Portfolio optimization under partially observed stochastic volatility

Frederi G. Viens

Department of Statistics and Department of Mathematics
1399 Mathematical Sciences Building
Purdue University
West Lafayette, IN 47907-1399

viens@stat.purdue.edu

Abstract. The stochastic volatility extension of the Black-Scholes model for one stock price is applied to the problem of stochastic portfolio optimization. The main assumption is that the portfolio manager has discrete access to the continuous-time stock prices. 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 [10], we propose a Monte-Carlo-type algorithm for solving the optimization problem.

1. Introduction

Many practicioners in today's financial industry believe that most stock prices and indices are best modeled by continuous-time stochastic processes, and in particular by diffusion processes. In the early 1970's Black, Scholes, and Merton ([3], [23]) were the first to acknowledge this fact, leading to the celebrated model that bears the first two authors' names. The third author's name is most often associated with his so-called Mutual Fund theorems ([21], [22]), 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, with the Black-Scholes model as the underlying stock price. Accordingly, Merton obtains answers by solving Hamilton-Jacobi-Bellman (HJB) equations. This approach contains two important drawbacks. The first one is the notorious fact that the Black-Scholes model's basic assumptions, that a stock's mean rate of return and volatility are constants, is not satisfied in many markets. We choose to use the Stochastic Volatility (SV) model, one of several extensions/corrections of the Black-Scholes model which have recent appeared. The second drawback is that the optimal portfolio given by the HJB equation makes changes continuously in time, based on stock-price information that arrives continuously in time, although for many investors, not only does the presence of transaction costs forbids such continuous trading, but information does not come in continuously, and is thus incomplete. 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.

In Section 2 we present a new form of optimal stochastic filtering, which shows how to estimate stochastic volatility optimally in the setting of incomplete information. In Section 3, we explain how this filtered volatility must be used for the selection of an optimal portfolio. In section 4, we detail an algorithm of the Monte-Carlo variety, for simulating the solution of the stochastic optimal control problem of section 3. An outline of the proof of convergence of the algorithm is provided. We begin in Paragraph 1.1 with a precise statement the problem we will solve, after which we review the relevant bibliography in Paragraph 1.2.

1.1. **Statement of the problem.** The SV model is

$$dX_t = X_t \mu dt + X_t \sigma(Y_t) dW_t, \qquad B_t = e^{rt} \tag{1}$$

where $t \in \mathbf{R}_+$, B is the non-risky asset (savings account), W is a Brownian motion,X is the risky asset price, 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 $corr(W, Z) = \rho$ with $0 \leq |\rho| < 1$, i.e.

$$dY_t = \alpha(Y_t) dt + \beta(Y_t) dZ_t. \tag{2}$$

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

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

For $i = 0, 1, \cdots, N$, let $\mathcal{F}_i^X$ be the information contained in the discrete sequence of *observed* asset prices $X_0, X_1, \cdots, X_i$. Note that this is not the commonly used "filtration of X", which contains much more information. For $\bar{x} = (x_0, \cdots, x_N)$ a fixed sequence of positive numbers, denote by $\mathcal{F}_i^{\bar{x}}$, the scenario (event) $\{X_0 = x_0, \cdots, X_i = x_i\}$. The stochastic volatility filtering problem is to estimate the conditional probability distribution

$$p_i(dy) := \mathbf{P}\left[Y_i \in dy | \mathcal{F}_i^X\right]. \tag{4}$$

This probability measure is random since it depends on the values $X_0, X_1, \cdots, X_i$. However at time i, the values $X_0 = x_0, \cdots, X_i = x_i$ 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 $p_i(dy)$ is non-random, depending only on the parameters $\bar{x}_i := (x_0, x_1, \cdots, x_i)$. We denote it by $p_i^{\bar{x}}(dy)$.

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 $s \in [i, i+1]$. 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

$$V(0, x_0, w_0) = \sup_a \mathbf{E}\left[U\left(\mathcal{W}_{N+1}^{a,b}\right) | X_0 = x_0, \mathcal{W}_0 = w_0\right] \tag{5}$$

for all $i = 0, \cdots, N$, 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 $w_0, x_0, x_1, \cdots, x_i$. 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).

1.2. **Significance of the problem.** Nonlinear stochastic filtering has a key role in partially observed stochastic control. We cite [13], [11], [12], [2] and recently [28]. Recent advances on finance-related aspect of this topic are still restricted to non-stochastic volatility: [29], [25], [20], [24], 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 $\sigma^2(Y_\cdot)$ 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 [19], [4], [14]). Dan Nelson ([26]; 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 $\delta \to 0$, leading many to believe the task is now to "bridge the gap to continuous time" (see [4]; [16],

[17], [18]). 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 [10] 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: [6] (a new *projection* filter); [7] (reduction to linear (Kalman) filtering in a special case); [27] (filtering w.r.t exogenous observation noise, not stochastic volatility).

2. Filtering with stochastic volatility

We establish an explicit recursion relation for the filter in (4). Note first that replacing X by $\tilde{X} = \ln X$ does not change $\mathcal{F}_i^X$, while in $\mathcal{F}_i^{\bar{x}}$ we simply need to replace x_i by $\tilde{x}_i = \ln x_i$. The advantage of $\tilde{X}$ is the explicit formula for $t \geq i$:

$$\tilde{X}_t = \tilde{x}_i + \int_i^t \left(\mu - \sigma^2 \left(Y_s\right)/2\right) ds + \int_i^t \sigma\left(Y_s\right) dW_s. \tag{6}$$

Since $corr\,(W, Z) = \rho$ for some $\rho \in (-1, 1)$, we decompose $W = \rho Z + \theta \tilde{W}$ where $\theta^2 + \rho^2 = 1$ and $\tilde{W}$ is independent of Z. Using the definition of conditional expectation, one can prove that the nonlinear stochastic filter $p_i^{\bar{x}}\left(dy_i\right)$ is given recursively by

$$p_{i+1}^{\bar{x}}\left(dy\right) = \frac{\int p_i^{\bar{x}}\left(dy_i\right) \mathbf{E}^Z \left[\mathbf{1}_{dy}\left(Y_{i+1}\right) g_{i+1}^Z\left(\tilde{x}_{i+1} - \tilde{x}_i\right) | Y_i = y_i\right]}{\int p_i^{\bar{x}}\left(dy_i\right) \mathbf{E}^Z \left[g_{i+1}^Z\left(\tilde{x}_{i+1} - \tilde{x}_i\right) | Y_i = y_i\right]}, \tag{7}$$

where $\mathbf{E}^Z$ is an expectation with respect to Z only, and where for each fixed realization of the increments $\{dZ_s : s \in [i, i+1]\}$, $g_{i+1}^Z\left(\tilde{x}\right)$ is the density at $\tilde{x}$, with respect to the randomness of $\tilde{W}$, of the random variable $\tilde{X}_{i+1} - \tilde{x}_i$, given $\tilde{X}_i = \tilde{x}_i$ and $Y_i = y_i$. Since Y and Z are non-random in the eyes of $\tilde{W}$, the explicit formula (6) easily yields

$$g_{i+1}^Z\left(\tilde{x}\right) = \left(2\pi\tau\right)^{-1/2} \exp\left(-\frac{\left(\tilde{x} - \zeta\right)^2}{2\tau}\right) \tag{8}$$

where the random variables τ and ζ are defined by

$$\begin{aligned} \zeta &= \int_i^{i+1} \left[\left(\mu - \frac{\sigma^2\left(Y_s\right)}{2}\right) ds + \rho\sigma\left(Y_s\right) dZ_s\right], \\ \tau &= \int_i^{i+1} \theta^2 \sigma^2\left(Y_s\right) ds. \end{aligned}$$

In view of the complexity of the iterative formula (7), there is currently no hope to evaluate $p_i^{\bar{x}}$ by any other method than the "smart"-Monte-Carlo algorithm recently established in [10], even for the simplest of examples. However, the proof of convergence of our Monte-Carlo method uses the explicit formulas above in a crucial way.

The algorithm of [10] (detailed in Section 5 therein), itself a bootstrapping extension of the genetic algorithm of [9], yields a good approximation (order $n^{-1/3}$) of $p_i^{\bar{x}}$ as the empirical distribution of a family of n interacting particles $\left(Y_i^k\right)_{k=1}^n$

$$\hat{p}_i^{\bar{x}}\left(dy\right) = \frac{1}{n} \sum_{j=1}^n \delta_{Y_i^j}\left(dy\right). \tag{9}$$

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 appropriate diffusions, with time step $m = n^{1/3}$.

The proof presented in [10] assumes that $\rho = 0$, and shows the convergence in L^1 of $\hat{p}_i^{\bar{x}}(f)$ to $p_i^{\bar{x}}(f)$ for deterministic bounded test functions f. We leave it to the reader to check that the following extension of their result is not any deeper.

Theorem 2.1. *Assume $\rho \in (-1,1)$. Let $x = (x_i : i = 1,2,\cdots,N)$ be a fixed sequence of positive numbers. The approximate filter $\hat{p}_i^{\bar{x}}$ of [10], which is a random probability measure in the probability space $\left(\hat{\Omega}, \hat{\mathcal{F}}, \hat{\mathbf{P}}\right)$ where the particles in (9) are defined, is such that for any random function f on $\hat{\Omega} \times \mathbf{R}$ satisfying $|f(\omega,y)| \leq c_f$ for some deterministic constant c_f and all ω, y:*

$$\hat{\mathbf{E}}\left[\left|\int \hat{p}_i^{\bar{x}}(dy)\, f(y) - \int p_i^{\bar{x}}(dy)\, f(y)\right|\right] \leq \frac{CB^i c_f}{n^{1/3}}.$$

The constants B and C depend only on $\mu, \sigma, \alpha, \beta, \rho$.

3. General portfolio optimization

For any scenario $\bar{x} := (x_0, x_1, \cdots, x_N)$, and any $i \leq N$, we define $\bar{x}_i = (x_0, \cdots, x_i)$. Our portfolio optimization problem can be imbedded in a dynamic one as follows: for all $w, x, \bar{x}$, for all $i = 1,2,\cdots,N$, for all $s \in [i, i+1]$, find

$$V(s,x,w) = V(s,x,w;\bar{x}_i) = \sup_{a \in A_0} \mathbf{E}\left[U\left(\mathcal{W}_{N+1}^a\right) | X_s = x, \mathcal{W}_s^a = w, \mathcal{F}_i^{\bar{x}}\right]. \tag{10}$$

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

$$\mathcal{W}_t = a_i X_t + (\mathcal{W}_i - X_i a_i) e^{r(t-i)} \tag{11}$$

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

Theorem 3.1. *For $s \in [i, i+1)$, V in (10) satisfies the Hamilton-Jacobi-Bellman (HJB) equation*

$$\frac{\partial V}{\partial s} + \sup_{a \in A_0}\left[(\mathcal{A}^a V)(s,x,w)\right] = 0. \tag{12}$$

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

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

Moreover, there exists an optimal control in A_0, i.e. the sup *in (12) is attained. This theorem also holds if A_0 is replaced by any proper subset of A_0.*

Proof. This is easily established using the classical proof of the HJB equation for stochastic control, and using the fact that because of self-financing (11), $\mathcal{W}$ is deterministic given X.

Here $Z^{i,X}$ 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 [30], [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 3.2, since the controls in A_0 can only change at times $i = 0,1,\cdots,N$, 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 (12), and is the key to our Monte-Carlo method.

Proposition 3.2. *Let I be any subset of $\mathbf{R}$,and replace A_0 by its restriction to I-valued sequences., For any $i = 0, 1, \cdots, N$, for any $f = f(x, w; \bar{x}_i)$, define*

$$\Phi^i(f)(x_i, w_i) = \Phi^i(f)_{\bar{x}_{i-1}}(x_i, w_i) := \sup_{a_i \in \mathbf{I}} E\left[f\left(X_{i+1}, \mathcal{W}^{a_i}_{i+1}; \bar{x}_i\right) | \mathcal{F}^{\bar{x}}_i, \mathcal{W}^{a_i}_i = w_i\right]. \tag{13}$$

Then we have

$$V(i, x_i, w_i)_{\bar{x}_{i-1}} := V(i, x_i, w_i; \bar{x}_i) = \Phi^i\left(\Phi^{i+1}\left(\cdots \Phi^N(U)\right)\right)_{\bar{x}_{i-1}}(x_i, w_i) \tag{14}$$

and the control $a^* = (a^*_0; \cdots, a^*_N)$ which is obtained by calculating an optimal a^*_i for the sup *in formula (13) is such that $a^* \in A_0$, and attains the* sup *in (10), i.e. is an optimal control.*

Proof. One notes that given $\{\mathcal{F}^{\bar{x}}_i, \mathcal{W}_i = w_i\}$, the pair $(X, \mathcal{W})$ is a Markov process because we are allowed to replace $\sigma(Y_t)^2$ by $Z^{i,X}_t(X_t)$. It is then easy to check that $\sup_{a_i \in I}$ can be replaced by $\sup_{a \in A_0}$ in the definition (13). One can then use Jensen's inequality to derive an upper bound on $V(N-1, w_{N-1}, x_{N-1})$. A reverse inequality is found by using the existence of an optimal control in Theorem 3.1, and the Proposition follows easily by iteration. ■

4. Monte-Carlo method

Since the only available numerical method to approximate $p^{\bar{x}}_i$ is the Monte-Carlo method of [10], it is natural to approximate our optimization problem using further Monte-Carlo techniques.

4.1. **Algorithm.** The following assumption, which is not restrictive in practice, is designed to simplify the proof of the convergence result (Theorem 4.1).

L1: The utility function U and its derivative U' are bounded. The volatility function σ^2 is bounded as $\sigma_0^2 \leq \sigma^2(y) \leq \sigma_1^2$ where σ_0 and σ_1 are positive constants.

We assume that $X_0 = x_0$ and $Y_0 = y_0$ are given. Without abandoning the hypothesis that (X, Y) have dynamics given by (1), (2), we impose that the only possible observations $\bar{x}$ satisfy

$$|\log x_{i+1} - \log x_i| \leq K_{m''} \tag{15}$$

where $K_{m''}$ is a constant depending only on an integer m''. We will see below that an optimal choice is $K_{m''} = C + C(\log m)^{1/2}$ for some C depending only on the SDEs' coefficients, where m is the number of Euler steps per unit interval of time. For large m, this choice makes it very unlikely (probability of order N/m) that an observed $\bar{x}$ would not satisfy the truncation condition (15). This condition implies that the observations are bounded: $\max_i |\log x_i| \leq |\log x_0| + NK_{m''}$ for all $\bar{x}$.

We also need to discretize the scenarii. For a fixed interger m', we impose that for all $\bar{x}$ and all i

$$x_i = \frac{k}{m'} \tag{16}$$

for some positive integer value k. As soon as m' is a multiple of 100, this is of course in accordance with the fact that stocks are traded in cents.

It is traditional to impose that $0 \leq w_i \leq e^{rN} w_{\max}$, for all i. This is equivalently a condition on the control set A_0, and can be achieved by saying that we pull out of the game if bankruptcy or a high level of wealth occurs. We may also impose the same truncation condition (15) on the w_i's as we did on the x_i's. Lastly we impose the same discretization condition (16) on the w_i's as on the stock prices.

We restrict our study to bounded strategies: assume there exists $A > 0$ such that for all $a \in A_0$,

$$\max_i |a_i| \leq A. \tag{17}$$

The discretization restrictions on wealth and stock prices imply that the strategy values a_i are automatically discretized. This defines a finite set A_0 of controls, with projection $A_{0,i}$ onto the i-th coordinate. By self-financing it is clear that the typical mesh size for A_0 is clearly of the order $1/m'$. However, if A_0 is the smallest set of strategies that is strictly necessary to achieve all possible scenarii under conditions (16) and (15), then this mesh size is not uniform, and in some locations, it is not even small. Therefore it is better to assume that A_0 is the union of this smallest set and of $((1/m')\mathbf{Z}\cap[-A,A])^N$, or, more generally:

L2: for all $a\in A_0$ there exists $a'\in A_0$ such that for all i, $|a_i-a'_i|\leq 1/m'$.

We introduce some notation.

- **Truncation.** For all $\bar{x}\in\mathbf{R}_+^N$, for all $i=0,\cdots,N$, if $|\log(x_{i+1}/x_i)|\leq K_m$ let $\{x_{i+1}\}=x_{i+1}$, if $\log(x_{i+1}/x_i)>K_m$ let $\{x_{i+1}\}=x_ie^{K_m}$, and if $\log(x_{i+1}/x_i)<-K_m$ let $\{x_{i+1}\}=x_ie^{-K_m}$. For a pair $(x,w)\in\mathbf{R}_+^2$, $\{(x,w)\}$ will denote either $(\{x\},\{w\})$ or $(\{x\},w)$ depending on whether we truncate w-values.
- **Discretization.** For all $x\in\mathbf{R}_+$, let $[x]$ be the largest element of $(1/m')\mathbf{N}$ that is smaller than x. For a pair $(x,w)\in\mathbf{R}_+^2$ we let $[(x,w)]=([x],[w])$.
- **Euler method**. For any d-dimensional Markov process $\mathcal{X}$ with infinitesimal generator $\mathcal{L}$ given by $(\mathcal{L}f)(x)=a_{kj}(x)\,\partial f/\partial x_k\partial x_j(x)+b_j(x)\,\partial f/\partial x_j(x)$, with σ a square root of the matrix a, the m-step Euler scheme for $\mathcal{X}$ on $[i,i+1]$ is the process defined by $\hat{\mathcal{X}}_i^m(0)=\mathcal{X}_i$ and
$$\hat{\mathcal{X}}_i^m(j+1)=\hat{\mathcal{X}}_i^m(j)+m^{-1}b\left(\hat{\mathcal{X}}_i^m(j)\right)+m^{-1/2}\sigma\left(\hat{\mathcal{X}}_i^m(j)\right)B_{i,j}$$
where $(B_{i,j})_{i=1,\cdots,N,j\in\mathbf{N}}$ is a family of independent standard d-dimensional Brownian motions. We denote $\hat{\mathcal{X}}_{i+1}^m=\hat{\mathcal{X}}_i^m(m)$.
- **Monte-Carlo approximation of Φ^i**. Let $\{\mathcal{X}(k):k=1,\cdots,n\}$.be a sequence of independent random variables sharing the same distribution as some integrable random variable $\mathcal{X}$. Let
$$\mathbf{M}[\mathcal{X}]=\frac{1}{n}\sum_{k=1}^n\mathcal{X}(k). \tag{18}$$
This empirical mean is the standard Monte-Carlo procedure for approximating $\mathbf{E}X$. For fixed $i\in\{1,\cdots,N\}$, for all bounded functions $f=f_{\bar{x}_i}(x_{i+1},w_{i+1})$, let
$$\hat{\mathbf{\Phi}}^i(f)_{\bar{x}_{i-1}}(x_i,w_i)=\sup_{a\in A_{0,i}}\mathbf{M}\left[f\left(\left[\left\{\widehat{(X,\mathcal{W}^a)}_{i+1}^m\right\}\right]\right)|\mathcal{W}_i=w_i,X_i=x_i,\mathcal{L}(Y_i)=\hat{p}_i^{\bar{x}}\right] \tag{19}$$
where $\widehat{(X,\mathcal{W}^a)}_{i+1}^m$ is the restriction to (x,w) of the basic Euler approximation of $(X,Y,\mathcal{W}^a)_{i+1}$. Here the conditioning under $\mathbf{M}$ means that the common starting point of the n independent copies of $\widehat{(X,\mathcal{W}^a)}_{i+1}^m$ is $(x_i,w_i,\hat{p}_i^x(dy))$ (or, more correctly, the starting distribution is $\delta_{x_i}\otimes\delta_{w_i}\otimes\hat{p}_i^{\bar{x}}$). The iteration of this procedure yields our Monte-Carlo method, denoted by
$$\hat{V}(i,\bar{x}_i,w_i)=\hat{\mathbf{\Phi}}^i\left(\hat{\mathbf{\Phi}}^{i+1}\circ\cdots\circ\hat{\mathbf{\Phi}}^N(U)\right)_{\bar{x}_{i-1}}(x_i,w_i).$$
- In the remainder of the paper, we write $\mathbf{E}[\cdot|\bar{x}_i,w_i,y]$ as a shorthand for $\mathbf{E}[\cdot|\mathcal{F}_i^{\bar{x}},\mathcal{W}_i=w_i,Y_i=y]$ when there is no risk of confusion.

The theorem below exhibits a value $\gamma_0>1$ such that that the optimal choice for n,n',m,m',m'' is $m=m''=n^{1/3}=(n')^{1/2}=(m')^{1/\gamma_0}$. The number of Monte-Carlo particles Y_i^k used for constructing $\hat{p}_i^{\bar{x}}$ is n, while only $n'=n^{2/3}$ Monte-Carlo particles are needed to construct $\hat{\mathbf{\Phi}}^i$ given $\hat{p}_i^{\bar{x}}$. It is more convenient, and may indeed be more efficient in practice, to use $n'=n$. That way, the k-th particle used to calculate $\hat{V}(i,\bar{x}_i,w_i)$ can be generated from the $(X,Y,\mathcal{W})$-dynamics starting at x_i,Y_i^k,w_i. Let S_i be the set of all possible sequences $\bar{x}_i=(x_j:j\leq i)$ that satisfy (15) and (16). Let T_i be the projection on the ith coordinate of S_i if one applies the truncation procedure in the variable w, otherwise let T_i be $(1/m')\mathbf{Z}\cap[-w_{\max},w_{\max}]$.

The Monte-Carlo procedure is as follows.

(1) **Initialization**. Let $X_0^k=x_0$, $\mathcal{W}_0^k=w_0$ and $Y_0^k=y_0$ for all $k=1,\cdots,n$. 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 $\bar{x} \in S$, 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 $\hat{V}$**. 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 $\hat{X}_{i+1}^m(k)$ using the Euler scheme with time step $1/m$ for the pair (X, Y) starting from $\left(x_i, Y_i^k(\bar{x})\right)$, 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
$$\widehat{\mathcal{W}^a}_{i+1}^m(k) = a_i \hat{X}_{i+1}^m(k) + a_i\left(\hat{X}_{i+1}^m(k) - x_i\right)e^r + w_i e^r,$$
   (b) calculate the Monte-Carlo average
$$\hat{F}(a_i, \bar{x}_i, w_i) = \frac{1}{n}\sum_{k=1}^n \hat{V}\left(i+1, \left[\left\{\widehat{(X, \mathcal{W}^a)}_{i+1}^m(k)\right\}\right], \bar{x}_i\right),$$
   (c) the i-th step optimal strategy and maximum expected yield are
$$\hat{V}(i, x_i, w_i, x_{i-1}) = \max_{a_i \in A_{0,i}} \hat{F}(a_i, \bar{x}_i, w_i), \qquad a_i^*(\bar{x}_i, w_i) = \underset{a_i \in A_{0,i}}{\arg\max} \hat{F}(a_i, \bar{x}_i, w_i).$$

4.2. **Usage and transaction costs.** For each $i = 0$ to N, a practicioner who observes stock value x_i and portfolio value w_i needs only to use the constant strategy $a_i^*(\bar{x}_i, w_i)$ in the interval $[i, i+1)$. At time i, this investor's expected maximal utility will then be given by $\hat{V}(i, x_i, w_i, x_{i-1})$, up to an error of order $1/m$, as our main Theorem 4.1 proves.

The power of this algorithm lies in the fact that, although its storage requirements are important (since at the very least, the values of a_i^* must be stored for each scenario $(\bar{x}, \bar{w})$), the computations only need to be performed once, and the practicioners may access the values of a_i^* directly from a calculated worksheet. This is particularly interesting for persons wishing to trade at high frequency, since the real-time computational demands are minimal.

Extending the algorithm to arbitrary non-random observation times $T_0, T_1, \cdots, T_N, T_{N+1} = T$ is straightforward. With equally spaced times, we let $\delta^{-1} = (T_{i+1} - T_i)^{-1}$ be the observation frequency. Clearly, higher δ^{-1} implies a higher expected yield. However, taking into account transaction costs implies that a higher δ^{-1} means a lower mean rate of return for the stock.

For illustration purposes, placing ourselves in the HARA case, we will assume it has been established that the uncertainty on the volatility, and its filtering at rate δ^{-1}, translate into a maximum expected utility in which the classical constant σ is replaced by a function $\sigma(\delta)$ that grows linearly in δ:
$$V(t, w) = \frac{1}{p} w^p \exp(pr(T-t)) \exp\left(\frac{2p}{1-p}\frac{(\mu - r)^2}{(\sigma_0 + c_0\delta)^2}(T-t)\right).$$
Incorporating symmetric proportional transaction costs is trivial: if one assumes that the purchase or sale of one dollar of stock costs c dollars, since there are δ^{-1} transactions in each unit of time, the mean rate of return μ needs simply to be replaced by $\mu - c/\delta$. Thus to optimize the value of V over all possible δ, we simply need to find the maximum of
$$R^2(\delta) := \left(\frac{\mu - r - c\delta^{-1}}{\sigma_0 + c_0\delta}\right)^2.$$
The maximum $R_{\max}$ occurs at $\delta_{\max} := c\left(1 + \sqrt{1 + \sigma_0(\mu - r)(cc_0)^{-1}}\right)/(\mu - r)$. This $R_{\max}$ has a numerator that ranges from $(\mu - r)/2$ to $\mu - r$ as $\mu - r$ increases, and a denominator of the same order

of magnitude as σ_0 as long as the original risk premium $\mu - r$ is not too small, and indeed very close to σ_0 for large $\mu - r$. In conclusion, when the risk premium is large, if one chooses the observation/transaction rate δ^{-1} appropriately, one can expect to do almost as well as with full observation.

4.3. **Convergence.**

Theorem 4.1. *Let*

$$\varepsilon > 0, \qquad \gamma_0 = 1 + \frac{2\sigma_1^2}{(1-\rho^2)\,\sigma_0^2}, \qquad \gamma > \gamma_0, \qquad K_0^* = \sup_{y \in \mathbf{R}} \left| \mu - \frac{\sigma^2(y)}{2} \right|, \qquad K > K_0^*.$$

Assume $n^{1/3} \geq m = m''$, $m' \geq m^\gamma$, and $n' \geq m^2$. Assume that the truncation procedure $\{\}$ in the definition of $\hat{\Phi}$ operates only on the random variable X. Assume (15) and (16) hold with

$$K_m = K^* + \sqrt{(2+\varepsilon)\,\sigma_1^2 \log m}, \tag{20}$$

$$N < \frac{K_m}{(1-\rho^2)\,\sigma_0^2}. \tag{21}$$

Then there exist constants $B, C > 0$ that depend only on the coefficients $\rho, \mu, \sigma, \alpha, \beta, r, A$, an integer m_0 depending also on γ, ε, and a constant $K^*_{\varepsilon,\rho,\sigma}$ depending only on $K^*, \varepsilon, \rho, \sigma$, such that for each $m > m_0$,

$$\sup_{\bar{x}, w_l} \hat{\mathbf{E}}\left[\left|V(i, \bar{x}_i, w_i) - \hat{V}(i, \bar{x}_i, w_i)\right|\right] \leq (N - i + 1)\frac{C}{m}\left(\|U\|\left(B^i + 1 + K^*_{\varepsilon,\rho,\sigma}\right) + \|U'\|\frac{e^{(N-i)\mu}}{m^{\gamma-1}}\right). \tag{22}$$

If the truncation procedure operates on both X and $\mathcal{W}$, then the statements above remain true with σ_1 replaced by $(1 \vee A)\,\sigma_1$ where A is the bound in (17). The supremum over w_i can be over any arbitrary set of values.

Strictly speaking, our theorem only guarantees a good approximation if m grows exponentially with N. It is possible to formulate a theorem in which this condition is not needed. The following is an informal discussion of what needs to be done.

For typical values of the parameters μ, σ, ρ (significantly less than 1), with Euler time step $m = 1000$, an admissible choice for N is 20, which may be uncomfortably small for people wishing to trade at high frequency. However, in practice (see [10]) the term B^i can be replaced by 1 when Y is an ergodic process, such as the oft used, mean-reverting O.U. process (3). Moreover, interpreting the term $e^{(N-i)\mu}$ as exponential in N is betraying the fact if N represents a period of T years, it equals $(1+R)^T$ where R is the effective annual rate of increase corresponding to μ, and thus does not depend on N.

The only problematic condition is (21). The presence of K^* makes it possible to choose N large, but the tradeoff is a larger value of γ; examining the proof of Theorem 4.1, one can see that for manageble values of m, γ must be linear in K^*, which means that a linear increase in N implies an exponential increase in the number of space-mesh points m'. A better solution is to require bounded observations $|x_i| \leq x_{\max}$ for all $i \leq N$. Then one can see that the restriction (21) is not needed. The probability of one's actual observations not being bounded is given by

$$\mathbf{P}\left[\sup_{t \in [0,N]} |X(i)| \geq x_{\max}\right] = \mathbf{P}\left[\sup_{t \in [0,N]} \left(\int_0^N \sigma(Y_s)\,dW_s + \int_0^N \left(\mu - \frac{1}{2}\sigma^2(Y_s)\right) ds\right) \geq \log(x_{\max}/x_0)\right].$$

Choosing $x_{\max} = x_0 \exp\left(NK_0^* + \sigma_1\sqrt{N \log m}\right)$ ensures that this probability no grater than $1/m$. If $\mu - \sigma^2/2$ is always negative (large volatility), the term NK_0^* can be deleted. For small volatility, $K_0^* < \mu$ so that $\exp NK_0^* < (1+R)^T$. The only remaining factor in $x_{\max}$ is $\exp \sigma_1 \sqrt{N \log m}$, which is the same order as the bound on x_i implied by the truncation condition (15). We conclude that one can generally ignore condition (21). The probability that the actual observations do not satisfy (15) is of order $1/m$.

To prove Theorem 4.1, we need several propositions on convergence speeds: the first result covers the standard Euler method for SDEs (see [1]), the second treats the standard Monte-Carlo method, the third investigates the operation $\hat{\Phi}$. The fourth establishes a type of contraction property for $\hat{\Phi}$.

Proposition 4.2. *Referring to the description of the Euler method above, let $q(x,x')$ and $\hat{q}^m(x,x')$ be the densities of the laws of $\mathcal{X}_{i+1}$ and $\hat{\mathcal{X}}_{i+1}^m$ respectively, given $\mathcal{X}_i = x$. There exist constants C and C' depending only on a,b such that $q(x,x') + \hat{q}^m(x,x') \le Ce^{-C'|x-x'|^2}$ and*

$$|x-x'| > 2/m \quad \Rightarrow \quad |q(x,x') - \hat{q}^m(x,x')| \le \frac{C}{m} e^{-C'|x-x'|^2}.$$

Proposition 4.3. *Let $\mathcal{X}$ be a bounded random variable (there exists a deterministic constant C such that $|\mathcal{X}| \le C$). Let $\{\mathcal{X}(k) : k = 1, \cdots, n\}$ be a family of independent copies of $\mathcal{X}$, under some other probability measure $\hat{\mathbf{P}}$ and let $\mathbf{M}[\mathcal{X}]$ be as in (18). Then*

$$\hat{\mathbf{E}}[|\mathbf{M}[\mathcal{X}] - \mathbf{E}[\mathcal{X}]|] \le K_u \frac{C}{n^{1/2}}$$

where K_u is a universal constant.

Proof. The proof is presumably well known. It uses the basic ideas of large deviations and subgaussian random variables. ■

Proposition 4.4. *Let the integers n, n', m, m', m'' be fixed and assume the values of $\bar{x}$ are limited by the condition*

$$|\log x_{i+1} - \log x_i| \le K_{m''}$$

where $K_{m''}$ satisfies (20) for some fixed $\varepsilon > 0$, with σ_1 replaced by $(1 \vee A)\,\sigma_1$ if the truncation function $\{\}$ operates on both x and w. For any bounded function $f = f_{\bar{x}_i}(x_{i+1}, w_{i+1})$,

$$\text{(23)}\quad \sup_{\bar{x}_l w_i} \hat{\mathbf{E}}\left[\left|\Phi^i(f)_{\bar{x}_i-1}(x_i,w_i) - \hat{\Phi}^i(f)_{\bar{x}_i-1}(x_i,w_i)\right|\right] \le \frac{CB^i\|f\|}{n^{1/3}} + \frac{C\|f\|}{m} + \frac{\|f'\|}{m'} + \frac{K_\varepsilon\|f\|}{m''} + \frac{K_u\|f\|}{(n')^{1/2}}$$

where $\|f\| = \sup_{\bar{x}_{i+1}, w_{i+1}} |f(\bar{x}_{i+1}, w_{i+1})|$ and $\|f'\| = \sup_{\bar{x}_{i+1}, w_{i+1}} \left(\left|\frac{\partial f}{\partial x_{i+1}}\right| + \left|\frac{\partial f}{\partial w_{i+1}}\right|\right)$, B and C depend only on $\rho, \mu, \sigma, \alpha, \beta, r, A$, K_ε depends only on ε, and K_u is a universal constant.

Proof. The first step is to notice that we have:

$$\text{(24)}\qquad \mathbf{E}\left[f\left(X_{i+1}, \mathcal{W}_{i+1}^a\right)|\mathcal{F}_i^{\bar{x}}, \mathcal{W}_i = w_i\right] = \int p_i^{\bar{x}}(dy)\,\mathbf{E}\left[f\left(X_{i+1}, \mathcal{W}_{i+1}^a\right)|\mathcal{F}_i^{\bar{x}}, \mathcal{W}_i = w_i, Y_i = y\right].$$

As a consequence we can rewrite

$$\begin{aligned}
&\left|\Phi^i(f)_{\bar{x}_i-1}(x_i,w_i) - \hat{\Phi}^i(f)_{\bar{x}_i-1}(x_i,w_i)\right| \\
&\le \left|\sup_{a\in A_{0,i}} \int p_i^{\bar{x}}(dy)\,\mathbf{E}\left[f\left((X,\mathcal{W}^a)_{i+1}\right)|\bar{x}_i, w_i, y\right] - \sup_{a\in A_{0,i}} \int \hat{p}_i^{\bar{x}}(dy)\,\mathbf{E}\left[f\left((X,\mathcal{W}^a)_{i+1}\right)|\bar{x}_i, w_i, y\right]\right| \\
&+ \sup_{a\in A_{0,i}} \int \hat{p}_i^{\bar{x}}(dy)\left|\mathbf{E}\left[f\left((X,\mathcal{W}^a)_{i+1}\right)|\bar{x}_i,w_i,y\right] - \mathbf{E}\left[f\left(\left\{(X,\mathcal{W}^a)_{i+1}\right\}\right)|\bar{x}_i,w_i,y\right]\right| \\
&+ \sup_{a\in A_{0,i}} \int \hat{p}_i^{\bar{x}}(dy)\left|\mathbf{E}\left[f\left(\left\{(X,\mathcal{W}^a)_{i+1}\right\}\right)|\bar{x}_i,w_i,y\right] - \mathbf{E}\left[f\left(\left\{\widehat{(X,\mathcal{W}^a)}_{i+1}^m\right\}\right)|\bar{x}_i,w_i,y\right]\right| \\
&+ \sup_{a\in A_{0,i}} \int \hat{p}_i^{\bar{x}}(dy)\left|\mathbf{E}\left[f\left(\left[\left\{\widehat{(X,\mathcal{W}^a)}_{i+1}^m\right\}\right]\right)|\bar{x}_i,w_i,y\right] - \mathbf{E}\left[f\left(\left\{\widehat{(X,\mathcal{W}^a)}_{i+1}^m\right\}\right)|\bar{x}_i,w_i,y\right]\right| \\
&+ \sup_{a\in A_{0,i}} \int \hat{p}_i^{\bar{x}}(dy)\left|\mathbf{E}\left[f\left(\left[\left\{\widehat{(X,\mathcal{W}^a)}_{i+1}^m\right\}\right]\right)|\bar{x}_i,w_i,y\right] - \mathbf{M}\left[f\left(\left[\left\{\widehat{(X,\mathcal{W}^a)}_{i+1}^m\right\}\right]\right)|\bar{x}_i,w_i,y\right]\right| \\
&= A_1 + A_2 + A_3 + A_4 + A_5.
\end{aligned}$$

A_5 is controlled uniformly in $a, \bar{x}_i, w_i, y$ using Proposition 4.3 with $\mathcal{X} = f\left(\left[\left\{\widehat{(X, \mathcal{W}^a)}_{i+1}^m\right\}\right]\right)$. Controlling A_4 is trivial using the fact that $|f(x,w) - f([(x,w)])| \leq \|f'\|/m'$. Proposition 4.2 for $\mathcal{X} = (X, Y, \mathcal{W})$ easily yields a bound for A_3. To control A_2, since f is bounded, it is sufficient to estimate

$$\mathbf{P}\left[\left|\int_i^{i+1} \sigma(Y_s)\, dW_s\right| > K_{m''} - \sup_{y'}\left|\mu - \frac{1}{2}\sigma^2(y')\right| |y\right].$$

Using Chebyshev's inequality $\mathbf{P}(Z > a) \leq \mathbf{E}\left[e^{\lambda Z^2}\right] e^{-\lambda a^2}$ one sees that to obtain a bound of $1/m$ it is sufficient to take $\lambda = (2+\varepsilon)^{-1} \sigma_1^{-2}$ and $K_{m''}$ as in (20). To finish the proof of the proposition, we must estimate

$$\sup_{\bar{x}_i} \sup_{w_i} \hat{\mathbf{E}} \sup_{a \in A_{0,i}} \left|\int \left(p_i^{\bar{x}}(dy) - \hat{p}_i^{\bar{x}}(dy)\right) F_{x_i, w_i, a}(y)\right|.$$

where $F_{x_i,w_i,a}(y) = \mathbf{E}\left[f\left((X, \mathcal{W}^a)_{i+1}\right) |\bar{x}_i, w_i, y\right]$. The presence of the supremum inside the expectation causes a problem, which is remedied using the Hahn-Jordan decomposition of the measure $\nu = p_i^{\bar{x}} - \hat{p}_i^{\bar{x}}$ into a difference of positive measures $\nu \mathbf{1}_H + \nu \mathbf{1}_{H^c}$, making it possible to bring the supremum inside the y-integration, so Theorem 2.1 can be applied to the function $\mathbf{1}_H(y) \sup_a |F_{x_i,w_i,a}(y)|$.■

Proposition 4.5. *For any bounded functions $f = f_{\bar{x}_i}(x_{i+1}, w_{i+1})$ and $g = g_{\bar{x}_i}(x_{i+1}, w_{i+1})$, which may moreover be random under $\hat{\mathbf{P}}$, we have*

$$\sup_{\bar{x}_i} \sup_{w_i} \hat{\mathbf{E}}\left[\left|\hat{\Phi}^i(g)_{\bar{x}_{i-1}}(x_i, w_i) - \hat{\Phi}^i(f)_{\bar{x}_{i-1}}(x_i, w_i)\right|\right] \leq \sup_{\bar{x}_{i+1}} \sup_{w_{i+1}} \hat{\mathbf{E}}\left[\left|f_{\bar{x}_i}(x_{i+1}, w_{i+1}) - g_{\bar{x}_i}(x_{i+1}, w_{i+1})\right|\right],$$

where the suprema in $\bar{x}_i$ and $\bar{x}_{i+1}$ are over the finite sets determined by conditions (15) and (16), while the suprema in w_i are over arbitrary sets.

Proof. Using the rule (24), Fubini's theorem, and the trivial fact $\mathbf{E}f(Z) \leq \sup f$, the proof is straightforward.■

The last and most technical step before proving Theorem 4.1 is to control $\|f'\|$ with $f = \Phi^i \circ \cdots \circ \Phi^N(U)$, which we do in the next proposition

Proposition 4.6. *Let $f_{\bar{x}_{i-1}}(x_i, w_i) = \Phi^i \circ \cdots \circ \Phi^N(U)_{\bar{x}_{i-1}}(x_i, w_i)$. There exists a constant $K_{\rho,\sigma,\mu}$ depending only on ρ, σ, μ, r, A, such that*

$$\sup_{w_i, \bar{x}_i}\left[\left|\frac{\partial f_{\bar{x}_{i-1}}}{\partial x_i}\right| + \left|\frac{\partial f_{\bar{x}_{i-1}}}{\partial w_i}\right|\right] \leq K_{\rho,\sigma,\mu,r,A}\left(\|U\| \exp\left(\frac{K_{m''}^2}{\theta^2 \sigma_0^2} + iK_{m''}\right) + \|U'\| e^{\mu(N-i)}\right)$$

Proof. From the expression $f_{\bar{x}_{i-1}}(x_i, w_i) = \sup_{a \in A_0} \int p_i^{\bar{x}}(dy)\, \mathbf{E}\left[U\left(\mathcal{W}_{N+1}^a\right) |\bar{x}_i, w_i, y\right]$, the explicit formula

$$\mathcal{W}_{N+1}^a = w_i e^r + \sum_{j=i}^N a_j\left(X_{j+1} - e^r X_j\right) e^{(N-j)r}, \tag{26}$$

which follows from the self-financing condition, the expression for X_j/x_i given by (6), and since $p_i^{\bar{x}}$ also depends explicitly on x_i via formula (7), $f_{\bar{x}_{i-1}}$ appears as an explicit function of (x_i, w_i). An upper bound on its derivatives is obtained using the general fact $\left|\frac{d}{dx}(\sup g_a)(x)\right| \leq \sup_a \left|\frac{dg_a}{dx}(x)\right|$ for any family $\{g_a\}$ of C^1 functions. It follows immediately that

$$\left|\frac{\partial}{\partial w_i} f_{\bar{x}_{i-1}}(x_i, w_i)\right| \leq \|U'\| e^r.$$

For some constant $K_{\rho,\sigma,\mu}$ depending only on ρ, σ, μ, the following upper bounds are straighforward, using conditions L1, (15), L2, and the martingale property for $j \mapsto X_j/x_i$:

$$\left|\frac{\partial}{\partial x_i}\mathbf{E}\left[U\left(w_i e^r + x_i \sum_{j=i}^N a_j\left(\mathcal{E}_{i,j+1}(y) - e^r \mathcal{E}_{i,j}(y)\right) e^{(N-j)r}\right)\right]\right| \leq 2A\|U'\| \frac{e^{\mu+r}}{e^{\mu-r} - 1} e^{\mu(N-i)},$$

$$\left|\frac{\partial}{\partial x_i}\int p_{i-1}^{\bar{x}_{i-1}}(dy')\,\mathbf{E}^Z\left[\mathbf{1}_{dy}(Y_i)\,g_i^Z(\tilde{x}_i-\tilde{x}_{i-1})\,|Y_{i-1}=y'\right]\right|$$
$$\leq K_{\rho,\sigma,\mu}\exp\left(iK_{m''}\right)\int p_{i-1}^{\bar{x}_{i-1}}(dy')\,\mathbf{P}^Z\left[Y_i\in dyY_{i-1}=y'\right]$$

$$\int p_{i-1}^{\bar{x}}(dy')\,\mathbf{E}^Z\left[g_i^Z(\tilde{x}_i-\tilde{x}_{i-1})\,|Y_{i-1}=y'\right]\geq K_{\rho,\sigma,\mu}\exp\left(-\frac{K_{m''}^2}{\theta^2\sigma_0^2}\right);$$

from which the proposition easily follows. ■

Proof of Theorem 4.1.

We prove the theorem by backwards induction on the value i. By definition we have $V(N+1,\bar{x},w)=U(w)=\hat{V}(N+1,\bar{x},w)$ and the theorem holds true for $i=N+1$. Now assume (22) holds true for some value $i\leq N+1$. By definition of $f_{\bar{x}_{i-1}},V,\hat{V}$, and Propositions 4.4, 4.5 and 4.6, we get:

$$\begin{aligned}
&\sup_{\bar{x}_{i-1}}\sup_{w_{i-1}}\hat{\mathbf{E}}\left|V(i-1,\bar{x}_{i-1},w_{i-1})-\hat{V}(i-1,\bar{x}_{i-1},w_{i-1})\right|\\
&\leq\sup_{\bar{x}_{i-1}}\sup_{w_{i-1}}\hat{\mathbf{E}}\left|\Phi^{i-1}\left(f_{\bar{x}_{i-1}}\right)_{\bar{x}_{i-2}}(x_{i-1},w_{i-1})\right)-\hat{\Phi}^{i-1}\left(f_{\bar{x}_{i-1}}\right)_{\bar{x}_{i-2}}(x_{i-1},w_{i-1})\right)\right|\\
&+\sup_{\bar{x}_i}\sup_{w_i}\hat{\mathbf{E}}\left[\left|V(i,\bar{x}_i,w_i)-V(i,\bar{x}_i,w_i)\right|\right]\\
&\leq\|U\|\left(\frac{CB^i}{n^{1/3}}+\frac{C}{m}+\frac{K_\varepsilon}{m''}+\frac{K_u}{(n')^{1/2}}\right)+\frac{K_{\rho,\sigma,\mu,r,A}}{m'}\left(\|U\|\exp\left(\frac{K_{m''}^2}{\theta^2\sigma_0^2}+iK_{m''}\right)+\|U'\|\,e^{\mu(N-i)}\right)\\
&+(N-i+1)\frac{C}{m}\left(\|U\|\left(B^i+1+K_{\varepsilon,\rho,\sigma}^*\right)+\|U'\|\frac{e^{(N-i)\mu}}{m^{\gamma-1}}\right).
\end{aligned}$$

It is optimal to choose $n^{1/3}=m=m''=(n')^{1/2}$. Using the definition of K_m in (20), and the bound on N in (21), for m large enough, we obtain, for a constant $K_{\theta,\sigma}^*$ depending only on K^*,θ,σ,

$$\exp\left(\frac{K_{m''}^2}{\theta^2\sigma_0^2}+iK_{m''}\right)\leq\frac{K_{\theta,\sigma}^*}{m'}\exp\left(\frac{2(2+\varepsilon)\sigma_1^2\log m}{\theta^2\sigma_0^2}\right)=\frac{K_{\theta,\sigma}^*}{m'}m^{\frac{2(2+\varepsilon)\sigma_1^2}{\theta^2\sigma_0^2}},$$

hence the choice for m' as announced in the Theorem, which finishes the proof. ■

We conclude this article by studying the effect of discretizing the set of all possible strategies. Assuming the control set A_0 in the original problem (5) takes bounded values as in (17), let $\tilde{A}_0$ be the set of strategies described in the the paragraph leading to condition L2. [The study for nonbounded A_0 requires a localization step which we omit.] The following proposition makes no claims regarding the effect on the optimal strategy, but from the practitioner's standpoint, since an optimal strategy exists for the discretized control set, only the effect on V is relevant.

Proposition 4.7. *Let $V(0,\cdot)$ be given by (10). Let*

$$\tilde{V}(0;x,w)=\sup_{a\in\tilde{A}_0}\mathbf{E}\left[U\left(\mathcal{W}_{N+1}^a\right)|X_0=x,\mathcal{W}_0^a=w\right].$$

If conditions L1 and L2 hold, then there is a constant K depending only on σ,μ,r such that

$$0\leq\sup_{x,w\in\mathbf{R}_+}\left[V(0;x,w)-\tilde{V}(0;x,w)\right]\leq\frac{1}{m'}(N+1)\,C^N\|U'\|_\infty.$$

Proof. The lower bound is trivial. The upper bound is easy if one writes $V(0)=\sup_{A_0}F$ and $\tilde{V}(0)=\sup_{\tilde{A}_0}F$ with $F(a)=\mathbf{E}\left[U\left(\mathcal{W}_{N+1}^a\right)|X_0=x,\mathcal{W}_0^a=w\right]$, and one uses the explicit expression (26) together with the fact that $V(0)-\tilde{V}(0)\leq F(a^*)-F(\tilde{a}^*)$ where a^* is the optimal control in A_0 and $\tilde{a}^*$ is its closest point in $\tilde{A}_0$. ■

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 corporate 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] Bouleau, N.; Lépingle, D. *Numerical Methods for Stochastic Processes.* John Wiley & Sons (1994).

[6] 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.

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

[8] 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.

[9] 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.

[10] Del Moral, P.; Jacod, J.; Protter, Ph. The Monte-Carlo method for filtering with discrete-time observations. (1999). To appear in *Annales de l'I.H.P.*

[11] 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.

[12] El Karoui, N.; Jeanblanc-Picqué, M. Martingale measures and partially observable diffusions. *Stochastic Anal. Appl.* **9** (1991), no. 2, 147–176.

[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] 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.

[17] Gallant, A. R., Hsieh, D. and Tauchen, G. E. Estimation of stochastic volatility models with diagnostics, *Journal of Econometrics* **81** 159-192 (1997).

[18] Gallant, A. R. and Tauchen, G. E. (1996), 'Which Moments to Match?', *Econometric Theory* **12** (4), 657-681.

[19] Ghysels, E.; Harvey, A.C.; Renault, E. Stochastic volatility. *Statistical methods in finance*, 119–191, Handbook of Statist. **14**, North-Holland, 1996.

[20] 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.

[21] Merton, R. Lifetime portfolio selection under uncertainty. The continuous time case. *Review of Economics and Statistics* **51** (1969), 247-257.

[22] Merton, R. Optimum consumption and portfolio rules in a continuous time model. *Journal of Economic Theory* **3** (1971), 373-413.

[23] Merton, R. The theory of rational option pricing. *Bell Journal of Economics and Management Science* **4** (1973), 141-183.

[24] 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.

[25] Nagai, H. Risk-sensitive dynamic asset management with partial information. Preprint.

[26] Nelson, D.B. ARCH models as diffusion approximations. *J. Econometrics* **45** (1990), no. 1-2, 7–38.

[27] 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.

[28] 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.

[29] 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.

[30] Wonham, W. M. On the separation theorem of stochastic control. *SIAM J. Control* **6** (1968) 312–326. 93.60