Efficient option pricing with transaction costs

Michael Monoyios

Department of Economics and Finance, Brunel University, Uxbridge, Middlesex UB8 3PH, UK

A fast numerical algorithm is developed to price European options with proportional transaction costs using the utility-maximization framework of Davis (1997). This approach allows option prices to be computed by solving the investor's basic portfolio selection problem without insertion of the option payoff into the terminal value function. The properties of the value function can then be used to drastically reduce the number of operations needed to locate the boundaries of the no-transaction region, which leads to very efficient option valuation. The optimization problem is solved numerically for the case of exponential utility, and comparisons with approximately replicating strategies reveal tight bounds for option prices even as transaction costs become large. The computational technique involves a discrete-time Markov chain approximation to a continuous-time singular stochastic optimal control problem. A general definition of an option hedging strategy in this framework is developed. This involves calculating the perturbation to the optimal portfolio strategy when an option trade is executed.

1. Introduction

This article develops an efficient optimal procedure for computing European option prices in the presence of transaction costs on trading the underlying stock. This issue arises because the Black–Scholes (1973) option pricing methodology relies on perfect replication of the option payoff by a continuously rebalanced hedging portfolio involving the underlying stock. It is therefore inapplicable in markets with transaction costs as the hedging costs would be ruinously expensive.

Attempts to circumvent this problem include the early work of Leland (1985) and Boyle and Vorst (1992), who used a fixed hedging time scale, which is not necessarily an optimal policy. Furthermore, the pricing bounds become wider as the hedging error is reduced. Bensaid *et al.* (1992) 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. For a call option this method reduces to the trivial strategy of buying the underlying asset and holding it to maturity, as proven by Soner, Shreve and Cvitanić (1995) and Cvitanić, Pham and Touzi (1999) following a conjecture of Davis and Clark (1994). This illustrates a fundamental feature of option hedging under transaction

costs, namely, that attempting to eliminate all risk results in unrealistically wide valuation bounds.

Hodges and Neuberger (1989) recognized that an optimal valuation method, incorporating a utility-maximization objective, is a more viable valuation programme. By comparing the utility achieved with and without the obligations of an option contract, they specified reservation bid and asking prices for an option by requiring that the same utility be 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 the ϵ-arbitrage approach of Bertsimas, Kogan and Lo (2001), and local risk minimization, studied by Mercurio and Vorst (1997) and Lamberton, Pham and Schweizer (1998).

The utility-maximization approach is promising, but one usually has to solve a formidable singular stochastic optimal control problem, further complicated by the insertion of the option payoff into the terminal value function. The search for a more efficient procedure to value options with market frictions, while retaining the optimality of the utility-maximization approach, is the aim of the present study.

In this spirit we implement 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. In this approach an investor fixes a "fair" price for an option by requiring that an infinitesimal diversion of funds into the purchase or sale of the option has no effect on the investor's maximum utility. This is essentially a marginal version of the valuation methods pioneered by Hodges and Neuberger.

Our methodology results in fast computation of option prices within bounds that are tight, even for large transaction costs. The advance in computation speed is achieved in two ways. First, the option prices are obtained directly from the investor's basic portfolio selection problem, without the presence of the option. This is a direct consequence of Davis' (1997) general pricing formula. It allows us to use properties of the value function and of the optimal trading strategy to drastically reduce the number of computations needed to locate the boundaries of the investor's no-transaction (NT) region. Second, we derive analytically the boundaries of the NT region one period prior to maturity of the option. Since the NT region narrows as we move closer to the present time, we obtain natural bounds on the state space over which the backward-recursive dynamic programming algorithm to locate the NT boundaries must be carried out.

The rest of the article proceeds as follows. In Section 2 we set up a portfolio selection scenario in which the optimal strategy to maximize expected utility of wealth at a finite horizon time is sought, both with and without some initial wealth diverted into the purchase or sale of European options, and we state Davis' (1997) general option pricing formula. We formulate a general definition of an option hedging strategy for such a utility-maximizing investor. In Section 3

we consider a specific market model with transaction costs and present the dynamic programming solution to the portfolio choice problem. In Section 4 we specialize to the case of an exponential utility function and develop a numerical algorithm based on a Markov chain approximation to the continuous-time dynamic programming problem. In Section 5 we present numerical solutions for option prices and hedging strategies, and make comparisons with the approximate replication approach of Leland (1985). Section 6 concludes and suggests directions for further research. An appendix contains a derivation of a result used in the implementation of the numerical algorithm.

2. Portfolio selection and option valuation

We utilize a finite time interval $[0, T]$, where T corresponds 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 $W^{\pi}_{t,S,x,y}(u)$ when using the trading strategy π and starting in the state (t, S, x, y). The wealth $W^{\pi}_{t,S,x,y}(u)$ consists of $X^{\pi}_{t,S,x,y}(u)$ dollars in cash and $Y^{\pi}_{t,S,x,y}(u)$ shares of stock whose price at time u is $S(u)$, so that

$$W^{\pi}_{t,S,x,y}(u) = X^{\pi}_{t,S,x,y}(u) + Y^{\pi}_{t,S,x,y}(u) S(u) \tag{1}$$

The investor's objective is to maximize the expected utility of wealth at time T. Denote the investor's maximum utility by

$$V(t, S, x, y) = \sup_{\pi} \mathbf{E}_t \left[U\left(W^{\pi}_{t,S,x,y}(T)\right) \right] \tag{2}$$

where $\mathbf{E}_t$ denotes the expectation operator conditional on the information at time t. The supremum in (2) is taken over a suitable set of admissable policies; these are described in the next section, when we specialize to a market with proportional transaction costs.

Consider the alternative optimization problem that 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

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

The superscript "o" 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. In (2) and (3) the quantities δ and p would be measurable with respect to the time-t information.

In Hodges and Neuberger (1989) option pricing bounds were derived by requiring that the same utility is achieved when an option is traded as when it is not. Another approach was proposed by Davis (1997). 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

$$\frac{\partial V^{(o)}}{\partial \delta}(t, S, x - \delta, y, \delta, \hat{p})\bigg|_{\delta=0} = 0 \qquad (4)$$

This results in the pricing formula

$$\hat{p}(t, S, x, y) = \frac{\mathbf{E}_t\left[U'\left(W^{\pi^*}_{t,S,x,y}(T)C(S(T))\right)\right]}{V_x(t, S, x, y)} \qquad (5)$$

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 that maximizes the expected utility in (2). This is the trading strategy which optimizes a portfolio *without* options, and the formula (5) for $\hat{p}$ shows no dependence on the optimization problem (3) containing embedded options. This is the key to the fast computation of option prices with transaction costs.

We write (5) as

$$\hat{p}(t, S, x, y) = \frac{F(t, S, x, y)}{V_x(t, S, x, y)} \qquad (6)$$

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

$$F(t, S, x, y) := \mathbf{E}_t\left[U'\left(W^{\pi^*}_{t,S,x,y}(T)C(S(T))\right)\right] \qquad (7)$$

Davis' pricing methodology reduces to Black–Scholes pricing in complete, frictionless markets. It has been studied theoretically in various contexts by a number of authors, including Rabeau (1996), Cvitanić and Karatzas (1996), Karatzas and Kou (1996), Bensoussan and Julien (2000), Frittelli (2000), Rouge and El Karoui (2000), Schäl (2000) and Stettner (2000). These works focus mainly on characterizing the martingale pricing measure associated with the fair price $\hat{p}$, on connections with no arbitrage, and with the minimal-entropy martingale measure. The focus of our article is on the numerical implementation of the pricing procedure under proportional transaction costs to illustrate the simplification it affords in calculating option prices.

An aspect of the utility-based approaches to derivative pricing that merits further attention is the effect on the underlying asset market of the introduction
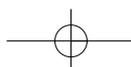
of an option. Detemple and Selden (1991) show that there are important general equilibrium effects when options are added to an incomplete market that can affect the price of the stock. In the above models such effects are implicitly assumed to be small, and this can be viewed as an approximation to simplify computation. In the case of the model examined here, the formula in (5) only involves the optimization problem in the absence of options (because its derivation involves infinitesimal diversions of wealth into options), and this is a potential justification for assuming that equilibrium effects are small.

The issue of what is an appropriate utility function and risk-aversion coefficient when implementing the above methods is an important one. Some method for calibrating risk aversion is necessary, and one approach is to examine the proportions of risky to risk-free asset holdings in empirical portfolios since these are directly related to risk-aversion characteristics. This is an important area for future research.

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 such as the local risk minimization approach in Lamberton, Pham and Schweizer (1998) or the ϵ-arbitrage approach of Bertsimas, Kogan and Lo (2001).

In the case of a utility-maximization approach to option pricing, the pricing problem is first embedded in the portfolio selection problem to determine a price. Then, if one computes the optimal trading strategy in the presence of the option trade, it 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 that ϵ options are written at price $\hat{p}$ given by (5). Then the investor's optimal trading strategy will be $\pi^{\dagger}$, maximizing

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

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 $\pi^{\dagger}$ 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

$$V^{(o)}(t, S, x - \delta, y, \delta, p) = \mathbf{E}_t \left[U \left(W_{t,S,x-\delta,y}^{\pi^\dagger}(T) + \frac{\delta}{p} C(S(T)) \right) \right] \qquad (8)$$

which defines the optimal trading strategy $\pi^\dagger$ for this utility-maximization 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*

$$X_{t,S,x,y}^{\pi^h}(u) = X_{t,S,x,y}^{\pi^\dagger}(u) - X_{t,S,x,y}^{\pi^*}(u) \qquad (9)$$

$$Y_{t,S,x,y}^{\pi^h}(u) = Y_{t,S,x,y}^{\pi^\dagger}(u) - Y_{t,S,x,y}^{\pi^*}(u) \qquad (10)$$

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

We shall see that the above definition of a hedging strategy is a correct one when we illustrate its numerical features in Section 5.

3. A market with transaction costs

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

$$dB(u) = rB(u)du \qquad (11)$$

$$dS(u) = S(u)[b\,du + \sigma\,dZ(u)] \qquad (12)$$

where $Z = \{Z(u), 0 \le u \le T\}$ is a one-dimensional standard Brownian motion defined on a complete probability space $(\Omega, \mathcal{F}, \mathbf{P})$. Denote by $\mathbf{F} = \{\mathcal{F}(u), 0 \le u \le T\}$ the $\mathbf{P}$-augmentation of the filtration $\mathcal{F}^Z(T) = \sigma(Z(u); 0 \le u \le T)$ generated by Z. The constant coefficients r, b, σ represent the risk-free 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 $\lambda(0 \le \lambda < 1)$ 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 $(0 \le \mu < 1)$ represents the proportional transaction cost rate associated with selling stock. In all other respects markets are assumed perfect.
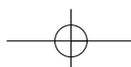
We shall also make use of a binomial approximation to the above market model for numerical computation. The bond and stock prices follow the discrete-time processes

$$B(u) + \delta B(u) \equiv B(u + \delta u) = \exp(r.\delta u)B(u) \tag{13}$$

$$S(u) + \delta S(u) \equiv S(u + \delta u) = \omega S(u) \tag{14}$$

where ω is a binomial random variable:

$$\omega = \exp\left[\left(b - \sigma^2/2\right)\delta u \pm \sigma\sqrt{\delta u}\right], \quad \text{each with probability } q = \frac{1}{2} \tag{15}$$

and δu is a small time interval.

Define $(L(u), M(u))$, a pair of **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

$$\mathrm{d}X(u) \equiv \mathrm{d}X_{t,S,x,y}^{L,M}(u) = rX(u)\,\mathrm{d}u - (1 + \lambda)\,S(u)\,\mathrm{d}L(u) + (1 - \mu)\,S(u)\,\mathrm{d}M(u)$$
$$\tag{16}$$

The number of shares held follows the process

$$\mathrm{d}Y(u) \equiv \mathrm{d}Y_{t,S,x,y}^{L,M}(u) = \mathrm{d}L(u) - \mathrm{d}M(u) \tag{17}$$

and the wealth of the investor is given by

$$W(u) \equiv W_{t,S,x,y}^{L,M}(u) = X_{t,S,x,y}^{L,M}(u) + Y_{t,S,x,y}^{L,M}(u)\,S(u) \tag{18}$$

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 maximize expected utility of wealth at time T. We introduce the set S, which defines the solvency region in the absence of an option trade, as

$$S = \left\{(S, x, y) \in \mathbf{R}^+ \times \mathbf{R}^2 \,\middle|\, x + (1 + \lambda)Sy \geq 0, \; x + (1 - \mu)Sy \geq 0\right\} \tag{19}$$

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

$$\left(S(u), X_{t,S,x,y}^{L,M}(u), Y_{t,S,x,y}^{L,M}(u)\right) \in S, \quad \text{almost surely, } \forall u \in [t, T] \tag{20}$$

For an investor who trades options at time t and then seeks to maximize expected utility of wealth, the set of admissable trading strategies is altered. For

example, when a contingent claim is written 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 optimization problem without the derivative security, though for computing hedging strategies this is not the case.

The value functions $V(t,S,x,y)$ and $V^{(o)}(t,S,x,y,\delta,p)$ satisfy the same dynamic programming equations but with different terminal boundary conditions. The function $F(t,S,x,y)$ of equation (7) is not necessarily a value function, but it 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 condition to be applied at the terminal time T for $V(t,S,x,y)$ is

$$V(T,S,x,y) = U(x+yS) \tag{21}$$

where it is assumed that no transaction costs are 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 (21) is replaced by

$$V(T,S,x,y) = U(x+c(y,S)) \tag{22}$$

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

$$c(y,S) = \begin{cases} (1+\lambda)yS, & \text{if } y<0 \\ (1-\mu)yS, & \text{if } y\geq 0 \end{cases} \tag{23}$$

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

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

$$V^{(o)}(T,S,x,y,\delta,p) = U\left(x+yS+\frac{\delta}{p}C(S)\right) \tag{24}$$

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

$$F(T,S,x,y) = C(S)U'(x+yS) \tag{25}$$

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 are presented below. They will be used in formulating a numerical algorithm in the next section.

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. 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) = $ 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 it therefore satisfies

$$V(t, S, x, y) = V\big(t, S, x - S(1 + \lambda)\delta L, y + \delta L\big) \quad \text{in BUY} \tag{26}$$

where δL, the number of shares bought, can take any positive value up to that required to take the portfolio to the boundary between the NT and BUY regions. Allowing $\delta L \downarrow 0$, (26) becomes

$$\frac{\partial V}{\partial y}(t, S, x, y) - (1 + \lambda)S\frac{\partial V}{\partial x}(t, S, x, y) = 0 \quad \text{in BUY} \tag{27}$$

Similarly, in the SELL region, the value function satisfies

$$V(t, S, x, y) = V\big(t, S, x + S(1 - \mu)\delta M, y - \delta M\big) \quad \text{in SELL} \tag{28}$$

where δM represents the number of shares sold. Letting $\delta M \downarrow 0$, (28) becomes

$$\frac{\partial V}{\partial y}(t, S, x, y) - (1 - \mu)S\frac{\partial V}{\partial x}(t, S, x, y) = 0 \quad \text{in SELL} \tag{29}$$

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

$$V(t, S, x, y) \geq V\big(t, S, x - S(1 + \lambda)\delta L, y + \delta L\big) \quad \text{in NT} \tag{30}$$

and

$$V(t, S, x, y) \geq V\big(t, S, x + S(1 - \mu)\delta M, y - \delta M\big) \quad \text{in NT} \tag{31}$$

which on expansion imply that the left-hand sides of (27) and (29) 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

$$V(t, S, x, y) = \mathbf{E}_{\delta t}V(t + \delta t, S + \delta S, x + \delta x, y) \quad \text{in NT} \tag{32}$$

where $\mathbf{E}_{\delta t}$ denotes expectation over the time interval δt. In the limit $\delta t \to 0$, δS and δx are given by (12) and (16), respectively (with $dL = dM = 0$ since we are

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

$$V_t + rxV_x + bSV_S + \frac{1}{2}\sigma^2 S^2 V_{SS} = 0 \quad \text{in NT} \tag{33}$$

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

These equations can be condensed into the PDE

$$\max\left[V_y - (1+\lambda)SV_x, \; -\left(V_y - (1-\mu)SV_x\right), \; V_t + rxV_x + bSV_S + \frac{1}{2}\sigma^2 S^2 V_{SS} \right] = 0 \tag{34}$$

The solution of the optimization problem is obtained by observing that if we can compute the value function in the NT region along with the boundaries of this region, we can calculate its value in the BUY and SELL regions using (26) and (28).

We use the equations (26), (28) and (32) and augment them with the specific properties of the optimal trading strategy to create a backward-recursive dynamic programming algorithm. Assume that the stock and bond prices evolve in discrete time according to (13)–(15). Then the discrete dynamic programming equation is

$$V(t,S,x,y) = \max_{(\delta L, \delta M)} \Big[\mathbf{E}_{\delta t} V\big(t + \delta t, \omega S, R(x - S(1+\lambda)\delta L), y + \delta L\big),$$

$$\mathbf{E}_{\delta t} V(t + \delta t, \omega S, Rx, y),$$

$$\mathbf{E}_{\delta t} V\big(t + \delta t, \omega S, R(x + S(1-\mu)\delta M), y - \delta M\big) \Big] \tag{35}$$

where $R \equiv \exp(r.\delta t)$ and the maximum is achieved by the first, second or third terms in (35) when the state (t,S,x,y) is in the BUY, NT and SELL regions, respectively.

Equation (35) expresses the value function at time t in terms of its counterpart at $t + \delta t$ by comparing the three possibilities: buying δL shares and allowing the stock to diffuse; not trading and allowing the stock to diffuse; or 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).
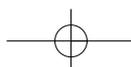
To implement the above algorithm we specialize to the case of exponential utility in the next section.

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

$$U(W) = -\exp(-\alpha W) \tag{36}$$

with constant risk-aversion parameter α. With this choice the investor's optimal trading strategy becomes independent of the wealth held in the risk-free asset. The assumption of exponential utility reduces the dimensionality of the optimization problem that we must solve. We defer to a later article the comparison of option prices generated by alternative choices of the utility function.

For exponential utility, the optimal trading strategy is characterized by a time-varying NT region with boundaries $\zeta_b(t) < \zeta_s(t)$, where $\zeta(t)$ represents the wealth held in the stock at time t. For logarithmic or power utility, we hypothesize that the results will be similar to those we present below for exponential utility, but $\zeta(t)$ will correspond to the ratio of wealth held in the stock to that held in the bond.

Define

$$H(t, S, y) := V(t, S, 0, y) \tag{37}$$

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

$$V(t, S, x, y) = H(t, S, y) \exp\left(-\alpha x \, e^{r(T-t)}\right) \tag{38}$$

The resulting reduction in dimensionality means that the discrete dynamic programming algorithm (35) reduces to

$$H(t, S, y) = \max_{(\delta L, \delta M)} \Big[\mathbf{E}_{\delta t} H(t + \delta t, \omega S, y + \delta L) \exp(\alpha S(1 + \lambda) \delta L . \beta(t)),$$

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

$$\mathbf{E}_{\delta t} H(t + \delta t, \omega S, y - \delta M), \exp\left(-\alpha S(1 - \mu) \delta M . \beta(t)\right) \Big] \tag{39}$$

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

For option price computations on a binomial tree, we need the boundaries of the NT region at each node of the tree. We chracterize these by the number of shares held at the NT boundaries, and these are therefore 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 $\lambda = \mu = 0$.
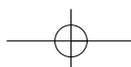
The optimal values of δL and δM, δL^* and δM^* satisfy

$$y + \delta L^* = y_b(t,S) \quad \text{and} \quad \delta M^* = 0, \qquad \text{if } y < y_b(t,S)$$

$$\delta L^* = \delta M^* = 0, \qquad \text{if } y_b(t,S) \le y \le y_s(t,S)$$

$$\delta L^* = 0 \quad \text{and} \quad y - \delta M^* = y_s(t,S), \qquad \text{if } y > y_s(t,S) \qquad (40)$$

Applying (40) and (38) to equations (26), (28) and (32), we obtain the following representations for $H(t,S,y)$ in the BUY, SELL and NT regions.

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

$$H(t,S,y) = H\big(t,S,y_b(t,S)\big) \exp\big(\alpha S(1+\lambda)(y_b(t,S) - y)\beta(t)\big) \qquad (41)$$

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

$$H(t,S,y) = H\big(t,S,y_s(t,S)\big) \exp\big(-\alpha S(1-\mu)(y - y_s(t,S))\beta(t)\big) \qquad (42)$$

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

$$H(t,S,y) = \mathbf{E}_{\delta t} H(t + \delta t, \omega S, y) \qquad (43)$$

Equations (41)–(43) give the value function $H(t,S,y)$ in the BUY, NT and SELL regions provided that 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 (39) 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 discretization 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 is accomplished by deriving analytically the NT boundaries at $T - \delta t$, as shown in the Appendix, and noting that the NT region is wider at this time than at any preceding time.[1] 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 maximization operator of (39) 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 maximization operator of (39) for increasing values of y in steps of h_y until the latter is greater

---

[1] This was confirmed by solving the problem without this assumption.
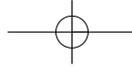
than or equal to the former at, say, y^s, which we assume satisfies $y^s = y_s(t, S)$, the boundary between the NT and SELL regions at the node (t, S).

3. Having located the NT boundaries for the node (t, S), the value function at all points outside the NT region is determined by assuming that the investor transacts to its boundaries (ie, applying equations (41) and (42)), while the function in the NT region is found by assuming that the investor does not transact and applying equation (43).

4. With exponential utility, the NT boundaries at any given time are characterized solely in terms of the wealth held in the stock.[2] 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 trivially. This property is not satisfied by the value function $V^{(o)}(t, S, x, y, \delta, p)$.

To summarize, the algorithm is very efficient because (1) it draws on the known properties of the value function $V(t, S, x, y)$ under exponential utility and (2) it restricts the state space over which we carry out the search for the NT boundaries by limiting the y-vector used to the interval $[y_b(T - \delta t, S), y_s(T - \delta t, S)]$, which is derived analytically. 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.

As an indication of the efficiency gains from the algorithm proposed in this article it is instructive to compare the computation times for option prices and optimal portfolios. For the numerical results under exponential utility presented in the next section, the solution of the basic portfolio problem without options, from which the fair price $\hat{p}$ is computed, takes approximately one-seventh of the computation time needed to solve the portfolio problem with options. To compute option prices in the latter case one must also conduct a search over different option prices to find a price which gives the same utility as without options, thus increasing the computation time still further.

4.1 Option prices

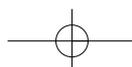To calculate option prices under exponential utility, the fair pricing formula (6) becomes

$$\hat{p}(t, S, y) = e^{-r(T-t)} \frac{G(t, S, y)}{H(t, S, y)} \tag{44}$$

where

$$G(t, S, y) \equiv \mathbf{E}_t\left[U\big(W^{\pi^*}_{t,S,0,y}(T)\big) C(S(T))\right] \tag{45}$$
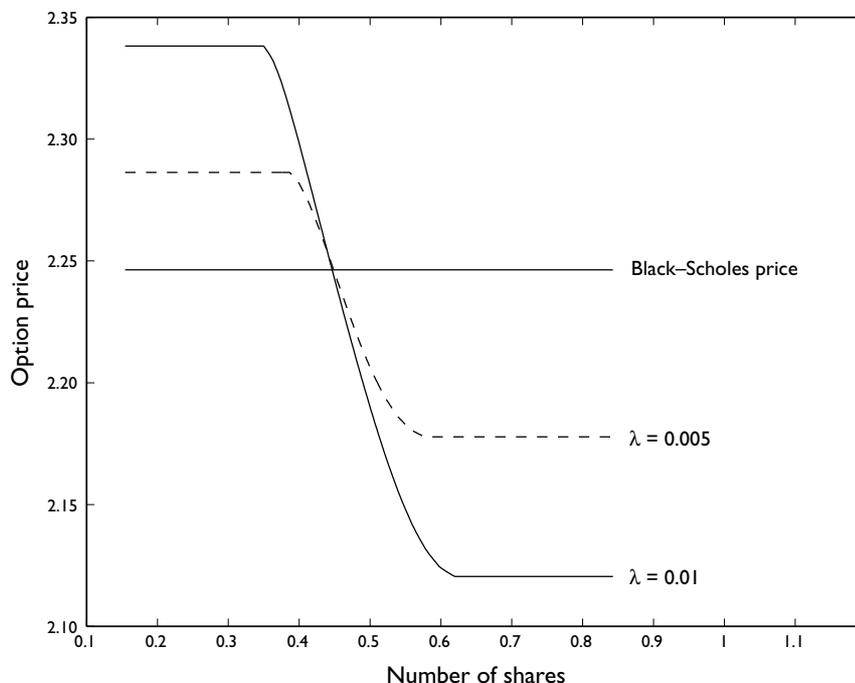
using the fact that $U'(w) = -\alpha U(w)$.

---

[2] This was confirmed by solving the problem without this assumption.

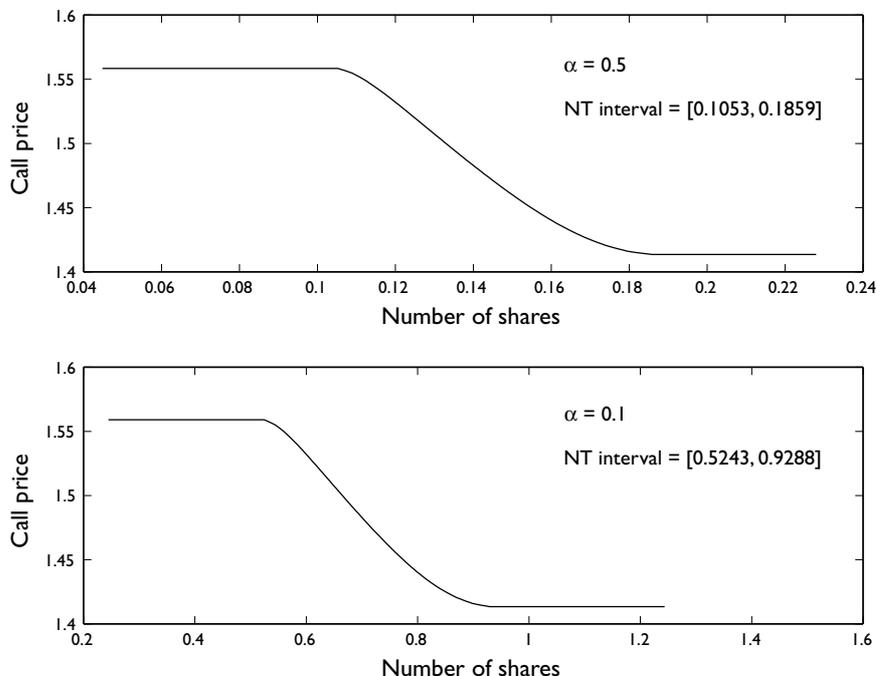**FIGURE 1**　At-the-money call option prices versus initial stock holding.

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. 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 time steps. We first confirmed some stylized facts about the investor's optimal trading strategy without options, which we summarize 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 and as we approach the horizon time T.
2. The NT region boundaries show a hyperbolic dependence on 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. An increase in risk aversion narrows the region of no transactions and shifts it to lower values of the stock holding.

FIGURE 2 Call option prices for different values of risk-aversion parameter α.

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 0.1.

Figure 1 shows at-the-money call option prices given by the general option pricing formula (44) plotted at time zero versus the investor's initial stock hold-ing, 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 to 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. It is interesting to observe that, depending on the investor's initial stock holding, the pricing method can produce a bid or ask price, or an intermediate price, which reflects the investor's current preferences for buying or selling the stock.

Figure 2 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

TABLE 1 Call bid and ask prices.

Strike	Ask price	Bid price	Leland ask	Leland bid
λ = 0.005, NT region = [0.3866, 0.5780]				
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
λ = 0.01, NT region = [0.3499, 0.6197]				
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
λ = 0.02, NT region = [0.2702, 0.7196]				
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
λ = 0.03, NT region = [0.1813, 0.8243]				
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	–

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 $\delta t = 0.02$.

spread becomes wider and is shifted to a higher 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 emphasizing. 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, which means that the investor has the *opportunity* to rehedge as frequently as the Leland strategy but chooses not to do so.

**FIGURE 3** Call option bid–ask values versus stock price
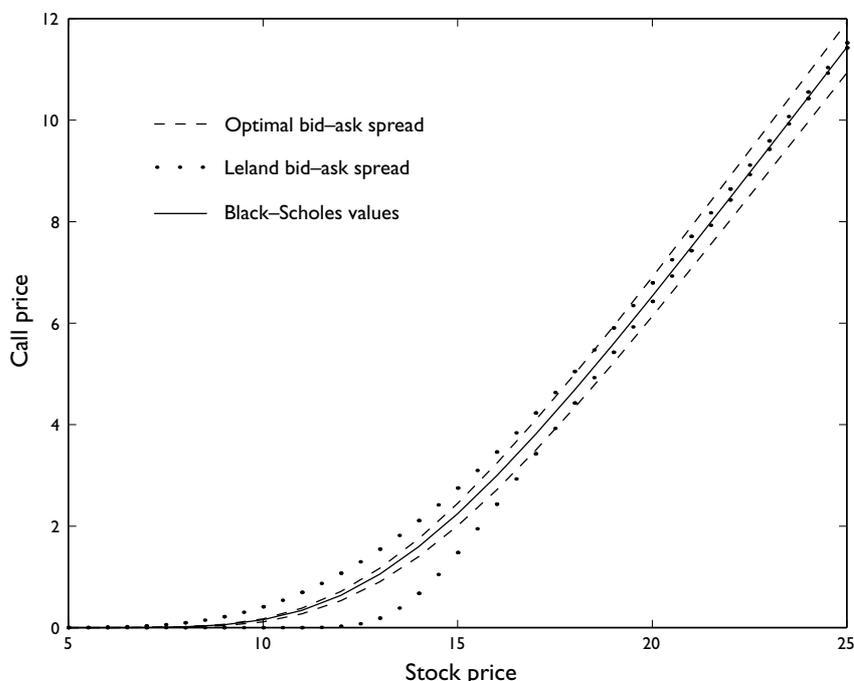
The parameters are $T = 1$ year, $r = 0.1$, $b = 0.15$, $\sigma = 0.25$, $\alpha = 0.1$, $K = 15$, $\lambda = 0.02$.

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 maximization.

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 3 is 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

FIGURE 4 Short call option hedging strategies versus stock price.
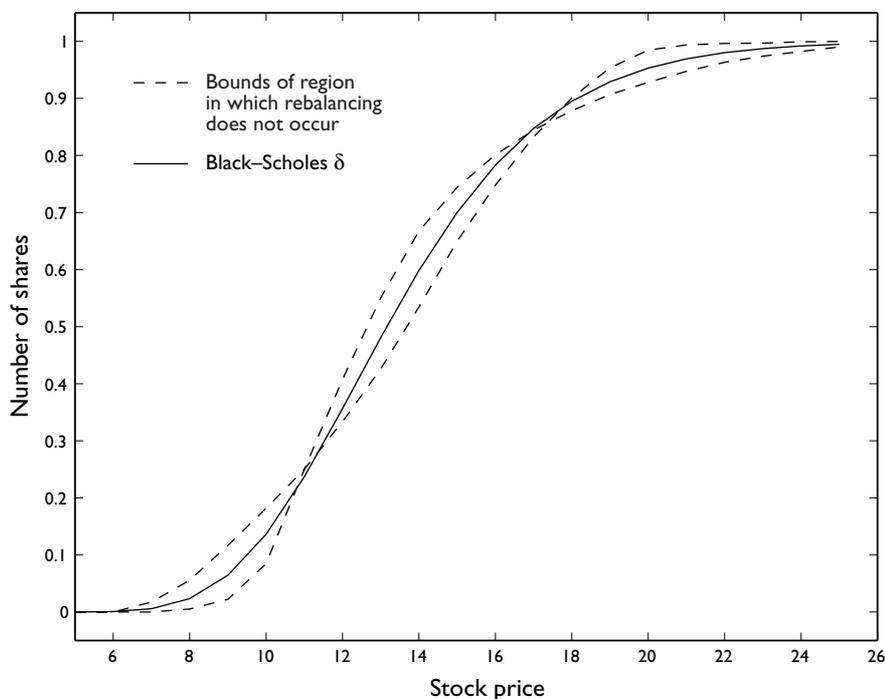
The call was sold at the "fair" asking price, and the other parameters are $T = 1$ year, $r = 0.1$, $b = 0.15$, $\sigma = 0.25$, $\alpha = 0.1$, $K = 15$, $\lambda = 0.005$.

tighter bounds on the option prices, except for the cases where the option is deep in-the-money, as before.

In Figure 4 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, while 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 those in Hodges and Neuberger (1989) and (for the limiting case of small transaction costs) Whalley and Wilmott (1997).

6. Conclusions

This article 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,

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 (39) for the value function $H(T - \delta t, S, y)$ reduces to

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

(46)

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

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

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

(47)

where ω_u and ω_d are the two possible realizations of the binomial random variable ω given in (15) (so that $q = \frac{1}{2}$), and the pseudo-probability q_+ is given by

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

(48)

Inserting the expression for δL^* into (46) gives the following representation for $H(T - \delta t, S, y)$ in the BUY region:

$$H(T - \delta t, S, y) = -\exp\left(-\alpha y RS(1 + \lambda)\right)\left(\frac{q}{q_+}\right)^{q_+}\left(\frac{1 - q}{1 - q_+}\right)^{(1 - q_+)}, \quad \text{in BUY} \quad (49)$$

We note that the value function's dependence on y and S enters via the product yS, the wealth held in the stock, as expected for an exponential utility function.

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

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

(50)

where the pseudo-probability q_- is given by

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

(51)

so that the value function in the SELL region is

$$H(T-\delta t, S, y) = -\exp\left(-\alpha y RS(1-\mu)\right)\left(\frac{q}{q_-}\right)^{q_-}\left(\frac{1-q}{1-q_-}\right)^{(1-q_-)}, \quad \text{in SELL} \quad (52)$$

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

$$H(T-\delta t, S, y) = -\left[q\exp(-\alpha y S\omega_u) + (1-q)\exp(-\alpha y S\omega_d)\right], \quad \text{in NT} \quad (53)$$

REFERENCES

Artzner, P., Delbaen, F., Eber, J. M., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance* **9**, 203–28.

Avellaneda, M., and Parás, A. (1994). Dynamic hedging portfolios for derivative securities in the presence of large transaction costs. *Applied Mathematical Finance* **1**, 165–93.

Bensaid, B., Lesne, J. P., Pagès, H., and Scheinkman, J. (1992). Derivative asset pricing with transaction costs. *Mathematical Finance* **2**, 63–86.

Bensoussan, A., and Julien, H. (2000). On the pricing of contingent claims with frictions. *Mathematical Finance* **10**, 89–108.

Bertsimas, D., Kogan, L., and Lo, A. W. (2001). Hedging derivative securities and incomplete markets: An ϵ-arbitrage approach. *Operations Research* **49**, 372–97.

Black, F., and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy* **81**, 637–54.

Boyle, P. P., and Vorst, T. C. F. (1992). Option replication in discrete time with transaction costs. *Journal of Finance* **47**, 271–93.

Clewlow, L., and Hodges, S. D. (1997). Optimal delta-hedging under transactions costs. *Journal of Economic Dynamics and Control* **21**, 1353–76.

Constantinides, G. M., and Zariphopoulou, T. (1999). Bounds on prices of contingent claims in an intertemporal economy with proportional transaction costs and general preferences. *Finance and Stochastics* **3**, 345–69.

Cvitanić, J., and Karatzas, I. (1996). Hedging and portfolio optimization under transaction costs: a martingale approach. *Mathematical Finance* **6**, 133–65.

Cvitanić, J., and Karatzas, I. (1999). On dynamic measures of risk. *Finance and Stochastics* **3**, 451–82.

Cvitanić, J., Pham, H., and Touzi, N. (1999). A closed-form solution to the problem of super-replication under transaction costs. *Finance and Stochastics* **3**, 35–54.

Davis, M. H. A. (1997). Option pricing in incomplete markets. In M. A. H. Dempster and S. R. Pliska (eds), *Mathematics of Derivative Securities*, pp. 216–26. Cambridge University Press, Cambridge, UK.

Davis, M. H. A., and Clark, J. M. C. (1994). A note on super-replicating strategies. *Philosophical Transactions of the Royal Society of London* A, **347**, 485–94.

Davis, M. H. A., and Norman, A. R. (1990). Portfolio selection with transaction costs. *Mathematics of Operations Research* **15**, 676–713.

Davis, M. H. A., and Zariphopoulou, T. (1995). American options and transaction fees. In M. H. A. Davis, D. Duffie, W. H. Fleming and S. E. Shreve (eds), *Mathematical Finance – IMA Volumes in Mathematics and its Applications* no. 65, pp. 47–61. Springer-Verlag, New York.

Davis, M. H. A., and Zervos, M. (1994). A problem of singular stochastic control with discretionary stopping. *Annals of Applied Probability* **4**, 226–40.

Davis, M. H. A., Panas, V. G., and Zariphopoulou, T. (1993). European option pricing with transaction costs. *SIAM Journal of Control and Optimization* **31**, 470–93.

Detemple, J., and Selden, L. (1991). A general equilibrium analysis of option and stock market interactions. *International Economic Review* **32**, 279–303.

Frittelli, M. (2000). The minimal entropy martingale measure and the valuation problem in incomplete markets. *Mathematical Finance* **10**, 39–52.

Hodges, S. D., and Neuberger, A. (1989). Optimal replication of contingent claims under transaction costs. *Review of Futures Markets* **8**, 222–39.

Karatzas, I., and Kou, S. (1996). On the pricing of contingent claims under constraints. *Annals of Applied Probability* **6**, 321–69.

Kushner, H. J. (1990). Numerical methods for stochastic control problems in continuous time. *SIAM Journal of Control and Optimization* **28**, 999–1048.

Kushner, H. J. (1997). Numerical methods for stochastic control problems in finance. In M. A. H. Dempster and S. R. Pliska (eds), *Mathematics of Derivative Securities*, pp. 504–27. Cambridge University Press, Cambridge, UK.

Kushner, H. J., and Martins, L. F. (1991). Numerical methods for stochastic singular control problems. *SIAM Journal of Control and Optimization* **29**, 1443–75.

Lamberton, D., Pham, H., and Schweizer, M. (1998). Local risk-minimization under transaction costs. *Mathematics of Operations Research* **23**, 585–612.

Leland, H. E. (1985). Option pricing and replication with transactions costs. *Journal of Finance* **40**, 1283–301.

Levental, S., and Skorohod, A. V. (1997). On the possibility of hedging options in the presence of transaction costs. *Annals of Applied Probability* **7**, 410–43.

Mercurio, F., and Vorst, T. C. F. (1997). Options pricing and hedging in discrete time with transaction costs. In M. A. H. Dempster and S. R. Pliska (eds), *Mathematics of Derivative Securities*, pp. 190–215. Cambridge University Press, Cambridge, UK.

Rabeau, N. M. (1996). Ph.D. Thesis, Imperial College, London.

Rouge, R., and El Karoui, N. (2000). Pricing via utility maximization and entropy. *Mathematical Finance* **10**, 259–76.

Rubinstein, M. (1998). Edgeworth binomial trees. *Journal of Derivatives* **5**, 20–7.

Schäl, M. (2000). Portfolio optimization and martingale measures. *Mathematical Finance* **10**, 289–303.

Stettner, L. (2000). Option pricing in discrete-time incomplete market models. *Mathematical Finance* **10**, 305–321.

Soner, H. M., Shreve, S. E., and Cvitanić, J. (1995). There is no nontrivial hedging portfolio for option pricing with transaction costs. *Annals of Applied Probability* **5**, 327–55.

Whalley, A. E., and Wilmott, P. (1997). An asymptotic analysis of an optimal hedging model for option pricing with transaction costs. *Mathematical Finance* **7**, 307–24.