Small Transaction Cost Asymptotics and Dynamic Hedging*

Claudio Albanese[†] Stathis Tompaidis [‡]

Keywords: Finance, Risk Management, Transaction Costs
Suggested Heading: O.R. Applications

ABSTRACT

Transaction costs are one of the major impediments to the implementation of dynamic hedging strategies. We consider an alternative to utility maximization, similar to the "good-deal" pricing framework in incomplete markets. We perform a dynamic risk-reward analysis for a family of non-self-financing strategies of practical importance: deterministic time hedging; i.e., hedging at predetermined, fixed, times. In the limit of small relative transaction costs, we carry out the asymptotic analysis and find that transaction costs affect the hedge ratios and that the time between trades is related in a simple way to the local sensitivities of the replication target.

*__Acknowledgments:__ We would like to acknowledge the helpful comments and suggestions of seminar participants at the University of Toronto, the Fields Institute, Université de Rennes I, Université de Paris VI, University of Texas at Austin, Swiss Federal Institutes of Technology in Zürich and Lausanne. We would also like to thank Peter Bossaerts, Peter Carr, Patrick Jaillet, Ehud Ronn, and Klaus Toft, for discussions and suggestions. Part of this work was completed when both authors were at the University of Toronto.

[†]Department of Mathematics, Imperial College, London, SW7 2AZ, United Kingdom, claudio.albanese@imperial.ac.uk

[‡]Corresponding author. IROM Department, McCombs School of Business, University of Texas at Austin, 1 University Station, B6500, Austin, TX 78712-1175, stathis.tompaidis@mccombs.utexas.edu

1. Introduction

Due to its practical relevance, the transaction cost problem has received much attention in the literature. Practitioners favor static hedging strategies because transaction costs are known at the contract inception. Dynamic hedging is however needed to manage derivatives books with non-linear instruments if the combination of all static hedge positions leaves a risky, nonlinear remainder unhedged.

When transaction costs are zero, Black and Scholes (1973) and Merton (1973) show that any payoff function can be replicated by means of a self-financing, dynamic, trading strategy. On the other hand, if proportional transaction costs are present, no matter how small, this dynamic hedging strategy, prescribed by the Black-Scholes model, becomes infinitely expensive. This result does not imply that all claims are unhedgeable. Consider for instance the case of a European call: by buying the stock and holding it until maturity, one can dominate the final payoff. Soner, Shreve, and Cvitanić (1995) show that covering the call by buying a unit of the underlying stock is actually the cheapest dominating hedging strategy (Edirisinghe, Naik, and Uppal (1993) studied the super-replication problem in discrete time, using a binomial model). Such a result provides the upper bound for the price of a call. Since it clearly overestimates observed call prices, the result indicates the need for models that allow the possibility of a loss.

Models where a loss is possible, and optimization of a trading strategy is sought instead, have been proposed by Hodges and Neuberger (1989), Davis and Norman (1990) and subsequently by Davis, Panas, and Zariphopoulou (1993). Defining optimization criteria and hedging an option on an underlying stock that incurs transaction costs is a non-trivial task. In the mathematical finance literature the attention focuses on self-financing strategies and minimization of a utility function of the mismatch between the payoff of the option and the value of the replicating portfolio at the expiration date of the option, for the simple reason that at that time the price of the option is unambiguous and no assumption needs to be made for the option price process. For the case of proportional transaction costs, the optimal trading strategy for a large class of utility functions has been shown to be described by a no-transaction interval and by infinitesimal transactions at the boundary of the interval, see Hodges and Neuberger (1989), Davis and Norman (1990), Davis, Panas, and Zariphopoulou (1993).

In an alternative approach, Leland (1985) removes the self-financing constraint and introduces trading at discrete times, imitating periodic marking-to-market. In Leland's framework transaction costs are compensated by systematic gains accumulated during the dynamic hedging process. To achieve these gains, the option price must be higher than the Black-Scholes price. A non-linear extension of the Black-Scholes equation can be used to calculate the modified option price, which leads to a hedging strategy that, on average, generates systematic gains. In Leland's original paper, the systematic gains offset transaction costs on average. Henrotte (1993), and Toft (1996) carry out an analysis of the variance of Leland's strategies, but do not incorporate the results in an optimization framework.

In this article, we extend Leland's framework and incorporate elements from the literature on "good deal" pricing in incomplete markets (see Cochrane and Saa-Requejo (2000) and Bernardo and Ledoit (2000)). In particular we perform a risk-reward analysis on trading strategies and require that optimal strategies lie on the efficient frontier, given that the risk-reward ratio is at a certain level. Due to the difficulty of solving the general problem we restrict ourselves to the case of a one-parameter family of hedging strategies where the parameter corresponds to the time between transactions, and where when trading occurs, the position is rebalanced to a position that reflects the market view.[1] This choice is motivated by both industry practice and practical constraints due to risk management concerns.[2]

[1] This choice is not necessarily optimal under the utility maximization framework. Indeed, for proportional transaction costs, the optimal solution when maximizing terminal time utility involves a no-trade region and small trades at the boundaries of that region. Our choice reflects the common practice of marking-to-market.

[2] The issue of hedging in discrete time when transaction costs prohibit continuous trading has been studied previously in the literature. Similar to our study, Boyle and Emanuel (1980), Leland (1985), Figlewski (1989), Gilster (1990), Gilster (1996), Bensaid, Lesne, Pagés, and Scheinkman (1992), Boyle and Vorst (1992), Edirisinghe, Naik, and Uppal (1993), Henrotte (1993), Toft (1996), and Bertsimas, Kogan, and Lo (2000), study time-based hedging strategies, in which trading occurs at an exogenous frequency. In addition, Henrotte (1993) and Toft (1996) study move-based hedging strategies, in which trading occurs when the price of the underlying asset changes by an exogenous amount.

We offer two alternative optimization criteria, corresponding to the points of view of a market-maker and a price-taker respectively. For the case of a market-maker, that sets the price competitively, we parameterize the efficient frontier by a risk-reward ratio. By constraining the risk-reward ratio the seller of the option is assured a certain level of compensation for the risk undertaken, where compensation is the expected gain from the trading strategy and risk is quantified in terms of the standard deviation of the profit and loss of the trading strategy over the time horizon.

An alternative point of view is that of a price-taker that may only observe the implied volatility at which an option trades. For that case we determine the strategy that maximizes the price-taker's compensation in terms of risk undertaken. The optimization problem can be formulated as: for a given, traded, implied volatility, and a family of strategies, find the strategy that maximizes the risk-reward factor over an investor-defined time horizon.[3]

In the limit of small transaction costs, we determine the asymptotically optimal strategy. Our analysis provides information on the scaling relations between the time interval between trades, the transaction cost, the surcharge over the Black-Scholes price and the sensitivities of the replication target. A similar asymptotic analysis for the case of utility maximization was carried out by Rogers (2000). Due to the difference in objectives, the scaling relationships are different. In particular, under utility maximization, the transaction cost of undertaking infinitesimal trades at the boundary of an interval of size x is of size $1/x$, while the utility cost of allowing an interval of finite size is of size x^2, leading to a choice that minimizes an expression of the form $ax^2 + b/x$. In our framework, we allow for positive expected gains, and also constrain the transactions to be of finite size, trading back to the position indicated by the market. In addition, we impose a minimum level of return for any risk undertaken; i.e., a minimum level of a risk-reward ratio. This results in the minimization of an expression of the form $a\sqrt{x} + b/\sqrt{x}$, leading to differences in the scaling relationships.

[3] Hodges and Neuberger (1989) and Clewlow and Hodges (1997) characterize optimally timed hedging strategies when the objective is to maximize the investor's utility function.

The remainder of the paper is organized as follows: in Section 2 we formulate the problem and motivate our choice of optimization criteria. In Section 3 we present results that describe the optimal trading strategy. In Section 4 we provide a numerical examination of the range of validity of the asymptotic analysis in terms of the size of the transaction cost, as well as comparative statics. Section 5 concludes the paper. All calculations and proofs are contained in the Appendices.

2. Formulation and optimization criteria

Consider a derivative security of European type whose payoff f_0 at expiration is contingent to the price of the stock S. The stochastic process for the stock is assumed to be geometric Brownian motion

$$\frac{dS_t}{S_t} = (\mu - \delta)dt + \sigma dW_t \tag{1}$$

where μ is the growth rate of the stock, δ is the dividend rate, σ is the volatility and W is a standard Wiener process. Pure discount bonds also trade with price

$$dB_t = rB_t dt \tag{2}$$

where r is the risk-free interest rate. The parameters μ, δ, σ and r are assumed to be known functions of S and t.

A trading strategy is described by a pair of adapted processes (a,b), where $a(t)$ is the number of shares and $b(t)$ is the number of bonds held in a portfolio at time t. The value of the portfolio at time t is equal to

$$\Pi(t) = a(t)S(t) + b(t)B(t)$$

Transaction costs are modeled as proportional to the number of shares exchanged. The cost of a trading strategy is the adapted process $c(t)$ where dc is the cost to transact from a portfolio (a,b) to a portfolio $(a+da, b+db)$, given by

$$dc = Sda + Bdb + \frac{k}{2}S\,|da|$$

where $k > 0$ is the relative, round-trip, transaction cost.

In the complete market framework, described by Equations (1), (2) when $k = 0$, any payoff f_0 satisfying mild integrability conditions, can be replicated by means of a self-financing trading strategy. The arbitrage-free price satisfies the Black-Scholes equation

$$\frac{\partial f_{BS}}{\partial t} + (r-\delta)S\frac{\partial f_{BS}}{\partial S} + \frac{\sigma^2 S^2}{2}\frac{\partial^2 f_{BS}}{\partial S^2} = r f_{BS} \tag{3}$$

with final condition $f_{BS}(S,T) = f_0(S)$. The payoff can be replicated by the self-financing strategy

$$a_{BS}(t) = \frac{\partial f_{BS}}{\partial S}(S(t),t), \qquad b_{BS}(t) = B(t)^{-1}\left(f_{BS}(S(t),t) - Sa_{BS}(t)\right).$$

If proportional transaction costs are present, no matter how small, the Black-Scholes dynamic hedging strategy is infinitely expensive, as shown by the following lemma:

Lemma 2.1. *Consider a price process $f_{BS}(S,t)$, $t \in [0,T]$, obeying the Black-Scholes Equation* (3)*. Let N be a positive integer and let $\Delta t = N^{-1}T$. Consider the dynamic hedging strategy for which the replicating portfolio is adjusted at times $t_j = j\Delta t$, $j = 1,\ldots,N$, in such a way that, at time t_j, the position in the stock consists of $a_{BS}(t_j)$ shares and $b_{BS}(t_j)$ bonds. If the relative transaction cost is greater than zero, $k > 0$, then the expected total transaction cost of the strategy is given, to leading order in N, by*

$$E\left[c(T)\right] = \frac{k\sigma}{\sqrt{2\Delta t}}\, E\left[\int_0^T S^2 \left|\frac{\partial^2 f_{BS}}{\partial S^2}(S,t)\right| dt\right] + O(1)$$

For a proof of Lemma 2.1 see the appendix. The lemma implies that in the continuous time limit; i.e., as $N \to \infty$ and $\Delta t \to 0$, the expected total transaction cost diverges. Continuous time hedging is thus unrealistic (similar results were obtained by Leland (1985), and Soner, Shreve, and Cvitanić (1995)).

Leland (1985) proposed to compensate for transaction costs through systematic gains accumulated during the dynamic hedging process. To achieve such systematic gains the process for the option price has to be different than the Black and Scholes process, since the discretely rebalanced Black-Scholes strategy does not lead to systematic gains or losses. Leland has introduced the following non-linear extension of the Black-Scholes equation for the modified price process:

$$\frac{\partial f_L}{\partial t} + (r-\delta)S\frac{\partial f_L}{\partial S} + \frac{\sigma^2 S^2}{2}\left(\frac{\partial^2 f_L}{\partial S^2} + \Lambda\left|\frac{\partial^2 f_L}{\partial S^2}\right|\right) = rf_L, \qquad f_L(S,T) = f_0(S) \tag{4}$$

where $\Lambda > 0$ is a positive parameter that we refer to as the *Leland volatility adjustment*.[4] The intuition behind the Leland volatility adjustment is the following: since transaction costs are proportional to the number of shares traded, and since the number of shares traded is proportional to how rapidly the number of shares, or Delta, Δ, of a replicating portfolio deviates from the Delta of the option, the magnitude of the transaction costs is proportional to the derivative of Delta; i.e., the Gamma, Γ, of a position.[5] Adjusting the volatility in the equation that determines the price of an option leads to systematic gains that are also proportional to Γ, making it possible to balance the expected gains with the expected transaction costs.

Suppose that at time t_0 the hedge ratios are chosen as follows:

$$a_L(t_0) = \frac{\partial f_L}{\partial S}(S(t_0),t_0), \qquad b_L(t_0) = B(t_0)^{-1}\left(f_L(S(t_0),t_0) - Sa_L(t_0)\right) \tag{5}$$

[4]Although the Leland equation is in general nonlinear, in the case of a European option with either concave or convex payoff the Leland equation becomes linear.

[5]Delta, or Δ, is the first derivative of the option price with respect to the price of the underlying asset. Gamma, or Γ, is the second derivative of the option price with respect to the price of the underlying asset.

and consider the difference at time $t_0 + \tau$ between the option value, given by Equation (4) and the portfolio value of a portfolio $(a_L(t_0), b_L(t_0))$, specified by Equation (5):

$$\delta\Pi(t_0+\tau) = a_L(t_0)S(t_0+\tau) + b_L(t_0)B(t_0+\tau) - f_L\left(S\left(t_0+\tau\right), t_0+\tau\right)$$

where τ is the — possibly random — stopping time when the next trade occurs. We have the following result:

Lemma 2.2. *To leading order for small expected values of the stopping time $E_{t_0}[\tau]$, we have that*

$$E_{t_0}\left[\delta\Pi(t_0+\tau)\right] = \frac{\Lambda}{2}|\Gamma_0|S_0^2\sigma^2 E_{t_0}\left[\tau\right] + O\left(E_{t_0}\left[\tau\right]^{3/2}\right)$$

where $\Gamma_0 = \frac{\partial^2 f_L}{\partial S^2}\left(S(t_0), t_0\right)$, $S_0 = S(t_0)$ and expectations are taken with respect to the real measure, and are conditional on information up to and including time t_0.

Lemma 2.2 suggests that the discounted price given by the Leland equation (4) leads to systematic gains for the strategy defined in (5). In Leland's original paper these systematic gains match transaction costs, on average. In this article, we take the point of view that the systematic gains should be such that the hedging strategy belongs to the efficient risk-reward frontier of feasible strategies.

It is useful to look at the Leland equation as a Black-Scholes equation with a modified volatility $\bar{\sigma}$, which depends on the sign of Γ, where

$$\bar{\sigma}^2 = \sigma^2\left(1 + \Lambda \operatorname{sgn}\left(\frac{\partial^2 f_L}{\partial S^2}\right)\right)$$

where sgn is the sign function:

$$\operatorname{sgn}(x) = \begin{cases} 1, & x > 0 \\ 0, & x = 0 \\ -1, & x < 0 \end{cases}$$

The risk-reward analysis is in the spirit of the literature on good-deal pricing in incomplete markets developed by Cochrane and Saa-Requejo (2000) and Bernardo and Ledoit (2000). Assume that at time t_0, the replicating portfolio is worth $f_L(S_0, t_0; \Lambda)$ given by the Leland equation with parameter Λ. In the following we will suppress the dependence on Λ from the notation for f_L. Risk and reward over a time horizon ΔT are measured in terms of the following functions:

$$R_{\Delta T}(S_0, t_0) = \sqrt{E_{t_0}\left[X(t_0)^2\right] - E_{t_0}\left[X(t_0)\right]^2}, \qquad G_{\Delta T}(S_0, t_0) = E_{t_0}\left[X(t_0)\right]$$

where $X(t_0)$ is the present value at time t_0 of the cumulative cash flows of the replicating strategy in the time interval $[t_0, t_0 + \Delta T]$ and t_i are the times at which trades occur:

$$X(t_0) \equiv \sum_{t_0 < t_i < \Delta T + t_0} e^{-r(t_i - t_0)} \delta\Pi(t_i) - \frac{k}{2} \sum_{t_0 < t_i < \Delta T + t_0} e^{-r(t_i - t_0)} \delta c(t_i)$$

Reward, $G_{\Delta T}$, is the expected gain of the strategy given in Equation (5) over the time interval $[t_0, t_0 + \Delta T]$. Risk, $R_{\Delta T}$, is the standard deviation of the gains over the same time interval. We note that the definition of Reward does not include transaction costs for setting up and for closing a hedging position. Conceptually, the inclusion of such costs would destroy the time scaling properties for the Reward, and the costs may be best covered through a separate charge, rather than through the trading strategy. In addition, it can be shown that such costs are of second order in the case of small relative roundtrip transaction costs.

The efficient frontier is parameterized by the ratios

$$J \equiv \frac{G_{\Delta T}(S, t)}{R_{\Delta T}(S, t)}, \qquad A \equiv \frac{G_{\Delta T}(S, t)}{\sqrt{\Delta T} \cdot R_{\Delta T}(S, t)} = \frac{J}{\sqrt{\Delta T}}, \tag{6}$$

that we hold fixed. A fixed risk-reward factor J signifies that the hedger wishes, on average, to achieve reward equal to J times the risk, over the hedger's time horizon; i.e., expected gains equal to J times the standard deviation of the expected gains. In contrast, a fixed risk-reward factor A is *independent* of the time horizon and signifies that the hedger wishes, on average, to achieve reward equal to A times the risk *per unit time*. The ratios J and A can also be thought of as Sharpe ratios corresponding to hedging the option following a prescribed strategy.

We offer two optimization criteria, appropriate for market participants with different points of view.

Optimization Criterion 1: *Given a family of trading strategies, a risk-reward factor A and a time horizon ΔT, find the cheapest strategy for which the given risk-reward ratio is achieved.*

Optimization Criterion 2: *Given a family of trading strategies, an implied volatility characterized by the value of the Leland volatility adjustment Λ and a time horizon ΔT, find the strategy that maximizes the risk-reward ratio A.*

Optimization criteria 1 and 2 are complimentary. Criterion 1 corresponds to the problem faced by a market-maker in a competitive options market. The market-maker wishes to offer the option at the lowest possible price that guarantees a certain level of compensation with respect to the risk undertook over the chosen time horizon. In the process, the market-maker identifies the strategy and parameters that minimize the offering price for the option.

On the other hand, a price-taker would employ optimization criterion 2. Given the option value, and hence the implied Leland volatility a price-taker would like to find the strategy and the parameters that maximize the risk-reward factor over the time horizon.

Instead of trying to solve optimization criteria 1 and 2 in the space of all allowed trading strategies, we limit our search to an one-parameter family of trading strategies that reflect common industry practice. In addition, the chosen strategies have the common characteristic that, in the limit of small transaction costs, the width of the no-transaction interval and the surcharge over the Black-Scholes price shrinks to zero, the frequency of trading tends to infinity, and the strategies converge to the Black-Scholes strategy. Such convergence is far from obvious in the space of all allowed strategies, as indicated by the paper by Soner, Shreve, and Cvitanić (1995). For example, given our objective, had we not constrained the strategies to trade back to the Black-Scholes position, it would have been possible, for some parameter value to optimally leave a derivative position unhedged, even with zero transaction costs. The family of strategies is characterized by the choice of times when trading occurs. At the times of trade, the portfolio is rebalanced back to the replication target as specified by the Leland equation. We denote the family as *deterministic time hedging*, where trading occurs at fixed time intervals; the parameter corresponds to the time between trades.

In addition to the deterministic time hedging strategy it is possible to apply our framework to families of strategies where the times of trade are random, and determined by the crossing of a threshold. Such families of strategies include the *delta hedging* strategies, where trading occurs whenever the difference in deltas; i.e., the number of shares in the portfolio between the replicating portfolio and the target, exceeds a certain threshold in absolute value; and the *portfolio-value-mismatch hedging* strategies, where trading occurs whenever the mismatch between the value of the replicating portfolio and the target exceeds a certain threshold.

3. Asymptotic analysis in the limit of small transaction costs

The asymptotic analysis in the limit of small transaction costs provides information on the scaling relations between the time interval between trades, the magnitude of the transaction cost, the Leland volatility adjustment and the sensitivities of the replication target such as the current Γ. The proofs of the theorems in this section are provided in the appendix. All our results are asymptotic in the limit of small transaction costs. The relation between the length of the time interval between trades and the time horizon of the hedger should be such that a large number of trades takes place within the time horizon:

$$\Delta T \gg \tau \tag{7}$$

where ΔT is the time horizon of the hedger and τ the time between successive trades. Combined with the results we describe below, Equation (7) is valid when

$$k \ll \sqrt{\pi} A \sigma$$

where k is the relative transaction cost, A is the risk-reward parameter defined in Equation (6), and σ the volatility defined in Equation (1).

For the case of deterministic time hedging we have the following result:

Theorem 3.1. *Under optimization criterion 1, the optimal length of the time interval between trades and the optimal Leland volatility adjustment, for the deterministic time hedging strategy, are given by*

$$\tau^* = \frac{k}{\sqrt{\pi}A\sigma} + o(k), \quad \Lambda^* = 2\sqrt{\frac{2Ak}{\sqrt{\pi}\sigma}} + o(k) \approx 2.13\sqrt{\frac{Ak}{\sigma}} + o(k)$$

Under optimization criterion 2, the optimal length of the time interval between trades and the maximum risk-reward ratio are given by

$$\tau^* = \frac{8k^2}{\pi\sigma\Lambda^2} + o\left(k^2/\Lambda^2\right), \quad A^* = \frac{\Lambda^2\sigma\sqrt{\pi}}{8k} + o(\Lambda^2/k) \approx 0.222\frac{\Lambda^2\sigma}{k} + o(\Lambda^2/k)$$

The above result holds for both puts and calls. It suggests that, from the point of view of a market-maker, to leading order in the roundtrip transaction cost the optimal length of the time interval between trades is proportional to the size of the roundtrip transaction cost and inversely proportional to the risk-reward factor and the volatility. The surcharge over the Black-Scholes price is proportional to the square root of the roundtrip transaction cost, a reminder that the price of an option is not an analytic function at $k = 0$.

For a price-taker, the optimal length of the time interval between trades is proportional to the square of the roundtrip transaction cost and inversely proportional to the volatility and the square of the Leland volatility adjustment.

In results we do not report, we have also performed the same asymptotic analysis for the case of delta hedging and portfolio value mismatch hedging families of strategies. In both cases, under our optimization framework, the size of the no-trade interval is proportional to the square root of the size of transaction costs, and only depends on the magnitude of the relative roundtrip transaction cost, the volatility and risk-reward factor.[6]

Theorem 3.1 provides intuition regarding upper and lower bounds for bid and ask option prices in a market with transaction costs. Hedging a short and a long position of the same option on an underlying stock with proportional transaction costs leads to corrections to the Black-Scholes price proportional to the square root of the relative roundtrip transaction cost but with different signs. This is due to the fact that if the Γ of a long position is positive, the Γ of a short position, for the same payoff, is negative. For example, hedging a long position in a put, or a call, would result in a price lower than the Black-Scholes price, while hedging a short position in the same put or call would result in a price higher than the Black-Scholes price. This observation seems to indicate that the bid-ask spread for option prices would be proportional to the square root of the relative roundtrip transaction cost. However the argument is incomplete, and justified only in illiquid option markets where hedging is necessary to cover one's position. In a liquid market the short and long positions can be matched against each other and only the outstanding position has to be covered. In this case the bid-ask spread may be unrelated to the size of the relative roundtrip transaction costs for the underlying stock, but instead depend on factors such as liquidity, inventory cost, etc. Nonetheless, the shift of the option price with respect to the Black-Scholes price is proportional to the square root of the relative roundtrip transaction cost and the sign of the correction is related to whether the actively managed outstanding positions in a particular option are short or long. In a typical situation, investment banks would be the marginal investor, and would be short Γ and hedge their positions. This would suggest that option prices would be higher than those predicted by the Black-Scholes model.

[6] Details are available from the authors.

4. Numerical Results

4.1. Base case

To determine the range of validity of the asymptotic analysis of Section 3, we perform numerical simulations around a base-case of a six-month European call option with strike of \$100 on a stock that pays no dividends, with current price \$100. We undertook Monte Carlo simulations following the optimal hedging strategy, prescribed by the results of the analysis in Section 3. The time horizon of the option writer was set, for the base case, to one month, $\Delta T = 1/12$. The underlying stock follows geometric Brownian motion with drift $\mu = 9\%$ per year and volatility $\sigma = 20\%$ per year. The interest rate of the riskless bond is set to $r = 4\%$ per year. The relative transaction costs are 0.1%, which correspond to a bid-ask spread of \$0.10 for the current stock price of \$100. For each simulation we used 10,000 paths. The risk-reward factor was set to $J_{\text{desired}} = 1$. This value corresponds to a hedger that wishes to achieve, on average, gains equal to the standard deviation of the cumulative profit/loss over the time horizon; i.e, one month. Equivalently, for a Gaussian distribution of cumulative profit/loss the hedger wishes to make a profit 84% of the time over the one month time horizon.

For the base case, the Black-Scholes price is \$6.63, while the adjusted volatility is 22.6% and the adjusted option price \$7.35; i.e., a surcharge over the Black-Scholes price of approximately 11%. The optimal number of transactions over the hedger's time horizon is 102, which is approximately 5 trades a day. The average realized gain is 5.8 cents, with a standard deviation of 6.2 cents. While the level of transaction costs is relatively low, there are indications of second order effects, beyond the asymptotic analysis of Section 3. In particular, the realized level of the risk-reward factor was 93% instead of 1, and the distribution of cumulative profit and loss was different from a Gaussian distribution, with negative skewness with a skewness coefficient of -0.3, and fat tails, with a kurtosis coefficient of 3.4. The standard error, for the number of simulation paths used, was approximately 1% for the expected gain and for the value of the realized risk-reward factor, and 0.1 for the skewness and kurtosis coefficients.

4.2. Range of validity of asymptotic analysis

Since the results from Section 3 are valid in the limit when the relative transaction cost tends to zero, $k \to 0$, to check whether the results are also valid for transaction costs of finite size, we use numerical simulation and vary the level of transaction costs. For each value of the level of transaction costs we monitor both the realized risk-reward factor, and the deviations of the distribution of the cumulative profit and loss from the trading strategy from the Gaussian distribution. The results are presented in Table I.

We note that, while the deviations from the asymptotic analysis are obvious even for low levels of relative transaction costs, these deviations are not too large, even when transaction costs reach 1%, corresponding to bid-ask spreads of \$1. From the table it is clear that only at very low levels the higher order effects disappear. The number of transactions for relative transaction costs of 0.01% correspond to trading approximately every 10 minutes, while for relative transaction costs of 1% trading occurs approximately every 2 days.

As expected, the second order effects become more obvious as the relative transaction cost increases. The difference between the realized risk-reward factor, J_{real}, and the desired one, J_{desired} (which was set to one) is an indication of the magnitude of the higher-order corrections with respect to the leading order term. We also note that even though the surcharge over the Black-Scholes price of the option is significant, and ranges from 4% to 31%, the optimal trading strategies produce very small expected cumulative gains over the period of a month, ranging from 2 to 15 cents.

4.3. Comparative statics

In addition to varying the magnitude of the relative roundtrip transaction costs, we have performed a study of comparative statics over the base case, by varying the drift and volatility, the strike price and the expiration date of the option, and the time horizon of the hedger. The results of this analysis are presented in Table II.

From the table we notice that the volatility and the strike of the option have the biggest impact, while the option expiration date, the time horizon of the hedger and the drift of the underlying stock have only a minor influence. The results indicate that the surcharge in the option price decreases as volatility increases (from a 13% surcharge when volatility is 10% to a 9% surcharge when volatility is 40%), while the magnitude of the risk and reward increase (from approximately 4 cents when volatility is 10% to approximately 8 cents when volatility reaches 40%). From varying the strike we notice that deep in the money options have the smallest surcharge (approximately 1% for a strike of 80) and the distribution of gains is closest to a normal distribution. As the option becomes at the money and out of the money, the surcharge in the option price increases (11% at a strike of 100, and 47% at a strike of 120), while the tails of the distribution of cumulative profit and loss become fatter. Following the discussion on price bounds for option prices in Section 3, we point out that if certain strikes are more liquid than others there would be an effect on the implied bid and ask volatilities, with the bid-ask spread widening for the illiquid strikes. On the other hand our results do not necessary imply a smile pattern in implied volatilities, since additional information on the identity and preferences of the marginal investor would be necessary.

From all the results in Table II, we notice that the size of the average gains increases with the Γ of the position, and is greatest for strikes at the money, and for larger volatilities.

5. Conclusions and Summary

We have presented an extension of the framework proposed by Leland (1985) for pricing derivatives in the presence of small transaction costs, and of the good-deal pricing literature for incomplete markets. Within our framework, we define optimization criteria appropriate for market-makers and for price takers. Our approach is an alternative to utility maximization and allows for the incorporation of common market practices, such as marking-to-market and investor-set time horizons.

In the case of small transaction costs we have carried out an asymptotic analysis and characterized the optimal length of the time interval between trades in terms of local sensitivity parameters (such as the option Γ), the volatility and the investor-specific parameters such as time horizon, and desired investor risk-reward ratio.

A. One period risk-reward analysis

We will assume that ΔT is small enough so that the leading order calculation for the gain and reward functions $G(S,t)$ and $R(S,t)$ is not affected by the variation of S and Γ over the time horizon.[7] Since returns over each hedging cycle are i.i.d. random variables, thanks to the central limit theorem, the risk and reward functions are proportional to the one period analogs $G_0(S,t)$ and $R_0(S,t)$. More precisely we have the following

Lemma A.1. *For any stopping time τ and to leading order in $E[\tau]$ as $E[\tau] \to 0$, we have*

$$G(S,t) = \left(\frac{\Delta T}{E[\tau]} + O\left(\sqrt{\frac{\Delta T}{E[\tau]}} \right) \right) \cdot G_1(S,t)$$

and

$$R(S,t) = \left(\sqrt{\frac{\Delta T}{E[\tau]}} + O(1) \right) \cdot R_1(S,t)$$

B. Proof of Lemma 2.1

Suppose that at time t_0 the replicating portfolio is hedged to the Black-Scholes point. To the leading order in τ as $\tau \to 0$, we have that

$$E\left[\frac{k}{2} S_0 \left| \delta a (t_0 + \tau) \right| \right] = \frac{k S_0^2}{2} \left| \frac{\partial^2 f}{\partial S^2}(S_0, t_0) \right| E\left[\left| W\left(\sigma^2 \tau\right) \right| \right] + O(\tau)$$

where $W(t)$ is the standard Wiener process. The probability distribution function for the random variable $w = W(\sigma^2 \tau)$ is

$$P(\tau, w) = \frac{1}{\sqrt{2\pi\sigma^2\tau}} \exp\left(-\frac{w^2}{2\sigma^2\tau} \right).$$

[7] In general, one needs to take into account the correlation between successive values of Γ. However, for small values of ΔT, Γ will be largely unchanged and the application of the central limit theorem goes through.

Hence

$$E\left[\left|W\left(\sigma^2\tau\right)\right|\right]=\sqrt{\frac{2\pi\sigma^2}{\tau}}\int_{-\infty}^{\infty}dw\,|w|\exp\left(-\frac{w^2}{2\sigma^2\tau}\right)=\sqrt{\frac{2\sigma^2\tau}{\pi}}$$

and we find

$$E\left[\frac{k}{2}S_0\left|\delta a\left(t_0+\tau\right)\right|\right]=\sqrt{\frac{\tau}{2}}\;\left(k\sigma S_0^2\right)\;\left|\frac{\partial^2 f}{\partial S^2}\left(S_0,t_0\right)\right|. \tag{8}$$

By summing over all the trades in a fixed time interval, we obtain the estimate in Lemma 2.1.

C. Proof of Lemma 2.2

We have that

$$\begin{aligned}\delta\Pi(\tau)&=a_L(t_0)S(t_1)+a_L(t_0)\int_{t_0}^{t_1}S(t)\delta dt+b_L(t_0)B(t_1)-f_L(S(t_1),t_1)\\&=b_L(t_0)B(t_0)r\tau+a_L(t_0)\delta S+a_L(t_0)S_0\delta\tau+f_L(S_0,t_0)-f_L(S_0+\delta S,t_0+\tau)+O(\tau^2+\delta S^3+\tau\delta S)\\&=\left(rf_L(S_0,t_0)-(r-\delta)a_L(t_0)S_0-\frac{\partial f_L}{\partial t}(S_0,t_0)\right)\tau-\frac{1}{2}\frac{\partial^2 f_L}{\partial S^2}(S_0,t_0)\delta S^2+O(\tau^2+\delta S^3+\tau\delta S)\end{aligned}$$

where $\delta S\equiv S\left(t_1\right)-S\left(t_0\right)$. Using the modified Black-Scholes Equation (4) for the price $f_L(S,t)$, we have

$$\delta\Pi(\tau)=\frac{1}{2}\Gamma_0S_0^2\left(\bar{\sigma}^2\tau-\left(\frac{\delta S}{S_0}\right)^2\right)+O(\tau^2+\delta S^3+\tau\delta S). \tag{9}$$

In the Black-Scholes limit, where $\Lambda=0$, the expected value of $\delta\Pi(\tau)$ is zero. The positive bias that a positive Λ introduces is given by

$$E\left[\delta\Pi\left(\tau\right)\right]=\frac{\Gamma_0S_0^2}{2}\Lambda\sigma^2E\left[\tau\right]+O\left(\tau^2+\delta S^3+\tau\delta S\right)$$

D. Deterministic time strategies and Theorem 3.1

Consider a class of strategies according to which one places a trade at fixed time intervals, which are determined based on information available either at the strategy inception date or at the date of the previous trade. From Equation (9) we have that

$$G_1(S_0,t_0) = \frac{\Lambda}{2} S_0^2 \Gamma_0 \sigma^2 \tau - \frac{k}{2} S_0 \left|\delta a(\tau)\right|$$

The expected transaction cost is given, to leading order, from Equation (8), as

$$E\left[\frac{k}{2} S_0 \left|\delta a(t_0+\tau)\right|\right] = \sqrt{\frac{\tau}{2}}\ \left(k\sigma S_0^2\right)\ \left|\frac{\partial^2 f}{\partial S^2}(S_0,t_0)\right|.$$

The risk is given by

$$R_1(S_0,t_0) = \sqrt{E\left[(\delta\Pi)^2\right] - E\left[\delta\Pi\right]^2}.$$

We have that

$$\begin{aligned}
E\left[\delta\Pi\right]^2 &= \frac{1}{4}\Gamma_0^2 S_0^4 E\left[\left(\bar{\sigma}^2\tau - \left(\frac{\delta S}{S_0}\right)^2 - \Lambda\sigma^2\tau\right)^2\right] \\
&= \frac{\Gamma_0^2 S_0^4}{2\sqrt{2\pi\sigma^2\tau}} \int_0^\infty dw \left(\bar{\sigma}^2\tau - w^2 - \Lambda\sigma^2\tau\right)^2 \exp\left(-\frac{w^2}{2\sigma^2\tau}\right) \\
&= \frac{1}{2}\Gamma_0^2 S_0^4 \sigma^4 \tau^2
\end{aligned}$$

Hence the risk; i.e., the standard deviation of gains, is given by

$$R_1(S_0,t_0) = \frac{\Gamma_0 S_0^2 \sigma^2 \tau}{\sqrt{2}}.$$

The risk-reward constraint amounts to the following equation:

$$A = \frac{G_1(S,t)}{\sqrt{\tau} R_1(S,t)} = \frac{\Lambda}{\sqrt{2\tau}} - \frac{k}{\sqrt{\pi}\sigma\tau}$$

Hence, we find

$$\Lambda = A\sqrt{2\tau} + \frac{k}{\sigma}\sqrt{\frac{2}{\pi\tau}}. \tag{10}$$

The optimal value of τ, under optimization criterion 1, is the one for which the Leland volatility adjustment is minimum. i.e.

$$\tau^* = \frac{k}{\sqrt{\pi}A\sigma}$$

Moreover, the optimal Leland volatility adjustment is

$$\Lambda^* = 2\sqrt{\frac{2Ak}{\sqrt{\pi}\sigma}} \approx 2.12\sqrt{\frac{Ak}{\sigma}}.$$

Under optimization criterion 2, instead of choosing the time interval τ in Equation (10) to minimize the Leland volatility adjustment, Λ, we fix Λ and choose τ to maximize the risk-reward factor A. Theorem 3.1 follows.

References

Bensaid, Bernard, Jean-Philippe Lesne, Henri Pagés, and José Scheinkman, 1992, Derivative Asset Pricing with Transaction Costs, *Mathematical Finance* 2, 63–86.

Bernardo, Antonio E., and Olivier Ledoit, 2000, Gain, Loss and Asset Pricing, *Journal of Political Economy* 108, 144–172.

Bertsimas, Dimitris, Leonid Kogan, and Andrew W. Lo, 2000, When is Time Continuous?, *Journal of Financial Economics* 55, 173–204.

Black, Fisher, and Myron Scholes, 1973, The Pricing of Options and Corporate Liabilities, *Journal of Political Economy* 81, 637–654.

Boyle, Phelim P., and D. Emanuel, 1980, Discretely Adjusted Option Hedges, *Journal of Financial Economics* 8, 259–282.

Boyle, Phelim P., and Ton Vorst, 1992, Option Replication in Discrete Time with Transaction Costs, *Journal of Finance* 47, 271–293.

Clewlow, Les, and Steward D. Hodges, 1997, Optimal Delta-Hedging under Transaction Costs, *Journal of Economic Dynamics and Control* 21, 1353–1376.

Cochrane, John H., and Jesus Saa-Requejo, 2000, Beyond Arbitrage: Good deal asset price bounds in incomplete markets, *Journal of Political Economy* 108, 79–119.

Davis, Mark H. A., and A. Norman, 1990, Portfolio Selection with Transaction Costs, *Math. Operations Research* 15, 676–713.

Davis, Mark H. A., V. G. Panas, and Thaleia Zariphopoulou, 1993, European Option Pricing with Transaction Costs, *SIAM J. Control and Optimization* 31, 470–493.

Edirisinghe, C., V. Naik, and R. Uppal, 1993, Optimal Replication of Options with Transaction Costs and Trading Restrictions, *Journal of Financial and Quantitative Analysis* 28, 117–138.

Figlewski, Steven, 1989, Options Arbitrage in Imperfect Markets, *Journal of Finance* 44, 1289–1311.

Gilster, John E., Jr., 1990, The Systematic Risk of Discretely Rebalanced Option Hedges, *Journal of Financial and Quantitative Analysis* 25, 507–516.

Gilster, John E., Jr., 1996, Optimal Pricing Theory: Is Risk Free Hedging Feasible?, Michigan State University working paper.

Henrotte, Philippe, 1993, Transaction Costs and Duplication Strategies, Working paper, Stanford University.

Hodges, S. D., and A. Neuberger, 1989, Optimal Replication of Contingent Claims under Transactions Costs, *Review of Futures Markets* 8, 222–239.

Leland, Hayne, 1985, Option Prices with Transaction Costs, *Journal of Finance* 40, 1283–1301.

Merton, Robert C., 1973, Theory of Rational Option Pricing, *Bell Journal of Economics and Management Science* 4, 141–183.

Rogers, L. C. G., 2000, Why is the effect of proportional transaction costs $O(\delta^{2/3})$?, preprint, University of Bath.

Soner, H. M., S. E. Shreve, and J. Cvitanić, 1995, There is no nontrivial hedging portfolio for option pricing with transaction costs, *Annals of Applied Probability* 5, 327–355.

Toft, Klaus B., 1996, On the Mean-Variance Tradeoff in Option Replication with Transactions Costs, *Journal of Financial and Quantitative Analysis* 31, 233–263.

Table I
Validity of Asymptotic Analysis

This table examines the range of validity of the asymptotic analysis of Section 3. The option priced is a six-month European call, with strike price \$100, current stock price \$100, volatility $\sigma = 20\%$, interest rate $r = 4\%$ per year, rate of growth $\mu = 9\%$, and time horizon of the hedger $\Delta T = 1$ month. Trans. is the number of transactions, Adj. Vol. is the adjusted volatility, BS Price is the Black-Scholes price, Adj. Price is the adjusted price, Gain is the average gain of the strategy, Risk is the standard deviation of the gains, Skew and Kurt. are the skewness and kurtosis of the distribution of the cumulative profit and loss from the trading strategy, and J_{real} the realized risk-reward ratio.

Trans. Cost	Trans.	Adj. Vol.	BS Price	Adj. Price	Gain	Risk	Skew	Kurt.	J_{real}
0.01%	1023	20.9%	\$6.63	\$6.87	\$0.020	\$0.020	-0.1	3.1	97%
0.05%	205	21.9%	\$6.63	\$7.15	\$0.042	\$0.045	-0.3	3.2	94%
0.10%	102	22.6%	\$6.63	\$7.35	\$0.058	\$0.062	-0.4	3.2	93%
0.25%	41	24.0%	\$6.63	\$7.74	\$0.083	\$0.096	-0.4	3.1	87%
0.50%	20	25.5%	\$6.63	\$8.15	\$0.113	\$0.132	-0.5	3.4	86%
0.75%	14	26.6%	\$6.63	\$8.44	\$0.129	\$0.157	-0.7	3.8	82%
1.00%	10	27.5%	\$6.63	\$8.69	\$0.152	\$0.185	-0.8	3.9	82%

Table II
Comparative Statics

This table presents comparative statics from a base case of a six-month European call option, with strike price \$100, current stock price \$100, volatility $\sigma = 20\%$, interest rate $r = 4\%$ per year, rate of growth $\mu = 9\%$, relative transaction costs of 0.1%, and time horizon of the hedger $\Delta T = 1$ month. Trans. is the number of transactions, Adj. Vol. is the adjusted volatility, BS Price is the Black-Scholes price, Adj. Price is the adjusted price, Gain is the average gain of the strategy, Risk is the standard deviation of the gains, Skew and Kurt. are the skewness and kurtosis of the distribution of the cumulative profit and loss from the trading strategy, and J_{real} the realized risk-reward ratio.

Volatility	Trans.	Adj. Vol.	BS Price	Adj. Price	Gain	Risk	Skew	Kurt.	J_{real}
10%	51	11.8%	\$3.89	\$4.38	\$0.037	\$0.043	-0.4	3.3	88%
20%	102	22.6%	\$6.63	\$7.35	\$0.058	\$0.062	-0.4	3.2	93%
30%	153	33.2%	\$9.39	\$10.29	\$0.072	\$0.076	-0.3	3.0	95%
40%	205	43.8%	\$12.15	\$13.19	\$0.082	\$0.089	-0.3	3.3	92%

Strike	Trans.	Adj. Vol.	BS Price	Adj. Price	Gain	Risk	Skew	Kurt.	J_{real}
80	102	22.6%	\$21.80	\$21.98	\$0.016	\$0.019	0.0	4.0	85%
90	102	22.6%	\$13.15	\$13.64	\$0.040	\$0.044	-0.3	3.4	90%
100	102	22.6%	\$6.63	\$7.35	\$0.058	\$0.062	-0.4	3.2	93%
110	102	22.6%	\$2.76	\$3.43	\$0.056	\$0.060	-0.3	3.2	94%
120	102	22.6%	\$0.96	\$1.40	\$0.038	\$0.044	-0.2	3.9	87%

Time horizon	Trans.	Adj. Vol.	BS Price	Adj. Price	Gain	Risk	Skew	Kurt.	J_{real}
0.5 months	72	23.1%	\$6.63	\$7.48	\$0.033	\$0.036	-0.4	3.2	91%
1.0 months	102	22.6%	\$6.63	\$7.35	\$0.058	\$0.062	-0.4	3.2	93%
1.5 months	125	22.4%	\$6.63	\$7.29	\$0.078	\$0.086	-0.4	3.3	90%
2.0 months	145	22.2%	\$6.63	\$7.24	\$0.099	\$0.106	-0.3	3.2	93%

Expiration	Trans.	Adj. Vol.	BS Price	Adj. Price	Gain	Risk	Skew	Kurt.	J_{real}
3 months	102	22.6%	\$4.49	\$5.00	\$0.082	\$0.091	-0.3	3.2	91%
6 months	102	22.6%	\$6.63	\$7.35	\$0.058	\$0.062	-0.4	3.2	93%
9 months	102	22.6%	\$8.38	\$9.26	\$0.047	\$0.050	-0.3	3.1	94%
1 year	102	22.6%	\$9.93	\$10.93	\$0.039	\$0.043	-0.3	3.1	92%

Drift	Trans.	Adj. Vol.	BS Price	Adj. Price	Gain	Risk	Skew	Kurt.	J_{real}
5%	102	22.6%	\$6.63	\$7.35	\$0.057	\$0.062	-0.3	3.2	93%
7%	102	22.6%	\$6.63	\$7.35	\$0.057	\$0.062	-0.3	3.1	93%
9%	102	22.6%	\$6.63	\$7.35	\$0.058	\$0.062	-0.4	3.2	93%
11%	102	22.6%	\$6.63	\$7.35	\$0.058	\$0.061	-0.3	3.2	95%
13%	102	22.6%	\$6.63	\$7.35	\$0.059	\$0.061	-0.3	3.1	96%
15%	102	22.6%	\$6.63	\$7.35	\$0.057	\$0.062	-0.4	3.4	93%