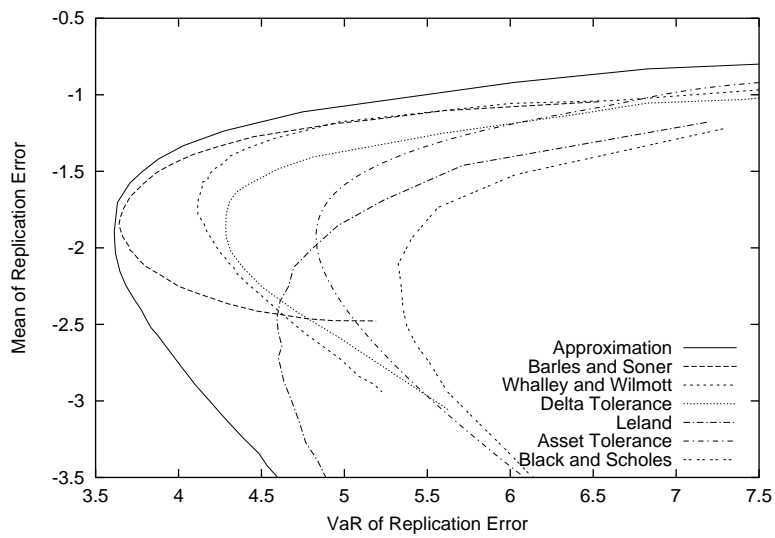


Figure 5: Comparison of hedging strategies in the mean-VaR framework.

converges to the Black-Scholes strategy. That is, in the limit, the amount of transaction costs tends to infinity. On the contrary, in the correct utility based strategy, as γ increases, the decrease in the size of the hedging bandwidth is largely compensated by the increase in the volatility adjustment. Roughly speaking, in the limit as $\gamma \rightarrow \infty$ the utility based strategy tends to the Leland's strategy where the rehedging interval approaches zero.

Despite the presence of the volatility adjustment, the Barles and Soner strategy performs worse than the Whalley and Wilmott strategy when we consider a highly risk averse hedger. The reason is that under realistic model parameters the Barles and Soner strategy tends to overvalue the volatility adjustment with respect to the size of the hedging bandwidth. That is, the Barles and Soner strategy showed to be unable to sustain a correct inter-relationship between the size of the hedging bandwidth and the volatility adjustment. Our simulations show that either undervaluation or overvaluation of the volatility adjustment results in a drastic deterioration of the performance of a hedging strategy.

In addition, we have carried on similar simulations for different values of r , μ , σ , $T - t$, and λ . Due to the space limitations, the results of these simulations are not presented. Qualitatively, the relative performance of different hedging strategies remains the same as that in Figures 4 and 5. Needless to say that the relative advantage of our approximation strategy over the asymptotic strategies decreases when the level of transaction costs decreases. When $\lambda < 0.001$, our approximation strategy and both the asymptotic strategies produce practically identical results in the given range of parameter γ .

It is important to note that the assumption $\mu = r$ is not very essential for the results of simulations and is made solely for simplicity. Our study shows that an increase in μ with respect to r results in some decrease in the advantage of the approximation strategy over the Leland and Black-Scholes strategies for the values of the risk aversion parameter close to the lower end of the chosen interval. The other important remark is that, according to our simulation analysis, the assumption on the absence of transaction costs at terminal time T is not very essential either. If we include transaction costs at T and conduct a similar simulation analysis, we will get essentially the same pictures as those on Figures 4 and 5 with the only difference: all the curves experience an (approximately) equal shift below from their original positions.

6 Dwelling on the Results

The approximation methodology presented in Section 4 produces quite an impressive result as applied to the problem of finding a closed-form expressions for the optimal hedging strategy for a short European call option. However, to arrive to this result we started with the numerical calculations of the optimal hedging strategy for a large set of the model parameters. To find the optimal hedging strategy for another type of a European option (for example, a long European call), we need again to start with the numerical computations of the optimal hedging strategy for this option. Since the numerical algorithm is cumbersome to implement and the calculation of the optimal hedging strategy is time consuming, the approximation methodology is unlikely to be commonly used by the practitioners.

From the other side, there is great practitioner resistance to the idea of utility theory. The most usual objections to the idea of utility theory from a typical practitioner are the following: (i) We do not know the utility function of a hedger, (ii) Even if we know the utility function we need to specify somehow the risk aversion parameter of a hedger. However, the delta tolerance hedging strategy is commonly accepted among the practitioners.

First of all we would like to convince the practitioners that the utility based hedging is very much alike the delta tolerance hedging. Ignoring for the moment the hedging bandwidth H_0 and the volatility adjustment H_σ , the optimal hedging strategy for a short European call option takes the form

$$\Delta = \frac{\partial V}{\partial S} \pm 1.08 \frac{\lambda^{0.31}}{\sigma^{0.25}} \left(\frac{\Gamma}{\gamma} \right)^{0.5}. \quad (23)$$

If we do not know the risk aversion parameter γ , the knowledge of the rest of the model parameters does not help us: we cannot determine the optimal hedging bandwidth. However, what we do know in this case is that the form of the optimal hedging strategy is given by

$$\Delta = \frac{\partial V}{\partial S} \pm h_w \Gamma^{0.5}, \quad (24)$$

where h_w is some unknown parameter which depends on the hedger's risk aversion. In addition, we know that h_w decreases when the hedger's risk aversion increases. Similarly, when the hedger's risk aversion decreases, h_w increases. Consequently, the parameter h_w is closely related to parameter H (see equation (11)) in the delta tolerance strategy and is a proxy for the measure of risk of the hedging portfolio. Without knowing the risk aversion parameter of a hedger we can do the following: we provide the simulation analysis, similarly as in Section 5, where we vary the parameter h_w and span all the possible combinations of risk and return. Then we present the results of the simulation analysis to a particular hedger and allow him to choose the best risk-return tradeoff according to his own preferences. This is similar to

what some investment companies are doing when they want to reveal the risk aversion of a prospective client: they propose him either to choose what fraction of his wealth to invest in the risky assets by giving him a sort of menu or to choose a fund from a list of funds reflecting their perspective risk-return tradeoffs.

Second, in this section we would like to present the intuition behind the other method to find the optimal hedging strategy with transaction costs. This method have a great advantage over the approximation method in that one does not need to implement the numerical calculations of the optimal hedging strategy. Instead, the other method is based on simulations, which are much easier to implement. However, this other method heavily relies on the knowledge of the form of the optimal strategy which is obtained using the asymptotic and approximation methods.

If we, in fact, use the hedging strategy given by (23), we will discover that its performance is similar to that of the Whalley and Wilmott hedging strategy given by (12). The introduction of the hedging bandwidth H_0 and the volatility adjustment H_σ helps to improve the performance of the hedging strategy given by (23). Quite generally, if we do not know the hedger's risk aversion, we can describe the optimal hedging strategy for a short European call option as

$$\begin{aligned}\Delta &= \frac{\partial V(\sigma_m)}{\partial S} \pm \left(h_w \Gamma^{0.5} + \frac{h_0}{S(T-t)} \right), \\ \sigma_m^2 &= \sigma^2 \left(1 + h_\sigma (S^2 \Gamma)^{0.15} \right).\end{aligned}\tag{25}$$

Similarly to (24) we have introduced two parameters, h_0 and h_σ , which depend on the hedger's risk aversion: as γ increases, h_0 decreases and h_σ increases.

Now we are going to present an important idea suggested by the results of our simulation analysis: if we choose a suitable risk-return framework, there are some optimal combinations of the triple of parameters (h_w, h_0, h_σ) that generate the efficient frontier. That is, suppose we do not know the parameters (h_w, h_0, h_σ) of the optimal hedging strategy. Using a simulation analysis we generate points in the chosen risk-return space for every possible combination of (h_w, h_0, h_σ) . Having fulfilled these simulations we can note that there is an efficient frontier in the risk-return space such that one cannot improve the risk-return tradeoff for any point belonging to this efficient frontier. That is, any other combinations of (h_w, h_0, h_σ) which produce a lower risk have a lower return, or, similarly, any other combinations of (h_w, h_0, h_σ) which produce a higher return have a higher risk. Figure 6 illustrate the efficient frontier in some risk-return space.

Consequently, the main idea behind the other method of finding the optimal hedging strategy is to use the simulation analysis in order to find the efficient frontier of a hedging strategy in some risk-return space. Note

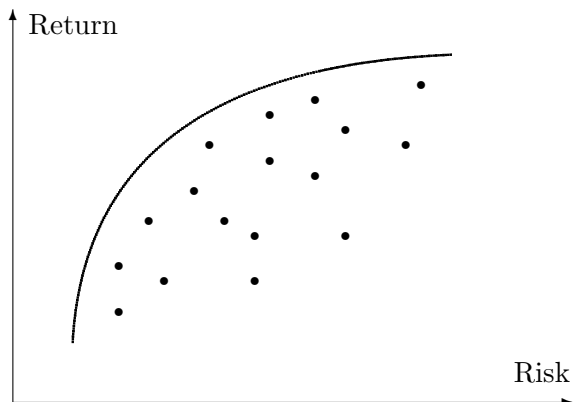


Figure 6: Possible combinations of the risk and return of a hedging strategy with the triple of parameters (h_w, h_0, h_σ) and the efficient frontier.

that the general form of the optimal hedging strategy is obtained using the asymptotic and approximation methods applied to the utility based hedging problem. At first sight, the risk-return optimization has nothing to do with the utility maximization. However, there is a common belief that these two problems are, in fact, deeply interrelated. Indeed, the mean-variance utility function presented by Markowitz was meant to reflect the risk-return tradeoff faced by an investor.

7 The Optimization Method

The detailed study of the utility based hedging strategy for a plain vanilla European option shows that the general description of the optimal hedging strategy could be given by

$$\begin{aligned} \Delta &= \frac{\partial V(\sigma_m)}{\partial S} \pm \left(h_w |\Gamma|^\alpha + \frac{h_0}{S(T-t)} \right), \\ \sigma_m^2 &= \sigma^2 \left(1 + h_\sigma \text{sign}(\Gamma) |S^2 \Gamma|^\beta \right), \end{aligned} \tag{26}$$

where V and Γ are the value and the gamma of the option, respectively, in a market with no transaction costs (the Black-Scholes market model). The rational under this general description is the following: Our study shows that the form of the optimal hedging bandwidth of any option resembles the form of the option gamma. That is, ceteris paribus, the hedging bandwidth increases when the (absolute value of the) option gamma increases. This is also true for any portfolio of options. Moreover, when the hedger's risk aversion is not high, it is optimal to introduce the hedging bandwidth H_0 . The latter increases when the time to maturity decreases. In addition,

the visual observations and our approximation analysis shows that H_0S is approximately constant with respect to S . Finally, when the hedger's risk aversion is rather high, it is optimal to hedge with the adjusted volatility. When the option gamma is positive, the adjusted volatility is higher than the original volatility. On the contrary, when the option gamma is negative, the adjusted volatility is lower than the original volatility. The form of the volatility adjustment also resembles the form of the option gamma: the volatility is adjusted more in regions where the (absolute value of the) option gamma is higher.

The purpose of this section is to describe our risk-return optimization method where we search for the combinations of the triple of parameters (h_w, h_0, h_σ) which belong to the efficient frontier or, at least, the efficient set. In addition, since we are intended to present a general method of finding the optimal hedging strategy for any plain vanilla European option, we need to specify somehow a single pair of parameters (α, β) in (26). Our study shows that $\alpha \in (0.3, 0.7)$ and $\beta \in (0, 0.3)$. Actually, the optimal value of α and β depends not only on a particular type of option, but also on the hedger's risk aversion. This results in the fact that the comparison of the performances of the optimal hedging strategies for different pairs of parameters (α, β) does not reveal any substantial difference among them. Consequently, the choice of the values of (α, β) can be made quite arbitrary, for example $\alpha = 0.5$ and $\beta = 0$. The choice $\beta = 0$ seems a bit surprising, but is completely equivalent to the Leland's volatility adjustment and could probably be justified by the following: the presence of the volatility adjustment is much more important than its (slight) dependence on the option gamma. If someone is not satisfied with this arbitrary choice, one can always modify the optimization method to find the optimal combinations of parameters $(\alpha, \beta, h_w, h_0, h_\sigma)$. There is no issue of principle here, just the increased computational load.

Now we turn on to the formal presentation of the optimization method in some risk-return space. We choose some triple of parameters of the hedging strategy, (h_w, h_0, h_σ) , perform path simulations, and estimate the return, $\eta = \eta(h_w, h_0, h_\sigma)$, and the risk, $\rho = \rho(h_w, h_0, h_\sigma)$, associated with this particular hedging strategy. By return we mean the expected replication error. By risk we mean some suitable risk metric, for example, the variance or the VaR of the replication error. The hedging strategy with parameters (h'_w, h'_0, h'_σ) is considered⁶ to be better than the strategy with parameters (h_w, h_0, h_σ) if

$$\eta(h'_w, h'_0, h'_\sigma) \geq \eta(h_w, h_0, h_\sigma) \text{ and } \rho(h'_w, h'_0, h'_\sigma) \leq \rho(h_w, h_0, h_\sigma). \quad (27)$$

That is, if the strategy with (h'_w, h'_0, h'_σ) provides either higher return with less or equal risk, or less risk with higher or equal return than the strategy with (h_w, h_0, h_σ) .

⁶Note that this is true for any risk averse hedger.

The optimization method we propose is based on a sequential improvement of the risk-return tradeoff of a hedging strategy. That is, starting with some (h_w, h_0, h_σ) we search for a new $(h'_w, h'_0, h'_\sigma) = (h_w + \Delta h_w, h_0 + \Delta h_0, h_\sigma + \Delta h_\sigma)$ such that the strategy with (h'_w, h'_0, h'_σ) is better than the strategy with (h_w, h_0, h_σ) . The most general procedure for finding the risk-return improvement step could be interpreted as finding a unit vector $\mathbf{u} = \langle a, b, c \rangle$ such that

$$D_{\mathbf{u}}\eta(h_w, h_0, h_\sigma) \geq 0, \quad (28)$$

and

$$D_{\mathbf{u}}\rho(h_w, h_0, h_\sigma) \leq 0, \quad (29)$$

where $D_{\mathbf{u}}\eta(h_w, h_0, h_\sigma)$ and $D_{\mathbf{u}}\rho(h_w, h_0, h_\sigma)$ are the directional derivatives of $\eta(h_w, h_0, h_\sigma)$ and $\rho(h_w, h_0, h_\sigma)$, respectively, in the direction of the unit vector \mathbf{u} . The reader is reminded that the directional derivative of some function $f(x, y, z)$ in the direction of the unit vector $\mathbf{u} = \langle a, b, c \rangle$ is defined as

$$\begin{aligned} D_{\mathbf{u}}f(x, y, z) &= \lim_{h \rightarrow 0} \frac{f(x + ha, y + hb, z + hc) - f(x, y, z)}{h} \\ &= f_x(x, y, z)a + f_y(x, y, z)b + f_z(x, y, z)c = \nabla f(x, y, z)\mathbf{u}. \end{aligned}$$

Consequently, to implement an improvement step from the point (h_w, h_0, h_σ) we need first to find the partial derivatives of η and ρ at this point and then to find a vector \mathbf{u} such that both (28) and (29) are satisfied. Note that the vector \mathbf{u} is not unique. However, since our goal is to find the points on the efficient frontier, the actual path from the benchmark $(h_w, 0, 0)$ to (h'_w, h'_0, h'_σ) belonging to the efficient frontier is not important, especially since this path is, in fact, more or less short.

In a practical realization of the optimization method the algorithm based on the partial derivatives is very time consuming, since in order to calculate a partial derivative with high precision we need to simulate a great number of stock paths. Below we describe the algorithm based on the implementation of an improvement step with respect to (h_w, h_0) and a consequent improvement step with respect to (h_w, h_σ) .

The starting point for our optimization method is the risk-return tradeoff given by the benchmark strategy $(h_w, 0, 0)$. The introduction of h_0 and h_σ helps to improve the risk-return tradeoff of the benchmark strategy. It is important to note that both h_0 and h_σ are positive. Suppose that some (h_w, h_0, h_σ) is an improvement of the benchmark strategy. We denote by A the point in the risk-return space with coordinates $(\rho(h_w, h_0, h_\sigma), \eta(h_w, h_0, h_\sigma))$ (see Figure 7). The further improvement of the risk-return tradeoff by means of $(\Delta h_w, \Delta h_0)$ consists in the further increase of h_0 by some Δh_0 . We denote by B the point in the risk-return space with coordinates $(\rho(h_w, h_0 + \Delta h_0, h_\sigma), \eta(h_w, h_0 + \Delta h_0, h_\sigma))$. Since by increasing h_0 we increase the total hedging bandwidth, the strategy

$(h_w, h_0 + \Delta h_0, h_\sigma)$ has a higher return and a higher risk than the strategy (h_w, h_0, h_σ) . To decrease the risk (and the return at the same time) as compared to the strategy $(h_w, h_0 + \Delta h_0, h_\sigma)$ we need now to decrease the hedging bandwidth H_w . This is achieved by decreasing h_w by some Δh_w . We try to find the size of Δh_w such that a point C with coordinates $(\rho(h_w - \Delta h_w, h_0 + \Delta h_0, h_\sigma), \eta(h_w - \Delta h_w, h_0 + \Delta h_0, h_\sigma))$ has a better risk-return tradeoff than the original point A . However, this is not always possible. If, for example, the transition between B and C goes along the line (B, C') (see Figure 7), then the improvement step by $(\Delta h_w, \Delta h_0)$ is feasible. Otherwise, if the transition between B and C goes along the line (B, C'') , it is not possible to improve the risk-return tradeoff by $(\Delta h_w, \Delta h_0)$.

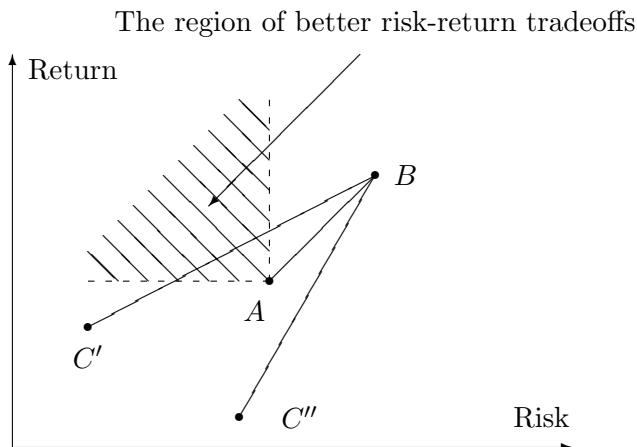


Figure 7: Illustration of the improvement step by means of $(\Delta h_w, \Delta h_0)$.

Now we turn on to the description of the improvement step by means of $(\Delta h_w, \Delta h_\sigma)$. As before, we denote by A the point in the risk-return space with coordinates $(\rho(h_w, h_0, h_\sigma), \eta(h_w, h_0, h_\sigma))$ (see Figure 8). The improvement of the risk-return tradeoff by means of $(\Delta h_w, \Delta h_\sigma)$ consists in the increase of h_σ by some Δh_σ . We denote by B the point in the risk-return space with coordinates $(\rho(h_w, h_0, h_\sigma + \Delta h_\sigma), \eta(h_w, h_0, h_\sigma + \Delta h_\sigma))$. This improvement step might be possible if $\rho(h_w, h_0, h_\sigma + \Delta h_\sigma) < \rho(h_w, h_0, h_\sigma)$. That is, if the increase of h_σ decreases the risk. If it is so, there are two typical situations: if the option gamma is positive, a new strategy $(\rho(h_w, h_0, h_\sigma + \Delta h_\sigma), \eta(h_w, h_0, h_\sigma + \Delta h_\sigma))$ has also a higher return. This situation is illustrated by the point B' . If the option gamma is negative, a new strategy with a lower risk has also a lower return. This situation is illustrated by the point B'' . In the latter case to increase the return (and the risk at the same time) as compared to the strategy $(h_w, h_0, h_\sigma + \Delta h_\sigma)$ we need now to increase the hedging bandwidth H_w . This is achieved by

increasing h_w by some Δh_w . We try to find the size of Δh_w such that a point C with coordinates $(\rho(h_w + \Delta h_w, h_0, h_\sigma + \Delta h_\sigma), \eta(h_w + \Delta h_w, h_0, h_\sigma + \Delta h_\sigma))$ has a better risk-return tradeoff than the original point A . Again, this is not always possible. If, for example, the transition between B'' and C goes along the line (B'', C') (again, see Figure 8), then the improvement step by $(\Delta h_w, \Delta h_\sigma)$ is feasible. Otherwise, if the transition between B'' and C goes along the line (B'', C'') , it is not possible to improve the risk-return tradeoff by $(\Delta h_w, \Delta h_\sigma)$. If it is not possible to implement the improvement step by either $(\Delta h_w, \Delta h_0)$ or $(\Delta h_w, \Delta h_\sigma)$, it means that A belongs to the efficient frontier.




Figure 8: Illustration of the improvement step by means of $(\Delta h_w, \Delta h_\sigma)$.

Finally we present the results (see Figures 9 and 10) of our risk-return optimization of the hedging strategy for a short European call option in the mean-variance and the mean-VaR frameworks. The option parameters were the same as in Section 5. We varied the parameter h_w of the hedging strategy in $[0.025, 2.5]$. We used fixed steps for Δh_0 and Δh_σ and adaptive steps for Δh_w . To implement the optimization and compute the risk and return of the hedging strategy with higher precision we generated 400,000 paths for each triple of parameters (h_w, h_0, h_σ) .

The optimized strategy is compared against the approximation strategy and the benchmark strategy (i.e., the strategy with $h_0 = h_\sigma = 0$). The risk-return tradeoff of the Leland strategy is also presented as a standard benchmark for comparison. It is interesting to note that irrespective of the chosen risk metric the optimization in either the mean-variance or the mean-VaR framework produces similar results to those of the approximation strategy. However, a closer look at the optimal values of h_0 and h_σ reveals that for the same value of h_w the optimization in the mean-VaR framework gives a slightly higher value of h_σ than the optimization in the mean-variance framework. Nevertheless, it is difficult to notice any substantial difference

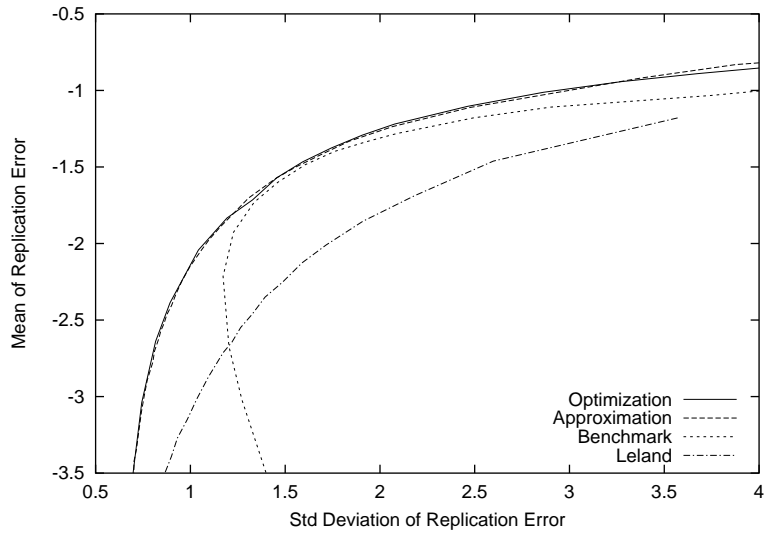


Figure 9: Comparison of hedging strategies in the mean-variance framework.

in the optimized and the approximation strategy in Figure 10.




Figure 10: Comparison of hedging strategies in the mean-VaR framework.

References

- Andersen, E. D. and Damgaard, A. (1999). “Utility Based Option Pricing with Proportional Transaction Costs and Diversification Problems: an Interior-Point Optimization Approach”, *Applied Numerical Mathematics*, *29*, 395–422.
- Barles, G. and Soner, H. M. (1998). “Option Pricing with Transaction Costs and a Nonlinear Black-Scholes Equation”, *Finance and Stochastics*, *2*, 369–397.
- Black, F. and Scholes, M. (1973). “The Pricing of Options and Corporate Liabilities”, *Journal of Political Economy*, *81*(3), 637–654.
- Clelow, L. and Hodges, S. (1997). “Optimal Delta-Hedging under Transaction Costs”, *Journal of Economic Dynamics and Control*, *21*, 1353–1376.
- Davis, M. H. A. and Panas, V. G. (1994). “The Writing Price of a European Contingent Claim under Proportional Transaction Costs”, *Computational and Applied Mathematics*, *13*, 115–157.
- Davis, M. H. A., Panas, V. G., and Zariphopoulou, T. (1993). “European Option Pricing with Transaction Costs”, *SIAM Journal of Control and Optimization*, *31*(2), 470–493.
- Henrotte, P. (1993). “Transaction Costs and Duplication Strategies”, Working Paper, Stanford University and HEC.

- Hodges, S. D. and Neuberger, A. (1989). “Optimal Replication of Contingent Claims under Transaction Costs”, *Review of Futures Markets*, 8, 222–239.
- Hoggard, T., Whalley, A. E., and Wilmott, P. (1994). “Hedging Option Portfolios in the Presence of Transaction Costs”, *Advances in Futures and Options Research*, 7, 21–35.
- Judd, K. L. (1998). *Numerical Methods in Economics*. The MIT Press.
- Leland, H. (1985). “Option Pricing and Replication with Transaction Costs”, *Journal of Finance*, 5, 1283–1301.
- Martellini, L. and Priaulet, P. (2002). “Competing Methods for Option Hedging in the Presence of Transaction Costs”, *Journal of Derivatives*, 9(3), 26–38.
- Mohamed, B. (1994). “Simulations of Transaction Costs and Optimal Re-hedging”, *Applied Mathematical Finance*, 1, 49–63.
- Whalley, A. E. and Wilmott, P. (1997). “An Asymptotic Analysis of an Optimal Hedging Model for Option Pricing with Transaction Costs”, *Mathematical Finance*, 7(3), 307–324.
- Zakamouline, V. I. (2005). “European Option Pricing and Hedging with both Fixed and Proportional Transaction Costs”, *Journal of Economic Dynamics and Control*, forthcoming.
- Zakamouline, V. I. (2004). “Efficient Analytic Approximation of the Optimal Hedging Strategy for a European Call Option with Transaction Costs”, Working paper, Bodø Graduate School of Business, Norway.