

Hedging Options under Transaction Costs and Stochastic Volatility

Jacek Gondzio, Roy Kouwenberg, Ton Vorst

31 January 2000

MS-00-004

For other papers in this series see <http://www.maths.ed.ac.uk/preprints>

Hedging Options under Transaction Costs and Stochastic Volatility*

Jacek Gondzio[†]

Department of Mathematics & Statistics,
The University of Edinburgh,
Mayfield Road, Edinburgh EH9 3JZ,
United Kingdom.

Roy Kouwenberg[‡] and Ton Vorst[§]

Econometric Institute,
Erasmus University Rotterdam,
P.O.Box 1738, 3000 DR Rotterdam,
The Netherlands.

January 31, 2000

*Research partially supported through contract "HPC-Finance" (no. 951139) of the European Commission. Partial support also provided by the "HPC-Finance" partner institutions: Universities of Bergamo (IT), Cambridge (UK), Calabria (IT), Charles (CZ), Cyprus (CY), Erasmus (ND), Technion (IL) and the "Centro per il Calcolo Parallelo e i Supercalcolatori" (IT).

[†]Email: gondzio@maths.ed.ac.uk, URL: <http://maths.ed.ac.uk/~gondzio/>

[‡]Corresponding author: Erasmus University Rotterdam, Econometric Institute, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands. Tel.nr.: +31-104082388, Fax.nr.: +31-104089162, Email: kouwenberg@few.eur.nl, URL: <http://www.few.eur.nl/few/people/kouwenberg/>

[§]Email: vorst@few.eur.nl, URL: <http://www.few.eur.nl/few/people/vorst/>

Hedging Options under Transaction Costs and Stochastic Volatility

Abstract

In this paper we consider the problem of hedging contingent claims on a stock under transaction costs and stochastic volatility. Extensive research has clearly demonstrated that the volatility of most stocks is not constant over time. As small changes of the volatility can have a major impact on the value of contingent claims, hedging strategies should try to eliminate this volatility risk. We propose a stochastic optimization model for hedging contingent claims that takes into account the effects of stochastic volatility, transaction costs and trading restrictions. Simulation results show that our approach could improve performance considerably compared to traditional hedging strategies.

JEL Classifications Codes: G10, G11, G13.

Keywords: option hedging, stochastic volatility, stochastic programming, computational finance, high-performance computing.

1 Introduction

In this paper we consider the problem of hedging contingent claims under transaction costs and stochastic volatility. Extensive research during the last two decades has demonstrated that the volatility of stocks is not constant over time (Bollerslev e.a. 1992). Engle (1982) and Bollerslev (1986) introduced the family of ARCH and GARCH models to describe the evolution of the volatility of the asset price in discrete time. Econometric tests of these model clearly reject the hypothesis of constant volatility and find evidence of volatility clustering over time. In the financial literature stochastic volatility models have been proposed to model these effects in a continuous-time setting (Hull and White 1987, Scott 1987, Wiggins 1987). Pricing methods for options on a stock with a stochastic volatility process are now widely available, both in the discrete-time and the continuous-time framework (Heston 1993, Finucane and Tomas 1997 and Ritchken and Trevor 1999). Practicable methods for hedging options under stochastic volatility are rare however.

Schweizer (1991, 1995) has proposed methods to minimize the replication error of contingent claims in general incomplete markets, including stochastic volatility as a special case. Schweizer (1995) only considers trading strategies involving the riskless bond and the underlying stock itself. As the bond and the underlying stock price are insensitive to changes of the volatility, these hedging schemes are deemed to be inefficient compared to strategies involving traded option contracts on the underlying stock (Frey and Sin 1999). Traded options on the underlying stock are sensitive to changes in the unobservable stock price volatility. This observation is used in the simple delta-vega hedging scheme: traded options are added to the portfolio of the investor in order to eliminate the exposure to small changes of the volatility. Unfortunately an effective delta-vega hedge has to be rebalanced frequently. As the bid-ask spreads on exchange traded options are considerable, frequent updating of a delta-vega hedge could result in losses due to transaction costs.

Static hedging methods try to compose a buy-and-hold portfolio of exchange traded options that replicate the payoff of the contingent claim under consideration (Derman, Ergener and Kani 1995, Carr, Ellis and Gupta 1998). The static hedging strategy does not require any rebalancing and is therefore quite efficient in avoiding transaction costs. Unfortunately the odds of coming up with a perfect static hedge for a particular over-the-counter product are small, as the number of (liquid) traded option contracts is limited. Avellaneda and Paras (1996) proposes an algorithm to construct a static portfolio of options that matches the desired payoff as closely as possible, while the residual is priced and hedged with a trading strategy involving the underlying stock. A disadvantage of this approach is that the static hedge can only be efficient if traded options are available with sufficiently similar maturity and moneyness as the over-the-counter product that has to be hedged.

In this paper we propose a stochastic optimization model to extend the simple delta-vega hedging scheme. The hedge portfolio in our model consists of the underlying stock and exchange traded options with sufficient liquidity. The model has a limited number of trading dates on which the hedge portfolio can be rebalanced (e.g. weekly), while transaction costs and trading restrictions are taken into account. The goal of the model is to minimize hedging errors by following an appropriate dynamic trading strategy. An important feature is that we only minimize the

hedging error at the first few trading dates and not until the final maturity of the contingent claim. We think that our specification of the hedging model is useful because: 1. The planning horizon of traders is shorter than the maturity of their contingent claims and they are usually more interested in overnight profits and losses. 2. The portfolio of liabilities of a trader might change frequently due to additional buying and selling. 3. Risk limits like ‘Value At Risk’ are often imposed with a relatively short horizon.

The stochastic optimization hedging model requires a set of scenarios as input. The scenarios are distinct paths for the prices of the assets and liabilities at each trading date. It is clear that the performance of a hedge constructed with the stochastic optimization model will crucially depend on the quality of the price scenarios. Without a good scheme for generating scenarios the stochastic optimization model is merely a theoretical concept, not a practicable hedging tool. Our contribution is that we propose reliable methods to construct scenarios for stochastic volatility models. Moreover, results of simulations with an illustrative example show that the stochastic optimization hedging strategy can really outperform a delta-vega hedging scheme in the presence of transaction costs. The strategy of the stochastic optimization model makes sense intuitively: it stays close to a delta-vega neutralized position but with some additional slack to avoid needless transaction costs.

We now shortly outline the contents of the paper. In Section 2 we first describe traditional hedging strategies like delta-vega hedging and static replication. Next we introduce the stochastic optimization hedging model for markets with stochastic volatility and transaction costs. The stochastic optimization hedging model requires a set of asset price scenarios as input. In Section 3 we propose methods to construct accurate scenarios of stock and option prices for a one-factor stochastic volatility process. Given the fine-grained approximation of the underlying distribution of the prices in Section 3, the stochastic optimization model might become huge and hard to solve however. We propose an aggregation algorithm that reduces a set of price scenarios to a smaller size, while preserving important properties like the absence of arbitrage. In Section 4 we discuss methods to solve the stochastic optimization model. Finally, we use simulations to investigate the relative performance of our approach for a specific hedging problem.

2 A Stochastic Optimization Approach to Hedging

In this section we will introduce a stochastic optimization model for hedging a contingent claim on a stock with stochastic volatility and under transaction costs. As an introduction, we will shortly describe delta-vega hedging and static hedging methods.

2.1 Delta-Vega Hedging

The cornerstone of the Black-Scholes option pricing formula (Black and Scholes 1973) is the construction of a dynamic investment strategy that perfectly replicates the payoff of a European call or put option at the maturity date of the contract. Let $S(t)$ denote the price of the underlying stock at time t and let $f(t, S(t))$ denote the value of a European call (or put) option with

maturity T . The Black-Scholes option pricing model assumes that the stock price follows a geometric Brownian motion:

$$(1) \quad dS(t) = \mu S(t)dt + \sigma S(t)dZ(t),$$

where μ is the continuously compounded drift rate of the stock price $S(t)$, σ is the instantaneous volatility and $dZ(t)$ is a standard Brownian motion. It follows from Ito's lemma that the stochastic process for the option price $f(t, S(t))$ is:

$$(2) \quad df = \frac{\partial f}{\partial S}dS + \frac{\partial f}{\partial t}dt + \frac{1}{2}\frac{\partial^2 f}{\partial S^2}(dS)^2 = \frac{\partial f}{\partial S}dS + \left(\frac{\partial f}{\partial t} + \frac{1}{2}\frac{\partial^2 f}{\partial S^2}\sigma^2 S(t)^2\right)dt$$

Under the assumption of continuous and frictionless trading, a riskless hedge for a short position in the option f can be constructed by buying $\Delta_f = \frac{\partial f}{\partial S}$ stocks. This trading strategy is self-financing: all purchases of stock are financed by borrowing at the constant interest rate r and the proceeds of selling stocks are deposited on a money market account at the same rate. The value of the investor's portfolio Π is $\Pi = \Delta_f S - f$, following the stochastic process:

$$(3) \quad d\Pi = \Delta_f dS - df = -\left(\frac{\partial f}{\partial t} + \frac{1}{2}\frac{\partial^2 f}{\partial S^2}\sigma^2 S(t)^2\right)dt$$

It is clear from (3) that the instantaneous change of the portfolio value is non-stochastic: the delta-hedging strategy has eliminated all risk due to price changes in the underlying stock. Because the portfolio Π is riskless it must earn the riskfree rate: $d\Pi = r\Pi dt$. If we substitute the last expression in (3), then we derive a partial differential equation for the option price f having the famous Black-Scholes textbook formula as its solution.¹

In practice the dynamic replication strategy for European options will only be perfect if all of the assumptions underlying the Black-Scholes formula hold. First of all, frictionless and continuous trading in the underlying stock will be impossible. Leland (1985) and Boyle and Vorst (1992) have derived a perfect replication strategy for European options under transaction costs. For general contingent claims on a stock (e.g. American-style) under market frictions the delta might still be used as first order approximation to set up a riskless portfolio. However, if the volatility of the underlying stock varies stochastically then the delta hedging method might fail severely.

A simple method to limit the volatility risk is to consider the volatility sensitivity *vega* of the contract f : $\Upsilon_f = df/d\sigma$. Suppose that an exchange traded option g is available with $\Upsilon_g = dg/d\sigma$. The instantaneous volatility exposure of a short position in f could be eliminated by taking a position of Υ_f/Υ_g in the traded option g . The delta of the entire portfolio now is: $(\Upsilon_f/\Upsilon_g)\Delta_g - \Delta_f$. By adding an offsetting position in the stock (having a delta of 1 and a vega of 0), the portfolio Π can be made locally insensitive to 'small' changes in the stock price and the volatility. Note that the delta and vega of options depend on factors like time t , the stock

¹Under the additional boundary constraint that $f(T, S(T))$ equals the payoff of the option at maturity

price $S(t)$, the interest rate r and the volatility $\sigma(t)$. The portfolio will have to be rebalanced frequently to ensure delta-vega neutrality. With transaction costs frequent rebalancing might result in considerable losses.

2.2 Static Hedging and Super-Replication

The goal of static hedging is to construct a buy-and-hold portfolio of exchange traded claims that perfectly replicates the payoff of a given over-the-counter product. The static hedge portfolio requires no readjustments, so the transaction costs are relatively small. Unfortunately the odds of coming up with a perfect static hedge for a given over-the-counter claim are small, given the limited number of exchange listed option contracts with sufficient trading volume. For example consider barrier options: the static hedge of Derman, Egener and Kani (1995) requires an infinite number of vanilla options, while the static hedge of Carr, Ellis and Gupta (1998) requires at least one option with a specific maturity and exercise price under the additional assumption that risk neutral drift-rate of the stock is zero. We would like to study the hedging problem under less restrictive assumptions.

Avellaneda and Paras (1996) propose a method to construct a *static* hedge portfolio of exchange traded options, combined with a *dynamic* trading strategy involving the underlying stock. As an example consider a European over-the-counter product with payoffs $F(S(T))$ at maturity T , where S is the price of the underlying stock. Suppose that H exchange traded vanilla options are available with payoff $G_h(S(T))$ at maturity T , for $h = 1, 2, \dots, H$. If the investor buys a static portfolio $(\lambda_1, \dots, \lambda_H)$ of the vanilla options to hedge a short position in the over-the-counter product then he faces a residual payoff (or payment) at maturity T of: $R(S(T)) = \sum_{h=1}^H \lambda_h G_h(S(T)) - F(S(T))$. Avellaneda and Paras (1996) assume that the residual is super-replicated with a self-financing trading strategy in the stock and the riskless bond. Super-replication means that the self-financing trading strategy must payoff at least $R(S(T))$ at maturity. Combined with the proceeds of the super-replicating strategy, the payoff of the static option portfolio will always meet the liability of the over-the-counter product.

Note that the dynamic super-replication strategy of Avellaneda and Paras (1996) could become rather expensive in a stochastic volatility context, because it only involves the stock and the riskless bond. For example, given a stock price process with stochastic volatility, Frey and Sin (1999) prove that the minimum cost super-replication strategy for a call option is buying one unit of the underlying stock. For general European-style claims Avellaneda, Levy and Paras (1995) derive a least cost super-replicating strategy involving the stock and the riskless bond, given that volatility lies between the bounds σ_{min} and σ_{max} . These strategies will tend to be conservative and expensive, as the stock price is not sensitive to changes in the volatility. A drawback of the approach of Avellaneda and Paras (1996) is that rebalancing of the option portfolio is impossible because there is no information about future option prices. The method is not suited for American-style claims and hedging problems where a long term liability is tracked with a portfolio of short term options.

2.3 A Stochastic Optimization Approach

In this section we propose a stochastic optimization model to extend the delta-vega hedging scheme and the approximate static hedge of Avellaneda and Paras (1996). We consider a discrete-time economy: the current time is denoted by t_0 and there is limited set of future trading dates $t = t_1, t_2, \dots, t_K$. The investor wishes to hedge a portfolio of contingent claims G with value $G(t, \omega)$ at each time t . The value of the portfolio of liabilities G depends on the uncertain state of the world ω , which will be fully revealed at the planning horizon t_K . The complete set $\Omega = \{\omega_1, \omega_2, \dots, \omega_Q\}$ of events could be thought of as a set of scenarios. Early exercise and other reasons for pre-mature termination of contracts (e.g. barrier options) could influence the value of the liabilities $G(t, \omega)$: these events should be modelled in the set of scenarios Ω . As information is revealed progressively through time, most scenarios will share some information at the first few dates and only become unique at the planning horizon. This information structure can be depicted as a familiar event tree (see figure 1).

The hedge portfolio of the investor consists of the underlying stock $S(t, \omega) = P_0(t, \omega)$, the exchange traded vanilla options $P_h(t, \omega)$ for $h = 1, 2, \dots, H$ and the money market account $M(t, \omega)$. We assume that the investor plans to rebalance his hedge at the trading dates $t = t_0, t_1, \dots, t_K$. The hedge portfolio is denoted by $(M(t, \omega), \lambda_0(t, \omega), \lambda_1(t, \omega), \dots, \lambda_H(t, \omega))$, where $M(t, \omega)$ denotes the position in the money market account, $\lambda_0(t, \omega)$ is the number of stocks in the portfolio and $\lambda_h(t, \omega)$ represents the number of plain vanilla options for $h = 1, 2, \dots, H$.

We assume that cash can be added to and withdrawn from the money market account without transaction costs and earns the constant interest rate r . However the stock has to be bought at the ask price $S(t, \omega)(1 + \gamma_0)$ and sold at the bid price $S(t, \omega)(1 - \gamma_0)$, while the proportional transaction costs for each option trade are γ_h for $h = 1, 2, \dots, H$. For ease of exposition we assume the stock does not pay dividends and that the exchange traded options will not be exercised prior to maturity. The cash payoff of the traded options is denoted by $D_h(t, \omega)$, which in most cases will only be non-zero at maturity.

We define the hedging error $R^G(t, \omega)$ as the value of the hedge portfolio minus the value of the liability G . The goal of the model is to find a trading strategy that minimizes these hedging errors at each future trading date $t = t_1, \dots, t_K$. Note that the final trading date t_K of the model could be prior to the maturity of the liabilities G . So we only minimize the hedging error at the first few trading dates and not at every trading date up to the final maturity of the liabilities. We think that this specification of the hedging model is useful, as the portfolio of liabilities of a trader might change frequently due to additional buying and selling. Traders are concerned about overnight profits and losses: their planning horizon will be shorter than the maturity of their liabilities. Moreover risk limits like ‘Value At Risk’ are usually imposed with a relatively short horizon.

We now specify the equations of the stochastic optimization hedging model. We assume that the investor maximizes a utility function $U_t(\cdot)$ over the hedging errors $R^G(t, \omega)$ at each trading date t , given an initial endowment M^{ini} and an initial portfolio $\{\lambda_h^{ini} : h = 0, 1, \dots, H\}$. The expected utility of the investor is maximized by choosing an optimal trading strategy $(M(t, \omega), \lambda_0(t, \omega), \lambda_1(t, \omega), \dots, \lambda_H(t, \omega))$ in the underlying stock and the exchange-traded options

(details about the mathematical definition of the model can be found in Appendix A):

$$(4) \quad \max_{M, \lambda} \sum_{k=1}^K E_{\omega} [U_{t_k}(R^G(t_k, \omega))]$$

Subject to:

$$(5) \quad M(t_1, \omega) = M^{ini} - \sum_{h=0}^H (\lambda_h(t_1, \omega) - \lambda_h^{ini}) P_h(t_0, \omega) - \sum_{h=0}^H \gamma_h |\lambda_h(t_1, \omega) - \lambda_h^{ini}| P_h(t_0, \omega)$$

$$(6) \quad M(t_{k+1}, \omega) = e^{r(t_k - t_{k-1})} M(t_k, \omega) + \sum_{h=0}^H \lambda_h(t_k, \omega) D_h(t_k, \omega) \\ - \sum_{h=0}^H (\lambda_h(t_{k+1}, \omega) - \lambda_h(t_k, \omega)) P_h(t_k, \omega) - \sum_{h=0}^H \gamma_h |\lambda_h(t_{k+1}, \omega) - \lambda_h(t_k, \omega)| P_h(t_k, \omega) \\ \text{for } k = 1, 2, \dots, K - 1$$

$$(7) \quad R^G(t_k, \omega) = e^{r(t_k - t_{k-1})} M(t_k, \omega) + \sum_{h=0}^H \lambda_h(t_k, \omega) (P_h(t_k, \omega) + D_h(t_k, \omega)) - G(t_k, \omega), \\ \text{for } \forall \omega \in p_{t_k}^G, \text{ otherwise } R^G(t_k, \omega) = 0$$

Equation (5) is the initial budget equation of the investor, while (6) represents the budget equation at the trading dates $t = t_1, t_2, \dots, t_{K-1}$. The hedging error is defined in (7) as the mid-price of the portfolio of traded assets (just before rebalancing) minus the value of the over-the-counter liability G . This definition of the hedging error is included as an example and could be adjusted for specific applications. The objective of model (4) is to maximize the expected utility $U_t(\cdot)$ over the hedging error R^G . Alternatively, $U_t(\cdot)$ could be defined as a loss-function like $\min\{R^G(t, \omega), 0\}$.

Note that we could easily add trading restrictions to the stochastic optimization model. For example, short selling and borrowing could be allowed up to a given fraction of the portfolio value. The stochastic optimization hedging (SOH) model has a number of clear advantages over the simple delta-vega hedge and the approach of Avellaneda and Paras (1996). The simple delta-vega hedge is extended by anticipating future transaction costs due to rebalancing of the hedge-portfolio. The approach of Avellaneda and Paras (1996) is extended as we allow for trading of options at a number of future dates. This feature allows the SOH model to track a long term liability with a portfolio of short term options (with sufficient liquidity).

A successful implementation of the SOH model requires two crucial elements:

1. a set of scenarios and probabilities for the stock and option prices and
2. efficient methods to solve the optimization problem.

Suitable decomposition methods and parallel implementations for solving the problem will be discussed in Section 4. In the next section we will concentrate on generating scenarios for the stochastic hedging model.

3 Generating Stochastic Volatility Scenarios

The SOH model proposed in the previous section requires an event tree to describe the stochastic evolution of the asset prices through time. First we apply a trinomial process to generate probability distributions on a fixed grid of stock price versus volatility. Option prices are next added to the grid, while we carefully avoid arbitrage opportunities. Finally we introduce an aggregation algorithm that constructs a sparse event tree for the SOH model. In this section we choose an asymmetric N-GARCH model as the underlying process for the stock price and its volatility, however the proposed methods can be applied in general for any stochastic volatility process.

3.1 Stochastic Volatility Processes

In the financial-economic literature a large number of continuous-time stochastic volatility processes have been proposed to overcome the limitations of constant-volatility models for option pricing (e.g. Hull and White 1987, Scott 1987 and Wiggins 1987). As an example, we consider the Hull and White (1987) specification for a stock price process with stochastic volatility:

$$(8) \quad \begin{aligned} dS(t) &= \mu S(t)dt + \sigma(t)S(t)dZ(t), \\ d\sigma(t) &= \kappa(\bar{\sigma} - \sigma(t))dt + \xi\sigma(t)^\alpha dW(t). \end{aligned}$$

where $\sigma(t)$ is the instantaneous stock volatility at time t , κ , $\bar{\sigma}$, ξ , α are parameters of the stochastic volatility process, $\sigma(0)$ is the initial volatility and $dW(t)$ is a standard Brownian motion. Let ρ define the correlation between $dZ(t)$ and $dW(t)$.

The process of Hull and White (1987) allows for mean-reversion of the stochastic volatility. Another important property of the model is the correlation ρ between the instantaneous shocks to the stock price $dZ(t)$ and the shocks to the volatility $dW(t)$. The correlation ρ allows the stochastic volatility model to overcome some of the ‘smile-problems’ of the Black-Scholes model in the plots of implied volatility v.s. exercise price (Bakshi, Cao and Chen 1997).

A drawback of continuous-time stochastic volatility models in general is that the coefficients are hard to estimate with timeseries of stock price observations. The family of discrete-time ARCH and GARCH stochastic volatility models do allow for relatively easy statistical estimation. As an example we consider the nonlinear asymmetric GARCH model (Engle and Ng 1993):

$$(9) \quad \begin{aligned} \ln S(t + \Delta t) - \ln S(t) &= (r + \lambda\sigma(t) - \frac{1}{2}\sigma(t)^2)\Delta t + \sigma(t)\sqrt{\Delta t}Z(t + \Delta t), \\ \sigma(t + \Delta t)^2 &= \beta_0 + \beta_1\sigma(t)^2 + \beta_2\sigma(t)^2(Z(t + \Delta t) - c)^2. \end{aligned}$$

where r is the riskless interest rate, Δt is the discrete time-step, β_0 , β_1 , β_2 and c are parameters of the asymmetric N-GARCH volatility specification. Let σ_0 denote the initial volatility.

The asymmetric N-GARCH volatility process could be mean-reverting depending on the values of β_0 and β_1 , while the model captures the correlation between return and volatility innovations

with the parameter c . Heston and Nandi (1997, 1998) have derived a closed-form solution for European option prices under an asymmetric N-GARCH volatility model. Moreover, numerically efficient lattice interpolation methods have been developed to compute both European and American option prices under discrete-time GARCH volatility specifications (Finucane and Tomas 1997 and Ritchken and Trevor 1999).

For the SOH model we will assume an asymmetric N-GARCH model as the underlying stochastic process for the stock and option prices, as it suits the discrete-time formulation of the stochastic optimization model. This assumption can be made without loss of generality as most continuous-time stochastic volatility models are limiting cases of the asymmetric N-GARCH model (see Nelson 1990, Duan 1997 and Ritchken and Trevor 1999).

3.2 Generating Scenarios with a Trinomial Process

Given the asymmetric N-GARCH model as underlying process, we now would like to generate scenarios of future stock and option prices to include into the stochastic optimization hedging model. A straightforward method for generating scenarios is random sampling: we simply draw a large number of paths randomly from the calibrated N-GARCH model. Additionally, option prices could be computed for each of the sampled paths with an appropriate pricing method (e.g. Heston and Nandi 1997 or Ritchken and Trevor 1999). This simulation approach for generating scenarios is straightforward at first sight, but raises a number of questions. How do we guarantee the absence of arbitrage opportunities in the set of simulated stock and option prices? Moreover, could we generate scenarios more efficiently in order to avoid numerous option price calculations?

We propose a method for generating scenarios that tries to resolve the problems mentioned above. Instead of simulation, we use a trinomial process to approximate the underlying N-GARCH process. Starting with the initial stock price $S(t)$ and volatility $\sigma(t)$, we assume that the stock price moves to one of the following three prices after a small timestep $\Delta t = 1$:

$$(10) \quad \begin{aligned} S_u(t+1) &= e^{+\eta\psi} S(t), & \text{with probability } p_u &= \frac{\sigma(t)^2}{2\eta^2\psi^2} + \frac{(r - \frac{1}{2}\sigma(t)^2 + \lambda\sigma(t))}{2\eta\psi}, \\ S_m(t+1) &= S(t), & \text{with probability } p_m &= 1 - \frac{\sigma(t)^2}{\eta^2\psi^2}, \\ S_d(t+1) &= e^{-\eta\psi} S(t), & \text{with probability } p_d &= \frac{\sigma(t)^2}{2\eta^2\psi^2} - \frac{(r - \frac{1}{2}\sigma(t)^2 + \lambda\sigma(t))}{2\eta\psi}. \end{aligned}$$

where p_u , p_m and p_d represent the probability of an upside, middle and downside move respectively. The parameter ψ is a fixed constant, chosen to be σ_0 . The jump-size parameter η is an integer and determined as: $\eta = \lfloor \sigma(t)/\sigma_0 \rfloor + 1$. With this parameter setting the probabilities are always positive and they sum up to one (Ritchken and Trevor 1999).

Given that the stock price has moved to $S_a(t+1)$ (with $a = u, m$ or d), the stochastic volatility

is next updated according to:

$$(11) \quad \begin{aligned} \sigma_a^2(t+1) &= \beta_0 + \beta_1 \sigma^2(t) + \beta_2 \sigma^2(t) (\epsilon_a(t) - c)^2, \\ \epsilon_a(t) &= \left[\ln(S^a(t+1)/S(t)) - (r_f - \frac{1}{2} \sigma^2(t) + \lambda \sigma(t)) \right] / \sigma(t). \end{aligned}$$

By applying the trinomial process k times recursively and keeping track of the probabilities, we can construct a joint distribution of the stock price and the volatility at the trading date $t+k$. Note however that the number of distinct points in the space of stock price and volatility is 3^k and will be huge if k is large. In order to increase the efficiency we will construct probability distributions on a fixed grid of stock price and volatility. Suppose that the stock price versus volatility grid consists of the points (S^i, σ^j) for $i = 1, 2, \dots, n_I$ and $j = 1, 2, \dots, n_J$ and that it includes the initial point (S_0, σ_0) . Let $p^{ij}(t)$ denote the probability of point (S^i, σ^j) at time t under the approximated N-GARCH process.

If we apply the trinomial process from a starting point *on* the grid at time t , then the three jump-points probably end up somewhere *off* the grid at time $t+1$. Note that if we use equal log-spacing of ψ between the stock prices (i.e. $\ln(S^{i+1}/S^i) = \psi$), then at least the three stock prices will always end up *on* the grid. Under this assumption we conclude that for a given jump to $(S_a(t+1), \sigma_a(t+1))$ there exists a point (S^i, σ^j) on the grid with $S^i = S_a(t+1)$ and $\sigma^j \leq \sigma_a(t+1) \leq \sigma^{j+1}$.

As an approximation, we now assign the probability mass p_a of each jump $a = u, m$ and d to points on the grid by using linear interpolation weights:

$$\begin{aligned} p_a \frac{(\sigma_a(t) - \sigma^j)}{(\sigma^{j+1} - \sigma^j)} & \text{ is assigned to } p^{i,j+1}(t+1) \text{ in point } (S^i, \sigma^{j+1}) \text{ and} \\ p_a \frac{(\sigma^{j+1} - \sigma_a(t))}{(\sigma^{j+1} - \sigma^j)} & \text{ is assigned to } p^{i,j}(t+1) \text{ in point } (S^i, \sigma^j). \end{aligned}$$

By applying the trinomial process (10) and the linear interpolation scheme in a forward recursion we can generate the joint distribution of the stock price and the volatility on the grid. Note that under the trinomial process multiple points at time t could assign probability mass to a particular point on the grid at time $t+1$ and that all these contribution should be summed. The resulting probabilities at time $t+1$ are always positive and sum up to one if the process never jumps outside the boundaries of the grid.²

Note that due to the interpolation scheme we construct an approximation of the true trinomial process. However we can easily guarantee the quality of this approximation. It turns out that our linear interpolation of *probabilities in a forward recursion* is exactly equivalent to a numerical option pricing scheme that interpolates *prices in a backward recursion*.³ Both the methods of

²It is easy to find appropriate boundaries for the grid, because a mean-reverting stochastic volatility process will have a minimum and a maximum volatility. Given the maximum volatility we can estimate the maximum and minimum stock price after k timesteps, using the jump-size η .

³The proof is available on request from the authors and skipped due to limited space.

Ritchken and Trevor (1999) and Finucane and Tomas (1997) for American option pricing under stochastic volatility are based on the interpolation of prices on a grid. As the option prices computed with these methods converge to their true values (Ritchken and Trevor 1999), it is clear that also our approximated probability distributions will converge.

The advantages of using a trinomial process recursively on a fixed grid of stock price and volatility are numerous. The approach is fast and efficient compared to simulation and we can easily improve the quality of the approximation by taking a n -nomial process with $n = 5, 7, 9, \dots$ instead of a trinomial process (Ritchken and Trevor 1999). Moreover we can construct a conditional transition density from a particular area of the grid at time t to another area at time $t + k$. By changing the drift of the N-GARCH process we can construct a risk neutral distribution on the grid. Under the risk neutral distribution non-zero probabilities are assigned to exactly the same points on the grid as under the actual distribution. These properties will be used in the next sections to eliminate arbitrage opportunities and to construct event trees with option prices.

3.3 Checking for Arbitrage Opportunities

We use a discrete-time trinomial process to approximate the joint distribution of the stock price and its volatility on a fixed grid. In order to generate scenarios for the stochastic optimization hedging model we also need to add option prices to the grid. This can be achieved if we calculate the option prices at each point of the grid with an appropriate option pricing method for the N-GARCH process (e.g. Heston and Nandi 1997 or Ritchken and Trevor 1999). Suppose that we have added option prices to our grid at time t_1 , representing the first future trading date of the SOH model, while t_0 is the initial date.

An important requirement for our set of option prices and stock prices on the grid at time t_1 and t_0 is that they should not contain arbitrage opportunities. Loosely speaking, if we could generate infinite profits by trading assets in our one-period economy without taking risk then arbitrage opportunities exist. In the context of a stochastic optimization model for financial planning arbitrage opportunities will lead to unbounded or severely biased solutions (Klaassen 1997). In this section we propose methods to guarantee that the scenarios for the SOH model are arbitrage-free.

Formally, arbitrage opportunities do not exist if there are no portfolios having a negative price that provide a non-negative payoff. Let $P_0(t_0)$ denote the initial stock price, while $P_h(t_0)$ for $h = 1, 2, \dots, H$ are the initial prices of the options that we include in the SOH model. The prices of these options in each point (S^i, σ^j) of the grid at the first trading date t_1 are denoted by $P_h^{ij}(t_1)$ and their payoffs by $D_h^{ij}(t_1)$. For ease of exposition we treat the stock as asset number 0: $P_0^{ij}(t_1) = S^i$ and $D_0^{ij}(t_1) = 0$. Harrison and Kreps (1979) and Naik (1995) prove that arbitrage opportunities do not exist if and only if we can find a number $\pi^{ij}(t_1)$ for each point on the grid such that conditions (12-14) are satisfied:

$$(12) \quad (1 - \gamma_h)P_h(t_0) \leq e^{-rk} \sum_{i=1}^{n_I} \sum_{j=1}^{n_J} \pi^{ij}(t_1) \left(P_h^{ij}(t_1) + D_h^{ij}(t_1) \right) \leq (1 + \gamma_h)P_h(t_0),$$

for $h = 0, 1, 2, \dots, H$

$$(13) \quad \sum_{i=1}^{n_I} \sum_{j=1}^{n_J} \pi^{ij}(t_1) = 1,$$

$$(14) \quad \pi^{ij}(t_1) \geq 0, \quad \pi^{ij}(t_1) > 0 \Leftrightarrow p^{ij}(t_1) > 0.$$

where r is the risk free interest rate and $p^{ij}(t_1)$ are the actual probabilities for the grid at time t_1 constructed with the trinomial process.

Condition (12) implies that the discounted expected value of the stock and option prices under the measure $\pi^{ij}(t_1)$ should lie in between the initial bid- and ask-price. Conditions (13) and (14) require $\pi^{ij}(t_1)$ to be a probability measure equivalent to the actual measure $p^{ij}(t_1)$, which has been constructed with the trinomial process. A solution of the no-arbitrage conditions is usually called a *risk neutral measure* or an *equivalent martingale measure*.

Note that we could construct a risk neutral measure for the stock price by changing the drift rate and the probabilities of the trinomial N-GARCH process (see Duan 1995 and Ritchken and Trevor 1999). This equivalent measure always satisfies the condition (12) *for the stock price*, but not necessarily *for the option prices*. If we would calculate all option prices with the risk neutral trinomial process in a backward recursion on the grid using linear interpolation, then condition (12) is satisfied by definition. However if we use another option pricing method, then we need to solve the equations above for a risk neutral measure in order to prove the absence of arbitrage opportunities explicitly.

The number of no-arbitrage conditions is equal to the number of assets (i.e. $H + 2$), while the number of decision variables (*non-zero* probabilities $\pi^{ij}(t_1)$) is equal to the number of points on the grid reached by the trinomial process. If the number of degrees of freedom in the no-arbitrage system is large, then usually there will be multiple feasible solutions. We can easily increase the degrees of freedom by refining the multi-nomial process and the price-volatility grid. If this approach fails then there probably exists a severe inconsistency in the initial prices and/or the option pricing scheme (e.g. a violation of put-call parity with transaction costs). The dual of the system (12-14) can be inspected to track the source of such an inconsistency.

3.4 Constructing Event Trees

For a successful implementation of the SOH model it is crucial that we represent the prices of the assets and the liabilities efficiently in the structure of an event tree. In the previous sections we introduced a method to construct (transition) distributions of stock and option prices on a grid and ways to avoid arbitrage opportunities. In this section we will use these results to construct sparse event trees without arbitrage opportunities for the SOH model. First we partition the fine-grained grid into a number of buckets, each of which corresponds to a node in the sparse event tree. We prove that if we represent the prices in each bucket by the conditional expectation under a risk neutral measure, then there are no arbitrage opportunities in the event tree.

Let $\pi^{ij}(t_k)$ denote the *unconditional* risk neutral probability of the point (S^i, σ^j) on the grid at time t_k , for $i, j \in \mathcal{X}$ with $\mathcal{X} = \{i, j | i = 1, 2, \dots, n_I \text{ and } j = 1, 2, \dots, n_J\}$. Let $\pi_{ij}^{uv}(t_k, t_{k+1})$ denote the *conditional* risk neutral transition probability from (S^i, σ^j) at time t_k to (S^u, σ^v) at time

t_{k+1} . For ease of exposition we denote the stock price by $P_0^{ij}(t_k) = S^i$. The prices $P_h^{ij}(t_k)$ and dividends $D_h^{ij}(t_k)$ on the grid satisfy the following no-arbitrage condition:

$$(15) \quad P_h^{*ij}(t_k) = e^{-r(t_{k+1}-t_k)} \sum_{u,v \in \mathcal{X}} \pi_{ij}^{uv}(t_k, t_{k+1}) (P_h^{uv}(t_{k+1}) + D_h^{uv}(t_{k+1})),$$

$$(1 - \gamma_h) P_h^{ij}(t_k) \leq P_h^{*ij}(t_{k+1}) \leq (1 + \gamma_h) P_h^{ij}(t_k), \text{ for } h = 0, 1, 2, \dots, H$$

To get a sparse representation of the distribution we would like to have N_k nodes in the event tree at the trading date t_k . Therefore we partition the price-volatility grid at time t_k into N_k disjoint and exhaustive sets \mathcal{A}_n , with $\bigcup_{n=1}^{N_k} \mathcal{A}_n = \mathcal{X}$ and $\mathcal{A}_n \cap \mathcal{A}_m = \emptyset$ for $n, m = 1, 2, \dots, N_k$. Suppose we want to construct a set of successors $m = 1, 2, \dots, N_{k+1}$ at time t_{k+1} for node n . Similarly, we partition the grid at time t_{k+1} into N_{k+1} disjoint and exhaustive sets \mathcal{B}_m , for $m = 1, 2, \dots, N_{k+1}$.

We would like to represent each bucket by just one representative vector of asset prices in the event tree. Another important requirement for the event tree is that it should not contain arbitrage opportunities, as this would lead to unbounded or biased solutions of the SOH model. We now prove that if we represent the asset prices in each bucket by the conditional expectation under the risk neutral measure, then there are no arbitrage opportunities in the event tree:

$$(16) \quad E[P_h(t_k) | \mathcal{A}_n] = \sum_{i,j \in \mathcal{A}_n} \frac{\pi^{ij}(t_k)}{\sum_{u,v \in \mathcal{A}_n} \pi^{uv}(t_k)} P_h^{*ij}(t_k) = \sum_{i,j \in \mathcal{A}_n} \pi^{ij}(t_k | \mathcal{A}_n) P_h^{*ij}(t_k) =$$

$$= \sum_{i,j \in \mathcal{A}_n} \pi^{ij}(t_k | \mathcal{A}_n) \left[e^{-r(t_{k+1}-t_k)} \sum_{u,v \in \mathcal{X}} \pi_{uv}^{ij}(t_k, t_{k+1}) (P_h^{uv}(t_{k+1}) + D_h^{uv}(t_{k+1})) \right]$$

$$= e^{-r(t_{k+1}-t_k)} \sum_{u,v \in \mathcal{X}} \sum_{i,j \in \mathcal{A}_n} \pi^{ij}(t_k | \mathcal{A}_n) \pi_{ij}^{uv}(t_k, t_{k+1}) (P_h^{uv}(t_{k+1}) + D_h^{uv}(t_{k+1}))$$

$$= e^{-r(t_{k+1}-t_k)} \sum_{u,v \in \mathcal{X}} \pi_{\mathcal{A}_n}^{uv}(t_k, t_{k+1} | \mathcal{A}_n) (P_h^{uv}(t_{k+1}) + D_h^{uv}(t_{k+1}))$$

$$= e^{-r(t_{k+1}-t_k)} \sum_{m=1}^{N_{k+1}} \sum_{u,v \in \mathcal{B}_m} \pi_{\mathcal{A}_n}^{uv}(t_k, t_{k+1} | \mathcal{A}_n) (P_h^{uv}(t_{k+1}) + D_h^{uv}(t_{k+1}))$$

$$= e^{-r(t_{k+1}-t_k)} \sum_{m=1}^{N_{k+1}} \pi_{\mathcal{A}_n}^{\mathcal{B}_m}(t_k, t_{k+1} | \mathcal{A}_n) \sum_{u,v \in \mathcal{B}_m} \frac{\pi_{\mathcal{A}_n}^{uv}(t_k, t_{k+1} | \mathcal{A}_n)}{\pi_{\mathcal{A}_n}^{\mathcal{B}_m}(t_k, t_{k+1} | \mathcal{A}_n)} (P_h^{uv}(t_{k+1}) + D_h^{uv}(t_{k+1}))$$

$$= e^{-r(t_{k+1}-t_k)} \sum_{m=1}^{N_{k+1}} \pi_{\mathcal{A}_n}^{\mathcal{B}_m}(t_k, t_{k+1} | \mathcal{A}_n) E_{\mathcal{A}_n}[P_h(t_{k+1}) + D_h(t_{k+1}) | \mathcal{B}_m]$$

In order to generate an entire event tree while preserving the no-arbitrage property we apply equation (16) recursively, starting from the initial point. During this forward recursion each conditional probability distribution is generated with the trinomial process and we construct equivalent risk neutral measures. We refer to appendix B for a detailed description of the risk neutral aggregation algorithm based on (16).

The aggregation algorithm preserves the property of no-arbitrage and derives a risk neutral probability measure for the aggregated event tree. The risk neutral probabilities are worthless however for financial planning purposes. For example, from the no-arbitrage condition (15) it follows that the expected return of each asset under the risk neutral measure is approximately equal to the risk free rate r . For a financial planning model we should additionally specify the *actual* probability distribution for the asset prices, which is given by the N-GARCH process (9) in this paper. It is clear from the N-GARCH process that the expected asset return under the actual probability distribution increases with the riskiness of the asset: there is a positive premium for risk λ .

To get a good approximation of the N-GARCH process we would prefer to aggregate the event tree by taking conditional expectations of the asset prices in the buckets under the *actual probability measure*. However in order to guarantee the absence of arbitrage opportunities, each bucket of nodes is represented by the conditional expectation of the prices under the *risk neutral probability measure*. We could try to eliminate this approximation error by representing the prices in each bucket by the conditional expectation under the actual measure and next we could check for arbitrage opportunities by solving the linear system (15). If the system is feasible, then we can use these prices in the event tree without any problem. However if the no-arbitrage system is infeasible, then we have to reject the set of prices and we are forced to accept some approximation error. We could still minimize the distance of the prices in the event tree to the conditional expectation under the actual measure by solving a small non-linear optimization model.

4 Solving and Testing the SOH Model

In this section we show how the SOH model can be solved numerically. First the nonlinear SOH model is reformulated as a multi-stage stochastic programming model. For multi-stage stochastic programming models a wide range of specialized solution methods and parallel decomposition techniques are available. Finally, we test the performance of the SOH model and other hedging strategies with a simulation for a specific hedging problem.

4.1 Solving the SOH Model with Stochastic Programming Methods

The number of variables and constraints of the SOH model grows exponentially with the number of trading dates, due to the path dependency of the decisions. Moreover, the SOH model includes a nonlinear restriction for modeling proportional transaction costs. Given the computational complexity of the problem, solving interesting instances of the model seems impossible at first sight. However as shown by Edirisinghe, Naik, and Uppal (1993), the nonlinear restriction can be linearized by introducing additional variables for the number of assets bought and the number of assets sold.

The linearized version of the SOH model has the structure of a linear multi-stage stochastic programming model, with a possibly non-linear objective (e.g. a concave utility function). For

Table 1: Estimated Coefficients of the Asymmetric N-GARCH Model.

Parameter	Value	St.Err.	T-Value
λ	0.0518	0.0205	2.528
β_0	0.0248	0.0043	5.774
β_1	0.8756	0.0120	72.886
β_2	0.0934	0.0091	10.231
c	0.2570	0.0472	5.442

Estimated coefficients of the asymmetric N-GARCH model (9), using a sample of 2584 daily AEX index returns from the period 2-1-1989 up to 26-11-1998. The model has been estimated by maximizing the likelihood function with the BHHH algorithm, without restrictions on the parameters.

the class of linear multi-stage stochastic programming problems, a wide range of specialized solution methods and decomposition techniques have been developed (see Birge and Louveaux 1997). A number of decomposition methods split the stochastic programming problem up into small parts which can be solved independently on a parallel machine. With a suitable parallel decomposition approach, models consisting of millions of variables and millions of constraints can be solved (see Vladimirou and Zenios 1997).

For the numerical example solved in this section we apply the primal-dual column generation method of Gondzio and Sarkissian (1996), which is a descendent of the general Dantzig-Wolfe (1961) and Benders (1962) decomposition methods. The algorithm splits up a multi-stage stochastic programming problem into a first stage master problem (corresponding to the first few trading dates) and a large number of relatively small second stage subproblems (representing later trading dates). The subproblems could be solved independently on the processors of a parallel machine with the interior point algorithm HOPDM (Gondzio 1995). With this state-of-the-art solution approach Gondzio and Kouwenberg (1999) solved a stochastic programming model for financial planning with 6 trading dates and 5 million scenarios at the planning horizon.

4.2 Numerical Example of a Hedging Problem

As an experiment we consider the hedging problem of a financial institution that has sold an at-the-money call option on the Amsterdam Exchanges Index (AEX index), with a maturity of one year (250 trading days). To hedge this short position the institution could use futures on the AEX index and five liquidly traded call options on the AEX index with a maturity of 3 months and moneyness equal to 0.950, 0.975, 1.000, 1.025 and 1.050 respectively. We estimate and forecast the prices of the options under the assumption that the AEX index follows the asymmetric N-GARCH model (9). We have estimated the N-GARCH model with daily returns from the period 2-1-1989 up to 26-11-1998, a sample of 2584 observations. Table 1 shows the estimated coefficients, which are all significant. The N-GARCH forecast of the conditional standard deviation of the AEX-index return at 27-11-1998 is 28.1%.

We calculate option prices with a trinomial process for the N-GARCH process and linear interpolation on a grid of stock price and volatility, as described in Section 3. The grid consists of layers with 1000 points for the stock price versus 100 points for the volatility, with 250 time-

steps of one trading day up to the maturity of the 1-year option. We assume that the financial institution hedges on a weekly basis (after 5 trading days each) and include the first three weeks in the SOH model. The event tree for the SOH model is constructed with the lattice aggregation algorithm of appendix B. Each 5-day conditional return distribution of prices on the grid is partitioned into 15 buckets with approximately equal probability. The asset prices in each bucket are represented by the conditional expectation under the risk neutral measure.

With a planning horizon of three weeks, the aggregated event tree for the SOH model has 15 nodes after one week, 225 nodes after two weeks and 3375 nodes after three weeks. We checked the consistency of the risk neutral measure and the prices of the aggregated event tree and found that the risk neutral pricing conditions (12-14) are satisfied in each node. We use the aggregated actual probabilities for the nodes, *not the risk neutral probabilities*, for calculating expected hedging errors in the SOH model.

For the objective function of the SOH model we choose a downside risk measure: the goal of the model is to minimize the average of the absolute negative hedging errors over all nodes in the event tree. We assume a transaction cost of 0.25% for buying or selling of the AEX index futures and a transaction cost of 2.5% for trading the three month call options on the AEX index. Money can be lent at a continuously compounded rate of 5% and borrowed at an additional spread of 0.5%. The initial wealth of the financial institution is set equal to the ask price of the 1-year at-the-money call option of fl 8.41. This means that the financial institution has sold the option at a premium of 2.5%.

4.3 Evaluation of Hedge Performance via Simulation

We test the performance of the SOH model and other hedging strategies with a simulation for the hedging problem described previously. First we solve the SOH model with three trading dates in order to get the SOH hedging strategy. Next we simulate paths from the calibrated GARCH model, using the trinomial process on the grid of stock price versus volatility. For each simulated path we implement the hedging strategy of the SOH model at the three weekly trading dates. Note that each node in the aggregated event tree represents a number of points on the grid of stock price versus volatility. We use this relationship to get the appropriate weekly rebalancing decisions from the SOH model on each simulated path.

We compare the SOH strategy with a delta hedging strategy and a delta-vega hedging strategy. Both strategies are rebalanced to a neutral position at the three weekly trading dates. For the delta-vega hedging strategy we use the at-the-money option to eliminate the volatility exposure. We sample a total of 100000 simulation paths, while applying antithetic sampling as a variance-reduction technique. Note that the performance of the SOH strategy could be improved by resolving the SOH model at each trading date of each simulation path, instead of selecting *future* decisions from the *initial* event tree. However, resolving the SOH model during simulations would require a tremendous computational effort.

Table 2 shows the results of simulations where short selling of assets and borrowing are unrestricted. Panel A reports the results of the delta hedging strategy, panel B shows the performance of the delta-vega hedging strategy and panel C represents the SOH strategy. The results clearly

show that the delta hedging strategy (panel A) performs very badly: its average negative hedging error is approximately 100 times worse than the hedging error of the SOH model. This poor performance comes as no surprise, given that a pure delta hedging strategy ignores volatility movements. The delta-vega hedging strategy in panel B performs much better, because it leads to lower hedging errors and it entails less trading. The optimal SOH strategy in panel C improves the results even further by reducing transaction costs. Note that in panel C the average absolute gap between the delta (vega) of the hedge portfolio and the delta (vega) of the 1-year option is quite small. The optimal SOH policy looks like a slightly ‘loose’ delta-vega hedge, with some slack to avoid needless transactions.

Another important property is that the initial SOH strategy only involves options with moneyness 0.950 and 1.050, while the options with moneyness 0.975, 1.000 and 1.025 are absent in the hedge-portfolio. This result probably follows from the observation that the vega of an at-the-money option is higher than the vega of both out-of-the money and in-the-money options. A graph of vega versus moneyness looks like a parabola, with the top approximately located at-the-money. Moreover, this effect is stronger (i.e. the parabola is steeper) for short maturity options than for long maturity options. Suppose that we set up a vega hedge with a short maturity at-the-money option. Then we might end up with a severe vega mismatch if the stock price increases or decreases considerably. However, if we use both in-the-money and out-of-the-money options to construct the vega hedge then the effects of a stock price change on the vega’s of the options could partly cancel. In this way less transactions might be needed to keep the portfolio vega-neutral.

In practice portfolio managers and traders face limits on the amounts they can borrow and sell short. In Table 3 we limit the short selling of each asset and borrowing to 50% of the portfolio value. Note that the delta hedging strategy and the delta-vega hedging strategy are both infeasible now: they require too much borrowing. Panel D shows the results for a pure static replication strategy: option trading is only allowed at the initial date and stock trading is never allowed. Though the transaction costs in this case are minimal, the hedging error is quite big, probably due to increasing delta and vega gaps. Panel E shows the optimal SOH strategy within the short selling bounds. The hedging error in panel E is relatively low, while the delta and vega gap are reduced compared to the static option portfolio.

As a concluding remark: at first sight it seems that the hedging strategies of the SOH model make sense both intuitively and economically. The optimal SOH solution is on average quite close to a delta-vega hedging strategy, but with a bit of slack in order to avoid needless trading. A drawback of the SOH model is that the path-dependency of the strategies limits the number of future trading dates that we can consider (otherwise both the event tree and the computational effort explode). Note however that the reduction of transaction costs by the optimal SOH policy seems considerable compared to the strict delta-vega hedging strategy, even though we included only three trading dates in the model.

5 Conclusions

In this paper we considered hedging of options under transaction costs and stochastic volatility. As traditional methods like delta-vega hedging and static hedging are not fully appropriate in this context, we introduced the stochastic optimization hedging (SOH) model. The SOH model takes account of transaction costs, stochastic volatility and trading restrictions. The model has a number of trading dates on which the hedge portfolio of stocks and traded options can be rebalanced. The goal of the model is to minimize the hedging errors at the first few trading dates by following an appropriate dynamic trading strategy.

The SOH model requires an event tree of asset prices as input. The performance of the SOH hedge depends crucially on the quality of this event tree, which is an approximation of the underlying price process. Our contribution is that we propose reliable methods to construct event trees for stochastic volatility models. Without loss of generality we focus on the asymmetric N-GARCH model as an underlying process. First we represent the N-GARCH model on a grid of stock price versus volatility using a trinomial process. Next we propose methods to avoid arbitrage opportunities in the set of stock and option prices on the grid. Finally we construct sparse event trees for the SOH model by applying an aggregation algorithm which preserves the no-arbitrage property.

We investigated the performance of the SOH model with a simulated hedging problem. Even though the SOH model for this example consists of just three trading dates, the solutions make sense both intuitively and economically. On average the hedging policy of the SOH model is quite close to a delta-vega neutral strategy, but with some slack to avoid needless transaction costs. The SOH hedging policy reduced transaction costs considerably compared to traditional delta hedging and delta-vega hedging strategies. Moreover the SOH strategy could easily incorporate restrictions on short selling and borrowing and outperformed the optimal static hedge portfolio.

References

- [1] Avellaneda M., A. Levy, and A. Paras (1995), Pricing and Hedging Derivative Securities in Markets with Uncertain Volatilities, *Applied Mathematical Finance*, vol. 2, 73-88.
- [2] Avellaneda M., and A. Paras (1996), Managing the Volatility Risk of Portfolios of Derivative Securities: the Lagrangian Uncertain Volatility Model, *Applied Mathematical Finance*, vol. 3, 21-52.
- [3] Bakshi G., C. Cao and Z. Chen (1997), Empirical Performance Of Alternative Option Pricing Models, *Journal of Finance*, vol. 52, 2003-2049.
- [4] Benders J.F. (1962), Partitioning Procedures for Solving Mixed-Variables Programming Problems, *Numerische Mathematik*, vol. 4, 238-252.
- [5] Birge J.R. and F. Louveaux (1997), *Introduction to Stochastic Programming*, Springer-Verlag, New York,.
- [6] Black F. and M. Scholes (1973), The Pricing of Options and Corporate Liabilities, *Journal of Political Economy*, vol. 81, 637-659.
- [7] Bollerslev T. (1986), Generalized Autoregressive Conditional Heteroscedasticity, *Journal of Econometrics*, vol. 31, 307-327.
- [8] Bollerslev T., R.Y. Chou and K.F. Kroner (1992), ARCH Modelling in Finance: a Review of the Theory and Empirical Evidence, *Journal of Econometrics*, vol. 52, 55-59.
- [9] Boyle P.P. and T. Vorst (1992), Option Replication in Discrete Time with Transaction Costs, *Journal of Finance*, vol. 47, 271-291.
- [10] Carr P., K. Ellis, and V. Gupta (1998), Static Hedging of Exotic Options, *Journal of Finance*, Vol 53, 1165-1190.
- [11] Dantzig G.B. and P. Wolfe (1961), The Decomposition Algorithm for Linear Programming, *Econometrica*, vol. 29, 767-778.
- [12] Derman E., D. Ergener and I. Kani (1995), Static options replication, *The Journal of Derivatives*, Vol. 2, 78-95.
- [13] Duan J.-C. (1995), The GARCH Option Pricing Model, *Mathematical Finance*, vol. 5, 13-32.
- [14] Duan J.-C. (1997), Augmented GARCH(p,q) Process and its Diffusion Limit, *Journal of Econometrics*, vol. 79, 97-127.
- [15] Edirisinghe C., V. Naik and R. Uppal (1993), Optimal Replication of Options with Transaction Costs and Trading Restrictions, *Journal of Financial and Quantitative Analysis*, vol. 28, 117-139.
- [16] Engle R.F. (1982), Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of U.K. Inflation, *Econometrica*, vol. 50, 987-1008.

- [17] Engle R.F. and V. Ng (1993), Measuring and Testing the Impact of News and Volatility, *Journal of Finance*, vol. 48, 1749-1779.
- [18] Finucane T.J. and M. Tomas (1997), American Stochastic Volatility Call Option Pricing: A Lattice Based Approach, *Review of Derivatives Research*, vol. 1, 183-201.
- [19] Frey R. and C.A. Sin (1999), Bounds on European Option Prices under Stochastic Volatility, *Mathematical Finance*, vol. 9, 97-116.
- [20] Gondzio J.(1995), HOPDM (version 2.12) - A Fast LP Solver Based on a Primal-Dual Interior Point Method, *European Journal of Operational Research*, vol. 85, 221-225.
- [21] Gondzio J. and R. Kouwenberg (1999), High Performance Computing for Asset Liability Management, Preprint MS-99-004, Department of Mathematics & Statistics, The University of Edinburgh, UK.
- [22] Gondzio J. and R. Sarkissian (1996), Column Generation with a Primal-Dual Method, Logilab Technical Report 96.6, Department of Management Studies, University of Geneva, Switzerland.
- [23] Harrison J.M. and D.M. Kreps (1979), Martingales and Arbitrage in Multiperiod Securities Markets, *Journal of Economic Theory*, vol. 20, 381-408.
- [24] Heston S.L. (1993), A Closed-Form Solution for Options with Stochastic Volatility, *Review of Financial Studies*, vol. 6, 327-344.
- [25] Heston S.L. and S. Nandi (1997), A Closed-Form GARCH Option Pricing Model, Working Paper 97-9, Federal Reserve Bank of Atlanta.
- [26] Heston S.L. and S. Nandi (1998), Preference-Free Option Pricing with Path-Dependent Volatility: A Closed-Form Approach, Working Paper 98-20, Federal Reserve Bank of Atlanta.
- [27] Hull J. and A. White (1987), The Pricing of Options on Assets with Stochastic Volatilities, *Journal of Finance*, vol. 42, 218-300.
- [28] Klaassen P. (1997), Discretized Reality and Spurious Profits in Stochastic Programming Models for Asset/Liability Management, *European Journal of Operational Research*, vol. 101, 374-392.
- [29] Leland H.E. (1985), Option Pricing and Replication with Transaction Costs, *Journal of Finance*, vol. 40, 1283-1301.
- [30] Naik V. (1995), Finite State Securities Market Models and Arbitrage, in *Handbooks in OR&MS*, Vol. 9, eds.: R. Jarrow et al, Elsevier Science, 31 - 64.
- [31] Nelson (1990), ARCH Models as Diffusion Approximations, *Journal of Econometrics*, vol. 45, 7-38.
- [32] Ritchken P. and R. Trevor (1999), Pricing Options Under Generalized GARCH and Stochastic Volatility Processes, *Journal of Finance*, vol. 54, 377-402.

- [33] Scott L. (1987), Option Pricing When the Variance Changes Randomly: Theory, Estimation and An Application, *Journal of Financial and Quantitative Analysis*, vol. 22, 419-438.
- [34] Schweizer M. (1991), Option Hedging for Semimartingales, *Stochastic Processes and their Applications*, vol. 37, 339-363.
- [35] Schweizer M. (1995), Variance-Optimal Hedging in Discrete Time, *Mathematics of Operations Research*, vol. 20, 1-32.
- [36] Vladimirov H. and S.A. Zenios (1997), Parallel Algorithms for Large-Scale Stochastic Programming, in *Parallel Computing in Optimization*, eds.: Migdals S., P. Pardalos and S. Storøy, Kluwer Academic Publishers, 413 - 469.
- [37] Wiggins J. (1987), Option Values under Stochastic Volatility: Theory and Empirical Evidence, *Journal of Financial Economics*, vol. 19, 351-372.

Appendix A: Mathematical Setup of the SOH Model

The current time is denoted by t_0 and we assume that the investor plans to rebalance his hedge at the trading dates t_1, t_2, \dots, t_K . The stochastic processes of the asset prices and the factors in the economy are described by a finite sample space Ω consisting of Q elements: $\Omega = \{\omega_1, \omega_2, \dots, \omega_Q\}$. A probability measure \mathcal{P} is defined on the sample space Ω . The information structure is described by the filtration $F = \{\mathcal{F}_t : t = t_0, t_1, \dots, t_K\}$ consisting of a nested sequence \mathcal{F}_t of algebras on Ω , with $\mathcal{F}_0 = \emptyset$, $\mathcal{F}_t \subset \mathcal{F}_{t+1}$ for $t = t_0, t_1, \dots, t_{K-1}$ and $\mathcal{F}_{t_K} = \Omega$. Corresponding to each algebra \mathcal{F}_t is a unique finest partition p_t of the sample space Ω . The partition p_t consists of N_t subsets p_{nt} of Ω with $\bigcup_{n=1}^{N_t} p_{nt} = \Omega$ and $\bigcap_{n=1}^{N_t} p_{nt} = \emptyset$.

The investor seeks to hedge the value of the general contingent claim $G(t, \omega)$ on a subset p_t^G of the partition p_t for $t = t_0, t_1, \dots, t_K$. The set of hedge instruments consists of the underlying stock $S(t, \omega) = P_0(t, \omega)$, the exchange traded vanilla options $P_h(t, \omega)$ for $h = 1, 2, \dots, H$ and the money market account $M(t, \omega)$. At each trading date t the investor should only have information about the subset of the partition p_t containing the realized sample path ω . Therefore the price processes $P_h(t, \omega)$ for $h = 0, 1, 2, \dots, H$ and $G(t, \omega)$ are adapted to the filtration F : the asset prices are constant on the finest partition p_t corresponding to the algebra \mathcal{F}_t at each trading date $t = t_0, t_1, \dots, t_K$.

The portfolio of the investor at time t just after trading is denoted by $\Lambda(t, \omega) = (M(t, \omega), \lambda_0(t, \omega), \lambda_1(t, \omega), \dots, \lambda_H(t, \omega))$, where $M(t, \omega)$ denotes the position in the money market account, $\lambda_0(t, \omega)$ is the position in the underlying stock and $\lambda_h(t, \omega)$ represents the position in the plain vanilla options for $h = 1, 2, \dots, H$. The trading strategy $\{\Lambda(t, \omega) : t = t_1, t_2, \dots, t_K\}$ is a predictable process with respect to the filtration F : e.g. $\Lambda(t, \omega)$ is constant on every subset of the partition p_{t-1} (i.e. measurable) for all $t = t_1, t_2, \dots, t_K$.

Appendix B: Risk Neutral Aggregation Algorithm

This appendix introduces a risk neutral aggregation algorithm based on equation (16). If the prices in the original event tree are arbitrage-free and we have a corresponding risk neutral measure, then the aggregated event tree will inherit this property by construction. From the set of dates $t = 0, 1, \dots, T$ we select a number of trading dates for inclusion in the stochastic optimization model. Let $\{t_0, t_1, \dots, t_K\}$ denote the K trading dates of the stochastic optimization model, where $t_0 = 0$, $t_K \leq T$ and $t_j < t_{j+1}$ holds. The dates $t = 0, 1, \dots, T$ that are not part of the set of trading dates will be aggregated, by applying the following algorithm:

1. Initialization:

set $k = 0$ and start with one partition $\mathcal{A}_1(t_0)$ and $N_{\mathcal{A}}(t_0) = 1$, consisting of the initial node at time t_0 .

2. Consider the first bucket:

$j(t_k) = 1$.

3. Partitioning:

partition the nodes at time t_{k+1} into buckets \mathcal{B}_l , for $l = 1, \dots, N_{\mathcal{B}}$.

4. Conditional probabilities:

The nodes at time t_k have been partitioned in buckets $\mathcal{A}_m(t_k)$ for $m = 1..N_{\mathcal{A}}(t_k)$. The active bucket is $\mathcal{A}_{j(t_k)}(t_k)$. Get the conditional risk neutral probabilities for the buckets \mathcal{B}_l at time t_{k+1} , given that we start in bucket $\mathcal{A}_{j(t_k)}(t_k)$ at time t_k .

5. Aggregation:

Use the conditional risk neutral probabilities to aggregate the prices at time t_{k+1} , by taking the conditional expectation in each bucket \mathcal{B}_l for $l = 1, \dots, N_{\mathcal{B}}$. Store these prices and probabilities for the aggregated event tree.

6. Next period:

if $k + 1 \leq K - 1$ then

(a) Store the buckets:

$N_{\mathcal{A}}(t_{k+1}) = N_{\mathcal{B}}$, $\mathcal{A}_m(t_{k+1}) = \mathcal{B}_m$ for $m = 1..N_{\mathcal{B}}$.

(b) $k = k + 1$

(c) Goto 2.

7. Next bucket or previous period:

(a) if $j(t_k) < N_{\mathcal{A}}(t_k)$ then

i. $j(t_k) = j(t_k) + 1$

ii. goto 3

(b) $k = k - 1$

(c) if $k = 0$ then the aggregation algorithm is finished: STOP.

(d) goto 7a

Table 2: Results of the Simulation I

<u>A. Delta Hedging.</u>				
Average Negative Hedging Error 0.1353				
Initial Portfolio Weights				
AEX Index	C095	C100	C105	Borrow
687.76 %	0	0	0	587.76 %
Date	T-costs	Turn-over	Delta Gap	Vega Gap
1	0.1418	688 %	0	100 %
2	0.0154	83 %	0	100 %
3	0.0153	84 %	0	100 %
<u>B. Delta-Vega Hedging.</u>				
Average Negative Hedging Error 0.0063				
Initial Portfolio Weights				
AEX Index	C095	C100	C105	Borrow
257.46 %	0	36.91 %	0	194.37 %
Date	T-costs	Turn-over	Delta Gap	Vega Gap
1	0.1284	292 %	0	0
2	0.0125	27 %	0	0
3	0.0273	48 %	0	0
<u>C. SOH Optimal Trading Strategy.</u>				
Average Negative Hedging Error 0.0014				
Initial Portfolio Weights				
AEX Index	C095	C100	C105	Borrow
200.07 %	34.18 %	0	11.29 %	145.64 %
Date	T-costs	Turn-over	Delta Gap	Vega Gap
1	0.1354	247 %	0.24 %	0.70 %
2	0.0070	24 %	0.87 %	5.59 %
3	0.0034	12 %	1.51 %	10.35 %

The table shows the results of simulations with different hedging strategies. Panel A represents delta-hedging with stock trading only. Panel B represents delta-vega-hedging with the at-the-money option and the stock. Panel C represents the optimal trading strategy of the SOH model, involving the five available traded options and the stock. For each model the objective value and the initial portfolio weights are reported, where C095 represents the call option with moneyness 0.95. Next for each trading date are shown: the average transaction costs in cents, the average portfolio turn-over, the average absolute delta gap and the average absolute vega gap. The portfolio turn-over is defined as the sum of the absolute changes of the money invested in each asset divided by two times the value of the portfolio. The absolute delta (vega) gap is the absolute difference between the delta (vega) of the hedge portfolio and the delta (vega) of the 1-year call option.

Table 3: Results of the Simulation II

<u>D. Static Replication Portfolio.</u>				
Average Negative Hedging Error 0.0138				
	Initial Portfolio Weights			
AEX Index	C095	C105	Borrow	Lend
0	67.99 %	2.21 %	0	29.80 %
	T-costs	Turn-over	Delta Gap	Vega Gap
Date	1	71 %	1.61 %	5.50 %
2	0	0	2.97 %	13.07 %
3	0	0	4.58 %	19.02 %
<u>E. SOH Optimal Trading Strategy.</u>				
Average Negative Hedging Error 0.0047				
	Initial Portfolio Weights			
AEX Index	C095	C105	Borrow	Lend
74.37 %	55.68 %	5.62 %	35.67 %	0
	T-costs	Turn-over	Delta Gap	Vega Gap
Date	1	137 %	0.68 %	3.72 %
2	0.0074	10 %	1.32 %	6.30 %
3	0.0053	12 %	1.72 %	13.31 %

The table shows the results of simulations with different hedging strategies. In panel D and E borrowing is restricted to 50% of the portfolio value and short selling of each asset is also restricted to 50% of the portfolio value. Panel D represents the static replication strategy, where option trading is only allowed at the initial date, while stock trading is never allowed. Panel E represents the optimal strategy of the SOH model, where both option and stock trading are unlimited within the short selling bound. For each model the objective value and the initial portfolio weights are reported, where C095 represents the call option with moneyness 0.95. Next for each trading date are shown: the average transaction costs in cents, the average portfolio turn-over, the average absolute delta-gap and the average absolute vega-gap.




Figure 1: Event tree for the stock price and its volatility