A NUMERICAL METHOD FOR FINANCIAL DECISION
PROBLEMS UNDER STOCHASTIC VOLATILITY

Enlu Zhou
Kun Lin
Michael C. Fu
Steven I. Marcus

Institute for Systems Research
University of Maryland, College Park
College Park, MD 20742, USA

ABSTRACT

In many financial decision problems, such as portfolio optimization or hedging, the goal is to compute an optimal investment strategy, in order to maximize expected utility or minimize expected risk. If the volatility of the risky asset(s) follows a stochastic process and is not observable, the problems usually do not have analytical solutions. Hence, we propose an efficient numerical method for these problems, based on a method developed recently for solving continuous-state partially observable Markov decision processes. Numerical applications are presented and discussed for a problem of hedging European put and call options.

1 INTRODUCTION

In many financial decision problems, such as portfolio optimization or hedging, an investor wants to find an optimal investment strategy to achieve a certain goal, for example, to maximize a utility function or to minimize a risk. He invests in the various assets available in market, including riskless assets such as bonds and deposits, and risky assets such as stocks and options. The price of a risky asset is usually modeled as a stochastic process, where the volatility is treated as a constant in some models, such as the classical Black-Scholes-Merton model (Black and Scholes 1973) (Merton 1973). However, stochastic volatility (SV) models are used heavily in financial engineering to capture the impact of time-varying volatility on financial markets and decision making (Shephard and Andersen 2008). Analytical results for SV models are derived mostly for option pricing, such as (Heston 1993) (Bates 1996) (Nicolato and Venardos 2003). For portfolio optimization or hedging, SV models in general do not have analytical results, and are hard to solve numerically due to computational difficulties.

Many financial decision problems under stochastic volatility fall into the framework of a stochastic control problem under partial observation, where the control is the investment strategy, the unobserved state is the risky asset volatility, and the observation is the risky asset price. A partially observed stochastic control problem can be transformed to a fully observed problem, but the transformation introduces a new state that is a continuous probability density function. Since a continuous density is generally infinite dimensional, the fully observable stochastic control problem has an infinite-dimensional state space. Efficient numerical methods for these problems are sparse, as opposed to a large number of existing methods for finite-dimensional stochastic control problems. Based on the idea of density projection for dimension reduction, we have recently developed an effective method for solving continuous-state partially observable Markov decision processes (Zhou, Fu, and Marcus 2008). In this paper, we will adapt this method to financial decision problems under stochastic volatility, and solve the problem of hedging European options in particular. A numerical method proposed in (Corsi, Pham, and Runggaldier 2007) solves the same problem through quantization of the observation and filtering processes, i.e., the price and the conditional density processes. The quantization of the conditional density still results in a relatively large dimension, compared to our density projection approach.

The rest of the paper is organized as follows. Section 2 sets up the general problem formulation of financial decision problems under stochastic volatility. Section 3 presents a numerical method for solving such problems. Section 4 applies the method to the hedging of European options under the quadratic risk and shortfall risk criteria, discusses the implication of the numerical results, and compares with other numerical methods. Section 5 concludes the paper.

2   FINANCIAL DECISION PROBLEMS UNDER STOCHASTIC VOLATILITY

2.1 Financial Example

For simplicity, we consider that an investor invests in a riskless asset with constant interest rate r and a risky asset whose price per share at time t is S_t. Suppose S_t is a geometric Brownian motion that satisfies the following stochastic differential equation (SDE):

$$dS_t = rS_t dt + X_t S_t dW_t, \tag{1}$$

where X_t is the (unobserved) volatility of the asset, and W_t is a Wiener process. The volatility X_t follows a mean-reverting process:

$$dX_t = \lambda(x_0 - X_t)dt + \eta dZ_t, \tag{2}$$

where the rate of reversion λ, mean reversion value x_0, and η are constant known parameters, and Z_t is a Wiener process that is independent of W_t.

Investment decisions take place at discrete times $t_0, t_1, \ldots, t_T$, simply denoted by $0, 1, \ldots, T$ in the following. Each time period lasts $\varepsilon = t_{k+1} - t_k$. Under the assumption that no additional capital is added (self-financing assumption), the investor observes the risky asset price S_k at time k, then decides the number of shares a_k invested in the risky asset, and the remainder of the capital is invested in the riskless asset. According to the analytical solution to (1), S_k evolves as

$$S_{k+1} = S_k \exp\left[(r - \frac{X_{k+1}^2}{2})\varepsilon + X_k\sqrt{\varepsilon}W_k\right],$$

where W_k is an i.i.d. sequence of standard Gaussian random variables. Let $Y_k = \ln S_k$, often referred to as the log-price. Then

$$Y_{k+1} = Y_k + (r - \frac{X_{k+1}^2}{2})\varepsilon + X_k\sqrt{\varepsilon}W_k. \tag{3}$$

The volatility process (2) can be discretized in time using the Euler's method (Milstein 1995), yielding

$$X_{k+1} = X_k + \lambda(x_0 - X_k)\varepsilon + \eta\sqrt{\varepsilon}Z_k,$$

where Z_k is an i.i.d. sequence of standard Gaussian random variables that are independent of W_k. Euler's method is the simplest discretization scheme, and is used here for simplicity. To achieve better approximation, we can use other higher order approximation methods in (Milstein 1995). The mean-reverting process (2) is an Ornstein-Uhlenbeck process that has an analytical solution, and hence can be discretized more precisely as

$$X_{k+1} = X_k e^{-\lambda\varepsilon} + x_0(1 - e^{-\lambda\varepsilon}) + \eta e^{-\lambda\varepsilon}\sqrt{\varepsilon}Z_k.$$

An investment strategy is denoted by $\pi = \{\rho_1, \ldots, \rho_T\}$, where each ρ_k is a function that maps the price history $Y_{0:k} \triangleq \{Y_0, \ldots, Y_k\}$ to an action a_k, i.e., the action a_k is selected based on the past observations of the risky asset price. Under the assumption of self-financing investment strategies, the wealth process evolves as

$$\begin{aligned} V_{k+1} &= (V_k - a_k S_k)e^{r\varepsilon} + a_k(S_{k+1} - S_k) \\ &= V_k e^{r\varepsilon} + \rho_k(Y_{0:k})(e^{Y_{k+1}} - e^{Y_k + r\varepsilon}). \end{aligned} \tag{4}$$

Let Π denote the set of all admissible investment strategies. In hedging, given a loss function $l : \mathbb{R} \to \mathbb{R}$, the hedging criterion for a derivative asset $h(Y_T)$ at maturity time T, is to minimize the expected loss over all admissible investment strategies:

$$\min_{\pi \in \Pi} E\left[l(h(Y_T) - V_T)\right].$$

In portfolio optimization, given a utility function $u : \mathbb{R} \to \mathbb{R}$, the goal is to maximize the expected utility over all admissible investment strategies:

$$\max_{\pi \in \Pi} E\left[u(V_T)\right].$$

2.2 Stochastic Control under Partial Observation

Hedging and portfolio optimization can be cast as stochastic control problems under partial observation as follows:

$$
\begin{align}
\min_{\pi \in \Pi} \quad J(x,v) \quad &= \quad E\left[g(Y_T,V_T)|X_0 = x, V_0 = v\right], \tag{5}\\
s.\,t. \quad X_{k+1} \quad &= \quad F(X_k,Z_k), \tag{6}\\
Y_{k+1} \quad &= \quad H(X_{k+1},X_k,Y_k,W_k), \tag{7}\\
V_{k+1} \quad &= \quad G(V_k,Y_{k+1},Y_k,\rho_k(Y_{0:k})), \tag{8}
\end{align}
$$

where functions g, F, H, and G are of suitable forms, X is the unobserved state, Y is the observation of X, V is the fully observed state, and ρ is the control. An optimal investment strategy, denoted by π^*, is an admissible investment strategy that achieves the minimum of (5). Please note that the above problem does not restrict W_k and Z_k to be Gaussian. Therefore, the numerical method we develop in the following applies to problems with any randomly distributed W_k and Z_k, as long as their distributions can be sampled.

Since X_k is unobserved, it needs to be estimated based on all the information available at time k, which includes the observation history $Y_{0:k}$, the action history $a_{0:k-1} = \{a_0,\ldots,a_{k-1}\}$, and the fully observed state history $V_{0:k} = \{V_0,\ldots,V_k\}$. With a little inspection, it is easy to see that for a fixed initial condition $V_0 = v$, a_{k-1} and V_k are both measurable with respect to $Y_{0:k}$. Hence, it is sufficient to use only $Y_{0:k}$ in the estimation of X_k. Define B_k as the conditional probability density of the X_k given $Y_{0:k}$,

$$B_k(x) \triangleq p(X_k = x|Y_{0:k}).$$

Since B_k is a sufficient statistic of $Y_{0:k}$, an investment strategy at time k, i.e., ρ_k, should be a function that maps B_k to an action a_k. Therefore, at each time k, the investor observes Y_k, estimates B_k, decides an action $a_k = \rho_k(B_k)$, and then reaches the wealth state V_{k+1} at the next time $k+1$.

3 NUMERICAL METHOD

3.1 Full Observation

In fact, by treating B_k as a state to replace the unobserved state X_k, the problem under partial observation (Q) can be transformed to an equivalent problem under full observation. The state B_k is observable, because it can be computed recursively via Bayes' rule using observed quantities:

$$
\begin{align}
B_{k+1}(x|Y_{0:k+1}) \quad &= \quad \int p(X_{k+1} = x, X_k | Y_{0:k}, Y_{k+1}) dX_k \\
&= \quad \frac{1}{p(Y_{k+1}|Y_{0:k})} \int p(Y_{k+1}|X_{k+1} = x, X_k, Y_{0:k}) p(X_{k+1} = x, X_k|Y_{0:k}) dX_k \\
&\propto \quad \int p(Y_{k+1}|X_{k+1} = x, X_k, Y_k) p(X_{k+1} = x|X_k, Y_{0:k}) p(X_k|Y_{0:k}) dX_k \\
&= \quad \int p(Y_{k+1}|X_{k+1} = x, X_k, Y_k) p(X_{k+1} = x|X_k) B_k(X_k|Y_{0:k}) dX_k, \tag{9}
\end{align}
$$

where $p(Y_{k+1}|X_{k+1},X_k,Y_k)$ is induced by (7), and $p(X_{k+1}|X_k)$ is induced by (6). The righthand side of (9) is a function of B_k, Y_k and Y_{k+1}. Hence, the evolution of B_k can be written as

$$B_{k+1} = \Psi(B_k,Y_k,Y_{k+1}). \tag{10}$$

Replacing X_k and (6) with B_k and (10), the resultant stochastic control problem is fully observable, and (B_k, Y_k, V_k) can be all viewed as states. Therefore, in principle we can apply dynamic programming to solve this problem:

$$
\begin{aligned}
J_T(b,y,v) &= g(y,v), \\
J_k(b,y,v) &= \min_{a \in A} E\left[J_{k+1}(B_{k+1}, Y_{k+1}, G(v, Y_{k+1}, y, a)) | B_k = b, Y_k = y\right], \quad k = 0, \dots, T-1.
\end{aligned}
$$

However, B_k is a continuous density, and hence is infinite dimensional in general. The above dynamic programming cannot be implemented in practice unless we can approximate B_k by some finite (preferably low) dimensional density. One naive way is to approximate B_k by a probability mass function (p.m.f.) supported on discretized points of X_k. This usually results in a very high dimension in order to achieve sufficient accuracy. An improvement is to adaptively change the support points in a certain way so that they concentrate on the area where B_k has most probability mass. To further reduce the dimensionality, we can use moments of the approximate p.m.f. to represent B_k. These intuitive ideas will be exploited and made rigorous in the following, through the idea of density projection with sequential Monte Carlo simulation.

3.2 Dimension Reduction

The idea of density projection provides an effective way for dimension reduction by orthogonally projecting an arbitrary density onto a parameterized family of densities. Hence, the density can be approximately represented by only a small number of parameters. The orthogonal projection is done through minimizing the Kullback-Leibler (KL) divergence of a density B from a parameterized family Ω, as follows:

$$
Proj_{\Omega}(B) \triangleq \arg\min_{f \in \Omega} D_{KL}(B\|f), \tag{11}
$$

where the KL divergence between B and f is defined as

$$
D_{KL}(B\|f) \triangleq \int f(x) \ln \frac{B(x)}{f(x)} dx.
$$

For an exponential family of densities, defined as

$$
\Omega = \{f(\cdot, \theta), \theta \in \Theta | f(x, \theta) = \exp[\theta^T c(x) - \varphi(\theta)]\},
$$

where $\varphi(\theta) = \ln \int \exp(\theta^T c(x)) dx$, $\theta = [\theta_1, \dots, \theta_m]$, and $c(x) = [c_1(x), \dots, c_m(x)]$. The minimization in (11) can be carried out analytically to yield the following equation related B with its projection $f(\cdot, \theta)$ (see (Zhou, Fu, and Marcus 2008) for a detailed derivation):

$$
E_{\theta}[c(X)] = E_B[c(X)], \tag{12}
$$

where E_B and E_{θ} denote the expectations with respect to B and $f(\cdot, \theta)$, respectively. Suppose B is a p.m.f., written as

$$
B(x) = \sum_{i=1}^{N} \alpha_i \delta(x - x_i),
$$

where $\{x_1, \dots, x_N\}$ are the support points, $\{\alpha_1, \dots, \alpha_N\}$ are the associated probabilities satisfying $\sum_{i=1}^{N} \alpha_i = 1$ and nonnegativity, and $\delta(\cdot)$ is a Kronecker delta function. Then (12) is reduced to

$$
E_{\theta}[c(X)] = \sum_{i=1}^{N} \alpha_i c(x_i).
$$

By projecting B_k to get its approximation $f(\cdot, \theta_k)$, the resultant stochastic control problem is finite dimensional, with states (θ, Y, V).

3.3 Simulated Dynamic Programming

The states (θ, Y, V) are continuous, so we first discretize (θ, Y, V): a set of M grid points on (θ, Y), denoted by Γ; and a set of M^V grid points on V, denoted by Γ^V. Then we estimate the transition probabilities between the grid points in Γ, and then apply dynamic programming on the discrete mesh $\Gamma \times \Gamma^V$. Algorithm 1 estimates the transition probability from a grid point $(\theta, y)^i \in \Gamma$ to a grid point $(\theta, y)^j \in \Gamma$, denoted as $P(i, j)$. It is based on the sequential Monte Carlo method (Cappé, Godsill, and Moulines 2007) and the density projection described above. In particular, the steps of "sampling state", "propagation", and "Bayes' updating" use Monte Carlo simulation to approximate the evolution of B_k (eqn. (9)), and yield a p.m.f. to approximate B_k; the step of "density projection" projects the p.m.f. to a specified exponential family; and the step of "sampling observation" simulates eqn. (7) of Y_k.

Algorithm 1   *Estimation of transition probabilities*

- *Input: grid point $(\theta, y)^i = (\theta^i, y^i)$, number of particles N; Output: transition probabilities $P(i, j), j = 1, \ldots, M$.*
- ***Sampling state****: Sample $\{x_1, \ldots, x_N\}$ from $f(\cdot, \theta^i)$.*
- ***Propagation****: Compute $\{\tilde{x}_1, \ldots, \tilde{x}_N\}$ according to $\tilde{x}_n = F(x_n, z_n)$ using $\{z_1, \ldots, z_N\}$ drawn from the distribution of Z.*
- ***Sampling observation****: Compute $\{y_1, \ldots, y_N\}$ according to $y_n = H(\tilde{x}_n, x_n, y^i, w_n)$ using $\{w_1, \ldots, w_N\}$ drawn from the distribution of W.*
- ***Bayes' updating****: For a fixed observation y_m, calculate the probability α_n^m associated with each x_n according to $\alpha_n^m \propto p(y_m | \tilde{x}_n, x_n, y^i), n = 1, \ldots, N$ and normalize. Do it for every $m = 1, \ldots, N$.*
- ***Density projection****: Project the p.m.f. $\sum_{n=1}^N \alpha_n^m \delta(x - x_n)$ onto the exponential family Θ, i.e., to compute θ_m according to $E_{\theta_m}(c) = \sum_{n=1}^N \alpha_n^m c(x_n)$. Do it for every $m = 1, \ldots, N$.*
- ***Estimation****: For $m = 1, \ldots, N$, find the nearest-neighbor of each (θ_m, y_m) in the set of grids, and count the frequency of each grid point. $P(i, j) = $ number of occurrences of the grid $(\theta^j, y^j)/M, j = 1, \ldots, M$.*

Algorithm 2 applies dynamic programming on the discrete mesh of $\Gamma \times \Gamma^V$ using the estimated transition probabilities.

Algorithm 2   *Dynamic Programming*

- *Input: $[P(i, j)]_{(i,j)}$; Output: $J_0(\theta, y, v), \forall (\theta, y, v) \in \Gamma \times \Gamma^V$, $\rho_0, \ldots, \rho_{T-1}$.*
- ***Initialization****: $J_T(\theta, y, v) = g(y, v), \forall (\theta, y, v) \in \Gamma \times \Gamma^V$. Set $k = T - 1$.*
- ***Iteration****: Compute $J_k(\theta, y, v)$ and find ρ_k according to*

$$J_k(\theta, y, v) = \min_{a \in A} \sum_{j=1}^M P(i, j) J_{k+1}\left((\theta, y)^j, Proj_{\Gamma^V} G(v, y^j, y, a)\right), \quad \forall (\theta, y, v) \in \Gamma \times \Gamma^V$$

    where i is the index of (θ, y) in Γ, and $Proj_{\Gamma^V}(\cdot)$ means to find the nearest-neighbor grid on Γ^V.
- ***Stopping****: If $k = 0$, stop; else, set $k := k - 1$ and go to the iteration step.*

4   EXAMPLES: HEDGING OF EUROPEAN OPTIONS

In this section, we apply the numerical methods stated in section 3 to hedging of European call and put options. Each will be hedged under the quadratic and shortfall cost criteria respectively, with the the risky asset price process (1) and volatility process (2). More specifically, the objective function for a put option under the quadratic risk criterion is

$$\min_{\pi \in \Pi} J(b, y, v) = E\left[\left((K - e^{Y_T})^+ - V_T\right)^2 | B_0 = b, Y_0 = y, V_0 = v\right],$$

for a call option under the quadratic risk criterion is

$$\min_{\pi \in \Pi} J(b, y, v) = E\left[\left((e^{Y_T} - K)^+ - V_T\right)^2 | B_0 = b, Y_0 = y, V_0 = v\right],$$

for a put option under the shortfall risk criterion is

$$\min_{\pi \in \Pi} J(b,y,v) = E\left[\left((K-e^{Y_T})^+ - V_T\right)^+ | B_0 = b, Y_0 = y, V_0 = v\right],$$

and for a call option under the shortfall risk criterion is

$$\min_{\pi \in \Pi} J(b,y,v) = E\left[\left((e^{Y_T} - K)^+ - V_T\right)^+ | B_0 = b, Y_0 = y, V_0 = v\right],$$

where K is the strike price at maturity time T, and $(\cdot)^+ = \max(\cdot,0)$.

4.1 Implementation

We project B_k to the Gaussian family, representing its density function via its mean μ_k and standard deviation σ_k, i.e. $\theta_k = (\mu_k, \sigma_k)$. Recall that Y_k is the log-price of the risky asset at time k. We choose an equally-spaced M grid points for (μ, σ, Y), and an equally-spaced M^V grid points for V. Using Algorithm 2, we obtain an estimated transition probability matrix $[P(i,j)]_{(i,j)}$. By plugging it into Algorithm 2, we solve an optimal hedging policy. In the case of quadratic put option, we can write the $J_T(\theta,y,v)$ and $G(v,y^j,y,a)$ in Algorithm 2 as:

$$
\begin{aligned}
J_T(\theta,y,v) &= ((K-e^y_T)_+ - v)^2 \\
G(v,y^j,y,a) &= ve^{r\varepsilon} + a(e^{y^j} - e^{y+r\varepsilon}).
\end{aligned}
$$

The above expressions for quadratic call option, shortfall call option, and shortfall put option can be written similarly.

4.2 Simulation Setup

We ran simulations on four hedging problems (e.g quadratic-call, quadratic-put, shortfall-call and shortfall-put) with the following parameters:

1. Strike price of the European Put/Call option: $K = 110$;
2. Riskless interest rate: $r = 0.05$;
3. Volatility: $x_0 = 0.15$, $\lambda = 0.1$, $\eta = 0.1$;
4. Number of periods: $T = 5$, time between periods: $\varepsilon = 0.2$;
5. Initial state: $\mu_0 = 0.1$, $\sigma_0 = 0.6$, $Y_0 = \ln 110 = 4.7$;
6. Number of particles in Algorithm 1: $N = 88$;
7. Discretization: $M = 396$ on $(\mu, \sigma) \in [0.1, 0.2] \times [0,1]$, $M^V = 100$ or 300 on $V \in [-10, 40]$.

The model parameters (K, r, x_0, λ, η, T ε, Y_0) are the same as in (Corsi, Pham, and Runggaldier 2007) for comparisons that we will discuss later.

4.3 Numerical Results

For each of the four risk-option combinations, we solve for a hedging policy, and plot its initial value function J_0 (i.e., the final risk of the hedging policy) and initial action a_0 (i.e., the number of shares held on the risky asset at time 0) against the initial wealth V_0. Figures 1-16 present these results.

4.4 Analysis of the Results: Quadratic Risk

We will first briefly describe the form of the optimal hedging strategy, a tuple of $(V_0, \{\rho_k\}_{\{0 \leq k \leq T-1\}})$, where V_0 is the initial wealth/captital, and $\{\rho_k\}_{\{0 \leq k \leq T-1\}}$ is the investment strategy. Next, we verify our simulation method by comparing the simulated relationship between initial wealth V_0 and initial action a_0 with their theoretical relationship. Finally, we will compare prices implied by our optimal hedging strategy against the arbitrage-free prices.
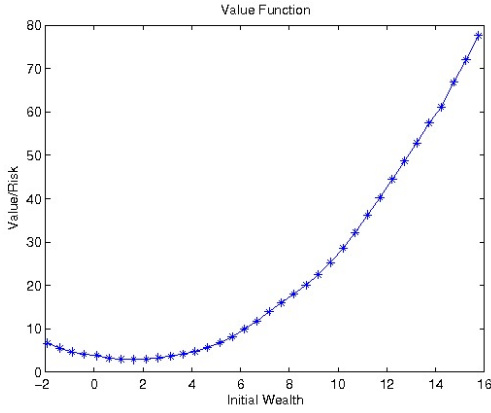
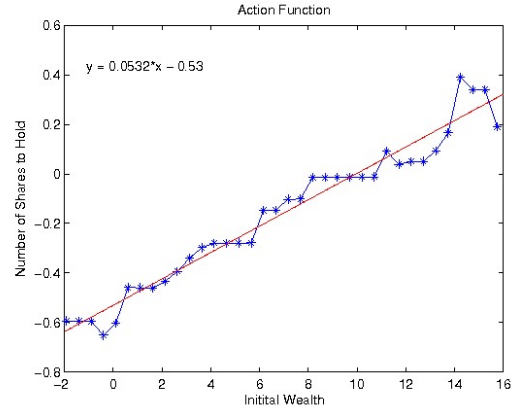Figure 1: Put Quadratic Value Function (M^V=100)

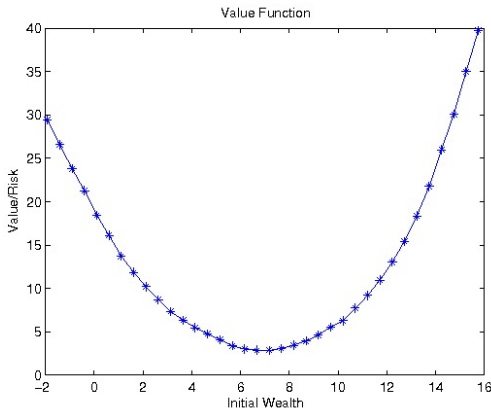Figure 2: Put Quadratic Initial Action (M^V=100)

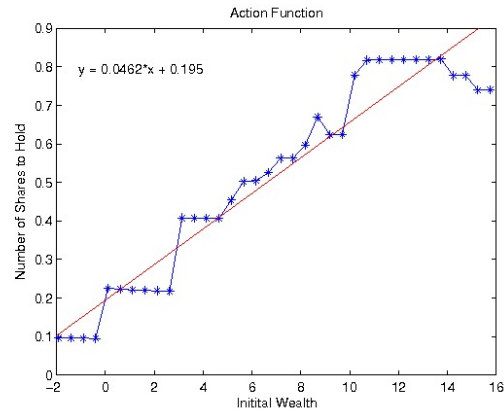Figure 3: Call Quadratic Value Function (M^V=100)

Figure 4: Call Quadratic Initial Action (M^V=100)

The theoretical analysis starts with recognizing the price process S_k is a semimartingale, and transforming it into a martingale by considering the discounted price process $\widetilde{S}_k = S_k e^{-rk\varepsilon}$. Applying the Kunita-Watanabe decomposition to the discounted option payoff $\tilde{h} = \sum_{k=0}^{T} h_k e^{-rk\varepsilon}$ gives

$$\tilde{h} = \tilde{h}_0 + \sum_{k=0}^{T-1} \rho_k^{\tilde{h}} \Delta \widetilde{S}_k + L_T^{\tilde{h}} \tag{13}$$

where $\tilde{h}_0 = E[\tilde{h}|\tilde{S}_0]$, $\rho_k^{\tilde{h}}$ is a process adapted to the filtration generated by the discounted price process $\widetilde{S}_k$, $\Delta \widetilde{S}_k = \widetilde{S}_{k+1} - \widetilde{S}_k$, $L_T^{\tilde{h}}$ is orthogonal to $\sum_{k=0}^{T-1} a_k^{\tilde{h}} \Delta \widetilde{S}_k$. If we choose to optimize over only the self-financing policies, the discounted wealth process is written as:

$$\widetilde{V}_T = \widetilde{V}_0 + \sum_{k=0}^{T-1} \rho_k \Delta \widetilde{S}_k \tag{14}$$

Comparing (13) with (14), it can be derived that the optimal strategy is:

$$\rho_k^* = \rho_k^{\tilde{h}}, \quad \widetilde{V}_0^* = \tilde{h}_0, \tag{15}$$
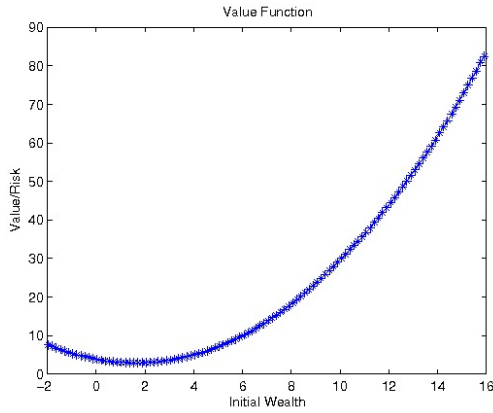
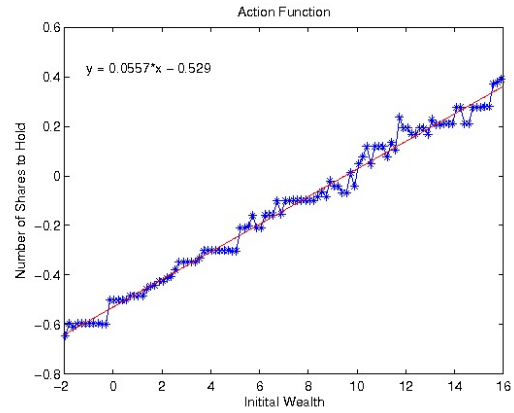Figure 5: Put Quadratic Value Function (M^V=300)

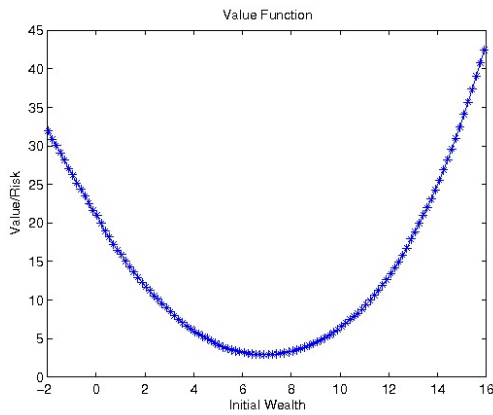Figure 6: Put Quadratic Action Function (M^V=300)

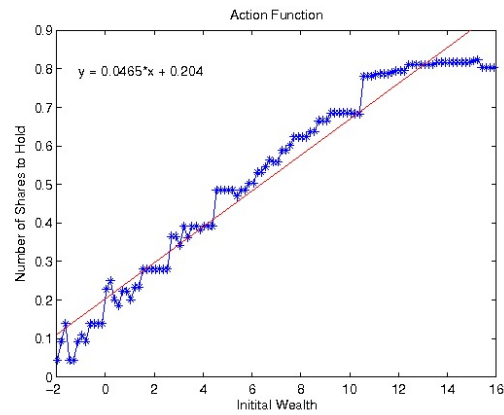Figure 7: Call Quadratic Value Function (M^V=300)

Figure 8: Call Quadratic Initial Action (M^V=300)

and the minimum risk achievable is $Var(L_T^{\tilde{h}})$. For a more detailed treatment of this topic, please refer to (Heath, Platen, and Schweizer 1998).

While (15) gives the optimal tuple $(V_0^*, \{\rho_k^*\}_{\{0 \leq k \leq T-1\}})$, we are also interested in the optimal number of shares to invest in the risky asset given an initial wealth and price. Looking at Figs. 2, 4, 6, and 8, we observe that a_0 appears to be an affine function of V_0, for a fixed Y_0. Indeed, it agrees with Theorem 5.1 from (Gourieroux, Laurent, and Pham 1998), which at time 0 can be written as

$$a_0^*(V_0) = \tau_1 V_0 + \tau_2.$$

where $a_0^*(V_0)$ is the optimal initial action for a given initial wealth V_0, and τ_1 and τ_2 are constants. While Theorem 5.1 from (Gourieroux, Laurent, and Pham 1998) gives the structure of the solution, the values of τ_1 and τ_2 are not directly computable. However, doing linear regressions on our numerical results (as shown by the lines in Figs. 2, 4, 6, and 8), we get estimates for τ_1 as shown in Table 1. Interpreting initial wealth as the amount the seller receives from writing the option, the seller allocates a constant proportion, τ_1, of it to the risky asset. We can think of this constant as the seller's relative risk exposure with respect to his initial wealth. A similar argument can be made for the European call case.

When an option is attainable in the sense that it can be written as (13) with $L_T^{\tilde{h}} = 0$, then $\widetilde{h}_0$ is its arbitrage-free price. If both the call and put options are attainable, then the following put-call parity (cf. pp.174, (Hull 2002)) should be satisfied:

$$C_0 = P_0 + S_0 + Ke^{-rT} \tag{16}$$

where C_0 and P_0 are arbitrage-free prices for European call and put options, respectively. However, if the option is not attainable, the smaller the difference from the arbitrage-free price, the better the option price. Our simulated option prices for the put and call options are the initial wealth corresponding the the minimum risk on Figs. 1, 3, 5, and 7. By plugging our simulated call price into (16), we compute the corresponding arbitrage-free put price, to compare with our simulated put price, as all shown in Table 2. We observe the differences between the arbitrage-free put prices and the simulated counterparts become smaller with an increasing M^V. We compare our results with other numerical results such as (Corsi, Pham, and Runggaldier 2007), and find them to be comparable.

In Table 3, we compare our minimum quadratic risk against the numerical results in (Corsi, Pham, and Runggaldier 2007), and find ours to be an improvement.

Table 1: Slope of Initial Action vs. Initial Wealth with Quadratic Cost

	$M^V = 100$	$M^V = 300$
Option	Simulated	Simulated
Call	0.0462	0.0465
Put	0.0532	0.0557

Table 2: Compare Arbitrage-Free Price with Simulated Price

Option	Loss Function	$M^V = 100$			$M^V = 300$		
		Arbitrage-Free$^\diamond$	Simulated	Diff	Arbitrage-Free$^\diamond$	Simulated	Diff
Put	Quadratic	-	1.6162	-	-	1.7057	-
Call	Quadratic	6.9810	7.1717	0.1907	7.0705	6.8896	0.1809
Put	Shortfall	-	3.1313	-	-	3.0435	-
Call	Shortfall	8.4961	9.1919	0.6958	8.4083	8.2274	0.1809

$\diamond$: obtained using (16)

Table 3: Minium Quadratic Value/Risk

Option	$M^V = 400$	$M^V = 300$
	(Corsi, Pham, and Runggaldier 2007)$^\diamond$	Ours
Put	~ 7.7	2.9248
Call	~ 8.0	2.9199

$\diamond$: the number is interpreted from graphs in (Corsi, Pham, and Runggaldier 2007)

4.5 Analysis of the Results: Shortfall Risk

Under the shortfall risk criterion, our simulation results (Figs. 9, 11, 13, and 15) show that a larger initial capital/wealth helps to reduce the risk of hedging until it reaches a threshold after which the risk stays at zero. Compared with the quadratic-risk induced option price, the shortfall-risk induced option price is always higher. These two observations are due to the fact that the shortfall risk function only penalizes losses, whereas quadratic hedging that penalizes both losses and gains.

4.6 Simulation Efficiency

We benchmark our results against the results obtained by (Corsi, Pham, and Runggaldier 2007). It has no analysis or data on the computation time; however, we can approximate the computational effort by the the number of grid points. Table 3 displays the comparison between their optimal values and ours. While their grid size is $1500 \times 400 = 600,000$, we only use $396 \times 300 = 118,800$ points. Our simulation obtains comparable results in terms of put-call parity and better results in terms of minimum value, with grid size less than one fifth of theirs.
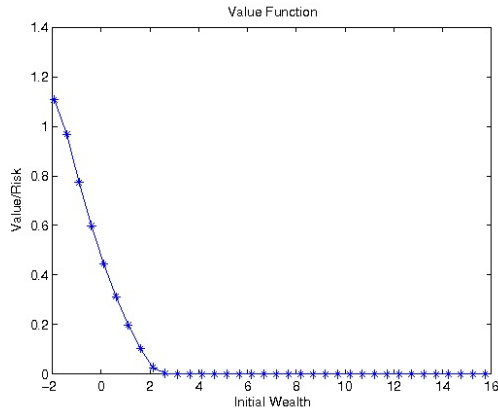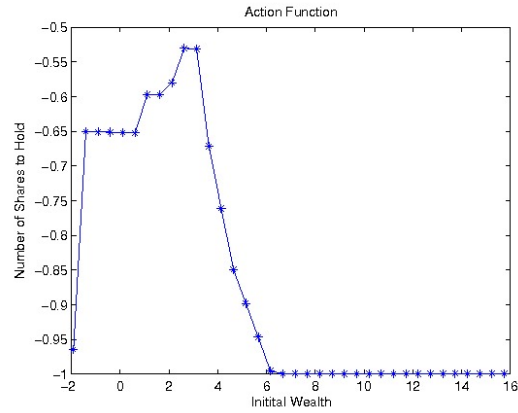
Figure 9: Put Shortfall Value Function (M^V=100)

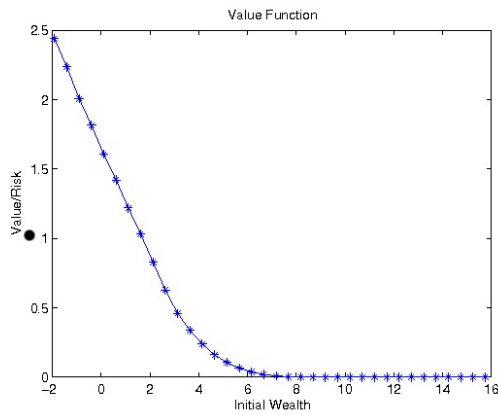Figure 10: Put Shortfall Initial Action (M^V=100)

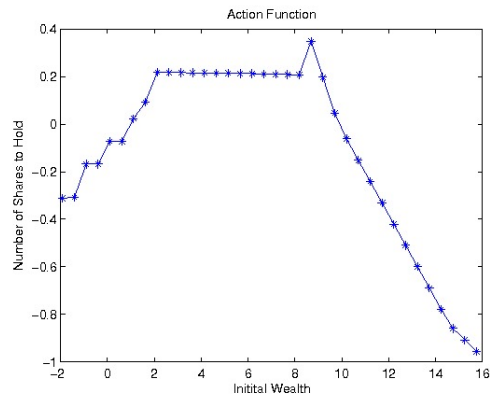Figure 11: Call Shortfall Value Function (M^V=100)

Figure 12: Call Shortfall Initial Action (M^V=100)

5 CONCLUSIONS

We have formulated several financial decision problems under stochastic volatility in the framework of a stochastic control problem under partial observation, for which we have developed an efficient numerical method. We have applied our method to the hedging of European options under quadratic risk and shortfall risk criteria, discussed the implications of the results, and compared our results with other analytical and numerical results. We believe the method presented in this paper has wide applications in financial engineering.

REFERENCES

Bates, D. S. 1996. Jumps and stochastic volatility: Exchange rate processes implicit in deutschemark option. *The Review of Financial Studies* 9 (1): 69–107.

Black, F., and M. Scholes. 1973. The pricing of options and corporate liabilities. *Journal of Political Economics* 81:637–659.

Cappé, O., S. J. Godsill, and E. Moulines. 2007. An overview of existing methods and recent advances in sequential Monte Carlo. *Proceedings of the IEEE* 95 (5): 899–924.
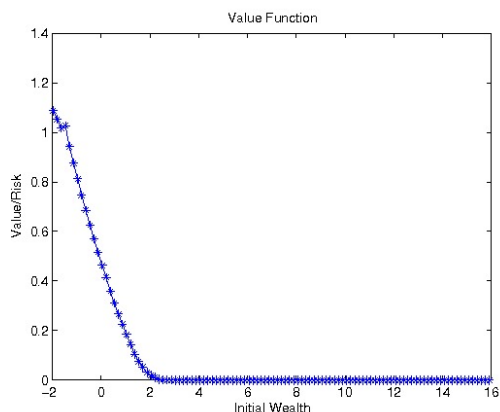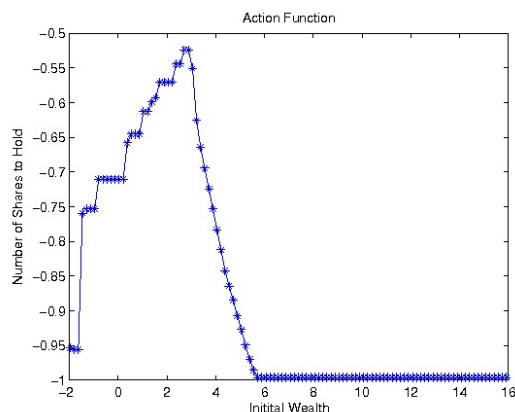
Figure 13: Put Shortfall Value Function (M^V=300)
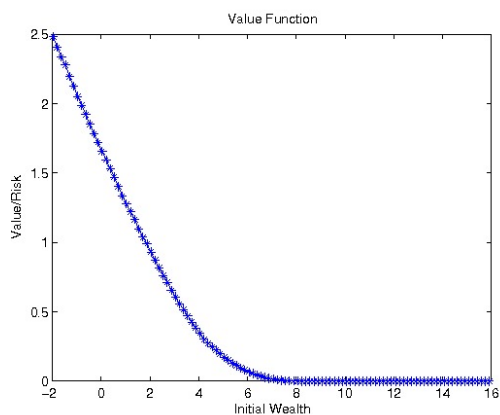

Figure 14: Put Shortfall Initial Action (M^V=300)
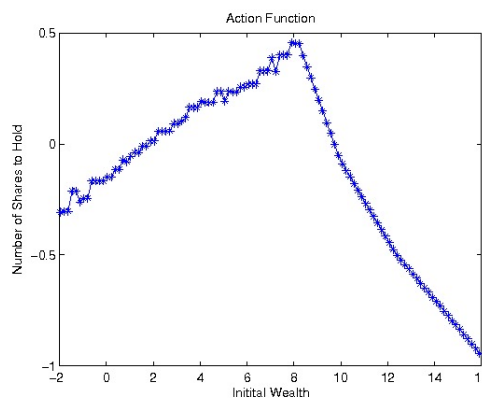

Figure 15: Call Shortfall Value Function (M^V=300)

Figure 16: Call Shortfall Initial Action (M^V=300)

Corsi, M., H. Pham, and W. J. Runggaldier. 2007. Numerical approximation by quantization of control problems in finance under partial observations. In *Mathematical modelling and numerical methods in finance. Special Volume of Handbook of Numerical Analysis*, ed. A. Bensoussan and Q. Zhang.

Gourieroux, C., J.-P. Laurent, and H. Pham. 1998. Mean-variance hedging and numéraire. *Mathematical Finance* 8 (3): 179 – 200.

Heath, D., E. Platen, and M. Schweizer. 1998. Comparison of some key approches to hedging in incomplete markets. Research Paper Series 1, Quantitative Finance Research Centre, University of Technology, Sydney.

Heston, S. L. 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6 (327-343).

Hull, J. C. 2002. *Options, futures, and other derivatives*. 5th ed. Prentice Hall.

Merton, R. C. 1973. Theory of rational option pricing. *Bell Journal of Economics and Management Science* 4 (1): 141–183.

Milstein, G. N. 1995. *Numerical integration of stochastic differential equations*. Mathematics and Its Applications. Kluwer Academic Publishers.

Nicolato, E., and E. Venardos. 2003. Option pricing in stochastic volatility models of the Ornstein-Uhlenbeck type. *Mathematical Finance* 13:445–466.

Shephard, N., and T. G. Andersen. 2008. Stochastic volatility: Origins and overview. Chapter prepared for the Handbook of Financial Time Series, Springer Verlag.

Zhou, E., M. C. Fu, and S. I. Marcus. 2008. Solving continuous-state POMDPs via density projection. *IEEE Transactions on Automatic Control*. Conditionally accepted.

AUTHOR BIOGRAPHIES

ENLU ZHOU is a Ph.D. candidate in the Department of Electrical and Computer Engineering, University of Maryland, College Park. She received the B.S. degree with highest honors in electrical engineering from Zhejiang University, China, in 2004. Starting from fall 2009, she will be an Assistant Professor in the Department of Industrial and Enterprise Systems Engineering, the University of Illinois at Urbana-Champaign. Her research interests lie in the areas of stochastic control and estimation, and simulation optimization, with applications towards financial engineering. Her email address is `<enluzhou@umd.edu>`.

KUN LIN is a PH.D candidate in Department of Electrical and Computer Engineering at University of Maryland, College Park. His research interest is financial engineering. His email address is `<klin@umd.edu>`.

MICHAEL C. FU is Ralph J. Tyser Professor of Management Science in the Robert H. Smith School of Business, with a joint appointment in the Institute for Systems Research and affiliate faculty appointment in the Department of Electrical and Computer Engineering, all at the University of Maryland, College Park. He received degrees in mathematics and EE/CS from MIT, and a Ph.D. in applied mathematics from Harvard University. His research interests include simulation optimization and applied probability, with applications in supply chain management and financial engineering. He served as Stochastic Models and Simulation Department Editor for *Management Science* 2006–2008 and Simulation Area Editor of *Operations Research* 2000–2005. He is a Fellow of INFORMS and IEEE. His email address is `<mfu@umd.edu>`.

STEVEN I. MARCUS is a Professor in the Department of Electrical and Computer Engineering and the Institute for Systems Research, University of Maryland. He received the B.A. degree in electrical engineering and mathematics from Rice University in 1971, and the S.M. and Ph.D. degrees in electrical engineering from MIT in 1972 and 1975, respectively. From 1975 to 1991, he was with the Department of Electrical and Computer Engineering, University of Texas, Austin, where he was the L.B. (Preach) Meaders Professor in Engineering. At the University of Maryland, he has served as Director of the Institute for Systems Research from 1991-1996, and as Chair of the Electrical and Computer Engineering Department from 2000-2005. Currently, his research is focused on stochastic control and estimation, with applications in manufacturing and telecommunication networks. He is a Fellow of IEEE and a member of the inaugural class of SIAM Fellows. His email address is `<marcus@umd.edu>`.