RC22810 (W0305-139) May 29, 2003
Mathematics

IBM Research Report

Calibrated Option Bounds

Alan J. King, Matti Koivu*, Teemu Pennanen*

IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

*Helsinki School of Economics
Finland

Calibrated option bounds

Alan J. King*

Matti Koivu†

Teemu Pennanen†

November 19, 2001
Revision December 10, 2002

Abstract

This paper proposes a numerical approach for computing bounds for the arbitrage-free prices of an option when some options are available for trading. Convex duality reveals a close relationship with recently proposed calibration techniques and implied trees. Our approach is intimately related to the uncertain volatility model of Avellaneda, Levy and Parás, but it is more general in that it is not based on any particular form of the asset price process and does not require the seller's price of an option to be a differentiable function of the cash-flows of the option. Numerical tests on S&P 500 options demonstrate the accuracy and robustness of the proposed method.

1 Introduction

In incomplete but arbitrage-free markets, the price of an option should lie somewhere between the least cost of super-replication (seller's price) and the greatest amount a hypothetical buyer of the option could pay for it without facing the risk of negative terminal wealth (buyer's price). When frictionless trading is possible, these bounds can be expressed as the supremum and infimum values over the set of all pricing measures of the discounted expected future cash-flows of the option. Ritchken and Kuo [26] proposed to compute such bounds numerically by solving two optimization problems over the set of martingale measures for a discrete market model. The numerical results in [26] were based on the trinomial tree of Cox and Rubinstein [7], and their bounds converged near the Black-Scholes value of the option as the number of trading stages was increased. However, it is well known that the true value of an option can be far from the BS-value and it may very well fall outside of the RK-bounds obtained with a trinomial tree. The

*IBM Thomas J. Watson Research Center, kingaj@us.ibm.com
†Helsinki School of Economics, [koivu, pennanen]@hkkk.fi

basic problem with RK-bounds is that they are completely determined by the set of possible scenarios in the market model, so these scenarios should contain all the information that might affect the claim's price. In practice, such models are hard to come up with.

This problem is avoided to a large extent in the approach developed in Avellaneda, Levy and Parás [2] and Avellaneda and Parás [3], where market-traded options may be used in the trading strategies of the seller and the buyer. The market prices of the claims are then a natural input to the model since they reflect the current market expectations. Much like the RK-bounds, the bounds of [3] can be expressed as the supremum and infimum values of the expected discounted cash-flows over a set of pricing measures of the discounted expected cash-flows. The essential difference is that now the pricing measures are required to be consistent with the observed market prices.

This view of the method of Avellaneda and Parás [3] shows its close relationship with calibration techniques which try to specify a single pricing measure that is optimal in some sense among those measures that are consistent with the observed market prices (see the survey of Jackwerth [15]). The approach of [2, 3] has the advantage that it does not attempt to select a single measure, so it is free of possible misspecifications in the optimization criterion. Instead, it gives for each option separately a "calibrated" interval where its price is bound to be. Accordingly, it seems reasonable to call the resulting bounds *calibrated option bounds.*

The bounds of [2, 3] are based on an uncertain volatility model, where the volatility process is assumed to stay within a given volatility band that is input to the model. The pricing measures that were considered in the definition of the bounds were the ones that correspond to volatility processes varying within this band. The proposed algorithm is a two-stage procedure where a quasi-Newton algorithm is used to minimize a function whose evaluation is done through a solution of a nonlinear partial differential equation. A requirement of this approach is that the seller's and buyer's prices of a contingent claim be differentiable functions of the cash-flows. Unfortunately, this is a strong condition on the underlying market model, and it fails, for example, in discrete incomplete market models; see Section 3.

The main contribution of this paper is a new approach for computing calibrated option bounds. It is based on the use of fully discrete market models, which are well adapted to numerical computations and allow for simplified derivations of various duality relations as in King [20]. The method is more general than that of Avellaneda and Parás [3] in the sense that it requires no assumptions on the particular form of the security price processes, nor do the writer's and buyer's prices need to be differentiable functions of the cash-flows of a claim. Also, the claims are allowed to be contingent on multiple securities. Numerical tests on S&P 500 options demonstrate that the method is fast, accurate and robust with respect to changes in the underlying market model.

The next section introduces the notation and the basic structures that will be used in the rest of this paper. Section 3 reviews the convex duality approach of King [20] in studying the relations between hedging and martingale measures in incomplete markets. This serves as an introduction to Section 4 where we incorporate market-traded options into the model and show how they give rise to the calibrated option bounds. Section 5 outlines the computational advantages of our approach, and Section 6 presents results of our experimentation on S&P500 options.

2 Discrete market models

We start by describing the finite-dimensional market model that the computational framework will be based on. There are many treatments of discrete market models in the literature ([13, 11, 21, 10, 24], etc). Our notation follows King [20].

The market consists of $J+1$ tradable securities with prices $S_t = (S_t^0, \ldots, S_t^J)$, and it is assumed that investors have no influence on these prices. The probability space is that of a finite set of scenarios (price paths) taken by S_t over discrete points in time $t = 0, \ldots, T$. The path histories of the security prices up to time t correspond one-to-one with a subset $\mathcal{N}_t$ of the nodes $\mathcal{N}$ of a *scenario tree.* The initial state at time $t = 0$ corresponds to the *root node* which will be denoted by 0. We do not assume that the tree is recombining, so for each $n \in \mathcal{N} \setminus \{0\}$ there is a unique node denoted by $a(n)$ preceding n at time $t-1$. This assumption is essential in incomplete markets where trading strategies are in general path dependent; see the discussion in [11, Section IIIA]. The set of nodes that can be reached from n at time $t+1$ is denoted by $\mathcal{C}(n)$. The price vector at node n will be denoted by S_n.

A probability measure P attaches a weight $p_n > 0$ to each leaf node $n \in \mathcal{N}_T$ with $\sum_{n\in\mathcal{N}_T} p_n = 1$, and weights $p_n = \sum_{m\in\mathcal{C}(n)} p_m$ to each intermediate node $n \in \mathcal{N}_t$ for $t = T-1, \ldots, 0$. The expected value of S under P at time t is

$$E^P S_t := \sum_{n\in\mathcal{N}_t} S_n p_n.$$

A probability measure $Q = \{q_n\}_{n\in\mathcal{N}}$ is called a *martingale measure for* S if the value of S at every node n is equal to its one-step ahead conditional expectation under Q, or equivalently,

$$q_n S_n = \sum_{m\in\mathcal{C}(n)} q_m S_m.$$

It is worth emphasizing that we do not assume any particular form for P. In particular, the price process S could be non-Markovian, it could have stochastic volatility or arbitrary large jumps, etc. This generality is one of the advantages of our computational approach.

These models can be viewed as discretizations of more realistic market models, where the securities take values in a continuum of real numbers. Convergence of such discretizations has been studied in Pennanen and Koivu [23]. The next section describes a weakly convergent (see Billingsley [4]) discretization of discrete-time geometric Brownian motion.

2.1 An example: Gauss-Hermite processes

Consider an asset whose price S_t follows a continuous time geometric Brownian motion with daily drift d and volatility σ. For $t = 1, \dots, T$ its logarithm $\xi_t = \ln S_t$ satisfies

$$\xi_t = \xi_{t-1} + d_t + e_t, \quad e_t \sim N(0, \sigma_t), \tag{2.1}$$

where

$$d_t = l_t d \quad \text{and} \quad \sigma_t = \sqrt{l_t}\sigma,$$

and l_t is the length of period t in days.

Given the parameters of (2.1) and the initial value ξ_0, we generate a scenario tree using the approach described in Omberg [22] and [23]. We use the Gauss-Hermite quadrature to obtain a sample $(e_1^{i_1})_{i_1=1}^{\nu_1}$ of size ν_1 of e_1 with associated probabilities $(\pi_1^{i_1})_{i_1=1}^{\nu_1} \subset (0, \infty)$. This gives an approximation of the possible values of the logarithmic index at time $t = 1$:

$$\xi_1^{i_1} = \xi_0 + d_1 + e_1^{i_1} \quad i_1 = 1, \dots, \nu_1.$$

We then generate a sample $(e_2^{i})_{i=1}^{\nu_2}$ of the second period innovations e_2, and the possible values of the logarithmic index at time $t = 2$ are

$$\xi_2^{i_1,i_2} = \xi_1^{i_1} + d_1 + e_2^{i_2} \quad i_1 = 1, \dots, \nu_1, \quad i_2 = 1, \dots, \nu_2.$$

Proceeding this way for $t = 2, \dots, T$, we obtain a scenario tree whose nodes $\mathcal{N}_t$ at time t are labeled by the t-tuples $(i_1, \dots, i_t)$. Defining

$$\mathcal{N} = \mathcal{N}_1 \cup \ldots \cup \mathcal{N}_T, \tag{2.2}$$

$$a(i_1, \dots, i_t) = (i_1, \dots, i_{t-1}), \tag{2.3}$$

$$C(i_1, \dots, i_t) = \{(i_1, \dots, i_{t+1}) \in \mathcal{N}_{t+1} \mid i_{t+1} \in \{1, \dots, \nu_{t+1}\}\}, \tag{2.4}$$

$$S_n = \exp(\xi_n) \quad \forall n \in \mathcal{N}, \tag{2.5}$$

$$p_{(i_1,\dots,i_t)} = \pi_1^{i_1} \cdots \pi_t^{i_t} \tag{2.6}$$

puts us in the discrete setting described in the previous section. Omberg [22] called such discrete processes *Gauss-Hermite processes*. As the number of branches increases, the GH process converges weakly to the discrete time geometric Brownian motion; see [23, Proposition 5].

As noted in [22], if the periods are of equal length and we choose $\nu_t = 2$ for all $t = 1, \dots, T$, then we obtain the binomial tree of Jarrow and Rudd [17]. Note that the probabilities in the Gauss-Hermite processes do not depend on the parameters μ, σ or the step lenght l_t. In particular, they cannot become negative like in some of the better known trinomial trees for large values of l_t. For $\nu > 3$, the ν-nomial Gauss-Hermite trees are not recombining since then the jumps are not integer multiples of each other. However, in our approach, the recombination property is not of interest.

An attractive feature of Gauss-Hermite processes is that the discretized one-step conditional probabilities of the logarithmic index match a maximum number of moments of the normal distribution. More precisely, with ν branches, the GH quadrature matches $2\nu - 1$ moments; see for example Stoer and Bulirsch [30, Section 3.6]. In particular, the trinomial Gauss-Hermite tree has the first five moments of the normal distribution (not just the expectation and volatility).

3 Option bounds of Ritchken and Kuo

By a European-style contingent claim we will mean a stochastic cash-flow F with pay-outs $\{F_n\}_{n\in\mathcal{N}}$ that depend on the prices of the underlying securities S. The existence of investors who may trade the securities S without restrictions or transaction costs leads naturally to the Ritchken-Kuo bounds for the price of F. The portfolio of securities held by an investor in state $n \in \mathcal{N}$ is denoted $\theta_n = (\theta_n^0, \dots, \theta_n^J)$, and its value is

$$S_n \cdot \theta_n := \sum_{j=0}^{J} S_n^j \theta_n^j.$$

A portfolio process $\{\theta_n\}$ is said to *super-replicate* the claim's cash-flows if all trades are self-financing and the portfolio always has non-negative value. The *writer's price* of the claim is the smallest amount of current cash required to begin a super-replication process. In a discrete market model, this is the optimum value of the optimization problem

$$\text{(W)} \qquad \begin{array}{lllll} \min_{V,\theta} & V & & & \\ \text{subject to} & S_0 \cdot \theta_0 & = & V & \\ & S_n \cdot (\theta_n - \theta_{a(n)}) & = & -F_n & (n \in \mathcal{N}_t,\ t \geq 1) \\ & S_n \cdot \theta_n & \geq & 0 & (n \in \mathcal{N}_T). \end{array}$$

The optimal solution θ is the corresponding super-replication strategy or the *writer's hedge.*

We will follow King [20] and analyze the problem (W) through convex programming duality; see Appendix. To derive the dual problem of (W), we write the Lagrangian in the form

$$\begin{aligned} l(V,\theta;x,y) &= V + y_0[S_0 \cdot \theta_0 - V] + \sum_{t=1}^{T} \sum_{n\in\mathcal{N}_t} y_n[S_n \cdot (\theta_n - \theta_{a(n)}) + F_n] - \sum_{n\in\mathcal{N}_T} x_n S_n \cdot \theta_n \\ &= \sum_{t=1}^{T} \sum_{n\in\mathcal{N}_t} y_n F_n + \sum_{n\in\mathcal{N}_T} [y_n - x_n] S_n \cdot \theta_n \\ &\quad + [1 - y_0]V + \sum_{t=0}^{T-1} \sum_{n\in\mathcal{N}_t} [y_n S_n - \sum_{m\in\mathcal{C}(n)} y_m S_m] \cdot \theta_n \end{aligned}$$

(here it is understood that $l(V,\theta;x,y)$ takes the value $-\infty$ if $x \not\geq 0$). The minimum of $l(V,\theta;x,y)$ with respect to (V,θ) gives the objective to be maximized in the dual. If $S_n \neq 0$ for all $n \in \mathcal{N}_T$, the dual of (W) becomes (after clearing x),

$$\begin{array}{ll} \underset{y}{\mathit{maximize}} & \displaystyle\sum_{t=1}^{T} \sum_{n\in\mathcal{N}_t} y_n F_n \\ \mathit{subject\ to} & y_n \geq 0 \quad n \in \mathcal{N}_T, \\ & y_0 = 1, \\ & \displaystyle\sum_{m\in\mathcal{C}(n)} y_m S_m = y_n S_n \quad n \in \mathcal{N}_t,\ t = 0,\ldots,T-1. \end{array}$$

This is essentially problem (P3) of Ritchken and Kuo [26]. If one of the assets, say S^0, is strictly positive in every node, we can use $\beta_n = S_0^0/S_n^0$ as a discount-factor, and then the last set of constraints in the above dual problem shows that the numbers

$$q_n = y_n/\beta_n \quad n \in \mathcal{N}$$

define a *martingale measure* for the discounted asset price process $(\beta_n S_n)_{n\in\mathcal{N}}$. We will denote the set of martingale measures by $\mathcal{M}$.

Theorem 3.1. *If S^0 is always strictly positive, then the writer's price of F can be expressed as*

$$\text{(W')} \qquad \sup_{Q\in\mathcal{M}} \quad E^Q \sum_{t=1}^{T} \beta_t F_t.$$

In particular, the writer's price is finite if and only if $\mathcal{M} \neq \emptyset$.

Proof. According to the above observations, the dual problem of (W) can be expressed as (W'). Since S_0 is always strictly positive, F can be super-replicated (take, for example, V large enough and invest everything in S^0). Thus, the optimum value in (W) is either finite or $-\infty$. In both cases, the result follows from Theorem 6.1 in the Appendix. □

Note that the writer's price is a nonlinear function of F, except in the exceptional case where the set $\mathcal{M}$ is a singleton, i.e. the market model is complete. An infinite-dimensional version of the above result with general probability spaces can be found in Delbaen, Kabanov and Valkeila [8].

The *buyer's price* of F is the maximum amount one could pay for it without the risk of having negative terminal wealth. In our setting, this is the optimum value in

$$\text{(B)} \qquad \begin{array}{llcll} \mathit{maximize}_\theta & V & & & \\ \text{subject to} & S_0 \cdot \theta_0 & = & -V & \\ & S_n \cdot [\theta_n - \theta_{a(n)}] & = & F_n & (n \in \mathcal{N}_t,\ t \geq 1) \\ & S_n \cdot \theta_n & \geq & 0 & (n \in \mathcal{N}_T). \end{array}$$

The optimal solution θ is the corresponding *buyer's hedge.* The derivation of the dual problem in this case shows that the buyer's price equals

$$\text{(B')} \qquad \inf_{Q\in\mathcal{M}} \quad E^Q \sum_{t=1}^{T} \beta_t F_t,$$

which is the same as (W') except that sup has been replaced by inf. The bounds (W') and (B') were introduced by Ritchken and Kuo [26]. The interval between the buyer's and seller's prices describes the possible range of arbitrage-free valuations of the claim in an incomplete market.

The above bounds depend on the measure P only through its support, i.e. the finite number of scenarios in the market model. Since these scenarios are the only input to the option bounds, they should contain all the relevant information that might affect the claims price. The numerical results of [26] were based on the trinomial tree of Cox and Rubinstein [7], and the resulting bounds converged near the Black-Scholes option value as the number of stages T was increased. However, it is well known that the true value of an option can be far from the BS-price and it may very well fall outside of the RK-bounds obtained with a trinomial tree.

In order to get more reasonable bounds with the RK-approach (or any other approach that is based only on modeling the underlying), one could try to use market models that better capture market expectations. There are, of course, many advances in this direction such as the jump-diffusion and stochastic volatility models.

Another possibility is to deduce market expectations from the prices of other market-traded securities. This is the idea behind model calibration techniques.

4 Calibrated option bounds

When there are options (other than F) available for trading, it is natural to try to use them as parts of a hedging strategy. If everything else remains unchanged, this can only improve the investors' situation. In particular, it can make the writer's price lower and buyer's price higher thus narrowing down the arbitrage interval.

Let G^k, $k = 1,\ldots,K$ be contingent claims with bid-ask prices $C_b^k \leq C_a^k$ and payoffs G_n^k. Allowing the writer to apply buy-and-hold strategies on these options leads to the following modification of (W).

$$\text{(WC)} \qquad \begin{array}{lll} \min_{V,\theta,\xi_+,\xi_-} & V & \\ \text{subject to} & S_0\cdot\theta_0 + C_a\cdot\xi_+ - C_b\cdot\xi_- = V & \\ & S_n\cdot(\theta_n-\theta_{a(n)}) = G_n\cdot(\xi_+-\xi_-) - F_n & (n\in\mathcal{N}_t,\ t\geq 1), \\ & S_n\cdot\theta_n \geq 0 & (n\in\mathcal{N}_T), \\ & \xi_+,\xi_- \geq 0, & \end{array}$$

where ξ_+^i and ξ_-^i are the amounts bought and shorted of G^k at time $t = 0$. Choosing $\xi_+ = \xi_- = 0$ reduces problem (WC) to (W), so its optimal value is at most that of (W). In general, the optimal hedge will be a mixture of a dynamic hedge in S and a buy-and-hold hedge in G.

Much like in Section 3, we will derive an expression for the optimal value of (WC) in terms of martingale measures. Again, we do this by deriving the dual problem. Writing the Lagrangian for (WC) as

$$\begin{aligned} l(V, \theta, \xi_+, \xi_-; x, y) &= V + y_0[S_0 \cdot \theta_0 + C_a \cdot \xi_+ - C_b \cdot \xi_- - V] - \sum_{n \in \mathcal{N}_T} x_n S_n \cdot \theta_n \\ &\quad + \sum_{t=1}^{T} \sum_{n \in \mathcal{N}_t} y_n [S_n \cdot (\theta_n - \theta_{a(n)}) - F_n + G_n \cdot (\xi_+ - \xi_-)] \\ &= \sum_{t=1}^{T} \sum_{n \in \mathcal{N}_t} y_n F_n + \sum_{n \in \mathcal{N}_T} [y_n - x_n] S_n \cdot \theta_n \\ &\quad + [1 - y_0]V + \sum_{t=0}^{T-1} \sum_{n \in \mathcal{N}_t} [y_n S_n - \sum_{m \in \mathcal{C}(n)} y_m S_m] \cdot \theta_n \\ &\quad + [C_a - \sum_{t=1}^{T} \sum_{n \in \mathcal{N}_t} y_n G_n] \cdot \xi_+ - [C_b - \sum_{t=1}^{T} \sum_{n \in \mathcal{N}_t} y_n G_n] \cdot \xi_-, \end{aligned}$$

we get the following dual

$$\begin{array}{ll} \underset{y}{maximize} & \displaystyle\sum_{t=1}^{T} \sum_{n \in \mathcal{N}_t} y_n F_n \\ subject\ to & y_n \geq 0 \qquad n \in \mathcal{N}_T, \\ & y_0 = 1, \\ & \displaystyle\sum_{m \in \mathcal{C}(n)} y_m S_m = y_n S_n \quad n \in \mathcal{N}_t,\ t = 1, \ldots, T-1, \\ & \displaystyle\sum_{t=1}^{T} \sum_{n \in \mathcal{N}_t} y_n G_n \leq C_a, \\ & \displaystyle\sum_{t=1}^{T} \sum_{n \in \mathcal{N}_t} y_n G_n \geq C_b. \end{array}$$

Note that this is exactly the same as the dual of (W) except for the two sets of additional constraints. Again, if S^0 is always strictly positive, we can write the dual in terms of martingale measures and we can apply Theorem 6.1 of Appendix to get the following.

Theorem 4.1. *If S^0 is always strictly positive, then the writer's price of F in the presence of the contingent claims G can be expressed as*

$$\text{(WC')} \qquad \sup_{Q \in \mathcal{M}_C} \quad E^Q \sum_{t=1}^{T} \beta_t F_t,$$

where

$$\mathcal{M}_C = \{Q \in \mathcal{M} \mid C_b \leq E^Q \sum_{t=1}^{T} \beta_t G_t \leq C_a\}.$$

In particular, the writer's price is finite if and only if $\mathcal{M}_C \neq \emptyset$.

The set $\mathcal{M}_C$ can be thought of as a set of martingale measures that have been calibrated to the observed market prices. Theorem 4.1 corresponds to Avellanedas and Parás [3, Proposition 2], where the set of martingale measures $\mathcal{M}$ was replaced by a set of measures corresponding to different volatility processes varying within a given band. In a sense, the above result is a nonparametric version of [3, Proposition 2] since it applies to arbitrary (but finite) asset price processes.

Much like above, the buyer's problem in the presence of market-traded claims becomes

$$\text{(BC)} \qquad \begin{aligned} &\max_{V,\theta,\xi_+,\xi_-} \quad V \\ &\text{subject to} \quad S_0 \cdot \theta_0 + C_a \cdot \xi_+ - C_b \cdot \xi_- = -V \\ &\qquad S_n \cdot (\theta_n - \theta_{a(n)}) = G_n \cdot (\xi_+ - \xi_-) + F_n \quad (n \in \mathcal{N}_t,\ t \geq 1), \\ &\qquad S_n \cdot \theta_n \geq 0 \quad (n \in \mathcal{N}_T), \\ &\qquad \xi_+, \xi_- \geq 0, \end{aligned}$$

and its dual can be written as

$$\text{(BC')} \qquad \inf_{Q \in \mathcal{M}_C} \quad E^Q \sum_{t=1}^{T} \beta_t F_t$$

Thus, when the claims G are available for trading, the arbitrage free-prices of a claim F are bounded by the optimal values of (BC') and (WC'). These problems are the same as (B') and (W') except that now the martingale measures are restricted by the constraints that require that they price the benchmark claims G consistently with their observed values. The additional constraints reduce the set of possible pricing measures, so the bounds obtained from (BC') and (WC') are never wider than the bounds of Ritchken and Kuo.

We will call the the minimum and maximum values in (BC') and (WC'), respectively, *calibrated option bounds.* This terminology is consistent with the literature on implied trees and various calibration techniques. The methods proposed in Rubinstein [29], Avellaneda et.al. [1], King et.al. [18], and Borwein et.al. [6] are designed to find a single measure that optimizes some criterion among all the measures that are consistent with the observed market prices. For instance, Jackwerth and Rubinstein [16] look for the consistent measure that maximizes the smoothness of its density function, and [1, 18, 6] try to find the one that minimizes the Kullback-Leibler entropy relative to a user-selected prior.

The calibrated option bounds (BC') and (WC') do not depend on any user specified inputs other than the discrete support for the security price process. Moreover, these bounds give for each option separately an interval where its price is bound to be. When applying pricing measures obtained by optimization of some user-specified criterion, as in [16, 1, 18, 6], it may be hard to tell how much the resulting valuations depend on the choice of the criterion. The calibrated option bounds could be used as an independent set of error bounds in that case. The numerical tests in Section 6 demonstrate that these bounds can be made quite tight by the available market information, and moreover do not depend strongly on the choice of support for the security price process. Another interesting feature of the calibrated option bounds is that their computation also yields hedging strategies for the option being priced.

Our approach to obtaining the calibrated option bounds can be shown to be strongly related to the *Lagrangian Uncertain Volatility Model* of Avellaneda and Parás [3] (see also Levy, Avellaneda and Parás [2]), which studies pricing and hedging of contingent claims under stochastic volatility when the volatility is assumed to stay within prespecified levels. To compare our approach with that of [3], let us write (WC) as a two-phase minimization problem

$$\underset{\xi_+,\xi_-\geq 0}{minimize} \quad C_a\cdot\xi_+ - C_b\cdot\xi_- + f(\xi_+,\xi_-), \tag{4.1}$$

where $f(\xi_+,\xi_-)$ is the writer's price of the residual cash-flow $F - G\cdot(\xi_+ - \xi_-)$. It is easy to check that a vector $(\bar{\xi}_+,\bar{\xi}_-,\bar{\theta})$ solves (WC) if and only if $(\bar{\xi}_+,\bar{\xi}_-)$ solves (4.1) and $\bar{\theta}$ is the writer's hedge for $F - G\cdot(\bar{\xi}_+ - \bar{\xi}_-)$. Using Theorem 3.1 we can write (4.1) as

$$\underset{\xi_+,\xi_-\geq 0}{minimize} \quad C_a\cdot\xi_+ - C_b\cdot\xi_- + \underset{Q\in\mathcal{M}}{max}\ E^Q\sum_{t=1}^{T}\beta_t[F_t - G_t\cdot(\xi_+ - \xi_-)]. \tag{4.2}$$

This is similar to problem (10) of [3], but there $\mathcal{M}$ is replaced by the set of probability measures corresponding volatility processes varying within given bounds. Our approach does not depend on the form of the asset price process and it can take into account multiple underlying securities.

5 Computational issues

The algorithm proposed in [3] is based on a quasi-Newton method that iteratively minimizes (over the option holdings) the writer's price of the residual cash-flow (function f in (4.1)). Quasi-Newton methods are a class of algorithms for solving unconstrained minimization problems with differentiable objective functions. They proceed by evaluating the objective and its gradient along the iterates; see for example Polak [25]. In [3], the evaluation of the objective (writer's price) was done through numerical solution of a nonlinear PDE. This can be a hard problem in itself, but a more serious problem is the computation of the gradient vectors. In general, in incomplete markets, there is no reason to assume that the writer's price is a differentiable function of the cash-flows. In such a case the gradient may not be everywhere defined, and quasi-Newton methods may fail to converge. For example, in a discrete probability space the writer's price is a piecewise linear function.

One could modify the approach of [3] by replacing the quasi-Newton algorithm with an algorithm for nondifferentiable convex minimization [14]. The writer's price is the pointwise supremum of linear functions of the cash-flow so it is convex. The *subgradient vectors* used by such algorithms would be the measures $Q \in \mathcal{M}$ attaining the maximum in the writer's price [28, Theorem 24(b)].

Our derivation of the calibrated option bounds suggests another approach: instead of a two-stage procedure for solving (4.2) we will solve problem (WC) all at once. This will not only give us the writer's price but also the optimal hedge in terms of the options and the underlying securities. Problem (WC) is a linear program, although it can be a very large one: it has $|\mathcal{N}|(2K+J+1)$ variables and $|\mathcal{N}|(J+1)+|\mathcal{N}_T|$ constraints, where $|\mathcal{N}|$ is the number of nodes in the scenario tree and $|\mathcal{N}_T|$ is the number of leaf nodes. If the size of the scenario tree increases, these numbers can become very large (recall that we are not working with recombining trees). However, today's generation of solvers and computers can solve linear programs with 1000's of variables in a matter of seconds.

The main problem in setting up and solving the model is the generation of scenario trees and communicating them and the optimization model into a solver. We have written a C-program that produces scenario trees with user specified features. The program takes as input a time series model to be discretized and produces a tree with given given period and branching structures. The tree can be incorporated into an optimization model which is then sent to an appropriate solver. Instead of the problems (BC') and (WC') we will set up and solve the corresponding primal problems (WC) and (BC). These turn out to be easier to set up and they are more natural from the point of view of hedging. From the optimal solution of the primal (the hedge) one can readily obtain the solution of the dual (the pricing measure). Moreover, the approach of numerically solving problems (WC) and (BC) allows various generalizations in the model. For example, portfolio constraints can be incorporated simply by restricting the variables (ξ_+, ξ_-, θ), and transactions costs by including bid-ask spreads on the underlying. Such modifications do not add much to the numerical complexity of the problem. But but they do affect the interpretation of solutions of the dual problem, which may no longer be martingale measures (cf. King [20]).

The above approach where one solves numerically a discrete version of a stochastic sequential decision making problem (hedging in our case) is known as *stochastic programming* in the optimization literature. There are many efficient computational approaches for that purpose (cf. [19], [5]). Other applications of stochastic programming to finance can be found in Gondzio, Kouwenberg and Vorst [12], Dempster and Thompson [9] and Edirisinghe, Naik and Uppal [11].

6 Numerical tests with S&P500 options

Table 1 displays the bid and ask closing prices of 48 European call and put options on the S&P500 index on September 10, 2002. The columns labeled STR and MAT give the strike prices and maturities, respectively.

Table 1: Options data

Call options				Put options			
STR	MAT	C_b	C_a	STR	MAT	C_b	C_a
890	17	31.5	33.5	750	17	0.4	0.6
900	17	24.4	26.4	790	17	1	1.3
905	17	21.2	23.2	800	17	1.3	1.65
910	17	18.5	20.1	825	17	2.5	2.85
915	17	15.8	17.4	830	17	2.6	3.1
925	17	11.2	12.6	840	17	3.4	3.8
935	17	7.6	8.6	850	17	3.9	4.7
950	17	3.8	4.6	860	17	5.5	5.8
955	17	3	3.7	875	17	7.2	7.8
975	17	0.95	1.45	885	17	9.4	10.4
980	17	0.65	1.15	750	37	5.5	5.9
900	37	42.3	44.3	775	37	6.9	7.7
925	37	28.2	29.6	800	37	9.3	10
950	37	17.5	19	850	37	16.7	18.3
875	100	77.1	79.1	875	37	23	24.3
900	100	61.6	63.6	900	37	31	33
950	100	35.8	37.8	925	37	41.8	43.8
975	100	26	28	975	37	73	75
995	100	19.9	21.5	995	37	88.9	90.9
1025	100	12.6	14.2	650	100	5.7	6.7
1100	100	3.4	3.8	700	100	9.2	10.2
				750	100	14.7	15.8
				775	100	17.6	19.2
				800	100	21.7	23.7
				850	100	33.3	35.3
				875	100	40.9	42.9
				900	100	50.3	52.3

We will compute calibrated bounds for each of the 48 options one at a time by using the remaining 47 options as the calibrating claims. In other words, we will solve problems (BC) and (WC) 48 times by letting F run over the 48 options and each time using the remaining 47 options as our "market-traded" claims G. The resulting values can then be compared with the actual market prices in Table 1.

6.1 A three-stage model

We let S^1 be the S&P500 index, and use $S = (1, S^1)$ as the dynamically traded securities. The period structure in the model is chosen according to the maturities of the options. That is, we assume that trading occurs at 0, 17, 37, and 100 days. A scenario tree is built by approximating the development of S^1 by the Gauss-Hermite process of Section 2.1. We choose the branching structure (50, 10, 10), or in the notation of Section 2.1, $\nu_1 = 50$, $\nu_2 = 10$ and $\nu_3 = 10$. This results in 5000 scenarios; see Figure 1.

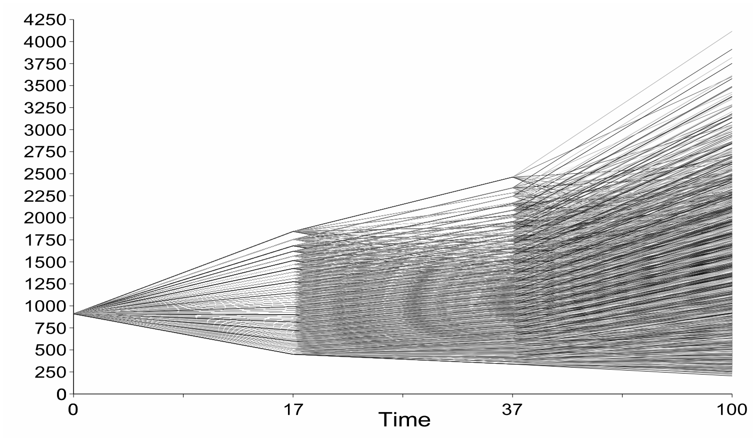

Figure 1: Scenario tree with branching structure (50, 10, 10).

We generated the optimization model with AMPL modeling language, and used CPLEX for the numerical solution. It took approximately 7 minutes for an Intel Pentium 4, 2.2 GHz processor with 512 RAM to set up and solve all the 96 problems (two for each option) as described in Section 5. The resulting calibrated bounds are given in Table 2.

Table 2: Calibrated bounds with a 3-period model

Call options						Put options					
STR	MAT	Cb	Ca	BP	WP	STR	MAT	Cb	Ca	BP	WP
890	17	31.5	33.5	30.98	32.00	750	17	0.4	0.6	0.00	1.15
900	17	24.4	26.4	24.67	26.05	790	17	1	1.3	0.87	1.44
905	17	21.2	23.2	21.79	23.08	800	17	1.3	1.65	1.16	1.74
910	17	18.5	20.1	18.93	20.30	825	17	2.5	2.85	2.18	2.86
915	17	15.8	17.4	16.08	17.51	830	17	2.6	3.1	2.74	3.15
925	17	11.2	12.6	10.43	12.77	840	17	3.4	3.8	3.20	3.90
935	17	7.6	8.6	7.68	9.09	850	17	3.9	4.7	4.40	4.80
950	17	3.8	4.6	3.39	4.79	860	17	5.5	5.8	4.61	5.94
955	17	3	3.7	2.99	3.89	875	17	7.2	7.8	6.80	7.98
975	17	0.95	1.45	0.65	1.66	885	17	9.4	10.4	9.97	10.78
980	17	0.65	1.15	0.66	1.44	750	37	5.5	5.9	3.80	6.64
900	37	42.3	44.3	40.58	42.58	775	37	6.9	7.7	6.33	7.95
925	37	28.2	29.6	26.38	28.38	800	37	9.3	10	7.90	11.23
950	37	17.5	19	13.82	18.98	850	37	16.7	18.3	13.49	19.53
875	100	77.1	79.1	75.48	77.48	875	37	23	24.3	21.77	25.65
900	100	61.6	63.6	59.88	61.88	900	37	31	33	32.72	34.05
950	100	35.8	37.8	32.27	39.29	925	37	41.8	43.8	43.62	45.02
975	100	26	28	23.71	28.74	975	37	73	75	72.24	76.91

995	100	19.9	21.5	17.73	22.48	995	37	88.9	90.9	87.09	94.08
1025	100	12.6	14.2	8.02	16.44	650	100	5.7	6.7	2.60	8.58
1100	100	3.4	3.8	0.00	12.80	700	100	9.2	10.2	6.65	11.25
						750	100	14.7	15.8	11.80	16.20
						775	100	17.6	19.2	16.95	19.75
						800	100	21.7	23.7	20.07	24.57
						850	100	33.3	35.3	32.74	36.50
						875	100	40.9	42.9	42.52	43.80
						900	100	50.3	52.3	52.02	54.02

Figure 2 plots these values along with the true bid/ask prices of Table 1. Except for some cases (most notably, some deep out of the money options) the calibrated bounds are close to the true bid/ask values. In some instances, one or both of the bounds even fall strictly between the true spread. In general, good bounds seem to result when there are many benchmark options with strikes close to the strike of the option we are trying to price. This seems rather natural in view of the fact that the bounds are obtained by trying to hedge the cash-flows of a given option using market-traded options. Good hedges are easier to come up when there are available options that are similar to the one being hedged.

6.2 Sensitivity analysis

Finite scenario trees only approximate the true asset price process. It is thus natural to ask to what extent the bounds depend on the choice of a tree. This can be studied by changing the tree and recomputing the bounds.

Table 3 displays the bounds obtained with trees that have the same period structure as above but with branching structures $(60, 10, 10)$ and $(70, 10, 10)$, respectively. For most options the resulting bounds are almost identical to those in Table 2.

The next experiment consists of adding one more trading period to the model. In addition to days 0, 17, 37 and 100, we now allow trading also in day 8, thus obtaining a four-period model. The branching structure is set at $(20, 10, 10, 10)$. The results are shown in Table 4. Again, for most options the changes are very small.

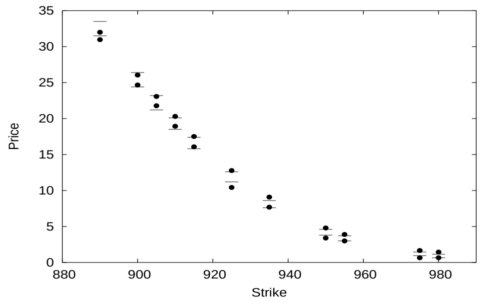

(a) Maturity 17 calls.

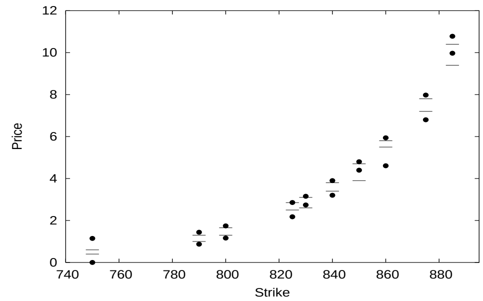

(b) Maturity 17 puts.

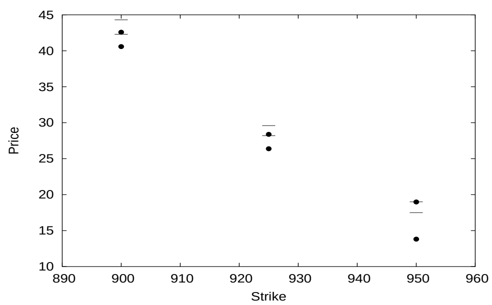

(c) Maturity 37 calls.

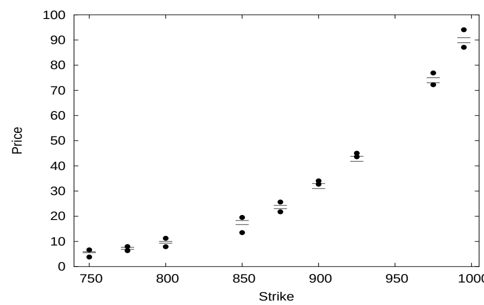

(d) Maturity 37 puts.

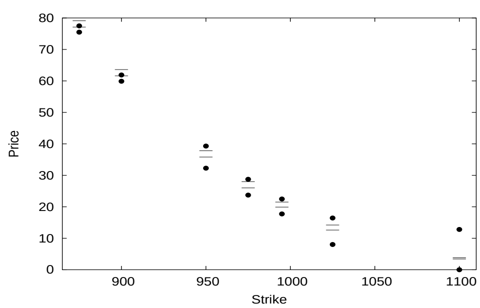

(e) Maturity 100 calls.

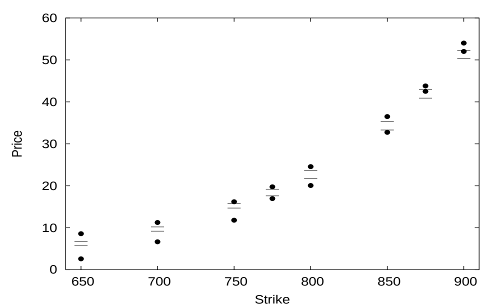

(f) Maturity 100 puts.

Figure 2: Calibrated bounds (•) and the true bid/ask prices (−).

Table 3: Bounds obtained with larger trees. The subscript 60 (resp. 70) refers to the tree with the branching structure $(60, 10, 10)$ (resp. $(70, 10, 10)$)

Call options						Put options					
STR	MAT	BC_{60}	WC_{60}	BC_{70}	WC_{70}	STR	MAT	BC_{60}	WC_{60}	BC_{70}	WC_{70}
890	17	30.79	31.91	30.79	31.73	750	17	0.00	1.16	0.00	1.17
900	17	25.10	26.01	25.32	25.91	790	17	0.70	1.44	0.81	1.44
905	17	22.03	23.05	22.23	23.01	800	17	1.10	1.74	1.14	1.74
910	17	18.99	20.30	19.14	20.27	825	17	2.09	2.86	2.25	2.86
915	17	15.96	17.57	16.04	17.60	830	17	2.80	3.17	2.67	3.17

925	17	10.71	12.72	11.09	12.72	840	17	3.01	3.90	3.33	3.90
935	17	7.22	9.25	6.74	9.40	850	17	4.31	4.77	4.13	4.77
950	17	3.52	4.89	3.80	4.88	860	17	4.35	5.94	4.35	5.71
955	17	2.67	3.97	2.20	3.97	875	17	7.42	8.25	7.65	8.46
975	17	0.72	1.66	0.88	1.60	885	17	9.77	10.29	9.77	10.06
980	17	0.42	1.44	0.45	1.34	750	37	3.80	6.65	3.80	6.65
900	37	40.58	42.58	40.58	42.58	775	37	6.32	7.95	6.31	7.95
925	37	26.68	28.38	26.85	28.38	800	37	7.90	11.23	7.90	11.23
950	37	13.85	18.98	13.96	18.95	850	37	13.58	19.53	13.73	19.53
875	100	75.48	77.48	75.48	77.48	875	37	21.99	25.65	22.26	25.65
900	100	59.88	61.88	59.88	61.88	900	37	32.72	34.05	32.72	34.05
950	100	31.96	39.29	32.18	39.29	925	37	43.62	45.02	43.62	45.02
975	100	23.83	28.74	24.19	28.74	975	37	72.53	76.91	72.80	76.91
995	100	17.73	22.48	18.06	22.48	995	37	86.87	94.08	87.31	94.08
1025	100	7.97	16.44	7.97	16.44	650	100	2.60	8.58	2.60	8.58
1100	100	0.00	12.82	0.00	12.87	700	100	6.65	11.25	6.65	11.25
						750	100	11.83	16.20	11.87	16.20
						775	100	16.95	19.75	16.95	19.75
						800	100	20.07	24.57	20.07	24.57
						850	100	32.74	36.50	32.74	36.50
						875	100	42.52	43.80	42.52	43.80
						900	100	52.02	54.02	52.02	54.02

Table 4: Bounds obtained with a four period model with branching structure $(20, 10, 10, 10)$

Call options				Put options			
STR	MAT	BC	WC	STR	MAT	BC	WC
890	17	30.10	32.00	750	17	0.00	1.07
900	17	24.59	26.05	790	17	0.74	1.44
905	17	21.06	23.08	800	17	1.10	1.74
910	17	18.46	20.30	825	17	2.08	2.86
915	17	15.20	17.60	830	17	2.68	3.17
925	17	10.70	13.00	840	17	3.07	3.90
935	17	6.73	9.40	850	17	3.97	4.80
950	17	3.53	4.93	860	17	4.37	5.94
955	17	2.47	3.97	875	17	7.17	8.35
975	17	0.66	1.66	885	17	9.77	11.32
980	17	0.41	1.43	750	37	3.80	6.64
900	37	40.58	42.58	775	37	6.33	7.95
925	37	26.38	28.38	800	37	7.90	11.23
950	37	13.82	18.98	850	37	13.25	19.53
875	100	75.48	77.48	875	37	21.71	25.65
900	100	59.88	61.88	900	37	32.72	34.05
950	100	31.97	39.29	925	37	43.62	45.02
975	100	23.71	28.74	975	37	72.27	76.91
995	100	17.50	22.48	995	37	86.49	94.08
1025	100	7.97	16.44	650	100	2.60	8.58
1100	100	0.00	12.90	700	100	6.65	11.25
				750	100	11.74	16.20
				775	100	16.95	19.75
				800	100	20.07	24.57

850	100	32.74	36.50
875	100	42.52	43.80
900	100	52.02	54.02

Appendix

Let f_i, $i = 1, \ldots, m$ be convex functions from $\mathbf{R}^n$ to $\mathbf{R}$, and consider the optimization problem

$$\text{(P)} \qquad \begin{array}{lll} \mathit{minimize}_x & f_0(x) & \\ \mathit{subject\ to} & f_i(x) \le 0 & i = 1, \ldots, r, \\ & f_i(x) = 0 & 1 = r+1, \ldots, m. \end{array}$$

Define the *Lagrangian*

$$l(x,y) = \begin{cases} f_0(x) + \sum_{i=1}^m y_i f_i(x) & \text{if } y_i \ge 0 \text{ for } i = 1, \ldots, r, \\ -\infty & \text{otherwise,} \end{cases}$$

and consider the *dual problem*

$$\text{(D)} \qquad \underset{y}{\mathit{maximize}} \quad g(y),$$

where the function g is defined by

$$g(y) = \inf_x l(x,y).$$

Combining Corollaries 28.3.1 and 27.3.2 of [27] we get the following.

Theorem 6.1. *The optimal value in* (P) *is greater than or equal to the optimal value in* (D)*. If all the functions f_i are of the form $f_i(x) = a_i \cdot x - b_i$, then the following are equivalent and imply that the optimum values of* (P) *and* (D) *are equal*

1. (P) *has a solution,*
2. *the optimum value of* (P) *is finite,*
3. (D) *has a solution,*
4. *the optimum value of* (D) *is finite.*

References

[1] M. Avellanedas, R. Buff, C. Friedman, N. Grandchamp, L. Kruk, and J. Newman. Weighted monte carlo: A new technique for calibrating asset-pricing models. *International Journal of Theoretical and Applied Finance*, 4:91–119, 2001.

[2] M. Avellanedas, A. Levy, and P. Parás. Pricing and hedging derivative securities in markets with uncertain volatilities. *Applied Mathematical Finance*, 2(2):73–89, 1995.

[3] M. Avellanedas and P. Parás. Managing the volatility risk of portfolios of derivative securities: the lagrangian uncertain volatility model. *Applied Mathematical Finance*, 3(1):21–52, 1996.

[4] Patrick Billingsley. *Convergence of probability measures.* Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons Inc., New York, second edition, 1999. A Wiley-Interscience Publication.

[5] John R. Birge. Stochastic programming computation and applications. *INFORMS Journal on Computing*, 9:111–133, 1997.

[6] J. Borwein, R. Choksi, and P. Marechal. Probability distributions of assets inferred from option prices via the principle of maximum entropy. *submitted*, 2001.

[7] J. Cox and M. Rubinstein. *Options markets.* Prentice Hall, New York, 1985.

[8] Freddy Delbaen, Yuri M. Kabanov, and Esko Valkeila. Hedging under transaction costs in currency markets: a discrete-time model. *Math. Finance*, 12(1):45–61, 2002.

[9] M.A.H. Dempster and G.W.P. Thompson. Dynamic portfolio replication via stochastic programming. In M.A.H. Dempster, editor, *Risk Management: Value at Risk and Beyond*, pages 100–128, Cambridge, 2002. Cambridge University Press.

[10] Darrell Duffie. *Dynamic Asset Pricing Theory.* Princeton University Press, Princeton, New Jersey, 1996.

[11] N. C. P. Edirisinghe, Vasant Naik, and R. Uppal. Optimal replication of options with transactions costs and trading restrictions. *Journal of Financial and Quantitative Analysis*, 28:117–138, 1993.

[12] J. Gondzio, K. Kouwenberg, and T. Vorst. Hedging options under transaction costs and stochastic volatility. *Journal of Economic Dynamics & Control*, to appear.

[13] J. Michael Harrison and Stanley R. Pliska. Martingales and stochastic integrals in the theory of continuous time trading. *Stochastic Processes and their Applications*, 11:215–260, 1981.

[14] Jean-Baptiste Hiriart-Urruty and Claude Lemaréchal. *Convex analysis and minimization algorithms. I*, volume 305 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences].* Springer-Verlag, Berlin, 1993. Fundamentals.

[15] J.C. Jackwerth. Option-implied risk-neutral distributions and implied binomial trees: a literature review. *Journal of Derivatives*, 7(2):66–82, 1999.

[16] J.C. Jackwerth and M. Rubinstein. Recovering probability distributions from option prices. *Journal of Finance*, 51:1611–1631, 1996.

[17] R. Jarrow and A. Rudd. *Option Pricing.* Dow Jones-Irwin, Homewood, IL, 1983.

[18] Alan King, Olga Streltchenko, and Yelena Yesha. Private valuation of contingent claims: Discrete time/state model. Technical report, University of Maryland, Baltimore, Maryland, 2001.

[19] Alan J. King, editor. *Approximation and Computation in Stochastic Programming*, volume 75 of *Mathematical Programming, Series B*, 1996.

[20] Alan J. King. Duality and martingales: a stochastic programming perspective on contingent claims. *Math. Program.*, 91(3, Ser. B):543–562, 2002. ISMP 2000, Part 1 (Atlanta, GA).

[21] Vasant Naik. Finite state securities market models and arbitrage. In R. A. Jarrow, V. Maksimovic, and W. T. Ziemba, editors, *Finance*, volume 9 of *Handbooks in Operations Research and Management Science*, pages 31–64, Amsterdam, 1995. Institute for Operations Research and the Management Sciences, North-Holland.

[22] E. Omberg. Efficient discrete time jump process models in option pricing. *Journal of Financial and Quantitative Analysis*, 23(2):161–174, 1988.

[23] T. Pennanen and M. Koivu. Integration quadratures in discretization of stochastic programs. *Stochastic Programming E-Print Series*, 2002.

[24] Stanley R. Pliska. *Introduction to Mathematical Finance: Discrete Time Models.* Blackwell Publishers, Malden, Massachusetts, 1997.

[25] Elijah Polak. *Optimization*, volume 124 of *Applied Mathematical Sciences.* Springer-Verlag, New York, 1997. Algorithms and consistent approximations.

[26] P. Ritchken and S. Kuo. Option bounds with finite revision opportunities. *Journal of Finance*, 43:301–308, 1988.

[27] R. Tyrrell Rockafellar. *Convex analysis.* Princeton Mathematical Series, No. 28. Princeton University Press, Princeton, N.J., 1970.

[28] R. Tyrrell Rockafellar. *Conjugate duality and optimization.* Society for Industrial and Applied Mathematics, Philadelphia, Pa., 1974. Lectures given at the Johns Hopkins University, Baltimore, Md., June, 1973, Conference Board of the Mathematical Sciences Regional Conference Series in Applied Mathematics, No. 16.

[29] M. Rubinstein. Implied binomial trees. *Journal of Finance*, 49(3):771–818, 1994.

[30] J. Stoer and R. Bulirsch. *Introduction to numerical analysis*, volume 12 of *Texts in Applied Mathematics*. Springer-Verlag, New York, second edition, 1993. Translated from the German by R. Bartels, W. Gautschi and C. Witzgall.