

Option Pricing with Transaction Costs Using a Markov Chain Approximation

Michael Monoyios*
Department of Economics and Finance
Brunel University
Uxbridge, Middx UB8 3PH, UK

June 5, 2000

Abstract

An efficient algorithm is developed to price European options in the presence of proportional transaction costs, using the optimal portfolio framework of Davis (1997). A fair option price is determined by requiring that an infinitesimal diversion of funds into the purchase or sale of options has a neutral effect on achievable utility. This results in a general option pricing formula, in which option prices are computed from the solution of the investor's basic portfolio selection problem, without the need to solve a more complex optimisation problem involving the insertion of the option payoff into the terminal value function. Option prices are computed numerically using a Markov chain approximation to the continuous time singular stochastic optimal control problem, for the case of exponential utility. Comparisons with approximately replicating strategies are made. The method results in a uniquely specified option price for every initial holding of stock, and the price lies within bounds which are tight even as transaction costs become large. A general definition of an option hedging strategy for a utility maximising investor is developed. This involves calculating the perturbation to the optimal portfolio strategy when an option trade is executed.

JEL classification: C61; G11; G13

Keywords: Option pricing; transaction costs; utility maximisation; singular stochastic control; Markov chain approximation

1 Introduction

The celebrated Black-Scholes (1973) option pricing methodology is not applicable in the presence of transaction costs on trading the underlying stock. This is

*e-mail: Michael.Monoyios@brunel.ac.uk. The author wishes to thank Bill Ziemba for helpful comments.

because it relies on the implementation of a continuously rebalanced portfolio strategy (to replicate the option payoff) that would be excessively costly in the face of any market frictions. Since perfect replication is not possible it is necessary to look for alternative criteria to determine fair valued derivative prices and, if possible, hedging strategies. There is an unavoidable trade-off between the transaction costs incurred in portfolio rebalancing and the tightness of the valuation bounds. Reducing the hedging error generally incurs large transaction costs and results in unrealistically wide valuation bounds. The question arises as to what is the optimal valuation and hedging policy, consistent with an investor's risk management objectives.

A number of methods have been suggested to solve this problem. The natural alternative to continuous hedging policies for options is to introduce a discrete timescale in which transactions take place. This is the course taken by Leland (1985) and Boyle and Vorst (1992). Both these models result in Black-Scholes type formulae for bounds on the value of the option, with an adjusted volatility which depends directly on the exogenously specified revision frequency. The hedging error is reduced if portfolio rebalancing occurs more frequently, but the pricing bounds become much wider. Moreover, restricting the form of the hedging strategy to portfolio rebalancing at fixed time intervals may not always be the optimum method of managing the risk from an option trade.

Bensaid, Lesne, Pagès and Scheinkman (1992), Edirisinghe, Naik and Uppal (1993) and Boyle and Tan (1994) replaced the replication strategy with a "super-replicating strategy" in which the hedging portfolio is only required to dominate, rather than replicate, the option payoff at maturity. In a discrete time setting it can sometimes be cheaper to dominate a contingent claim than to replicate it. However, as conjectured by Davis and Clark (1994) and subsequently proven by Soner, Shreve and Cvitanić (1995) and Cvitanić, Pham and Touzi (1999), in the continuous time limit the cheapest super-replicating strategy for a European call option is the trivial strategy of buying one share of the underlying stock and holding it to maturity. This illustrates a fundamental feature of option hedging under transaction costs, namely that eliminating all risk results in unrealistically wide valuation bounds.

Hodges and Neuberger (1989) recognised that a valuation method which incorporates some element of optimality, in the form of the agent's utility maximisation objective, is perhaps the most promising path to follow in designing effective option hedging policies. This point was also made by Dumas and Luciano (1991). By comparing the utility achieved with and without the obligations of an option contract, Hodges and Neuberger specified reservation bid and asking prices for an option by requiring that the same utility is achieved whether an option trade has been entered into or not. This approach was further developed by Davis, Panas and Zariphopoulou (1993), Clewlow and Hodges (1997) and Constantinides and Zariphopoulou (1999). Alternative criteria for determining an option hedging policy include quadratic criteria such as local risk minimisation, studied by Mercurio and Vorst (1997) and Lamberton, Pham and Schweizer (1998).

Davis (1997) suggested a utility maximisation approach to valuing options

in a general situation in which replication is either impossible or unfavourable. Davis derived a general option pricing formula for a “fair” option value, determined by balancing the marginal utility from diverting an infinitesimal fraction of the initial wealth into the options market with that from refraining from trading options and instead investing optimally in the underlying asset market.

In this paper we numerically study the optimal pricing procedure suggested by Davis (1997). We develop an associated definition of a hedging strategy implied by the new method, and compute option prices and hedging strategies under proportional transaction costs. Davis’ formula, which is reminiscent of a classical representative agent asset pricing formula, has not (until now) been tested numerically to see what implications it has for option values and hedging strategies in the presence of market frictions. Karatzas and Kou (1996) have applied this approach in frictionless markets to the situation in which there are constraints on the investor’s portfolio choices, whereas we consider markets that are otherwise complete, but option replication is rendered unfavourable by trading costs.

Our numerical results indicate that the new methodology inherits realistic and appealing features from the solution of the investor’s basic portfolio selection problem. The presence of frictions in a financial market qualitatively changes the nature of the optimal trading strategy, in that it specifies a region of the portfolio state space in which it is optimal not to transact, whilst outside this region the optimal action is to either buy or sell the risky asset. When the investor’s portfolio holdings lie outside the no transaction region, the derived option value is either the optimal bid or ask price. Alternatively, if the agent’s stock-bond holdings lie within the no transaction region, then the option value lies strictly within the bid-ask spread, and is uniquely specified given the initial portfolio holdings. Trading an option at any price within the optimal bid-ask interval results in a negligible effect on the investor’s utility. However, trading at a price outside the bid-ask interval drastically reduces (or increases, if the investor is on the favourable side of the trade) the achievable utility compared with the scenario in which an option is not traded.

We also make comparisons with the replication approach of Leland (1985). In general, we find that the optimal valuation method produces much tighter bounds on option prices than can be obtained by any strategy which attempts to eliminate risk by replication, and this effect is most marked for large transaction costs. In essence, the option trader is able to incorporate some risk in his options portfolio, which is unavoidable in the face of market frictions.

The pricing methodology has the desirable feature that one only has to solve the investor’s basic portfolio selection problem to price options. This is in contrast to previous attempts to embed the pricing problem in a utility maximisation framework, which require the solution of a much more difficult optimisation problem involving the credit or debit of an option payoff to the investor’s portfolio at the option maturity date. We develop an efficient algorithm to price options in our framework using a Markov chain approximation technique of the type pioneered by Kushner (1990).

The paper is organised as follows. In Section 2 we set up a portfolio selection

scenario, in which the optimal strategy to maximise expected utility of wealth at a finite horizon time is sought. We also formulate the alternative optimisation problem which results if a small amount of the initial wealth is diverted into the purchase or sale of European options. We illustrate the various ways that such stochastic optimal control problems can be used to specify an option price, and we state Davis' (1997) general option pricing formula. Then we give a general definition of an option hedging strategy for a utility maximising investor who has the choice of diverting some initial wealth into the options market. In Section 3 we consider a specific market model with transaction costs. We illustrate the properties of the solution to the portfolio choice problem, and of Davis' pricing formula. In Section 4 we specialise to the case of an exponential utility function, and develop an efficient numerical option pricing algorithm. In Section 5 we present numerical solutions for option prices and hedging strategies. Section 6 concludes and suggests directions for further research. An Appendix contains a derivation of some results used in the implementation of our numerical algorithm.

2 Portfolio Selection and Option Valuation

We shall utilise a finite time interval $[0, T]$, where T will correspond to the maturity of a European option. Consider an investor with concave utility function U , starting at time $t \in [0, T]$ with cash endowment x , and holding y shares of a stock whose price is S . The investor trades a dynamic portfolio whose value at time $u > t$ is $W_{t,S,x,y}^\pi(u)$ when he or she uses the trading strategy π and starts in the state (t, S, x, y) . The wealth $W_{t,S,x,y}^\pi(u)$ consists of $X_{t,S,x,y}^\pi(u)$ dollars in cash and $Y_{t,S,x,y}^\pi(u)$ shares of stock, whose price at time u is $S(u)$, so that

$$W_{t,S,x,y}^\pi(u) = X_{t,S,x,y}^\pi(u) + Y_{t,S,x,y}^\pi(u)S(u). \quad (1)$$

The investor's objective is to maximise expected utility of wealth at time T . We assume that the investor does not consume any wealth, which is a reasonable assumption for a trader in a financial institution. Denote the investor's maximum utility by

$$V(t, S, x, y) = \sup_{\pi} \mathbf{E}_t[U(W_{t,S,x,y}^\pi(T))], \quad (2)$$

where \mathbf{E}_t denotes the expectation operator conditional on the time- t information. The supremum in (2) is taken over a suitable set of admissible policies, to be described in the next section, when we specialise to a market with proportional transaction costs.

Consider the alternative optimisation problem which results if a small amount of the initial wealth is diverted into the purchase or sale of a European option whose payoff at time T is some non-negative random variable $C(S(T))$. To be precise, if the option price at time t is p and an amount of cash δ is diverted at this time into options, we define

$$V^{(o)}(t, S, x - \delta, y, \delta, p) = \sup_{\pi} \mathbf{E}_t \left[U \left(W_{t,S,x-\delta,y}^\pi(T) + \frac{\delta}{p} C(S(T)) \right) \right]. \quad (3)$$

The “o” superscript denotes that the investor’s portfolio at time T incorporates the option payoff. The value function in (3) is evaluated for the initial cash endowment $x - \delta$ to signify that the funds to buy (or sell, if $\delta < 0$) the options have come from (or been credited to) the initial wealth. The value function for initial cash x is

$$V^{(o)}(t, S, x, y, \delta, p) = \sup_{\pi} \mathbf{E}_t \left[U \left(W_{t,S,x,y}^{\pi}(T) + \frac{\delta}{p} C(S(T)) \right) \right], \quad (4)$$

namely the maximum utility that is achieved if the investor’s account is credited with the payoff of δ/p options at time T , without any corresponding adjustment to the initial cash position. In (2)–(4) the quantities δ and p would be measurable with respect to the time- t information.

Clearly $V^{(o)}(t, S, x, y, 0, \cdot) = V(t, S, x, y)$, and the reader will note the dependence in (3) on the volume of option contracts traded. There are a number of ways that the above optimisation problems can be used to specify valuation bounds or a price for the option. A reservation writing price $p_w(t, S, x, y)$ for the initial holdings (x, y) is defined as the minimum value at which the investor is prepared to write the claim, and so satisfies

$$V^{(o)}(t, S, x + p_w(t, S, x, y), y, -p_w(t, S, x, y), p_w(t, S, x, y)) = V(t, S, x, y), \quad (5)$$

since the same utility is achieved when selling the option for $p_w(t, S, x, y)$ as is achieved by not writing the option. We further define the *universal reservation writing price* $\bar{p}_w(t, S)$ as the maximum of reservation write prices across all initial holdings (x, y) . Therefore $\bar{p}_w(t, S)$ satisfies

$$V^{(o)}(t, S, x + \bar{p}_w(t, S), y, -\bar{p}_w(t, S), \bar{p}_w(t, S)) \geq V(t, S, x, y). \quad (6)$$

Inequality (6) guarantees that the writer will be willing to write the option at any price higher than $\bar{p}_w(t, S)$, independently of his current portfolio position.

Similarly, a reservation buying price $p_b(t, S, x, y)$ for the initial holdings (x, y) is

$$V^{(o)}(t, S, x - p_b(t, S, x, y), y, p_b(t, S, x, y), p_b(t, S, x, y)) = V(t, S, x, y). \quad (7)$$

The universal reservation buying price $\underline{p}_b(t, S)$ is the minimum of reservation buying prices across all initial holdings, and so satisfies

$$V^{(o)}(t, S, x - \underline{p}_b(t, S), y, \underline{p}_b(t, S), \underline{p}_b(t, S)) \geq V(t, S, x, y). \quad (8)$$

As pointed out by Constantinides and Zariphopoulou (1999), transaction prices of the option must lie in the interval $[\underline{p}_b, \bar{p}_w]$. For, if an option trader writes the option for a price $p > \bar{p}_w$ then the buyer is acting suboptimally, as he could have found a willing writer of the option at a price as low as \bar{p}_w . Likewise, if a transaction occurs at a price $p < \underline{p}_b$, then the writer is acting suboptimally.

Hodges and Neuberger (1989) initiated the above approach to finding option valuation bounds. They considered the special case of a risk-neutral world where

the stock return rate is equal to the riskless interest rate. In this case, and in the absence of any option transaction, the investor would choose not to invest in the stock at all. Therefore, when an option is written or purchased, the optimal trading strategy specified by (5) or (7) would consist of the incremental trades generated by the option transaction, i.e. the option hedging strategy. Hodges and Neuberger also restricted their model to the case where the agent's initial holdings in stock and cash are zero. Their method is an ingenious way of focusing attention on the hedging issue. In this paper we relax their assumptions to allow for a richer interaction between the hedging strategy and the agent's initial portfolio. Davis, Panas and Zariphopoulou (1993) and Constantinides and Zariphopoulou (1999) have also considered the issue of finding reservation prices for European options, without the assumption of risk-neutrality, and the latter derive analytic bounds on option values by embedding the pricing problem in an infinite horizon portfolio selection framework.

One way of determining a unique price for an option under market imperfections was proposed by Davis (1997). Davis suggests that an agent will be willing to trade the option at a "fair" price \hat{p} , such that there is a neutral effect on the investor's utility if an infinitesimal fraction of the initial wealth is diverted into the purchase or sale of the option at price \hat{p} . That is, \hat{p} is given by the solution of

$$\left. \frac{\partial V^{(o)}}{\partial \delta}(t, S, x - \delta, y, \delta, \hat{p}) \right|_{\delta=0} = 0. \quad (9)$$

This results in the pricing formula

$$\hat{p}(t, S, x, y) = \frac{\mathbf{E}_t[U'(W_{t,S,x,y}^{\pi^*}(T))C(S(T))]}{V_x(t, S, x, y)}, \quad (10)$$

where U' is the derivative of U , $V_x(t, S, x, y)$ denotes the partial derivative with respect to x , and π^* denotes the trading strategy which maximises the expected utility in (2). This is the trading strategy which optimises a portfolio *without* options, and the formula (10) for \hat{p} shows no dependence on the optimisation problems (3) and (4) containing embedded options.

The formula in (10) is the central subject of investigation in this article. Its form is reminiscent of a classical representative agent asset pricing formula (see, for example, Duffie (1996)). Our goal is to see how (10) translates into numerical evaluations of option prices under transaction costs. It is evident from (10) that, at the price \hat{p} , the investor is balancing the marginal gain from diverting wealth into options with that from refraining from option trading and instead investing optimally in the stock market.

We write (10) as

$$\hat{p}(t, S, x, y) = \frac{F(t, S, x, y)}{V_x(t, S, x, y)}, \quad (11)$$

where the function $F(t, S, x, y)$ is defined by

$$F(t, S, x, y) \equiv \mathbf{E}_t[U'(W_{t,S,x,y}^{\pi^*}(T))C(S(T))]. \quad (12)$$

Because the optimal wealth vector in (12) and (2) are the same, one only needs to solve the investor’s basic portfolio selection problem to calculate $F(t, S, x, y)$ and $\hat{p}(t, S, x, y)$. When transaction costs are charged on trading the stock this is a highly nontrivial simplification of the option pricing problem compared to the “reservation price” approach of (5)–(8), because the latter requires the solution of the optimisation problem with the option payoff for a large range of possible option prices.

Remark 1 *In the absence of transaction costs the prices p_w , p_b and \hat{p} all reduce to the Black-Scholes option price.*

2.1 Hedging

As well as finding sensible derivative prices under transaction costs, any feasible pricing methodology should say something concerning the risk management of an option position. In the case of zero transaction costs the answer to this question is automatic, in that the Black-Scholes methodology sets option prices by a hedging argument. Such comments also apply to imperfectly replicating approaches like that of Leland (1985), and to quadratic approaches like the local risk minimisation approach in Lamberton, Pham and Schweizer (1998).

In the case of a utility maximisation approach to option pricing, the situation is somewhat different, in that the pricing problem is first embedded into the utility maximisation problem to determine a price according to one of the above methods. Then one computes the optimal trading strategy in the presence of the option trade, which will be altered compared to the situation without the option, with the adjustment measuring the effect of the option trade. This adjustment will correspond to what is usually meant by an “option hedging strategy”.

Suppose ϵ options are written at price \hat{p} given by(10). Then the investor’s optimal trading strategy will be π^\dagger maximising

$$\mathbf{E}_t \left[U \left(W_{t,S,x+\epsilon\hat{p},y}^\pi(T) - \epsilon C(S(T)) \right) \right].$$

In the absence of the option trade, the investor’s optimal trading strategy is π^* to achieve the supremum in (2). Since the option trade has altered the investor’s optimal stock trading strategy, a natural definition of the “option hedging strategy” is the incremental trades generated by the option trade, that is, the difference of the trading strategies π^\dagger and π^* . This motivates the definition which follows below.

Let an amount δ be paid (or received, for the case when options are written) to trade options at time t for a given price p . We then write the value function in (3) as

$$V^{(o)}(t, S, x - \delta, y, \delta, p) = \mathbf{E}_t \left[U \left(W_{t,S,x-\delta,y}^{\pi^\dagger}(T) + \frac{\delta}{p} C(S(T)) \right) \right], \quad (13)$$

which defines the optimal trading strategy π^\dagger for this utility maximisation problem. If we compare the optimal portfolio in the presence of the option position

with that in the absence of the options, we obtain a measure of the additional holdings brought about by the option trade, which is a natural candidate for the option hedging strategy.

Definition 1 *The hedging strategy π^h for δ/p options traded at time t , each at price p , is one whose holdings $X_{t,S,x,y}^{\pi^h}(u)$, $Y_{t,S,x,y}^{\pi^h}(u)$ at time $u \in [t, T]$ satisfy*

$$X_{t,S,x,y}^{\pi^h}(u) = X_{t,S,x,y}^{\pi^\dagger}(u) - X_{t,S,x,y}^{\pi^*}(u), \quad (14)$$

$$Y_{t,S,x,y}^{\pi^h}(u) = Y_{t,S,x,y}^{\pi^\dagger}(u) - Y_{t,S,x,y}^{\pi^*}(u). \quad (15)$$

The hedging strategy can be written as $\pi^h = \pi^\dagger - \pi^$.*

In practice, a utility maximising investor would simply calculate an optimal trading strategy π^\dagger including the option trade, and would not directly calculate the hedging strategy π^h . Nevertheless, we shall see that the above definition of a hedging strategy is a correct one, when we illustrate its features in our numerical results.

3 A Market with Transaction Costs

We consider a market consisting of a riskless bond and a risky stock whose prices $B(u)$ and $S(u)$ at time $u \in [0, T]$ satisfy, in continuous time

$$dB(u) = rB(u)du, \quad (16)$$

$$dS(u) = S(u)[bdu + \sigma dZ(u)], \quad (17)$$

where $Z = \{Z(u), 0 \leq u \leq T\}$ is a one-dimensional standard Brownian motion defined on a complete probability space $(\Omega, \mathcal{F}, \mathbf{P})$. Denote by $\mathbf{F} = \{\mathcal{F}(u), 0 \leq u \leq T\}$ the \mathbf{P} -augmentation of the filtration $\mathcal{F}^Z(T) = \sigma(Z(u); 0 \leq u \leq T)$ generated by Z . The constant coefficients r, b, σ represent the riskless interest rate, stock growth rate, and stock volatility respectively. The stock is assumed to pay no dividends. Trading in the stock incurs proportional transaction costs, such that the purchase of ν shares of stock at price S reduces the wealth held in the bond by $(1 + \lambda)\nu S$, where λ ($0 \leq \lambda < 1$) represents the proportional transaction cost rate associated with buying stock. Similarly, the sale of ν shares of stock increases the wealth in the bond by $(1 - \mu)\nu S$, where μ ($0 \leq \mu < 1$) represents the proportional transaction cost rate associated with selling stock. In all other respects assume markets are “perfect”. Securities are infinitely divisible, the volatility σ is known, there is no limit on borrowing or lending at the same riskless rate, and there are no taxes or constraints on short selling with full use of the proceeds.

We shall also make use of a binomial approximation of the above market model (using a Cox-Ross-Rubinstein (1979) stock price tree, as modified by He (1990)), for numerically computing option prices and hedging strategies. The bond and stock prices follow the discrete time processes

$$B(u) + \delta B(u) \equiv B(u + \delta u) = \exp(r\delta u)B(u), \quad (18)$$

$$S(u) + \delta S(u) \equiv S(u + \delta u) = \omega S(u), \quad (19)$$

where ω is a binomial random variable:

$$\omega = \exp[(b - \sigma^2/2)\delta u \pm \sigma\sqrt{\delta u}], \quad \text{each with probability } q = \frac{1}{2}, \quad (20)$$

and δu is a small time interval.

Define $(L(u), M(u))$, a pair of \mathbf{F} -adapted, right-continuous, non-decreasing processes, such that $L(u)$ (respectively $M(u)$) is the cumulative number of shares of stock bought (respectively, sold) up to time u . Then in continuous time the wealth held in the bond, for an investor who begins trading in the state (t, S, x, y) , evolves as

$$dX(u) \equiv dX_{t,S,x,y}^{L,M}(u) = rX(u)du - (1+\lambda)S(u)dL(u) + (1-\mu)S(u)dM(u). \quad (21)$$

The number of shares held follows the process

$$dY(u) \equiv dY_{t,S,x,y}^{L,M}(u) = dL(u) - dM(u), \quad (22)$$

and the wealth of the investor is given by

$$W(u) \equiv W_{t,S,x,y}^{L,M}(u) = X_{t,S,x,y}^{L,M}(u) + Y_{t,S,x,y}^{L,M}(u)S(u). \quad (23)$$

The pair $(L, M) \equiv \{(L(u), M(u)), t \leq u \leq T\}$ constitutes a trading strategy for an investor in this financial market, who seeks to maximise expected utility of wealth at time T . We introduce the set \mathcal{S} , which defines the solvency region in the absence of an option trade, as

$$\mathcal{S} = \{(S, x, y) \in \mathbf{R}^+ \times \mathbf{R}^2 \mid x + (1+\lambda)Sy \geq 0, x + (1-\mu)Sy \geq 0\}. \quad (24)$$

A trading strategy (L, M) is said to be *admissible* (for the problem without options) if the corresponding holdings satisfy the solvency constraint

$$(S(u), X_{t,S,x,y}^{L,M}(u), Y_{t,S,x,y}^{L,M}(u)) \in \mathcal{S}, \quad \text{almost surely,} \quad \forall u \in [t, T]. \quad (25)$$

For an investor who trades options at time t and then seeks to maximise expected utility of wealth the set of admissible trading strategies is altered. For example, when writing a contingent claim, the work of Soner, Shreve and Cvitanić (1995) and Levental and Skorohod (1997) shows that, in order to keep the wealth of the writer non-negative, it is imperative to keep at least one share of the stock at all trading times. This issue does not enter our pricing methodology as it only requires us to solve an optimisation problem without the derivative security, though for computing hedging strategies this is not the case.

The value functions $V(t, S, x, y)$, and $V^{(o)}(t, S, x, y, \delta, p)$ will satisfy the same dynamic programming equations, but with different terminal boundary conditions. The function $F(t, S, x, y)$ of equation (12) is not necessarily a value function, but satisfies a similar recursive equation with the choice of control (the trading strategy) determined from the dynamic programming algorithm for $V(t, S, x, y)$.

The boundary equation to be applied at the terminal time T for $V(t, S, x, y)$ is

$$V(T, S, x, y) = U(x + yS), \quad (26)$$

where it is assumed that there are no transaction costs charged on cashing out the final portfolio (in keeping with much of the existing literature on transaction costs). Assuming that costs are charged on liquidating the portfolio then (26) is replaced by

$$V(T, S, x, y) = U(x + c(y, S)), \quad (27)$$

where $c(y, S)$ is the cash value of y shares of stock, each of price S , and is defined by

$$c(y, S) = \begin{cases} (1 + \lambda)yS, & \text{if } y < 0, \\ (1 - \mu)yS, & \text{if } y \geq 0. \end{cases} \quad (28)$$

Our results are not qualitatively altered if there are no costs on liquidation.

The terminal boundary condition for the optimisation problem involving options is (with the same remarks as above about liquidation costs)

$$V^{(o)}(T, S, x, y, \delta, p) = U\left(x + yS + \frac{\delta}{p}C(S)\right), \quad (29)$$

whilst the terminal boundary condition for $F(t, S, x, y)$ is

$$F(T, S, x, y) = C(S)U'(x + yS). \quad (30)$$

3.1 Dynamic Programming Equations

The dynamic programming equations satisfied by the function $V(t, S, x, y)$ in a market with proportional transaction costs have been derived by Hodges and Neuberger (1989) and Davis, Panas and Zariphopoulou (1993). We present these equations below along with a simple (and financially intuitive) justification, and we describe the investor's optimal trading strategy. We highlight how the characteristics of the value function $V(t, S, x, y)$ and the associated optimal portfolio strategy impinge on the investor's option valuation and hedging policy.

The investor's optimal strategy at any state (t, S, x, y) is as follows. The state space (t, S, x, y) is split into three distinct regions: the BUY, SELL and no transaction (NT) regions, from which it is optimal to buy stock, sell stock and not to trade, respectively. Moreover, with proportional transaction costs, the optimal trade when outside the NT region is to transact to the nearest boundary of the NT region. The optimal strategy can be summarised as *minimal trading to keep the portfolio in the NT region*. We denote the boundaries between the NT region and the BUY and SELL regions by y_b and y_s , respectively. In general, y_b and y_s will be functions of (t, S, x) , and will represent the number of shares held at the NT boundaries. If the state is in the NT region it drifts under the influence of the diffusion driving the stock price, on a surface defined by $Y(u) = \text{constant}$. If the state is in the BUY or SELL regions, an immediate transaction occurs taking the state to the nearest boundary of the NT region.

In the BUY region the value function remains constant along the path of the state dictated by the optimal trading strategy, and therefore satisfies

$$V(t, S, x, y) = V(t, S, x - S(1 + \lambda)\delta L, y + \delta L) \text{ in BUY,} \quad (31)$$

where δL , the number of shares bought, can take any positive value up to the one required to take the portfolio to y_b . Allowing $\delta L \downarrow 0$, (31) becomes

$$\frac{\partial V}{\partial y}(t, S, x, y) - (1 + \lambda)S \frac{\partial V}{\partial x}(t, S, x, y) = 0 \text{ in BUY.} \quad (32)$$

Equation (32) indicates that in the BUY region the marginal utilities of the risky and riskless assets are in the ratio $(1 + \lambda)S$, the cost of purchasing one share of the stock. It defines a direction in the state space in which the value function is constant in the BUY region. Such properties of the value function will carry over to the investor's option valuation and hedging strategy. Specifically, the general option pricing formula (10) yields an option price which does not vary with the investor's initial portfolio outside the NT region.

Similarly, in the SELL region, the value function satisfies

$$V(t, S, x, y) = V(t, S, x + S(1 - \mu)\delta M, y - \delta M) \text{ in SELL,} \quad (33)$$

where δM represents the number of shares sold. Letting $\delta M \downarrow 0$, (33) becomes

$$\frac{\partial V}{\partial y}(t, S, x, y) - (1 - \mu)S \frac{\partial V}{\partial x}(t, S, x, y) = 0 \text{ in SELL.} \quad (34)$$

Equations (32) and (34) are the well known "value matching" conditions satisfied by the value function of (2). (See, for example, Dixit (1991), Dumas (1991, 1992), Dumas and Luciano (1991), or the exposition by Harrison (1985) of the theory of optimally regulated Brownian motion.) They hold throughout the BUY and SELL regions, and (of course) at the boundaries y_b, y_s , regardless of whether these are chosen optimally or not.

Finally, in the NT region, since it is sub-optimal to carry out any stock trades, for any stock purchase δL or sale δM :

$$V(t, S, x, y) \geq V(t, S, x - S(1 + \lambda)\delta L, y + \delta L) \text{ in NT} \quad (35)$$

and

$$V(t, S, x, y) \geq V(t, S, x + S(1 - \mu)\delta M, y - \delta M), \text{ in NT,} \quad (36)$$

which on expansion imply that the left hand sides of (32) and (34) are non-positive and non-negative respectively in NT. Bellman's optimality principle for dynamic programming gives the value function at time t in terms of its counterpart at time $t + \delta t$ as

$$V(t, S, x, y) = \mathbf{E}_{\delta t} V(t + \delta t, S + \delta S, x + \delta x, y) \text{ in NT,} \quad (37)$$

where $\mathbf{E}_{\delta t}$ denotes expectation over the time interval δt . In the limit $\delta t \rightarrow 0$, δS and δx are given by (17) and (21) respectively (with $dL = dM = 0$ since we are

in the NT region). Applying Ito's lemma yields the Hamilton-Jacobi-Bellman equation for the value function in the NT region:

$$V_t + rxV_x + bSV_S + \frac{1}{2}\sigma^2 S^2 V_{SS} = 0 \text{ in NT}, \quad (38)$$

where the arguments of the value function have been omitted for brevity.

These equations can be condensed into the PDE

$$\max \left[V_y - (1 + \lambda)SV_x, -(V_y - (1 - \mu)SV_x, V_t + rxV_x + bSV_S + \frac{1}{2}\sigma^2 S^2 V_{SS}) \right] = 0. \quad (39)$$

The solution of the optimisation problem is obtained by observing that if we can compute the value function in the NT region along with the boundaries of this region, then we can calculate its value in the BUY and SELL regions using (31) and (33).

A closed form solution for the value function $V(t, S, x, y)$ is elusive. This feature is (notoriously) common to models with transaction costs. Constantinides (1986), Davis and Norman (1990), Dumas and Luciano (1991) and Shreve and Soner (1994) do find, in the infinite horizon (and hence, time independent) case, for HARA utility functions, an analytic form for the value function. However, the boundaries of the NT region must still be located numerically. We shall, therefore, in our finite horizon problem, go directly to the development of a numerical dynamic programming algorithm to obtain option prices and hedging strategies.

We use the equations (31), (33) and (37), and augment them with the specific properties of the optimal trading strategy to create a backward recursive dynamic programming algorithm. This yields the value function at time t , provided it is known at time $t + \delta t$, along with the location of the NT boundaries. Assume that the stock and bond prices evolve in discrete time according to (18)-(20). Then the discrete dynamic programming equation is

$$V(t, S, x, y) = \max_{(\delta L, \delta M)} \left[\mathbf{E}_{\delta t} V(t + \delta t, \omega S, R(x - S(1 + \lambda)\delta L), y + \delta L), \right. \\ \left. \mathbf{E}_{\delta t} V(t + \delta t, \omega S, Rx, y), \right. \\ \left. \mathbf{E}_{\delta t} V(t + \delta t, \omega S, R(x + S(1 - \mu)\delta M), y - \delta M) \right], \quad (40)$$

where $R \equiv \exp(r\delta t)$ and the maximum is achieved by the first, second or third terms in (40) when the state (t, S, x, y) is in the BUY, NT and SELL regions, respectively.

Equation (40) expresses the value function at time t in terms of its counterpart at $t + \delta t$ by comparing the three possibilities: (i) buying δL shares and allowing the stock to diffuse or (ii) not trading and allowing the stock to diffuse or (iii) selling δM shares and allowing the stock to diffuse.

The algorithm is an example of the Markov chain approximation technique for the numerical solution of continuous time stochastic control problems, pioneered by Kushner (1990); see also Kushner (1997) for a review of applications

in finance. The state variables and controls are approximated by discrete time, discrete state Markov chains, in such a manner that the solution to the discrete problem converges to the solution of the continuous time problem. The application here is to a singular control problem, along the lines of Kushner and Martins (1991). For the optimal portfolio problem studied here, the necessary proofs of convergence of the discrete time problem to the continuous one are supplied by Davis, Panas and Zariphopoulou (1993), who proved that the value function $V(t, S, x, y)$ is a viscosity solution of the variational inequality (39).

To implement the above algorithm we specialise to the case of exponential utility in the next section.

4 Option Prices and Hedging Strategies Under Exponential Utility

Following Hodges and Neuberger (1989) and Davis, Panas and Zariphopoulou (1993) we set the investor's utility function to be the negative exponential:

$$U(W) = -\exp(-\alpha W), \quad (41)$$

with constant risk aversion parameter α . With this choice the investor's optimal trading strategy becomes independent of the wealth held in the riskless asset. The assumption of exponential utility is made primarily to reduce the dimensionality of the optimisation problems that we must solve, as our goal is to compute numerical estimates of the option prices and hedging strategies implied by the general option pricing formula (10). We defer to a later article the comparison of option prices generated by alternative choices of the utility function. For example, with HARA utility functions, such as logarithmic or power utility, the optimal trading strategy is characterised by a time varying no transaction region with boundaries $\zeta_b(t) < \zeta_s(t)$, where $\zeta(t)$ represents the ratio of wealth held in the stock to that held in the bond at time t . We hypothesise that the results for option prices will be similar to those we present below for exponential utility, but the relevant state space variable will be $\zeta(t)$, instead of the amount of money invested in the stock.

Below we show how the optimisation problem for $V(t, S, x, y)$ simplifies under exponential utility.

Define

$$H(t, S, y) \equiv V(t, S, 0, y), \quad (42)$$

then since, with exponential utility, the optimal portfolio through time is independent of the wealth held in the bond, we have that

$$V(t, S, x, y) = H(t, S, y) \exp\left(-\alpha x e^{r(T-t)}\right). \quad (43)$$

The resulting reduction in dimensionality means that the discrete dynamic programming algorithm (40) reduces to

$$\begin{aligned}
H(t, S, y) = \max_{(\delta L, \delta M)} & \quad [\mathbf{E}_{\delta t} H(t + \delta t, \omega S, y + \delta L) \exp(\alpha S(1 + \lambda)\delta L \cdot \beta(t)) , \\
& \quad \mathbf{E}_{\delta t} H(t + \delta t, \omega S, y), \\
& \quad \mathbf{E}_{\delta t} H(t + \delta t, \omega S, y - \delta M), \exp(-\alpha S(1 - \mu)\delta M \cdot \beta(t))] ,
\end{aligned} \tag{44}$$

where $\beta(t) = \exp(r(T - t))$.

For exponential utility, the boundaries of the NT region become functions of t and S only. Denote them by $y_b(t, S)$ and $y_s(t, S)$, with $y_b(t, S) \leq y_s(t, S)$, and equality holding only in the case where $\lambda = \mu = 0$.

The optimal values of δL and δM , δL^* and δM^* , satisfy

$$\begin{aligned}
y + \delta L^* = y_b(t, S) \text{ and } \delta M^* = 0, & \quad \text{if } y < y_b(t, S) \\
\delta L^* = \delta M^* = 0, & \quad \text{if } y_b(t, S) \leq y \leq y_s(t, S) \\
\delta L^* = 0 \text{ and } y - \delta M^* = y_s(t, S), & \quad \text{if } y > y_s(t, S).
\end{aligned} \tag{45}$$

Applying (45) and (43) to equations (31), (33) and (37), we also obtain the following representation for $H(t, S, y)$ in the BUY, SELL and NT regions.

If $y < y_b(t, S)$, then

$$H(t, S, y) = H(t, S, y_b(t, S)) \exp(\alpha S(1 + \lambda)(y_b(t, S) - y)\beta(t)). \tag{46}$$

If $y > y_s(t, S)$, then

$$H(t, S, y) = H(t, S, y_s(t, S)) \exp(-\alpha S(1 - \mu)(y - y_s(t, S))\beta(t)). \tag{47}$$

If $y_b(t, S) \leq y \leq y_s(t, S)$, then

$$H(t, S, y) = \mathbf{E}_{\delta t} H(t + \delta t, \omega S, y). \tag{48}$$

Equations (46)-(48) give the value function $H(t, S, y)$ in the BUY, NT and SELL regions, provided we know $H(t, S, y)$ at and within the boundaries $y_b(t, S)$ and $y_s(t, S)$, along with the location of these boundaries. These are located by implementing the algorithm in (44) in the manner described below.

We create a large vector to represent possible values of y at each node of the stock price tree, with discretisation step h_y . The range of this vector must be large enough to locate $y_b(t, S)$ and $y_s(t, S)$ for all (t, S) on the binomial stock price tree. This can be accomplished by deriving analytically the NT boundaries at $T - \delta t$ (see the Appendix for this derivation), and noting that the NT region is wider at this time than at any preceding time. Then the following sequence of steps is performed.

1. Suppose we know the value function at $t + \delta t$ for all stock prices on the binomial tree at this time, and for all values of y in our discrete vector. Then starting at a time- t node of the stock price tree, (t, S) say, and from the minimum value of y in this vector, we compare the first and second

terms in the maximisation operator of (44) for increasing values of y in steps of h_y , until the latter is greater than or equal to the former, at say y^b , which we assume satisfies $y^b = y_b(t, S)$, the boundary between the NT and BUY regions at the node (t, S) .

2. We continue, comparing the second and third terms in the maximisation operator of (44) for increasing values of y in steps of h_y , until the latter is greater than or equal to the former, at say y^s , which we assume satisfies $y^s = y_s(t, S)$, the boundary between the NT and SELL regions at the node (t, S) .
3. Having located the boundaries of the NT region at the node (t, S) , the value function at all points outside this region is determined by assuming the investor transacts to its boundaries (i.e. applying equations (46) and (47)), whilst the function in the NT region is found by assuming the investor does not transact, and applying equation (48).

We further enhance the speed of the above algorithm for $V(t, S, x, y)$ because, for exponential utility, the NT boundaries at any time are characterised by the wealth held in the stock being constant. (This was confirmed by solving the problem without this assumption.) Therefore, having located the boundaries at a single node of the binomial tree at time t , the boundaries at all other time- t nodes are given easily. This property is not satisfied by the value function $V^{(0)}(t, S, x, y, \delta, p)$.

The other noteworthy feature of our algorithm is the fact that the y -vector we use is bounded to the interval $[y_b(T - \delta t, S), y_s(T - \delta t, S)]$, the NT boundaries one period prior to the option expiry. This allows us to restrict the state space over which we carry out the search for the NT boundaries at all earlier times, and makes for fast, accurate computation. These features can be exploited for option pricing because the “fair” pricing methodology only requires the solution of the investor’s optimal portfolio problem in the absence of options.

4.1 Option Prices

To calculate option prices, under exponential utility, the fair pricing formula (11) becomes

$$\hat{p}(t, S, y) = e^{-r(T-t)} \frac{G(t, S, y)}{H(t, S, y)}, \quad (49)$$

where

$$G(t, S, y) \equiv \mathbf{E}_t[U(W_{t,S,0,y}^{\pi^*}(T))C(S(T))], \quad (50)$$

using the fact that $U'(w) = -\alpha U(w)$.

4.2 Effect of Option Trading on Utility

To compare the investor’s maximum utility with and without the diversion of funds into an option trade, we must compare, in general, the quantities

$V^{(o)}(t, S, x - \delta, y, \delta, p)$ with $V(t, S, x, y)$. With exponential utility this becomes a comparison of

$$H^{(o)}(t, S, y, \delta, p) \exp(\alpha \delta \beta(t)) \quad \text{with} \quad H(t, S, y), \quad (51)$$

where $H^{(o)}(t, S, y, \delta, p)$ is defined by

$$H^{(o)}(t, S, y, \delta, p) \equiv V^{(o)}(t, S, 0, y, \delta, p). \quad (52)$$

5 Numerical Results

For our numerical results we used the following parameters as a base case: $T = 1$ year, $r = 0.1$, $b = 0.15$, $\sigma = 0.25$, $\alpha = 0.1$, and we took the transaction cost rates for buying and selling stock to be equal ($\mu = \lambda$). We used a stock price tree with at least 50 timesteps. First we confirmed some stylised facts about the investor's optimal trading strategy without options, which we summarise below, and which verify the robustness of our numerical algorithm. The optimal trading strategy has the following properties.

1. The boundaries of the NT region lie either side of the optimal portfolio without transaction costs, and the NT region widens with the transaction costs.
2. The NT region boundaries show a hyperbolic dependence versus the stock price, just as in the frictionless markets case, indicating that with exponential utility, and at a fixed time, the wealth in the stock is constant at the boundaries of the NT region.
3. As we approach the horizon time T the NT region widens considerably. This is as expected, indicating that portfolio rebalancing becomes less advantageous as the time to expiry lessens. This feature carries over to the hedging strategy for an option position.
4. An increase in risk aversion narrows the region of no transactions and shifts it to lower values of the stock holding.

Figure 1 shows at-the-money call option prices given by the general option pricing formula (49), plotted at time zero versus the investor's initial stock holding, y , for transaction cost parameters $\lambda = 0.005$ and $\lambda = 0.01$. The graphs are flat outside a certain range of y , which corresponds exactly with the width of the NT region for the particular transaction cost parameter. We see the widening of the option pricing bounds as the transaction costs are increased. The fair option price is higher when the investor's stock inventory is in the BUY region for the basic portfolio selection problem, then falls as we enter the NT region, and is at its lowest when the current inventory position is in the investor's SELL region. This is intuitively correct, since a buyer of shares will value a call option more highly than someone who wishes to sell stock. Of course, the opposite pattern is obtained for put options.




Figure 1: At-the-money call option prices versus initial stock holding. The parameters are $T = 1$ year, $r = 0.1$, $b = 0.15$, $\sigma = 0.25$, $\alpha = 0.1$, $S = K = 15$, $\lambda = 0.01$ and 0.005 .

5.1 The Effect on Utility of Trading Options

We analyse the case of the investor writing a call option at price p_w , the maximum price in Figure 1, which we label the “fair asking price”. If we compare $H^{(0)}(t, S, n, -p_w, p_w) \exp(-\alpha p_w \beta(t))$ with $H(t, S, y)$ (making the comparison in (51) with $\delta = -p_w$ and $p = p_w$), we find that the investor’s utility is virtually *unchanged* by the option sale. However, trading an option at a price outside the bid-ask spread always results in a reduction in utility (if one buys above the asking price or sells below the bid price) or an increase in utility (if one sells above the asking price or buys below the bid price).

To quantify the above points, in Figure 2 (upper graph) we have plotted the difference between the investor’s maximum utility when selling one call option at the fair asking price and the maximum utility in the absence of the option sale. If the investor starts out with a low stock inventory (i.e. he would, acting optimally, buy shares) then selling a call option reduces his utility (albeit




Figure 2: The upper graph shows the difference between the investor’s maximum utility when selling a call option at the “fair” asking price and the maximum utility without selling the option. The lower graph plots the utility with the option sale as a fraction of the utility without the option sale. The parameters are $T = 1$ year, $r = 0.1$, $b = 0.15$, $\sigma = 0.25$, $\alpha = 0.1$, $S = K = 15$, $\lambda = 0.005$ and $\hat{p} = p_w = 2.2864$.

by a small amount). At some value, y^* , of the initial stock holding y , the investor finds his utility is unchanged by the diversion of funds into the option transaction, and for $y > y^*$ the utility is increased by the sale of the option. The increase in utility reaches a maximum, then tails off toward zero for large values of y , reflecting the fact that when the investor wishes to sell large amounts of stock, selling a call option increases utility, but this effect diminishes as the initial stock holding increases. For all values of y , the percentage change in utility is bounded and small, as shown in the lower graph. Similar observations hold for any trade of an option within the “fair” bid-ask spread.

To summarise, the fair price given by the general option pricing formula is essentially a reservation price. A nice feature of the pricing method used in this paper is that the option pricing bounds can be found by solving the investor’s basic portfolio selection problem, and without having to solve the more complicated problem involving the purchase or sale of options.




Figure 3: Call option prices for different risk aversion parameters. The parameters are $T = 1$ year, $r = 0.1$, $b = 0.15$, $\sigma = 0.25$, $S = K = 10$, $\lambda = 0.01$, $\alpha = 0.5$ and $\alpha = 0.1$.

Figure 3 shows at the money call option prices for two different risk aversion parameters. The bid-ask spread is independent of risk aversion, but the range of values of the initial stock holding for which the fair price lies within the bid-ask spread becomes narrower and is shifted to a lower value, as the risk aversion increases.

In Table 1 we present call option prices for various strikes and transaction cost parameters, and for comparison we show the bid and ask prices generated by Leland's (1985) approximately replicating strategy, with a revision interval of $\delta t = 0.02$, which corresponds to approximately weekly portfolio rebalancing. A number of points are worth emphasising. First, in general, the optimal pricing approach places tighter bounds on the option price, particularly for large transaction costs. The intuition behind this feature is natural: Leland's strategy insists on portfolio rebalancing (thus incurring transaction costs) in situations where the optimal pricing procedure may not. We used a binomial tree with the same time step as the Leland revision interval to generate the prices in Table 1,

$\lambda = 0.005$, NT Region=[0.3866,0.5780]				
Strike	Ask Price	Bid Price	Leland Ask	Leland Bid
10	6.0471	5.8980	5.9940	5.9610
13	3.5841	3.4503	3.6190	3.4348
15	2.2864	2.1775	2.3869	2.0915
17	1.3419	1.2641	1.4878	1.1481
20	0.5423	0.5048	0.6724	0.3949
$\lambda = 0.01$, NT Region=[0.3499,0.6197]				
Strike	Ask Price	Bid Price	Leland Ask	Leland Bid
10	6.1199	5.8248	6.0187	5.9537
13	3.6476	3.3837	3.7088	3.3458
15	2.3376	2.1212	2.5164	1.9171
17	1.3788	1.2210	1.6336	0.9398
20	0.5613	0.4805	0.8010	0.2470
$\lambda = 0.02$, NT Region=[0.2702,0.7196]				
Strike	Ask Price	Bid Price	Leland Ask	Leland Bid
10	6.2675	5.6716	6.0775	5.9516
13	3.7798	3.2463	3.8807	3.2374
15	2.4475	2.0073	2.7502	1.4800
17	1.4612	1.1361	1.8940	0.3057
20	0.6063	0.4348	1.0401	0.0034
$\lambda = 0.03$, NT Region=[0.1813,0.8243]				
Strike	Ask Price	Bid Price	Leland Ask	Leland Bid
10	6.4068	5.5242	6.1450	-
13	3.9070	3.1159	4.0421	-
15	2.5556	1.9012	2.9590	-
17	1.5445	1.0589	2.1242	-
20	0.6537	0.3948	1.2592	-

Table 1: Call bid and ask prices. The parameters are $T = 1$ year, $r = 0.1$, $b = 0.15$, $\sigma = 0.25$, $\alpha = 0.1$, $S = 15$. For $\lambda = 0.03$ the Leland bid price is undefined for a revision interval of $\delta t = 0.02$.

which means that the investor has the *opportunity* to re hedge as frequently as the Leland strategy, but chooses not to do so.

The implication of these results is that the investor is prepared to bear more risk than the Leland strategy allows, and the size of this risk is determined by the utility function. The only exception to this feature is for options which are deep in the money. In this case Leland's bounds are tighter. The intuition here is as follows: for a deep in the money option, with very high probability of exercise, the optimal policy is to be (almost) fully hedged, and this is in accordance with Leland's strategy, which is designed to eliminate risk in a Black-Scholes type manner. Therefore, in these situations, Leland's strategy *is* optimal and falls within the spread given by utility maximisation.

For very large transaction costs Leland's strategy fails to produce a bid price for the option, as the effective volatility is no longer a real number. This point has also been made by Avellaneda and Parás (1994), who provided a solution to this problem using the notion of imperfectly dominating policies. The optimal pricing procedure never fails to produce a sensible option price, regardless of the level of transaction costs.

Figure 4 shows a plot of the call bid-ask spread specified by the optimal pricing formula, versus the stock price. We have also shown Leland's bid-ask spread with a revision interval $\delta t = 0.02$, equal to the time step of the binomial tree, and the Black-Scholes call values. We see how the optimal pricing procedure places tighter bounds on the option prices, except for the cases where the option is deep in the money, as before.

In Figure 5 we plot the hedging strategy for a short call position versus the initial stock price, produced using Definition 1. The dashed curves indicate the region in which the hedging portfolio is not rebalanced, whilst the solid curve is the Black-Scholes delta hedging strategy. The replacement of the unique Black-Scholes delta by a hedging bandwidth is in accordance with intuition and with previous results on optimal hedging under transaction costs, such as Hodges and Neuberger (1989) and (for the limiting case of small transaction costs) Whalley and Wilmott (1997).

6 Conclusions and Suggestions

This paper has developed a procedure for optimally valuing options in the presence of proportional transaction costs. The method involves treating an option transaction as an alternative investment to optimally trading the underlying stock. Option prices are determined by requiring that, at the margin, the diversion of funds into an option trade has no effect on an investor's achievable utility. Thus, the option trade is treated as a small perturbation on the investor's initial portfolio of assets. The methodology can therefore be extended to situations in which the basic portfolio contains many assets, including possibly other derivatives.

Option prices are computed by solving a singular stochastic optimal control problem via an efficient algorithm. We only need to solve the investor's




Figure 4: Call option bid-ask values versus stock price. The parameters are $T = 1$ year, $r = 0.1$, $b = 0.15$, $\sigma = 0.25$, $\alpha = 0.1$, $K = 15$, $\lambda = 0.02$.

fundamental portfolio selection problem to derive option prices, as opposed to other optimal procedures which require the solution of an optimisation problem containing an embedded option. The method is therefore a relatively tractable way of producing optimal pricing bounds.

The method places tight bounds on option prices and generates prices which can lie anywhere within these bounds, depending on the investor's initial holding of stock. The investor's utility is hardly affected when trading options at the optimal prices, whereas trading outside the bounds can lead to a drastic loss in utility.

The final conclusion is that a general approach to option pricing in the presence of transaction costs should be based on an optimal portfolio approach. Approximate replication schemes work well for small levels of transaction costs and are useful because they can lead to closed form approximate solutions, but they cannot cope with larger values of transaction costs. In these cases the optimal models outperform them as the frequency of trading is reduced.




Figure 5: Short call option hedging strategies versus stock price. The call was sold at the “fair” asking price, and the other parameters are $T = 1$ year, $r = 0.1$, $b = 0.15$, $\sigma = 0.25$, $\alpha = 0.1$, $K = 15$, $\lambda = 0.005$. The dashed curve indicates the region in which portfolio rebalancing does not take place, whilst the solid curve is the Black-Scholes delta.

There are a number of directions in which this work could be extended:

American Options The pricing of American options with transaction costs presents some interesting problems because one not only has to compute an optimal hedging strategy, but also an optimal exercise policy. This will involve a problem in singular control with optimal stopping, which has been studied by Davis and Zervos (1994). There is a further complication for the writer of an option in that it is not he, but the buyer of the option, who controls the exercise policy. Some preliminary ideas on this topic have been provided by Davis and Zariphopoulou (1995).

Equilibrium The option pricing method described in this paper selects a “fair” price within a closed interval provided by the reservation buying and selling prices. It therefore gives some hope of providing a framework in which

the buyer and seller of an option might agree on a price.

Different Preferences An analysis of how the risk preferences of the investor affect option prices would be interesting. Optimal portfolios for HARA utility functions are usually determined by selecting an optimal ratio of wealth in the risky and riskless assets, as opposed to the exponential function used in this paper, in which the wealth held in the risky asset is the important variable. We hypothesise that such patterns would transfer to the option valuation problem. One could also consider quadratic preferences, such as risk minimisation, as in Mercurio and Vorst (1997) or Lamberton, Pham and Schweizer (1998). Whilst this can be criticised on the grounds that it gives the same weighting to downside and upside risk, this approach does merit further study. Another possibility is to consider “coherent” measures of risk, introduced by Artzner, Delbaen, Eber and Heath (1999), and extended to a dynamic setting by Cvitanic and Karatzas (1999).

Appendix

We derive analytic formulae for the value function $H(T - \delta t, S, y)$ and the boundaries of the NT region $y_b(T - \delta t, S)$, $y_s(T - \delta t, S)$, one time period prior to the final time T , under exponential utility.

In the BUY region ($y < y_b(T - \delta t, S)$) the Bellman equation (44) for the value function $H(T - \delta t, S, y)$ reduces to

$$H(T - \delta t, S, y) = \max_{\delta L} \mathbf{E}_{\delta t} H(T, \omega S, y + \delta L) \exp(\alpha RS(1 + \lambda)\delta L), \quad \text{in BUY} \quad (53)$$

where $R = \exp(r\delta t)$ and $H(T, S, y) = -\exp(-\alpha yS)$.

We write out the above expectation explicitly, differentiate with respect to δL , and set the result to zero. This yields, after some tedious algebra, that the optimal number of shares to buy, δL^* , satisfies

$$y + \delta L^* \equiv y_b(T - \delta t, S) = \frac{1}{\alpha S(\omega_u - \omega_d)} \log \left(\frac{q(1 - q_+)}{(1 - q)q_+} \right), \quad (54)$$

where ω_u and ω_d are the two possible realisations of the binomial random variable ω given in (20), (so that $q = \frac{1}{2}$), and the pseudo-probability q_+ is given by

$$q_+ = \frac{R(1 + \lambda) - \omega_d}{\omega_u - \omega_d}. \quad (55)$$

Inserting the expression for δL^* into (53) gives the following representation for $H(T - \delta t, S, y)$ in the BUY region:

$$H(T - \delta t, S, y) = -\exp(-\alpha yRS(1 + \lambda)) \left(\frac{q}{q_+} \right)^{q_+} \left(\frac{1 - q}{1 - q_+} \right)^{(1 - q_+)}, \quad \text{in BUY.} \quad (56)$$

A similar analysis in the SELL region gives the optimal number of shares to sell, δM^* , as

$$y - \delta M^* = y_s(T - \delta t, S) = \frac{1}{\alpha S(\omega_u - \omega_d)} \log \left(\frac{q(1 - q_-)}{(1 - q)q_-} \right), \quad (57)$$

where the pseudo-probability q_- is given by

$$q_- = \frac{R(1 - \mu) - \omega_d}{\omega_u - \omega_d}, \quad (58)$$

so that the value function in the SELL region is

$$H(T - \delta t, S, y) = -\exp(-\alpha y R S(1 - \mu)) \left(\frac{q}{q_-} \right)^{q_-} \left(\frac{1 - q}{1 - q_-} \right)^{(1 - q_-)}, \quad \text{in SELL.} \quad (59)$$

Finally, in the NT region, the value function is given analytically by

$$H(T - \delta t, S, y) = -[q \exp(-\alpha y S \omega_u) + (1 - q) \exp(-\alpha y S \omega_d)]. \quad \text{in NT.} \quad (60)$$

References

- Artzner, Philippe, Freddy Delbaen, Jean-Marc Eber and David Heath, 1999, Coherent measures of risk, *Mathematical Finance* 9, 203–228.
- Avellaneda, Marco, and Antonio Parás, 1994, Dynamic hedging portfolios for derivative securities in the presence of large transaction costs, *Applied Mathematical Finance* 1, 165–193.
- Bensaid, Bernard, Jean-Philippe Lesne, Henri Pagès, and José Scheinkman, 1992, Derivative asset pricing with transaction costs, *Mathematical Finance* 2, 63–86.
- Black, Fischer and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–654.
- Boyle, Phelim P., and Ken S. Tan, 1994, Lure of the linear, *Risk* 7 ((4)), 43–46.
- Boyle, Phelim P., and Ton Vorst, 1992, Option replication in discrete time with transaction costs, *Journal of Finance* 47, 271–293.
- Clewlow, Les and Stewart D. Hodges, 1997, Optimal delta-hedging under transactions costs, *Journal of Economic Dynamics and Control* 21, 1353–1376.
- Constantinides, George M., 1986, Capital market equilibrium with transaction costs, *Journal of Political Economy* 94, 842–862.
- Constantinides, George M., and Thaleia Zariphopoulou, 1999, Bounds on prices of contingent claims in an intertemporal economy with proportional transaction costs and general preferences, *Finance and Stochastics* 3, 345–369.

- Cox, John C., Stephen A. Ross, and Mark Rubinstein, 1979, Option pricing: A simplified approach, *Journal of Financial Economics* 7, 229–263.
- Cvitanic, Jakša, and Ioannis Karatzas, 1999, On dynamic measures of risk, *Finance and Stochastics* 3, 451–482.
- Cvitanic, Jakša, Huyên Pham, and Nizar Touzi, 1999, A closed-form solution to the problem of super-replication under transaction costs, *Finance and Stochastics* 3, 35–54.
- Davis, Mark H.A., 1997, Option pricing in incomplete markets, in Michael A.H. Dempster and Stanley R. Pliska, eds.: *Mathematics of Derivative Securities* (Cambridge University Press, Cambridge, UK).
- Davis, Mark H.A., and J. Martin C. Clark, 1994, A note on super-replicating strategies, *Philosophical Transactions of the Royal Society of London A* 347, 485–494.
- Davis, Mark H.A., and Andrew R. Norman, 1990, Portfolio selection with transaction costs, *Mathematics of Operations Research* 15, 676–713.
- Davis, Mark H.A., Vassilios G. Panas, and Thaleia Zariphopoulou, 1993, European option pricing with transaction costs, *SIAM Journal of Control and Optimization* 31, 470–493.
- Davis, Mark H.A., and Thaleia Zariphopoulou, 1995, American options and transaction fees, in Mark H.A. Davis, Darrell Duffie, Wendell H. Fleming, and Steven E. Shreve, eds.: *Mathematical Finance - IMA Volumes in Mathematics and its Applications no.65* (Springer-Verlag, New York).
- Davis, Mark H.A., and Mihail Zervos, 1994, A problem of singular stochastic control with discretionary stopping, *Annals of Applied Probability* 4, 226–240.
- Dixit, Avinash, 1991, A simplified treatment of the theory of optimal regulation of Brownian motion, *Journal of Economic Dynamics and Control* 15, 657–673.
- Duffie, Darrell, 1996, *Dynamic Asset Pricing Theory, 2nd ed.* (Princeton University Press, Princeton, N.J.).
- Dumas, Bernard, 1991, Super contact and related optimality conditions, *Journal of Economic Dynamics and Control* 15, 675–685.
- Dumas, Bernard, 1992, Dynamic equilibrium and the real exchange rate in a spatially separated world, *Review of Financial Studies* 5, 153–180.
- Dumas, Bernard, and Elisa Luciano, 1991, An exact solution to a dynamic portfolio choice problem under transaction costs, *Journal of Finance* 46, 577–595.

- Edirisinghe, Chanaka, Vasanttilak Naik, and Raman Uppal, 1993, Optimal replication of options with transactions costs and trading restrictions, *Journal of Financial and Quantitative Analysis* 28, 117–128.
- Harrison, J. Michael, 1985, *Brownian Motion and Stochastic Flow Systems* (Wiley, New York, N.Y.).
- He, Hua, 1990, Convergence from discrete- to continuous-time contingent claims prices, *Review of Financial Studies* 3, 523–546.
- Hodges, Stewart D., and Anthony Neuberger, 1989, Optimal replication of contingent claims Under transaction costs, *Review of Futures Markets* 8, 222–239.
- Kushner, Harold J., 1990, Numerical methods for stochastic control problems in continuous time, *SIAM Journal of Control and Optimization* 28, 999–1048.
- Kushner, Harold J., 1997, Numerical methods for stochastic control problems in finance, in Michael A.H. Dempster and Stanley R. Pliska, eds.: *Mathematics of Derivative Securities* (Cambridge University Press, Cambridge, UK).
- Kushner, Harold J. and Luiz Felipe Martins, 1991, Numerical methods for stochastic singular control problems, *SIAM Journal of Control and Optimization* 29, 1443–1475.
- Karatzas, Ioannis, and Shing-gang Kou, 1996, On the pricing of contingent claims under constraints, *Annals of Applied Probability* 6, 321–369.
- Lamberton, Damien, Huyên Pham, and Martin Schweizer, 1998, Local risk-minimization under transaction costs, *Mathematics of Operations Research* 23, 585–612.
- Leland, Hayne E., 1985, Option pricing and replication with transactions costs, *Journal of Finance* 40, 1283–1301.
- Levental, Shlomo, and Anatolii V. Skorohod, 1997, On the possibility of hedging options in the presence of transaction costs, *Annals of Applied Probability* 7, 410–443.
- Mercurio, Fabio, and Ton C.F. Vorst, 1997, Options pricing and hedging in discrete time with transaction costs, in Michael A.H. Dempster and Stanley R. Pliska, eds.: *Mathematics of Derivative Securities* (Cambridge University Press, Cambridge, UK).
- Shreve, Steven E., and Halil M. Soner, 1994, Optimal investment and consumption with transaction costs, *Annals of Applied Probability* 4, 609–692.

Soner, Halil M., Steven E. Shreve, and Jakša Cvitanić, 1995, There is no non-trivial hedging portfolio for option pricing with transaction costs, *Annals of Applied Probability* 5, 327–355.

Whalley, A. Elizabeth, and Paul Wilmott, 1997, An asymptotic analysis of an optimal hedging model for option pricing with transaction costs, *Mathematical Finance* 7, 307–324.