

A Monte-Carlo method for portfolio optimization under partially observed stochastic volatility

Rahul Desai, Tanmay Lele*, and Frederi Viens[‡]

Department of Mathematics and School of Elec. and Comp. Engr., Purdue University, West Lafayette, IN 47907, U.S.A

*School of Chemical Engineering, Purdue University, West Lafayette, IN 47907, U.S.A

[‡]Department of Statistics and Department of Mathematics, Purdue University, West Lafayette, IN 47907, U.S.A

Emails: desai@ecn.purdue.edu, tanmay@ecn.purdue.edu, viens@purdue.edu

Abstract: In this paper we implement an algorithm for the optimal selection of a portfolio of stock and risk-free asset under the stochastic volatility (SV) model with discrete observation and trading. The SV model extends the classical Black-Scholes model by allowing the noise intensity (volatility) to be random. The main assumption is that the portfolio manager has discrete access to the continuous-time stock prices; as a consequence the volatility is not observed directly. In this partial information situation, one cannot hope for an arbitrarily accurate estimate of the stochastic volatility. Using instead a new type of optimal stochastic filtering, and its associated particle method due to del Moral, Jacod, and Protter [9], our algorithm, of the “smart” Monte-Carlo-type, approximates the new Hamilton-Jacobi-Bellman equation that is required for solving the stochastic control problem that is defined by the portfolio optimization question.

Keywords: *Stochastic portfolio optimization, stochastic volatility, particle filtering, Monte-Carlo methods.*

1 INTRODUCTION

The celebrated Black-Scholes model for stock prices, which was introduced and used by Black, Scholes, and Merton in the early 1970’s ([3], [24]), is still immensely popular, as a vast majority of financial practitioners in today’s financial industry believe that most stock prices and indices are best modeled by continuous-time stochastic processes driven by Brownian-like noise, and the Black-Scholes model is the simplest and best understood model in this class. Merton’s name is most often associated with his so-called Mutual Fund theorems ([22], [23]), which cast the problem of optimal selection of a portfolio of stock and risk-free asset in the framework of stochastic optimal control of diffusion processes. Merton originally used the Black-Scholes model as the underlying stock price. He showed that if the question is to maximize an expected future utility of the portfolio, the answer is obtained by solving a Hamilton-Jacobi-Bellman (HJB) partial differential equation (PDE).

Portfolio optimization using this HJB approach has been a popular one ever since. It has to be said that it appealed greatly to the world’s probability community because of the highly non-trivial use of stochastic analysis that is required

to prove the HJB equations, and because of the excitement created by finding such a beautiful application of stochastic theories to such a useful and popular area as finance. Further, the stochastic theory of HJB equations appears as a strong connection between a young and vigorous theory of stochastic analysis and the industrial engineering point of view of dynamic programming problems. A casual visitor in the land of the stochastic theory of finance will notice that there is an even more popular topic in finance than portfolio optimization that has benefitted from being cast into a theory based on stochastic analysis: the question of option pricing, introduced by Black and Scholes themselves. In reality, the financial industry today is in greater need of mathematical results on portfolio selection than on option pricing. One reason for the imbalanced distribution between the two topics in stochastic finance may be that option pricing is mathematically easier than portfolio optimization; indeed the HJB equation is a nonlinear parabolic PDE, while the fundamental equation of option pricing is the Black-Scholes equation, which is a linear parabolic PDE. We cite as our main practical motivation for the present work our desire to correct the imbalance. From a mathematical standpoint, we also have a preference for portfolio optimization precisely because, to paraphrase Leon Tolstoy, linear PDEs form in some sense a unique big happy family, while each nonlinear PDE is an individual unhappy family, and why would one turn one’s attention to a family that needs no help to achieve happiness...?

Among the hypotheses in Merton’s original work, two of the most important ones are still widely used in stochastic portfolio optimization: the use of a constant, or at the least, non-random volatility, and the assumption that observations and trading occur in continuous time. Both assumptions are far from being satisfied in most real financial markets. Today, it is widely recognized that constant volatility is the single most important deficiency of the Black-Scholes model. We choose to use the Stochastic Volatility (SV) model, one of several extensions/corrections of the Black-Scholes model which have recently appeared. Relaxing the second assumption is no less crucial. Indeed for many investors, not only does the presence of transaction costs forbid continuous trading, but information does not come in continuously, and is thus incomplete. In our view, this problem is just as concerning as the first one, but has received much less press. We tackle it by requiring that portfolio

selection and rebalancing occur only at a set of discrete observation times, based solely on the available observations. Our main philosophy is that, going against a natural mathematical tendency, one should not then revert to discrete time models; the underlying stocks should still be considered as continuous-time Brownian-motion-driven stochastic processes. The simple fact of having continuous-time models with discrete time information captures, in our view, the essence of the incompleteness of information that all portfolio managers are faced with.

2 THE MATHEMATICAL PROBLEM

2.1 Portfolio optimization with stochastic volatility

We will work exclusively in a model in which the risk-free asset B is assumed to have a constant interest rate r : $B_t = e^{rt}$ for all $t \geq 0$. Under the simplest SV model, the evolution of the price of a single risky asset X is given by the following stochastic differential equation:

$$dX_t = X_t \mu dt + X_t \sigma(Y_t) dW_t. \quad (1)$$

Here $t \in \mathbf{R}_+$, W is a Brownian motion, the mean rate of return μ is assumed to be constant for simplicity, and the stochastic volatility $\sigma(Y_t)$ is a deterministic function σ of a stochastic process Y that satisfies a diffusion equation driven by another Brownian motion Z such that $\text{corr}(W, Z) = \rho$ with $0 \leq |\rho| < 1$, i.e.

$$dY_t = \alpha(Y_t) dt + \beta(Y_t) dZ_t. \quad (2)$$

Typically (see [15], and their statistical study of the Standard & Poor 500 index), practitioners take $\sigma = \exp$ and $Y = a$ a fast-mean-reverting process such as the Ornstein-Uhlenbeck process with large α :

$$dY_t = \alpha(m - Y_t) dt + \sqrt{\alpha} dZ_t. \quad (3)$$

For $i = 0, 1, \dots, N$, let \mathcal{F}_i^X be the information contained in the discrete sequence of *observed* asset prices X_0, X_1, \dots, X_i . Note that for notational simplicity we assume here, and in the remainder of the paper unless otherwise specified, that the observation times are the integers. Note that \mathcal{F}_i^X is not the commonly used “filtration of X ”, which contains much more information. For $\bar{x} = (x_0, \dots, x_N)$ a fixed sequence of positive numbers, denote by $\mathcal{F}_i^{\bar{x}}$, the scenario (event) $\{X_0 = x_0, \dots, X_i = x_i\}$.

We consider self-financing portfolios $a = (a_i)_{i=0}^N$ with wealth $\mathcal{W}_s = \mathcal{W}_s^{a_i, b_i} = a_i X_s + b_i B_s$ for $s \in [i, i+1]$. This ensures that the strategy be constant in each interval $[i, i+1)$. Using wealth as a state variable is a standard choice, and thus we can reduce the number of control variables, by letting $b_i = (w_i - x_i a_i) e^{-ri}$. Assume $\mathcal{W}_0 = w_0$ is given. The basic portfolio maximization problem with horizon $N+1$ is to find a portfolio a^* that attains the supremum

$$V(0, x_0, w_0) = \sup_a \mathbf{E} \left[U \left(\mathcal{W}_{N+1}^{a, b} \right) \mid X_0 = x_0, \mathcal{W}_0 = w_0 \right] \quad (4)$$

for all $i = 0, \dots, N$, where the supremum is over all (a, b) that are non-anticipating, i.e. such that (a_i, b_i) are functions depending only on $w_0, x_0, x_1, \dots, x_i$. Other restrictions on (a, b) may be placed, such as requiring that \mathcal{W} be bounded below (no ruin), or that the possible values for (a_i, b_i) be bounded and/or discrete. Here U is some utility function. A typically choice is $U(w) = w^p/p$ for some $p \in (0, 1)$ (the so-called Hyperbolic Absolute Risk Averse (HARA) case).

We will see below that one of the most important concepts needed for solving this problem is how to find the best estimate for the volatility of X given only the discrete observations of X . This is called the stochastic volatility filtering problem, and can be written as the following conditional probability distribution:

$$p_i(dy) := \mathbf{P} [Y_i \in dy \mid \mathcal{F}_i^X]. \quad (5)$$

Before we explain what mathematical techniques are required for solving the portfolio optimization problem and implementing a numerical method for it using the stochastic volatility filter, we review the current literature related to the problem.

2.2 Position of the problem in the current literature

Nonlinear stochastic filtering has a key role in partially observed stochastic control. We cite [13], [10], [11], [2] and recently [29]. Recent advances on finance-related aspect of this topic are still restricted to non-stochastic volatility: [30], [26], [21], [25], in which the linear-quadratic and integral-quadratic models are considered, but only using standard linear filtering.

There is no literature on filtering of stochastic volatility in continuous time. The reason for this gap is that probabilists’ work on filtering of continuous-time processes have concentrated on continuous-time observation; in that situation, the volatility $\sigma^2(Y)$ is, in principle, obtainable exactly from the information in X (measurable w.r.t. the filtration of X), as the so-called quadratic variation $\langle X \rangle$ of X . However evaluating $\langle X \rangle_t$, a problem of estimation, rather than filtering, is treacherous in practice. The financial industry contains notorious stories of investment firms whose bankruptcy can be traced to a poorly estimated volatility.

The popular ARCH/GARCH models are designed to estimate stochastic volatility in a stable way (see [20], [4], [14]). Dan Nelson ([27]; see also [4]) showed that ARCH/GARCH models are in fact an approximate filter, since they converge to the full information SV as the observation time step $\delta \rightarrow 0$, leading many to believe the task is now to “bridge the gap to continuous time” (see [4]; [17], [18], [19]). But the quality of the ARCH/GARCH “filter” is only guaranteed for high observation frequency δ^{-1} . We adopt a different angle, seeking not an estimation but the *optimal* filter when δ is *fixed*. The very recent work [9] gives a numerical method for discrete-observation filtering of diffusions under stochastic volatility, which opens the way to numerically solving the stochastic volatility control problem with discrete information, as we detail below.

Recent work on volatility filtering that departs from the ARCH/GARCH framework, but differs from the optimal filtering approach includes: [5] (a new *projection* filter); [6] (reduction to linear (Kalman) filtering in a special case); [28] (filtering w.r.t exogenous observation noise, not stochastic volatility). [12], [16] study models with stochastic volatility, but in the first case, the issue is to hedge an option, while in the other, the SV is filtered out of unobserved noise in random observation times, reducing the problem to standard filtering.

3 FILTERING WITH STOCHASTIC VOLATILITY

The probability measure $p_i(dy)$ is random since it depends on the values X_0, X_1, \dots, X_i . However at time i , the values $X_0 = x_0, \dots, X_i = x_i$ are known to us (they constitute the *observation*, while Y is the *signal*) and therefore \mathcal{F}_i^X can be replaced by $\mathcal{F}_i^{\bar{x}}$, and $p_i(dy)$ can be considered as being non-random, depending only on the parameters $\bar{x}_i := (x_0, x_1, \dots, x_i)$. To make the deterministic dependence on \bar{x} appear clearly, we will systematically denote p by $p_i^{\bar{x}}(dy)$.

Using the concept of Bayes' formula, it is not difficult to establish an explicit recursion relation for the filter in (5). It can be found in [31]. However, in view of the complexity of this iterative formula, there is currently no hope to evaluate $p_i^{\bar{x}}$ by any other method than the "smart"-Monte-Carlo algorithm recently established in [9], even for the simplest of examples.

The algorithm of [9] (detailed in Section 5 therein), itself a bootstrapping extension of the genetic algorithm of [8], yields a good approximation (order $n^{-1/3}$) of $p_i^{\bar{x}}$ as the empirical distribution of a family of n interacting particles $(Y_i^k)_{k=1}^n$

$$\hat{p}_i^{\bar{x}}(dy) = \frac{1}{n} \sum_{j=1}^n \delta_{Y_i^j}(dy). \quad (6)$$

The particles evolve according to the iteration of a two-step (selection/mutation) process. In the mutation process, they evolve independently according to the Euler approximations of the diffusion Y defined by the original dynamics of (X, Y) , with time step $m = n^{1/3}$. To take the selection step, the particles, which are created jointly with observations simulated one time-step into the future, use these simulated observations to evaluate each particle's fitness relative to a bootstrapping version of the maximum likelihood estimator of how they should be distributed. The particles then rearrange their positions according to their fitnesses, which is the selection step. This algorithm is explicitly given on page 16 of [9].

4 MATHEMATICAL TOOLS

All the results quoted in this section are established in [31]. For any scenario $\bar{x} := (x_0, x_1, \dots, x_N)$, and any $i \leq N$, we define $\bar{x}_i = (x_0, \dots, x_i)$. Using a standard idea, we embed our portfolio optimization problem into a dynamic one as follows: for all w, x, \bar{x} , for all $i = 1, 2, \dots, N$, for

all $s \in [i, i+1]$, find

$$\begin{aligned} V(s, x, w) &= V(s, x, w; \bar{x}_i) \\ &= \sup_{a \in A_0} \mathbf{E} [U(\mathcal{W}_{N+1}^a) | X_s = x, \mathcal{W}_s^a = w, \mathcal{F}_i^{\bar{x}}]. \end{aligned} \quad (7)$$

Recall that the control set A_0 is the set of all sequences $(a_j)_{j=0}^N$ of the form $a_j = a_j(w_0, \bar{X}_j)$. It should be clear from the self-financing condition

$$\mathcal{W}_t = a_i X_t + (\mathcal{W}_i - X_i a_i) e^{r(t-i)} \quad (8)$$

that this is just as general as allowing a_j to be of the form $a_j = a_j(\bar{X}_j, \bar{\mathcal{W}}_j)$.

Theorem 1 For $s \in [i, i+1]$, V in (7) satisfies the Hamilton-Jacobi-Bellman (HJB) equation

$$\frac{\partial V}{\partial s} + \sup_{a \in A_0} [(\mathcal{A}^a V)(s, x, w)] = 0. \quad (9)$$

where for any fixed $a \in A_0$, \mathcal{A}^a is the infinitesimal generator of (\bar{X}, \mathcal{W}^a) in $[i, i+1]$ with σ replaced by $\sqrt{Z_s^{i,X}(x; \bar{x}_i)}$, where

$$Z_s^{i,X}(x; \bar{x}_i) := \mathbf{E} [\sigma^2(Y_s) | X_s = x, \mathcal{F}_i^{\bar{x}}].$$

Moreover, there exists an optimal control in A_0 , i.e. the sup in (9) is attained. This theorem also holds if A_0 is replaced by any proper subset of A_0 .

Here $Z_s^{i,X}$ appears naturally as the filtered expected value of the squared Stochastic Volatility $\sigma^2(Y_s)$. In this sense, the dynamics of V follow a so-called *separation principle* (see [32], [2]), i.e. the fact that the unobserved SV parameter $\sigma^2(Y_s)$ can be replaced by its filtered value at time s , given the current information, and all past discrete information. Note that in the calculation of this filtered value, although the current stock price may be invoked, one is not allowed to use any continuous flow of information for any non-zero length of time. This makes it impossible to estimate the SV using formulas such as $\sigma^2(Y_t) = \langle X \rangle_t$, notwithstanding the fact that the current stock price may be used. However, as we are about to see in Proposition 2, since the controls in A_0 can only change at times $i = 0, 1, \dots, N$, the optimal strategy only makes use of the information $\mathcal{F}_i^{\bar{x}}$ at those times, a fact which is arguably intuitively obvious.

We now present an iterative formula, which reduces the complexity of the HJB equation (9), and is the key to our Monte-Carlo method.

Proposition 2 Let I be any subset of \mathbf{R} , and replace A_0 by its restriction to I -valued sequences., For any $i = 0, 1, \dots, N$, for any $f = f(x, w; \bar{x}_i)$, define

$$\begin{aligned} \Phi^i(f)(x_i, w_i) &= \Phi^i(f)_{\bar{x}_{i-1}}(x_i, w_i) \\ &:= \sup_{a_i \in I} E [f(X_{i+1}, \mathcal{W}_{i+1}^{a_i}; \bar{x}_i) | \mathcal{F}_i^{\bar{x}}, \mathcal{W}_i^{a_i} = w_i]. \end{aligned} \quad (10)$$

Then we have

$$\begin{aligned} V(i, x_i, w_i)_{\bar{x}_{i-1}} &:= V(i, x_i, w_i; \bar{x}_i) \\ &= \Phi^i(\Phi^{i+1}(\dots \Phi^N(U)))_{\bar{x}_{i-1}}(x_i, w_i) \end{aligned} \quad (11)$$

and the control $a^* = (a_0^*, \dots, a_N^*)$ which is obtained by calculating an optimal a_i^* for the sup in formula (10) is such that $a^* \in A_0$, and attains the sup in (7), i.e. is an optimal control.

5 ALGORITHM

Theorem 1 shows that in principle, by solving a partial differential equation, one can find the optimal strategy for problem (7). However it should be clear that there is no hope of solving this equation explicitly, or perhaps even approximate it by an analytic method, because the SV filter is analytically intractable, and the PDE is inextricably linked to the SV filter via its coefficients.

Proposition 2 proves that the full optimization problem (7) can be replaced by an iteration of one-step optimization problems in which, at each step, the set of admissible strategies is one-dimensional, since between two observation times, we are not allowed to change our strategy. One can abandon any attempt of solving a PDE for the one-step optimization, by simply calculating the expectation that is to be maximized for a certain range of strategies, and investigating which strategy yielded the maximum expectation. Repeating this work for all scenarios that one wishes to consider yields a function indexed by the set of all scenarios. This function is used as the starting point of the next step. This is one of three basic ideas which leads to our algorithm.

The second idea is as follows. In order to calculate the Monte-Carlo expectations to be maximized, one needs to understand the dynamics of (X, Y) on the interval $[i, i + 1]$ under the conditioning $\{\mathcal{F}_i^{\bar{x}}, \mathcal{W}_i^{a_i} = w_i\}$. It should be clear, and is in fact true, that these dynamics are equivalent to the regular dynamics of (X, Y) given by (1), (2), started from the observed value x_i for X and the measure $p_i^{\bar{x}}(dy)$ for Y . Further, in order to simulate these dynamics, one can simply use an approximation of $p_i^{\bar{x}}(dy)$; naturally, one is lead to use the del Moral-Jacod-Protter approximation alluded to in section 3.

The third and last idea is that since one is using a Monte-Carlo-type algorithm for generating an approximation for the SV filter, and since there seems to be no other way to calculate the expectations to be maximized than by more Monte-Carlo methods, one should seriously consider arranging things so that the two Monte-Carlo methods work hand in hand. This is easily done if one is willing to use the same number of Monte-Carlo generations in the expectation calculation as there are particles in the filter approximation. Then in order to simulate the pair (X, Y) on $[i, i + 1]$ starting from the measure $(\delta_{x_i} \otimes p_i^{\bar{x}})(dx \otimes dy)$, it should be clear that it is sufficient to do one of the n Monte-Carlo runs started from each of the n points $(x_i, Y_i^k)_{k=1}^n$ where $(Y_i^k)_{k=1}^n$ is the set of the n particles in the SV filter.

These ideas lead to the general algorithm that follows. Let S_N and T_N denote the set of all strategies and the set of all possible wealths that one is willing or able to consider. The subscript N denotes the fact that these sets should depend on the number of observations N . More specifically, the notation S_i and T_i indicates that at the end of i time steps, each strategy in S_i is an i -dimensional vector, while the range of wealth values in T_i presumably grows with i in the same fashion that a binomial tree grows. The notation $A_{0,i}$ for the set of all possible strategies one is willing to consider at time i , has a similar interpretation, although in this case, there is

no reason to think that one should allow this set of scalars to depend on i . We will comment more on these point further down. Also, for simplicity, since we are using inter-observation times equal to 1, we assume that each time the Euler method is implemented, since it is implemented in an interval of length 1, if one wishes to use m Euler steps, then the Euler time step should be $1/m$. Lastly, since for a computer implementation, the sets S_i, T_i, A_i have to be discrete and finite, we will need some truncation and discretization procedures to ensure that our Euler methods, which in principle can yield arbitrary values, are then projected back to the closest values in S_i, T_i, A_i . We denote the truncation and discretization of an arbitrary value χ by $\{\{\chi\}\}$.

1. **Initialization.** Let $X_0^k = x_0, \mathcal{W}_0^k = w_0$ and $Y_0^k = y_0$ for all $k = 1, \dots, n$. For all $(\bar{x}, w) \in S_N \times T_N$ let $\hat{V}(N + 1, \bar{x}, w) = U(w)$.
2. **Calculation of the filter.** For each $\bar{x} \in S_N$, use the del Moral-Jacod-Protter method with Euler time step $1/m$ to calculate the particles $Y_i^k = Y_i^k(\bar{x})$ for all $i \leq N, k \leq n$.

Repeat step 3 for $i = N$ down to 0:

3. **Calculation of the approximate control solution \hat{V} and its corresponding approximate optimal strategy.** We assume that we know $\hat{V}(i + 1, x_{i+1}, w_{i+1}, \bar{x}_i)$ for all $\bar{x}_{i+1} \in S_{i+1}$ and $w_{i+1} \in T_{i+1}$, as well as the corresponding optimal strategy $a_{i+1}^*(\bar{x}_{i+1}, w_{i+1})$. From step 2, we also know $Y_i^k(\bar{x})$ for all $\bar{x}_i \in S_i, k \leq n$. For each $\bar{x}_i \in S_i, w_i \in T_i, a_i \in A_{0,i}$:

(a) independently for each $k \leq n$

- i. **simulate an observation one step into the future** $\hat{X}_{i+1}^m(k)$ using the Euler scheme with time step $1/m$ for the pair (X, Y) starting from $(x_i, Y_i^k(\bar{x}))$, over $[i, i + 1]$ [Note that it is necessary to simulate \hat{Y}_{i+1}^k also, but this value can be discarded],

- ii. calculate the corresponding **simulated wealth**, one step into the future

$$\begin{aligned} & \widehat{\mathcal{W}}_{i+1}^m(a_i) \\ &= a_i \hat{X}_{i+1}^m(k) + a_i \left(\hat{X}_{i+1}^m(k) - x_i \right) e^r + w_i e^r, \end{aligned}$$

- (b) calculate the **Monte-Carlo average** of these simulations

$$\begin{aligned} & \hat{F}(a_i, \bar{x}_i, w_i) \\ &= \frac{1}{n} \sum_{k=1}^n \hat{V} \left(i + 1, \left[\left\{ \left(\widehat{\mathcal{W}}_{i+1}^m(a_i) \right) \right\} \right], \bar{x}_i \right), \end{aligned}$$

- (c) the **approximate i -th step optimal strategy and maximum expected yield** are

$$\begin{aligned} \hat{V}(i, x_i, w_i, x_{i-1}) &= \max_{a_i \in A_{0,i}} \hat{F}(a_i, \bar{x}_i, w_i), \\ a_i^*(\bar{x}_i, w_i) &= \arg \max_{a_i \in A_{0,i}} \hat{F}(a_i, \bar{x}_i, w_i). \end{aligned}$$

6 IMPLEMENTATION

According to the del Moral-Jacod-Protter algorithm, for a one-dimensional situation such as our, it is preferable to have the relation $n = m^3$. The algorithm detailed above is proved to converge in [31], as long as the price discretization step, which we call $1/m'$ is related to m as $m' \geq m^{1+c(\sigma)}$ where $c(\sigma)$ is a constant that depends on the non-degeneracy and boundedness of σ . Similarly, the truncation condition needed for the convergence of the algorithm can be written in the form

$$|\log(x_{i+1}/x_i)| < K_m = K^* + \|\sigma\| \sqrt{\log m}$$

The convergence is proved [31] to occur roughly at the speed $1/m$. While the constants involved above are typically not large, the conditions for convergence still imply that even for very small values of m , the algorithm we propose is still extremely memory intensive. Indeed, it requires storing values of \hat{V} and a^* for all times, for all possible wealths, and most greedily, for all possible scenarios.

On the other hand, interacting particle filters in one dimension are known to require only a very small number of particles in order to achieve an excellent level of accuracy and stability; typically 10 particles are sufficient. Our simulations for the del Moral-Jacod-Protter algorithm appear to confirm this fact. This means that if one even only wishes to use $m = 3$ Euler steps per observation interval, there should be 27 particles, which is amply sufficient for filtering.

Although this level for the Euler approximation may seem uncomfortably small for purists of numerical methods for stochastic differential equation, we have chosen to use $m = 3$ in order to show that our algorithm can be implemented on a modest platform. Increasing m by only an order of magnitude would require supercomputer storage capacities. In the same vein, we decided to restrict the strategies to only 5 integer values for the number of stock holdings, from -2 to $+2$.

Because of the very low value for m , we have also decided to restrict the set S_N of scenarios to a set given by a the values of a binomial model. In order to be consistent with the parameters of the (X, Y) -dynamics, we chose to use the so-called Cox-Ross-Rubenstein correspondence, given as follows: at each time step, a stock price has the option of going up or down by a factor of u or d respectively, with probabilities both equal to $1/2$. The values of u and d are given by

$$u = \exp\left(\left(\mu - \sigma^2/2\right) \Delta t + \sigma\sqrt{\Delta t}\right), \quad (12)$$

$$d = \exp\left(\left(\mu - \sigma^2/2\right) \Delta t - \sigma\sqrt{\Delta t}\right). \quad (13)$$

Here Δt is the Euler time step, which is $1/m = 1/3$, while μ is the value determines in the model for X . For σ , some interpretive creativity is required. Our whole work is aimed at being able to use random values for the volatility, but now we much choose a representative single value σ . In order to stay within the confines of the conditions needed for convergence of the algorithm, it is best to choose the function $\sigma(y)$

to be bounded. People working in SV models usually take $\sigma(y) = \exp(y)$. Instead, we will use $\sigma(y) = H(\exp(y))$ where $H(\varepsilon)$ is the continuous piecewise linear function that is equal to the identity function for $\varepsilon \in [0.05; 0.2]$, and is constant outside this interval. This σ takes on the values between 0.05 and 0.2 fairly uniformly. By using a mean-reverting model (2) for Y with mean in that range, we can ensure that this Ornstein-Uhlenbeck process does indeed allow $\sigma(Y(t))$ to scan the values in $[0.05; 0.2]$ fairly uniformly. It is then consistent to take the scalar value of σ in the correspondence (12), (13) to be 0.125. We take the mean rate of return to be $\mu = 0.1$. To give the portfolio a good shot at making more money than the bank, we choose the risk-free rate to be $r = 0.02$.

We run this simulation for 4 time intervals, with initial stock price equal to 100 and initial wealth equal to 1000, using the so-called Hyperbolic Absolute Risk Averse (HARA) $U(w) = w^p/p$ for the utility function, with $p = 0.6$. This type of risk aversion means that one is not much more satisfied by becoming immensely rich than by having a decent return on one's portfolio, while one is extremely dissatisfied by very low returns. The results are displayed in Figures 1, 2, and 3. Each of these three figures are relative to a given realization of the stock price (a fixed scenario); the solid line denotes the evolution of the wealth of the optimal portfolio. The family of dashed lines below the solid line denote the evolutions of 100 randomly chosen other portfolio positions from the strategy set for the same scenario. The pictures show that the optimal portfolio exceeds the other strategies, and only a few strategies come close, which means that one would have to be very lucky to do as well as the optimal portfolio.

It is worth noting that the algorithm we propose actually outputs, in backwards induction in time, optimal strategy values and optimized expected terminal utility values as functions of the following independent parameters: wealth w , scenario \bar{x} . In order to obtain the figures below, one needs to work forward in time in the algorithm's output: for a fixed scenario \bar{x} , one uses self-financing to calculate the wealth at time $i+1$ that corresponds to the previously determined values of the optimal strategy up to time i , which in turn can be used to calculate the optimal strategy to hold in the interval $[i+1, i+2]$, allowing the forward calculation to continue. This is exactly what the practitioner will want to do with the algorithm's output. However, in the end, since this operation can be performed offline, the practitioner only needs to be handed a table containing the evolution of the optimal strategy values for each scenario that one wishes to consider. The associated evolutions of portfolio wealth and expected future utility can be given for each scenario as well, for the purpose of comparing with the actual outcome, or for commercial purposes, to show what can be expected from the market.

These last remarks give us hope that even if the offline calculations present a tremendous storage challenge if a high degree of accuracy wants to be guaranteed, storage can be reduced to a much smaller file by the above forward procedure, and that file can be reduced further still by throw-




Figure 1: An optimal strategy for a typical scenario compared to other randomly chosen strategies for the same scenario.




Figure 3: An optimal strategy for a third typical scenario compared to other randomly chosen strategies for the same scenario.




Figure 2: An optimal strategy for another typical scenario compared to other randomly chosen strategies for the same scenario.




Figure 4: A 3D picture of the particle filter's evolution for a typical path (dark solid line) of the stochastic volatility.

ing away a large proportion of the scenarios, keeping only a small proportion of the most representative ones, since this last, wasteful operation does not effect the precision with which the calculations were performed originally.

Figure 4 and 5 are graphical illustrations of the particle filter for given typical scenarios. At each instant, the particle filter is comprised of a certain number of particle positions. These are arranged in a histogram, which is linearly interpolated and, when time varies, yields a surface in 3D. The narrower the spread at each time, the more efficient the filter is, and if the mean of the filter actually follows the signal, one can conclude that the filter is working well. This is what is observed in the figures.

REFERENCES

[1] Bally, V.; Talay, D. The law of the Euler scheme for stochastic differential equations: II. Aproximation of the density. *Monte Carlo Methods and Applications*, **2** (1996), 93-128.

[2] Bensoussan, A. *Stochastic control of partially observable systems*. Cambridge University Press, Cambridge, 1992.

[3] Black, F.; Scholes, M. The pricing of options and cor-

porate liabilities. *J. Political Economy* **81** (1973), 635-654.

[4] Bollerslev, T.; Rossi, P.E. Introduction. In *Modelling stock market volatility – Bridging the gap to continuous time*. Peter E. Rossi, Ed. Academic Press, 1996.

[5] Brigo, D.; Hanzon, B. On some filtering problems arising in mathematical finance. The interplay between insurance, finance and control (Aarhus, 1997). *Insurance Math. Econom.* **22** (1998), no. 1, 53–64.

[6] Chib, S.; Nardari, F.; Shephard, N. Markov Chain Monte Carlo methods for Generalized Stochastic Volatility Models. Preprint (Nov. 1998).

[7] Crisan, D.; Gaines, J.; Lyons, T. Convergence of a branching particle method to the solution of the Zakai equation. *SIAM J. Appl. Math.* **58** (1998), no. 5, 1568–1590.

[8] Del Moral, P.; Guionnet, A. On the stability of interacting processes with applications to filtering and genetic algorithms. *Ann. Inst. H. Poincaré Probab. Statist.* **37** (2001), no. 2, 155–194.

[9] Del Moral, P.; Jacod, J.; Protter, Ph. The Monte-Carlo




Figure 5: Another 3D picture of the particle filter's evolution for another typical path (dark solid line) of the stochastic volatility.

method for filtering with discrete-time observations. (1999). To appear in *Annales de l'I.H.P.*

- [10] El Karoui, N.; Du, H. N.; Jeanblanc-Picqué, M. Existence of an optimal Markovian filter for the control under partial observations. *SIAM J. Control Optim.* **26** (1988), no. 5, 1025–1061.
- [11] El Karoui, N.; Jeanblanc-Picqué, M. Martingale measures and partially observable diffusions. *Stochastic Anal. Appl.* **9** (1991), no. 2, 147–176.
- [12] Fischer, P.; Platen, E.; Runggaldier, W. Risk minimizing hedging strategies under partial observation. *Seminar on Stochastic Analysis, Random Fields and Applications (Ascona, 1996)*, 175–188, Progr. Probab., **45**, Birkhäuser, Basel, 1999.
- [13] Fleming, W. H.; Pardoux, E. Optimal control for partially observed diffusions. *SIAM J. Control Optim.* **20** (1982), no. 2, 261–285.
- [14] Fornari, F.; Mele, A. *Stochastic volatility in financial markets – Crossing the bridge to continuous time*. Kluwer A.P., 2000.
- [15] Fouque, J.-P.; Papanicolaou, G.; Sircar, K.R. *Derivatives in Financial Markets with Stochastic Volatility*. Cambridge, 2000.
- [16] Frey, R.; Runggaldier, W. Risk-minimizing hedging strategies under restricted information: the case of stochastic volatility models observable only at discrete random times. *Financial optimization. Math. Methods Oper. Res.* **50** (1999), no. 2, 339–350.
- [17] Gallant, A.R., Hsu, C.-T. and Tauchen, G. Using High/Low Data to Calibrate Volatility Diffusions and Extract the Forward Integrated Variance. Unpublished manuscript, 1998.
- [18] Gallant, A. R., Hsieh, D. and Tauchen, G. E. Estimation of stochastic volatility models with diagnostics, *Journal of Econometrics* **81** 159-192 (1997).

- [19] Gallant, A. R. and Tauchen, G. E. (1996), ‘Which Moments to Match?’, *Econometric Theory* **12** (4), 657–681.
- [20] Ghysels, E.; Harvey, A.C.; Renault, E. Stochastic volatility. *Statistical methods in finance*, 119–191, Handbook of Statist. **14**, North-Holland, 1996.
- [21] Lim, A. E. B.; Moore, John B. A quasi-separation theorem for LQG optimal control with IQ constraints. *Systems Control Lett.* **32** (1997), no. 1, 21–33.
- [22] Merton, R. Lifetime portfolio selection under uncertainty. The continuous time case. *Review of Economics and Statistics* **51** (1969), 247-257.
- [23] Merton, R. Optimum consumption and portfolio rules in a continuous time model. *Journal of Economic Theory* **3** (1971), 373-413.
- [24] Merton, R. The theory of rational option pricing. *Bell Journal of Economics and Management Science* **4** (1973), 141-183.
- [25] Moore, J. B.; Zhou, Xun Yu; Lim, A. E. B. Discrete time LQG controls with control dependent noise. *Systems Control Lett.* **36** (1999), no. 3, 199–206.
- [26] Nagai, H. Risk-sensitive dynamic asset management with partial information. Preprint.
- [27] Nelson, D.B. ARCH models as diffusion approximations. *J. Econometrics* **45** (1990), no. 1-2, 7–38.
- [28] Nielsen, J.N.; Vestergaard, M. Estimation in continuous time stochastic volatility models using nonlinear filters. *Int. J. Theor. Appl. Finance* **3** (2000), no. 2, 279–308.
- [29] Rishel, R.. A strong separation principle for stochastic control systems driven by a hidden Markov model. *SIAM J. Control Optim.* **32** (1994), no. 4, 1008–1020.
- [30] Rishel, R. Optimal portfolio management with partial observation and power utility function. In *Stochastic analysis, control, optimization and applications*. Eds: W. M. McEneaney, G. Yin, Q. Zhang. (1999) Birkhauser.
- [31] Viens, F. Portfolio optimization under partially observed stochastic volatility. *COMCON 8. The 8th International Conference on Advances in Communication and Control, 2001*. W. Wells, Ed. 1-12. Optim. Soft., Inc, Pub. Div., 2002.
- [32] Wonham, W. M. On the separation theorem of stochastic control. *SIAM J. Control* **6** (1968) 312–326. 93.60