

THE RELAXED INVESTOR WITH PARTIAL INFORMATION

NICOLE BÄUERLE*, SEBASTIAN P. URBAN*, AND LUITGARD A. M. VERAART‡

ABSTRACT. We consider an investor in a financial market consisting of a riskless bond and several risky assets. The price processes of the risky assets are geometric Brownian motions where either the drifts are modelled as random variables assuming a constant volatility matrix or the volatility matrix is considered random and drifts are assumed to be constant. The investor is only able to observe the asset prices but not all the model parameters and hence information is only partial. A Bayesian approach is used with known prior distributions for the random model parameters.

We assume that the investor can only trade at discrete time points which are multiples of $h > 0$ and investigate the loss in expected utility of terminal wealth which is due to the fact that the investor cannot trade and observe continuously.

It turns out that in general a discretization gap appears, i.e., for $h \rightarrow 0$ the expected utility of the h -investor does not converge to the expected utility of the continuous investor. This is in contrast to results under full information in (Rogers, L.C.G. 2001. The relaxed investor and parameter uncertainty. Finance and Stochastics, 5(2), 131-154).

We also present simple asymptotically optimal portfolio strategies for the discrete-time problem. Our results are illustrated by some numerical examples.

KEY WORDS: Optimal investment, partial information, Bayesian approach, Markov decision problem, discrete versus continuous trading, discretization gap.

AMS SUBJECT CLASSIFICATIONS: 91G10, 90C40, 93E20, 93E11.

1. INTRODUCTION

This paper combines two lines of research which have attracted attention over the last years: Portfolio optimization under partial information and approximation of continuous trading strategies by strategies which adjust portfolio weights only in discrete time. It will turn out that some surprising effects occur.

More precisely, we consider a standard Black-Scholes market with d risky assets and one money market account with constant interest rate. In the first part of the paper we assume that the drift of the risky assets is not known to the investor. Her aim is to maximize the expected log-utility of her wealth at terminal time $T > 0$. We follow a Bayesian approach here. Clearly, when the drift rates μ are chosen according to a multivariate normal prior, then the conditional distribution of μ given the observation of the stock prices is again normal and there exists a simple sufficient statistic. This is useful since we can apply the *principle of estimation and control* to solve the portfolio problem. If the investor is able to trade in continuous-time this problem is well-understood and can e.g. be found in Karatzas & Zhao (2001). Now let us assume the investor is only able to observe the stock prices and adjust the portfolio at discrete points in time which are multiples of $h > 0$. In practice, trading continuously is impossible and the investor also might want to avoid transaction cost or relax and therefore reduces the trading frequency. As in Rogers (2001), we will call her the h -investor and the optimization problem

* Karlsruhe Institute of Technology (KIT), Institute for Stochastics, Kaiserstr. 89, 76133 Karlsruhe, Germany
nicole.baeuerle@kit.edu and sebastian.urban@kit.edu.

‡ London School of Economics, Department of Mathematics, Columbia House, Houghton Street, London WC2A 2AE, UK. Tel: +44 (0)207 107 5062, l.veraart@lse.ac.uk .

The second and third named authors gratefully acknowledge funding by the Ministry of Science, Research and Arts, Baden-Württemberg, Germany.

h -optimization problem in contrast to the continuous investor and the continuous problem. In the case of complete information these two portfolio problems have been compared by Rogers (2001) for power utilities. Rogers (2001) shows that in general the difference between these two problems is small.

We will show that this is not true anymore in the case of partial observation. In general, there remains a *discretization gap*, i.e., for $h \rightarrow 0$ the expected utility of the h -investor does not converge to the expected utility of the continuous investor. This is due to the fact that for the h -investor it is never optimal to short-sell stocks or the bond. Surprisingly, the h -investor does not lose information when she *observes* stock prices only in discrete time, i.e., the estimators of the drift of the h -investor and the continuous investor coincide at the discrete time points.

Furthermore, we show that the optimal trading strategy for the h -investor can be approximated by the very simple to implement “plug-in” Merton-ratio strategy, i.e., we take the Merton-ratio and replace the unknown μ by its estimate. In Taksar & Zeng (2007) the authors studied discrete-time portfolio problems where the drift and the volatility are driven by an unobservable discrete-time Markov chain and used Taylor-series expansion of the utility to derive simple approximations of the optimal portfolio strategy. We instead use the knowledge of the continuous problem to approximate the portfolio strategy.

In the second part of the paper we consider the same problem with known stock price drift but unknown volatility matrix. The continuous investor can immediately estimate the volatility matrix from a tiny period of observed stock prices, but this is not true for the h -investor who obtains information only at discrete time points. In this case the h -investor loses relevant information by discrete observation. We solve the h -optimization problem assuming that the prior distribution of $(\sigma\sigma^\top)^{-1}$ is a Wishart-distribution. Again, for $h \rightarrow 0$, there remains a discretization gap between the continuous expected utility and the limit of the h -expected utilities. We propose an easy to implement portfolio strategy for the h -investor which is asymptotically optimal for $h \rightarrow 0$.

By now there are quite some papers investigating portfolio problems with unknown, not necessarily constant drift. In most papers it is either assumed that the drift is a linear, mean-reverting diffusion or a function of a finite state Markov chain. Both cases lead to well-known, finite-dimensional filters. Lakner (1995, 1998) uses a martingale approach to represent the optimal terminal wealth and trading strategy. In Lakner (1995) explicit results in case of log and power utility with constant but random drift are given, whereas in Lakner (1998) the case of a linear diffusion is treated. Rishel (1999) uses a dynamic programming approach to solve the linear diffusion case with power utility. Brendle (2006, 2008) slightly extends this problem to a multivariate setting where the linear diffusion may be correlated with the stock prices and gives results for power and exponential utility. In Zohar (2001) inverse Laplace transforms are used to represent the optimal strategy and value function for general utility problems. Karatzas & Zhao (2001) present a comprehensive approach via the martingale method. Moreover, there are papers dealing with the hidden Markov-modulated drift problem. Honda (2003) investigates the problem for two states. The general case is considered in Sass & Hausmann (2004) with the help of the martingale approach and Malliavin calculus and in Rieder & Bäuerle (2005) by dynamic programming focusing on power and log-utility. The latter paper also compares the optimal portfolio strategies in the case of complete and partial information. The recent paper Björk et al. (2010) takes a general point of view where the drift is allowed to be an arbitrary semimartingale. Nevertheless, fairly explicit expressions for the optimal terminal wealth and the optimal portfolio strategy are derived in the cases of log, power and exponential utility.

Pham & Quenez (2001) treat portfolio problems with partially observed stochastic volatility. In discrete-time the portfolio problem with unobservable drift and volatility is investigated in Taksar & Zeng (2007). Frey & Runggaldier (1999) consider hedging problems in discrete-time with unknown volatility. An unknown jump intensity is treated in Bäuerle & Rieder (2007).

The question of approximating the (often) complicated filter and thus also the portfolio strategy is tackled in Pham et al. (2005) and Corsi et al. (2008) by means of quantization techniques.

The relation between continuous- and discrete-time filters is investigated in James et al. (1996). In Gandy & Veraart (2012+) the authors study the effect of estimating parameters in optimal investment strategies when the number of assets gets large.

For the approximation of continuous-time financial models and portfolio strategies by discrete-time models and strategies see e.g. He (1991), Duffie & Protter (1992) or Prigent (2003). The latter two sources consider among others general questions of convergence of stochastic integrals. Most of the results are positive in the sense that convergence of the wealth process is obtained. This is in contrast to our findings. In Duffie & Protter (1992) two counterexamples are given where the processes do not convergence, but these examples are quite academic. In most cases considered in these sources, the financial market is also approximated.

Our paper is organized as follows: In the next section we introduce the financial market with unknown drift. In Section 3 we review the continuous-time optimization problem, followed by the h -optimization problem in Section 4. In both cases the aim is to maximize the expected log-utility of the investor's wealth. Afterwards we establish in Section 5 the comparison between the continuous and the h -optimization problem. We show that there remains a discretization gap between the optimal value of the problems for $h \rightarrow 0$ and construct a simple but asymptotically optimal portfolio strategy for the h -investor. In Section 6 we investigate the same optimization problem with known drift but unknown volatility. We proceed in the same way as before, i.e., we first show that there exists a discretization gap and construct a simple but asymptotically optimal portfolio strategy for the h -investor. This case differs from the first one since the h -investor with unknown volatility loses information by the discrete observation of the stock prices. In Section 7 we provide some numerical results and give additional insight in the size and sensitivity of the discretization gap. We also discuss in the case of unknown volatility to which extend the loss of information can be blamed for the distance to the continuous investor. The paper ends with a conclusion.

2. A FIRST FINANCIAL MARKET MODEL

We consider a financial market with finite time horizon $0 < T < \infty$. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be the underlying probability space on which we consider a d -dimensional Brownian motion $W = (W_t^1, \dots, W_t^d)$. We suppose that there are d risky assets and one riskless bond with the following dynamics for $t \in [0, T]$. The price process S_t^0 of the riskless bond is given by

$$S_t^0 := e^{rt},$$

where $r \geq 0$ denotes the deterministic interest rate. The price process S_t^i of the risky asset $i, i = 1, \dots, d$, satisfies the stochastic differential equation

$$dS_t^i = S_t^i \left(\mu_i dt + \sum_{j=1}^d \sigma_{ij} dW_t^j \right),$$

where $\mu = (\mu_1, \dots, \mu_d)^\top \in \mathbb{R}^d$ and $\sigma = (\sigma_{ij}) \in \mathbb{R}_+^{d \times d}$ is assumed to be non-singular and we set $S_0^i = 1$.

In the first part of the paper we assume that the drift-vector μ is not known to the investor, however, the initial distribution Q_0 of μ is known. We choose a d -variate normal distribution $Q_0 := \mathcal{N}_d(\mu_0, \Sigma_0)$ and assume the drift to be independent of W . By $\mathcal{F}^S := (\sigma(S_u, 0 \leq u \leq t))_{t \in [0, T]}$ we denote the filtration obtained by observing the stock prices $S_t = (S_t^1, \dots, S_t^d)^\top$, a sub-filtration of

$$\mathcal{F}^{\mu, S} := (\sigma(\mu, S_u, 0 \leq u \leq t))_{t \in [0, T]} = (\sigma(\mu, W_u, 0 \leq u \leq t))_{t \in [0, T]}.$$

3. CONTINUOUS-TIME OPTIMIZATION

In a continuous-time setup, an investor with initial capital $x > 0$ chooses a trading strategy $\pi = (\pi_t)_{t \in [0, T]}$. We allow short-selling the stock and the bond, thus $\pi_t = (\pi_1(t), \dots, \pi_d(t))^\top \in \mathbb{R}^d$,

where $\pi_i(t)$ denotes the fraction of wealth invested in stock i at time t and $\mathbf{1} := (1, \dots, 1)^\top \in \mathbb{R}^d$. The wealth process satisfies the stochastic differential equation

$$dX_t^\pi = X_t^\pi \left(r + \pi_t^\top (\mu - r\mathbf{1}) \right) dt + X_t^\pi \pi_t^\top \sigma dW_t, \quad X_0^\pi = x. \quad (3.1)$$

We only consider *admissible* trading strategies, i.e., strategies from the set

$$\mathcal{A} := \{ \pi : [0, T] \times \Omega \rightarrow \mathbb{R}^d : \pi \text{ is measurable, self-financing, } \mathcal{F}^S\text{-adapted} \\ \text{and satisfies } \int_0^T \|\pi(s)\|^2 ds < \infty \text{ a.s.} \}.$$

The investor aims to maximize her expected logarithmic utility of terminal wealth. We define

$$J_\pi(x) := \mathbb{E}_x \log (X_T^\pi),$$

and the aim is to find

$$J(x) := \sup_{\pi \in \mathcal{A}} J_\pi(x), \quad (3.2)$$

where \mathbb{E}_x is the conditional expectation given $X_0 = x$. Note that π in (3.1) has to be adapted w.r.t. \mathcal{F}^S , but not w.r.t. $\mathcal{F}^{S,\mu}$ which makes the problem “non-standard”. However, it is well-known that such a problem can be solved by the *principle of estimation and control* which works as follows: First, note that for $t \in [0, T]$

$$\log S_t^i = \log S_0^i + \left(\mu_i - \frac{1}{2} \sum_{j=1}^d \sigma_{ij}^2 \right) t + \sum_{j=1}^d \sigma_{ij} W_t^j.$$

Hence, when we denote $Z_t := \mu t + \sigma W_t$ for $t \geq 0$ and $\sigma^i := (\sigma_{i1}, \dots, \sigma_{id})$, $i = 1, \dots, d$, we can write

$$\log S_t^i = \log S_0^i + Z_t^i - \frac{1}{2} \|\sigma^i\|^2 t$$

and obtain that $\mathcal{F}^S = \mathcal{F}^Z$, i.e., the filtration generated by the stock prices is the same as the filtration generated by the process $Z = (Z_t)$. Next, we define for $t \in [0, T]$ the estimator

$$\hat{\mu}_t := \mathbb{E} [\mu | \mathcal{F}_t^S],$$

which is the conditional expectation of μ given the stock prices until time t . The evolution of $(\hat{\mu}_t)$ is given by the so-called *Kalman-Bucy filter* on the “signal” μ and “observation” $\sigma^{-1} Z_t$. When we define $\Sigma := (\sigma \sigma^\top)$, then it can be shown that

$$\hat{\mu}_t = (\Sigma_0^{-1} + t \Sigma^{-1})^{-1} (\Sigma_0^{-1} \mu_0 + \Sigma^{-1} Z_t). \quad (3.3)$$

More details on the derivation can be found in Bain & Crisan (2009) and Fristedt et al. (2007). We can now introduce the innovation process (V_t) by

$$dV_t = \sigma^{-1} (dZ_t - \hat{\mu}_t dt).$$

Using Lévy’s characterization of the Brownian motion it is possible to show that (V_t) is an \mathcal{F}^S -Brownian motion under the given probability measure \mathbb{P} . Moreover, the stochastic differential equation for the wealth can be rewritten as

$$\begin{aligned} dX_t^\pi &= X_t^\pi \left(r + \pi_t^\top (\mu - r\mathbf{1}) \right) dt + X_t^\pi \pi_t^\top \sigma dW_t \\ &= X_t^\pi \left(r + \pi_t^\top (\hat{\mu}_t - r\mathbf{1}) \right) dt + X_t^\pi \pi_t^\top \sigma dV_t, \end{aligned} \quad (3.4)$$

which reduces the stochastic control problem with partial observation to one with complete observation, since all processes in (3.4) are \mathcal{F}^S -adapted. An explicit solution of (3.4) is given by

$$X_t^\pi = x \exp \left(\int_0^t \left(r + \pi_s^\top (\hat{\mu}_s - r\mathbf{1}) - \frac{1}{2} \pi_s^\top \Sigma \pi_s \right) ds + \int_0^t \pi_s^\top \sigma dV_s \right)$$

and, given that the local martingale $\int \pi^\top \sigma dV$ is a true martingale, we obtain for an investor with logarithmic utility that

$$J_\pi(x) = \mathbb{E}_x \log X_T^\pi = \log(x) + \mathbb{E}_x \left[\int_0^T \left(r + \pi_s^\top (\hat{\mu}_s - r\mathbf{1}) - \frac{1}{2} \pi_s^\top \Sigma \pi_s ds \right) \right]. \quad (3.5)$$

It is now easy to see and has e.g. been shown by Karatzas & Zhao (2001) that the simple strategy

$$\pi_t^* = \Sigma^{-1}(\hat{\mu}_t - r\mathbf{1}) \quad (3.6)$$

is an optimal portfolio strategy for the investor with logarithmic utility.

When short-selling is not allowed, we maximize (3.5) over all π which are elements of

$$\underline{\mathcal{A}} := \{\pi \in \mathcal{A} \mid \pi_t \in D \ \forall t \in [0, T]\}, \text{ where } D := \{x \in [0, 1]^d : x^\top \mathbf{1} \leq 1\}.$$

I.e., we consider the value function

$$J^{nss}(x) := \sup_{\pi \in \underline{\mathcal{A}}} \mathbb{E} \log(X_T^\pi),$$

where *nss* denotes *no short-selling*. The optimal portfolio strategy is here denoted by $\underline{\pi} = (\underline{\pi}_t)$, where $\underline{\pi}_t$ is the (unique) solution of

$$\begin{cases} a^\top (\hat{\mu}_t - r\mathbf{1}) - \frac{1}{2} a^\top \Sigma a \rightarrow \max \\ a \in D. \end{cases} \quad (3.7)$$

In dimension $d = 1$, if $\pi_t^* \notin [0, 1]$, we simply have to shift π_t^* from (3.6) to the nearest boundary of the interval $[0, 1]$. For $d \geq 2$ this is in general not true. For a comprehensive treatment of constraint portfolio problems see Cvitanić & Karatzas (1992) and Karatzas & Zhao (2001). Also note that the pointwise maximization in (3.7) leads to an \mathcal{F}^S -progressively measurable strategy (cf. Cvitanić & Karatzas (1992)).

4. THE h -OPTIMIZATION PROBLEM

Here we suppose that the investor observes the stock prices only at discrete times which are multiples of $h > 0$ and rebalances her portfolio only at these time points. We refer to this investor as the h -investor. To simplify the presentation we require that $N := T/h \in \mathbb{N}$. We define

$$\tilde{R}_n^i := \exp \left(\mu_i h - \frac{1}{2} \sum_{j=1}^d \sigma_{ij}^2 h + \sum_{j=1}^d \sigma_{ij} (W_{nh}^j - W_{(n-1)h}^j) \right) \quad (4.1)$$

for $n = 1, \dots, N$ and $i = 1, \dots, d$ and obtain $S_{nh}^i = \prod_{k=1}^n \tilde{R}_k^i$. In what follows we write

$$\begin{aligned} R_n^i &:= \frac{\tilde{R}_n^i}{e^{rnh}} - 1, \\ \tilde{Z}_n^i &:= \mu_i h + \sum_{j=1}^d \sigma_{ij} (W_{nh}^j - W_{(n-1)h}^j) \stackrel{d}{=} \mu_i h + \sum_{j=1}^d \sigma_{ij} \varepsilon_n^j \sqrt{h}, \end{aligned}$$

where $\varepsilon_n = (\varepsilon_n^1, \dots, \varepsilon_n^d)^\top$ and $\varepsilon_1, \varepsilon_2, \dots, \varepsilon_N$ are i.i.d. random vectors from a multivariate normal distribution $\mathcal{N}_d(0, I)$. Thus, we have $\tilde{Z}_n = (\tilde{Z}_n^1, \dots, \tilde{Z}_n^d)^\top \stackrel{d}{=} \mu h + \sigma \varepsilon_n \sqrt{h}$, where $\stackrel{d}{=}$ denotes equality in distribution. The conditional distribution of \tilde{Z}_n given $\mu = m$ is denoted by $Q^{\tilde{Z}}(\cdot | m)$. Obviously, we have $Q^{\tilde{Z}}(\cdot | m) = \mathcal{N}_d(mh, \Sigma h)$.

Let $\mathcal{F}^R := (\sigma(R_1, \dots, R_n))_{n \in \{0, 1, \dots, N\}}$ be the filtration generated by (R_1, \dots, R_n) . The portfolio strategies for the h -investor are now \mathcal{F}^R -adapted stochastic processes $\alpha = (\alpha_n) = ((\alpha_n^1, \dots, \alpha_n^d)^\top)_{n \in \{0, \dots, N-1\}}$. In what follows the interpretation of α_n^i will always be the fraction

of wealth invested in the i -th stock at time step n . Once the h -investor has chosen a portfolio strategy, the wealth at time nh under this strategy satisfies

$$X_{hn}^\alpha = e^{rh} X_{h(n-1)}^\alpha (1 + \alpha_{n-1}^\top R_n).$$

The set of strategies of the h -investor is given by

$$\mathcal{A}^h := \{\alpha = (\alpha_n)_{n \in \{0,1,\dots,N-1\}} \mid \alpha_n : \Omega \rightarrow \mathbb{R}^d, \alpha_n \text{ is } \mathcal{F}_n^R \text{-measurable } \forall n, \alpha \text{ is self-financing}\}.$$

Note that we set $\log(x) := -\infty$ if $x \leq 0$. With this agreement it is obvious that an optimal portfolio strategy is found among the strategies α with $\alpha_n \in D$.

The h -investor's value function is then defined by

$$J_0^h(x) := \sup_{\alpha \in \mathcal{A}^h} \mathbb{E}_x \log(X_T^\alpha), \quad (4.2)$$

where $x > 0$ is the initial wealth. It is now well-known that problem (4.2) can be solved by using the theory of (Bayesian) Markov Decision Problems, see e.g. (Bäuerle & Rieder, 2011, Chapter 5). In order to set up the filtered Markov Decision Problem we have to enlarge the state space by the conditional distribution of the unknown parameter given the observation of the stock prices up to this time. If we denote the conditional distribution of μ given the observations $(\tilde{z}_1, \dots, \tilde{z}_n)$ by $\mu(\cdot | \tilde{z}_1, \dots, \tilde{z}_n)$, it is a classical result of statistics, see e.g. (DeGroot, 1970, p. 175) that $Q_0 = \mathcal{N}_d(\mu_0, \Sigma_0)$ is a conjugate prior distribution to $Q^{\tilde{Z}}(\cdot | m) = \mathcal{N}_d(mh, \Sigma h)$ and we obtain with $\tilde{z}_n := \frac{1}{n} \sum_{k=1}^n \tilde{z}_k$:

$$\mu(\cdot | \tilde{z}_1, \dots, \tilde{z}_n) = \mathcal{N}_d\left(\left(\Sigma_0^{-1} + nh\Sigma^{-1}\right)^{-1} \left(\Sigma_0^{-1} \mu_0 + n\Sigma^{-1} \tilde{z}_n\right), \left(\Sigma_0^{-1} + hn\Sigma^{-1}\right)^{-1}\right).$$

Obviously, $(\sum_{k=1}^n \tilde{z}_k, n)$ is a sufficient statistic for $\mu(\cdot | \tilde{z}_1, \dots, \tilde{z}_n)$. Thus we will write $\mu(\cdot | s, n)$ where $s = \sum_{k=1}^n \tilde{z}_k$. For $n = 0, \dots, N$ and a portfolio strategy $\alpha = (\alpha_n) \in \mathcal{A}^h$, we define the n -th step value function

$$\begin{aligned} J_{n,\alpha}^h(x, s) &:= \mathbb{E} \left[\log(X_{Nn}^\alpha) \mid X_{nh} = x, \sum_{k=1}^n \tilde{Z}_k = s \right], \\ J_n^h(x, s) &:= \sup_{\alpha \in \mathcal{A}^h} J_{n,\alpha}^h(x, s). \end{aligned} \quad (4.3)$$

Note that the probability measure underlying the expectation in (4.3) is induced by Q_0 and the transition probabilities of the Markov Decision Process $(X_{nh}, \sum_{k=1}^n \tilde{Z}_k)$, i.e., it is essentially induced by Q_0 and $Q^{\tilde{Z}}(\cdot | m)$.

Problem (4.2) can now be solved recursively via the Bellman equation

$$J_n^h(x, s) = \sup_{a \in D} \left\{ \iint J_{n+1}^h \left(e^{rh} x \left(1 + a^\top \hat{z} \right), (s + \tilde{z}) \right) Q^{\tilde{Z}}(d\tilde{z} | m) \mu(dm | s, n) \right\},$$

where $\hat{z}^i := \exp(\tilde{z}^i - \frac{1}{2} \|\sigma^i\|^2 h - rh) - 1$.

Theorem 4.1. a) *The value function for the h -investor is for $n = 0, 1, \dots, N$ given by*

$$J_n^h(x, s) = \log(x) + d_n(s), \quad (x, s) \in \mathbb{R}_{>0} \times \mathbb{R},$$

where the d_n satisfy the following recursion

$$\begin{aligned} d_N(s) &:= 0, \\ d_n(s) &= rh + \sup_{a \in D} \left\{ \int \int \log(1 + a^\top \hat{z}) Q^{\tilde{Z}}(d\tilde{z} | m) \mu(dm | s, n) \right\} \\ &\quad + \int \int d_{n+1}(s + \tilde{z}) Q^{\tilde{Z}}(d\tilde{z} | m) \mu(dm | s, n). \end{aligned} \quad (4.4)$$

The value function of the problem is given by $J_0^h(x) = J_0^h(x, 0)$ where $\mu(\cdot | 0, 0) = Q_0(\cdot)$.

- b) *The optimal amounts which are invested in the stocks at time step n when $X_{nh} = x$ and $\tilde{Z}_1 + \dots + \tilde{Z}_n = s$ are given by*

$$f_n^*(x, s) = \alpha_n^*(s)x, \quad x > 0,$$

where $\alpha_n^*(s)$ is the maximizer of (4.4) and denotes the optimal proportions. The optimal portfolio strategy is constructed through the application of the maximizers.

This can be proved by induction along the lines of (Bäuerle & Rieder, 2011, Theorem 6.1.1) and we therefore skip the proof.

The optimal strategy is given by the maximum points in (4.4) and denoted by α^* but is not available in closed form. In the next section we propose a portfolio strategy which is suboptimal however “simple” and performs for $h \rightarrow 0$ asymptotically as well as the optimal strategy.

5. COMPARISON BETWEEN THE CONTINUOUS AND THE h -INVESTOR

In this section we compare the performance of the continuous investor with the performance of the h -investor.

5.1. Comparison of Information. First we compare the relevant information of both investors. The continuous investor uses the estimator $(\hat{\mu}_t)$ whereas the h -investor uses $(\mu(\cdot|\tilde{z}_1, \dots, \tilde{z}_n))$. We obtain for $t = nh$, that

$$\hat{\mu}_t = (\Sigma_0^{-1} + t\Sigma^{-1})^{-1}(\Sigma_0^{-1}\mu_0 + \Sigma^{-1}Z_t) = \int m\mu\left(dm \middle| \sum_{k=1}^n \tilde{Z}_k, \frac{t}{h}\right).$$

Thus, the h -investor has the same estimator at the rebalancing points as the continuous investor. For general results concerning discretization of filter processes see James et al. (1996).

Let us now discuss the different results for the situations with and without short-selling constraints.

5.2. No Short-selling. We consider first the problem where the continuous investor is not allowed to short sell the stocks or the bond. Obviously, we have $J_0^h(x) \leq J^{nss}(x)$. On the other hand consider the portfolio strategy $\underline{\alpha} = (\alpha_n)_{n \in \{0,1,\dots,N-1\}}$ given by $\alpha_n := \underline{\pi}_{nh}$, recalling the optimal strategy $\underline{\pi}$ for the continuous investor obliging short-selling constraints. It is admissible for the h -investor since it depends on the history only through $\hat{\mu}_{nh}$. We obtain, that for $h \rightarrow 0$ the expected log-utility of the h -investor using the optimal strategy of the continuous investor with short-selling constraints only at discrete points in time, converges to the expected log-utility of the continuous investor with short-selling constraints:

Theorem 5.1. *With $\underline{\pi}$ and $\underline{\alpha}$ as above it holds*

$$\lim_{h \rightarrow 0} J_{0,\underline{\alpha}}^h(x, 0) = J_{\underline{\pi}}^{nss}(x) = J^{nss}(x), \quad x > 0.$$

The proof is given in the Appendix. Theorem 5.1 immediately leads to the following corollary:

Corollary 5.2. *For $x > 0$ we obtain: $\lim_{h \rightarrow 0} J_0^h(x) = J^{nss}(x)$.*

We see that the h -investor can perform almost as well as the continuous investor with short selling constraints with an error which tends to zero if $h \rightarrow 0$. This is in line with the results in Rogers (2001).

The optimal portfolio strategy for the h -investor is complicated and the solution of problem (4.4) has to be computed numerically. However, it can be shown that for small h the optimal portfolio strategy is close to the discretization of $(\underline{\pi}_t)$. To show this let again $\alpha^* = (\alpha_n^*(s))$ be the optimal strategy for the h -investor, i.e., the solution of problem (4.4) and define for $t \in [0, T]$ the process

$$\underline{\pi}_t^h := \alpha_{\lfloor \frac{t}{h} \rfloor}^* \left(\sum_{k=1}^{\lfloor \frac{t}{h} \rfloor} \tilde{Z}_k \right). \quad (5.1)$$

This is an admissible strategy for the continuous investor and at time points $t = nh$, π_t^h corresponds to the optimal strategy of the h -investor. This piecewise constant strategy converges a.s. to the optimal strategy of the continuous investor with short-selling constraints:

Theorem 5.3. *With the preceding definitions it holds*

$$\lim_{h \rightarrow 0} \pi_t^h = \pi_t \quad \mathbb{P} - a.s., \quad t \in [0, T].$$

Again, the proof can be found in the Appendix.

5.3. With Short-selling. Let us now return to the original problem (3.2) where the continuous investor can do short-selling. As outlined before we cannot allow the h -investor to short-sell assets (this would lead to an expected utility of $-\infty$) which implies that a simple discretization of the optimal portfolio strategy (3.6) in the continuous setting is not admissible for the h -investor. Corollary 5.4 shows that the h -investor cannot in general do as well as the continuous investor, even if $h \rightarrow 0$. There still remains a gap between the maximal expected utilities.

Corollary 5.4. *The h -investor faces a discretization gap in the sense that*

$$\lim_{h \rightarrow 0} \left(J(x) - J_0^h(x) \right) = J(x) - J^{nss}(x) > 0, \quad x > 0.$$

In the case of an unknown drift for the stock price and no short-selling constraints, the situation which has been excluded in Rogers (2001), namely that $\pi_t^* \notin [0, 1]^d$ cannot be avoided here. Even if the true ratio $\Sigma^{-1}(\mu - r\mathbf{1})$ lies within $[0, 1]^d$, there is no guarantee that its estimator lies in $[0, 1]^d$ as well.

Hence, we encounter a somewhat surprising result where we do *not* have convergence of the maximal expected utility of the h -investor to the maximal expected utility of the continuous investor. The size of the discretization gap is discussed in Section 7.

6. UNKNOWN VOLATILITY

So far we have assumed that the drift μ is unknown and the volatility matrix is known. In this section we consider the reverse situation where the drift μ is known but the volatility matrix is unknown.

More precisely, the investor cannot observe $\Sigma^{-1} = (\sigma\sigma^\top)^{-1}$ directly but knows the initial distribution Q_0^{vol} of Σ^{-1} . We choose here a d -dimensional *Wishart-distribution* with ν degrees of freedom and parameter matrix Σ_0^{-1} , denoted by $Q_0^{vol} = \mathcal{W}_d(\nu, \Sigma_0^{-1})$, i.e., there exist i.i.d. vectors $A_i \sim \mathcal{N}_d(0, \Sigma_0^{-1})$, $i = 1, \dots, \nu$, such that $\Sigma^{-1} \stackrel{d}{=} \sum_{i=1}^{\nu} A_i A_i^\top$. The Wishart-distribution is a distribution on the set of all symmetric, positive definite matrices with dimension d and is a multivariate generalization of the χ^2 -distribution.

We denote by $\det(\cdot)$ the determinant and by $\text{tr}(\cdot)$ the trace of a matrix. Then the density of the d -dimensional Wishart distribution with ν degrees of freedom and parameter matrix Σ_0^{-1} is, see e.g. (DeGroot, 1970, p. 57),

$$q_0^{vol}(\tau) := c(\det(\Sigma_0^{-1}))^{-\frac{\nu}{2}} (\det(\tau))^{\frac{\nu-d-1}{2}} \exp\left(-\frac{1}{2}\text{tr}(\Sigma_0\tau)\right), \quad (6.1)$$

where c is the normalizing constant.

6.1. The Continuous Investor with Unknown Volatility. Even if we assume that the continuous investor does not know $\sigma\sigma^\top$, then she is able to estimate it perfectly from a tiny observation of the stock prices. To explain this, consider the quadratic covariation of $\log S_t$ which is given by

$$\langle \log S_t^i, \log S_t^j \rangle_t = \sum_{k=1}^d \sigma_{ik} \sigma_{jk} t = \sigma^i (\sigma^j)^\top t = \Sigma_{ij} t.$$

This quadratic covariation can be observed since for $h \rightarrow 0$ it is the limit in probability of

$$\sum_{k:kh \leq t} (\log S_{hk}^i - \log S_{h(k-1)}^i)(\log S_{hk}^j - \log S_{h(k-1)}^j).$$

Note that convergence in probability implies that there exists a subsequence which converges \mathbb{P} -a.s.. The optimization problem for the continuous investor is therefore

$$J^{vol}(x) := \sup_{\pi \in \mathcal{A}} \mathbb{E}_x \log(X_T^\pi) = \log(x) + \sup_{\pi} \mathbb{E}_x \left[\int_0^T r + \pi_s^\top (\mu - r\mathbf{1}) - \frac{1}{2} \pi_s^\top \Sigma \pi_s ds + \int_0^T \pi_s^\top \sigma dW_s \right].$$

One can easily see that in the case where short-selling is allowed, the optimal strategy denoted by $\pi^{vol} = (\pi_t^{vol})$ is constant in time and given by $\pi_t^{vol} := \Sigma^{-1}(\mu - r\mathbf{1}) \forall t > 0$. At $t = 0$, the value of Σ^{-1} is not yet known, but it makes no difference at all, which value one chooses and hence $\pi_0^{vol} := \tilde{c} \in \mathbb{R}$. The maximal expected utility is

$$\begin{aligned} J^{vol}(x) &= \log(x) + rT + T \sup_{a \in \mathbb{R}^d} \mathbb{E}_x \left[a^\top (\mu - r\mathbf{1}) - \frac{1}{2} a^\top \Sigma a \right] \\ &= \log(x) + rT + \frac{T}{2} \mathbb{E}_x \left[(\mu - r\mathbf{1})^\top \Sigma^{-1} (\mu - r\mathbf{1}) \right] \\ &= \log(x) + rT + \frac{\nu T}{2} (\mu - r\mathbf{1})^\top \Sigma_0^{-1} (\mu - r\mathbf{1}), \end{aligned} \quad (6.2)$$

since for $\Sigma^{-1} \sim \mathcal{W}_d(\nu, \Sigma_0^{-1})$, the expectation is given by $\mathbb{E} \Sigma^{-1} = \nu \Sigma_0^{-1}$ and for a vector $a \in \mathbb{R}^d$ one obtains $a^\top \Sigma^{-1} a \sim \mathcal{W}_1(\nu, a^\top \Sigma_0^{-1} a)$, see e.g. (DeGroot, 1970, p. 58).

If short-selling is excluded, the optimal strategy, denoted by $\underline{\pi}^{vol} = (\underline{\pi}_t^{vol})$, is still constant in time and given by the solution of

$$\begin{cases} a^\top (\mu - r\mathbf{1}) - \frac{1}{2} a^\top \Sigma a \rightarrow \max \\ a \in D. \end{cases} \quad (6.3)$$

We denote by

$$J^{vol, nss}(x) := \sup_{\pi \in \mathcal{A}} \mathbb{E}_x \log(X_T^\pi)$$

the corresponding maximal expected logarithmic utility.

We will see in the next section that the h -investor loses relevant information in this setting and is unable to estimate the volatility as perfectly as the continuous investor.

6.2. The h -Investor with Unknown Volatility. We consider now the h -investor in the market with unknown volatility matrix and known drift vector μ . We again require that $N = T/h \in \mathbb{N}$. Since the h -investor is only able to observe the stock prices at time points $\{0, h, \dots, T\}$, she observes the random variable

$$\begin{aligned} Y_n^i &:= \log(\tilde{R}_n^i) = \mu_i h - \frac{1}{2} \sum_{j=1}^d \sigma_{ij}^2 h + \sum_{j=1}^d \sigma_{ij} (W_{nh}^j - W_{(n-1)h}^j) \\ &\stackrel{d}{=} \mu_i h - \frac{1}{2} \sum_{j=1}^d \sigma_{ij}^2 h + \sum_{j=1}^d \sigma_{ij} \varepsilon_n^j \sqrt{h}, \end{aligned}$$

where we used the notation from Section 4 with $n = 1, \dots, N, i = 1, \dots, d$ and $\varepsilon_1, \varepsilon_2, \dots, \varepsilon_N$ are i.i.d. $\mathcal{N}_d(0, I)$ -distributed random vectors. The h -investor's value function is then defined by

$$J_0^{h, vol}(x) := \sup_{\alpha \in \mathcal{A}^h} \mathbb{E}_x \log(X_T^\alpha), \quad (6.4)$$

where the wealth X_T^α is again determined recursively in terms of the Y_n .

We are interested in the posterior distribution of the unknown matrix Σ^{-1} . The conditional distribution of the vector Y_n given $\sigma\sigma^\top = \tau$ is denoted by $Q^Y(\cdot|\tau)$ and clearly

$$Q^Y(\cdot|\tau) = \mathcal{N}_d(m(\tau), \tau h), \quad m(\tau) := \left(\mu - \frac{1}{2}(\tau_{11} \dots, \tau_{dd})^\top\right)h.$$

Given observations (y_1, \dots, y_n) , the corresponding likelihood function is therefore

$$l(y_1, \dots, y_n|\tau^{-1}) := \frac{1}{((2\pi)^d \det(h\tau))^{n/2}} \exp\left(-\frac{1}{2h} \sum_{i=1}^n (y_i - m(\tau))^\top \tau^{-1} (y_i - m(\tau))\right).$$

The exponential can be simplified by taking the trace:

$$\begin{aligned} \sum_{i=1}^n (y_i - m(\tau))^\top \tau^{-1} (y_i - m(\tau)) &= \text{tr} \left(\sum_{i=1}^n (y_i - m(\tau))^\top \tau^{-1} (y_i - m(\tau)) \right) \\ &= \text{tr} \left(\sum_{i=1}^n (y_i - m(\tau)) (y_i - m(\tau))^\top \tau^{-1} \right) = \text{tr} \left(\left(\sum_{i=1}^n y_i y_i^\top - 2 \sum_{i=1}^n y_i m(\tau)^\top + n m(\tau) m(\tau)^\top \right) \tau^{-1} \right). \end{aligned}$$

Hence we see that $(\sum_{i=1}^n y_i y_i^\top, \sum_{i=1}^n y_i, n)$ is a sufficient statistic for (y_1, \dots, y_n) . The posterior density of Σ^{-1} given (y_1, \dots, y_n) is then proportional to $l(y_1, \dots, y_n|\Sigma^{-1})q_0^{\text{vol}}(\Sigma^{-1})$, where the prior density q_0^{vol} was specified in (6.1). We denote by $\Sigma^{-1}(\cdot|y_1, \dots, y_n)$ the posterior distribution of Σ^{-1} given (y_1, \dots, y_n) . We will also use the notation $\Sigma^{-1}(\cdot|u, v, n)$ where $u = \sum_{i=1}^n y_i y_i^\top$, $v = \sum_{i=1}^n y_i$.

For $n = 0, \dots, N$ and a portfolio strategy $(\alpha_n) \in \mathcal{A}^h$ we can now define the n -th step value function with unknown volatility similarly to Section 4 by

$$\begin{aligned} J_{n,\alpha}^{h,\text{vol}}(x, (u, v)) &:= \mathbb{E} \left[\log(X_{Nn}^\alpha) \mid X_{nh} = x, \sum_{i=1}^n Y_i Y_i^\top = u, \sum_{i=1}^n Y_i = v \right], \\ J_n^{h,\text{vol}}(x, (u, v)) &:= \sup_{\alpha \in \mathcal{A}^h} J_{n,\alpha}^{h,\text{vol}}(x, (u, v)). \end{aligned}$$

The value function can again be determined recursively:

$$J_n^{h,\text{vol}}(x, (u, v)) = \sup_{a \in D} \left\{ \iint J_{n+1}^{h,\text{vol}} \left(e^{rh} x \left(1 + a^\top \hat{z} \right), (u + yy^\top, v + y) \right) Q^Y(dy|\tau^{-1}) \Sigma^{-1}(d\tau|u, v, n) \right\},$$

where $\hat{z}^i = \exp(y^i - hr) - 1$. The next theorem follows again from (Bäuerle & Rieder, 2011, Th. 6.1.1) and we will therefore skip the proof.

Theorem 6.1. a) *The value function for the h -investor is given by*

$$J_n^{h,\text{vol}}(x, (u, v)) = \log(x) + d_n^{\text{vol}}(u, v), \quad (u, v) \in \mathbb{R}_+^{d \times d} \times \mathbb{R}^d, n \in \mathbb{N}, x > 0,$$

where the d_n^{vol} satisfy the following recursion

$$\begin{aligned} d_N^{\text{vol}}(u, v) &:= 0, \\ d_n^{\text{vol}}(u, v) &= rh + \sup_{a \in D} \left\{ \int \int \log(1 + a^\top \hat{z}) Q^Y(dy|\tau^{-1}) \Sigma^{-1}(d\tau|u, v, n) \right\} \\ &\quad + \int \int d_{n+1}^{\text{vol}}(u + yy^\top, v + y) Q^Y(dy|\tau^{-1}) \Sigma^{-1}(d\tau|u, v, n). \end{aligned} \tag{6.5}$$

The value function of the problem is given by $J_0^{h,\text{vol}}(x) = J_0^{h,\text{vol}}(x, (0, 0))$, where $\Sigma^{-1}(\cdot|0, 0, 0) = Q_0^{\text{vol}}(\cdot)$.

b) *The optimal amounts which are invested in the stocks are given by*

$$f_n^{\text{vol}}(x, (u, v)) = \alpha_n^{\text{vol}}(u, v)x, \quad x \geq 0,$$

where $\alpha_n^{\text{vol}}(u, v)$ is the maximizer of (6.5) and denotes the optimal proportions. The optimal portfolio strategy is constructed through the application of the maximizers.

6.3. Comparison between Continuous and h -Investor with Unknown Volatility.

6.3.1. *Comparison of Information.* As we have discussed before, the continuous investor can estimate Σ perfectly from a very short observation of the stock price and from this time onwards works under full information. For the h -investor, however, this is not the case. The h -investor updates her estimator of Σ at discrete time points but does not work under full information.

This is a remarkable contrast to the situation where the drift was unknown and where the h -investor did not suffer a loss in information compared to the continuous investor.

6.3.2. *A Simple and Almost Optimal Strategy for the h -Investor.* We have seen before that the optimal strategy of the h -investor can theoretically be described, however the practical implementation is difficult. So it is reasonable to look for an admissible “simple” strategy which is almost optimal. In contrast to the case of unknown drift, however, the h -investor is unable to implement the strategy of the continuous investor (in discrete-time) since the continuous investor works almost immediately under full information.

To simplify the presentation we assume that $T = 2\tilde{N}h$, $\tilde{N} \in \mathbb{N}$, i.e., $N = 2\tilde{N}$. We propose now to update the strategy at even multiples of h only and use the difference of the two observations during the time interval $((2n - 2)h, 2nh]$ given by

$$\tilde{Y}_n := Y_{2n} - Y_{2n-1} = \sigma(\varepsilon_{2n} - \varepsilon_{2n-1})\sqrt{h}$$

for $n = 1, \dots, \tilde{N}$. The advantage of doing this is that $\tilde{Y}_1, \dots, \tilde{Y}_{\tilde{N}}$ are now i.i.d. and obviously the distribution of \tilde{Y}_n given $\sigma\sigma^\top = \tau$ is $Q^{\tilde{Y}}(\cdot|\tau) := \mathcal{N}_d(0, 2\tau h)$. Now we are in a classical situation with normally distributed observations with unknown covariance matrix. In this case the Wishart-distribution is a conjugate prior and hence, if we assume that the initial distribution of $\Sigma^{-1} = (\sigma\sigma^\top)^{-1}$ is $Q_0^{vol} = \mathcal{W}_d(\nu, \Sigma_0^{-1})$, then the conditional distribution of Σ^{-1} given the observations $(\tilde{y}_1, \dots, \tilde{y}_n)$ is also a Wishart-distribution, namely

$$\mathcal{W}_d\left(\nu + n, \left(\Sigma_0 + \frac{1}{2h} \sum_{k=1}^n \tilde{y}_k \tilde{y}_k^\top\right)^{-1}\right),$$

see (DeGroot, 1970, p.176–177). Obviously, $(\sum_{k=1}^n \tilde{y}_k \tilde{y}_k^\top, n)$ is a sufficient statistic for Σ^{-1} . Moreover, the conditional expectation of Σ^{-1} given $\sum_{k=1}^n \tilde{Y}_k \tilde{Y}_k^\top = s$ is

$$\widehat{\Sigma^{-1, h}}(s, n) := \mathbb{E}\left[\Sigma^{-1} \mid \sum_{k=1}^n \tilde{Y}_k = s\right] = (\nu + n)\left(\Sigma_0 + \frac{1}{2h}s\right)^{-1} \quad (6.6)$$

for $n = 1, \dots, \tilde{N}$. For $n = 0$ we set $\widehat{\Sigma^{-1, h}}(s, 0) := \nu\Sigma_0^{-1}$.

Similarly, we can also derive the conditional expectation of Σ given $\sum_{k=1}^n \tilde{Y}_k \tilde{Y}_k^\top = s$, see e.g. (von Rosen, 1988, Theorem 3.1):

$$\hat{\Sigma}^h(s, n) := \mathbb{E}\left[\Sigma \mid \sum_{k=1}^n \tilde{Y}_k = s\right] = \frac{1}{\nu + n - d - 1}\left(\Sigma_0 + \frac{1}{2h}s\right)$$

for $n = 1, \dots, \tilde{N}$ and $\hat{\Sigma}^h(s, 0) := \frac{1}{\nu - d - 1}\Sigma_0$. This conditional expectation can be used to define an admissible strategy for the h -investor by replacing Σ with $\hat{\Sigma}^h$ in the optimal strategy for the continuous investor.

6.3.3. *No Short-selling.* Consider now the optimization problem

$$\begin{cases} a^\top(\mu - r\mathbf{1}) - \frac{1}{2}a^\top \hat{\Sigma}^h(s, n)a \rightarrow \max \\ a \in D \end{cases} \quad (6.7)$$

and denote the maximizer by $\alpha_n^{h, vol}(s)$. We define a portfolio strategy by

$$\underline{\pi}_t^{h, vol} := \alpha_{\lfloor \frac{t}{2h} \rfloor}^{h, vol}\left(\sum_{k=1}^{\lfloor \frac{t}{2h} \rfloor} \tilde{Y}_k \tilde{Y}_k^\top\right).$$

When we consider $(\underline{\pi}_t^{h,vol})$ only at time points nh , then we obtain an admissible strategy for the h -investor. For general $t \in [0, T]$ it is admissible for the continuous investor and converges to the optimal strategy of the continuous investor with no short-selling restriction:

Theorem 6.2. *For $t \in [0, T]$ we obtain $\lim_{h \rightarrow 0} \underline{\pi}_t^{h,vol} = \underline{\pi}^{vol}$.*

The proof can be found in the Appendix.

Finally, when we denote by $\underline{\alpha}^{h,vol} = (\underline{\alpha}^{h,vol})$ the discretization of $\underline{\pi}_t^{h,vol}$, we obtain the following theorem.

Theorem 6.3. *For $x > 0$ we obtain $\lim_{h \rightarrow 0} J_{0, \underline{\alpha}^{h,vol}}^{h,vol}(x, (0, 0)) = J^{vol, nss}(x)$.*

A proof is given in the Appendix. The theorem tells us that for small time steps h , the h -investor can choose the simpler strategy $\underline{\alpha}^{h,vol}$ which is close to being optimal.

Again we obtain convergence of the h -investor's value function to the value function of the continuous investor with short-selling restriction as $h \rightarrow 0$:

Corollary 6.4. *For $x > 0$ we obtain $\lim_{h \rightarrow 0} J_0^{h,vol}(x) = J^{vol, nss}(x)$.*

We give a short proof in the Appendix.

6.3.4. *With Short-selling.* If we allow short-selling for the continuous investor we again observe a discretization gap between the expected utilities of the h -investor and the continuous investor:

Corollary 6.5. *For $x > 0$ we obtain $\lim_{h \rightarrow 0} (J^{vol}(x) - J_0^{h,vol}(x)) = J^{vol}(x) - J^{vol, nss}(x) > 0$.*

7. NUMERICAL EXAMPLES

In this section we illustrate our theoretical results by some numerical examples and investigate the speed of convergence of some of our convergence results. We also give examples for the discretization gap. We consider a market with one stock, i.e., $d = 1$, time horizon $T = 1$ year and interest rate $r = \log(1.03)$. The investor's initial capital is set to $x = 1$.

7.1. **Unknown Drift.** We first consider the situation where the drift μ is unknown and sampled from a $\mathcal{N}_1(\mu_0, \sigma_0^2)$ -distribution with $\mu_0 = \log(1.05)$, $\sigma_0^2 = \frac{1}{3}\mu_0$ unless stated otherwise. The known volatility is chosen to be $\sigma = 0.3$.

The expected log-utility of the continuously trading investor is known analytically, see Karatzas & Zhao (2001), and for our choice of market parameters it is $J(1) = \mathbb{E}[\log(X_1^*)] \approx 0.0389$. This can be used as an upper bound on the expected utility of the continuously trading investor with short-selling constraints and on that of the h -investor.

In order to get an impression about the loss which is due to the incomplete information we derive the expected utility of an investor who knows μ and can therefore implement the strategy $(\mu - r)/\sigma^2$. This strategy is obviously not \mathcal{F}^S -adapted and therefore not admissible. The corresponding expected utility is known analytically and given by

$$J^{drift, known}(x) := \log(x) + rT + \frac{T}{2\sigma^2}(\sigma_0^2 + (\mu_0 - r)^2).$$

For our choice of parameters we obtain $J^{drift, known}(1) \approx 0.1220$. We see that the difference in expected utility between an investor with full and one with partial information is severe. In the following we focus on admissible strategies.

7.1.1. *Comparison of Expected Utilities.* For different values of h , we investigate the difference between the expected utility of an h -investor using her optimal strategy $J_0^h(1)$ and $J_{0, \underline{\alpha}}^h(1, 0)$ which is the expected utility of an h -investor using the discretization of the optimal solution of the continuous investor with short selling constraints (denoted by $\underline{\alpha}_n$ in Theorem 5.1). We computed the expected utilities by Monte Carlo methods and solved the optimization problem (4.4) by numerical optimization. We used 1.500.000 replications to derive $J_{0, \underline{\alpha}}^h(1, 0)$, 300.000 replications for $N = 2, 4, 6$ and 85.000 replications for $N = 8$ to derive $J_0^h(1)$. The smaller

TABLE 1. Results of MC simulations of the expected terminal utility of the h -investor for different values of $N = \frac{1}{h}$ using first the lower-bound strategy $\underline{\alpha}$ and second the optimal strategy α^* .

N	$J_{0,\underline{\alpha}}^h(1,0)$	$J_0^h(1)$
2	0.0351	0.0352
4	0.0361	0.0361
6	0.0363	0.0365
8	0.0365	0.0366

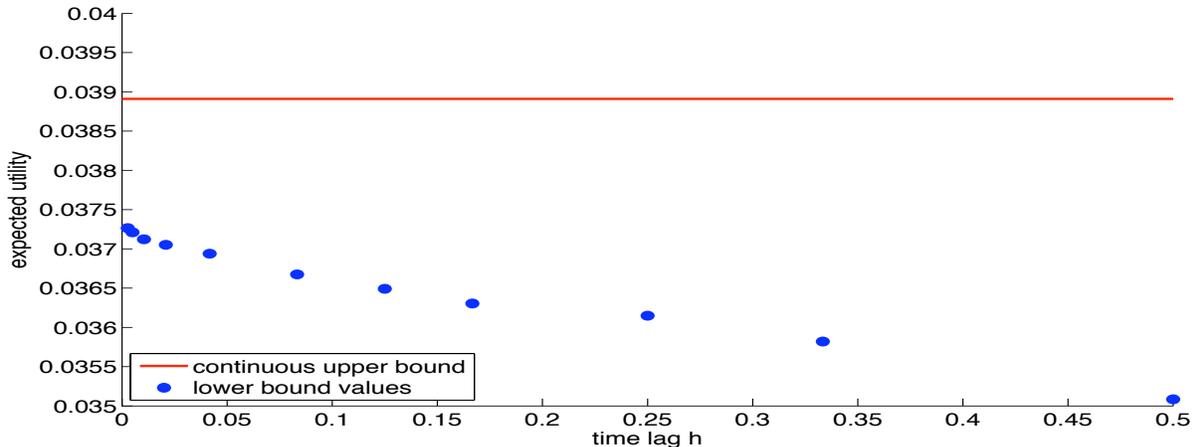


FIGURE 1. Expected utilities $J_{0,\underline{\alpha}}^h(1,0)$ (lower bound values) and $J(1)$ (continuous upper bound) for various time lags h .

number of replications in the latter case was necessary, since a numerical optimization has to be carried out at each time step, which resulted in a longer running time. The results are presented in Table 1. We observe that the expected utility increases with N , i.e., more frequent portfolio adjustments increase the expected utility. The simple strategy $\underline{\alpha}$ which just discretizes the optimal continuous solution with short-selling constraints results only in a slightly lower expected utility than using the optimal solution.

From Theorem 5.1 we know, that $\lim_{h \rightarrow 0} J_{0,\underline{\alpha}}^h(1,0) = J^{nss}(1)$ and obviously $J^{nss}(1) \leq J(1)$. Figure 1 shows $J_{0,\underline{\alpha}}^h(1,0)$ for different values of h and compares it to $J(1)$. We clearly find the discretization gap as stated in Corollary 5.4. The h -investor's value function only converges to $J^{nss}(1)$ and not to $J(1)$.

Figure 2 illustrates the sensitivity of the discretization gap to the time horizon T and the variance of the prior distribution. We plot $J(1)$ and $J^{nss}(1)$. One clearly finds a loss in terminal utility when having to obey short-selling restrictions, its magnitude increasing with a longer time horizon or a higher variance of the prior.

7.1.2. Comparison of Strategies. In Figure 3 we plot the h -investor's optimal strategy $\alpha_n^*(s)$ for fixed n and various values of $s = \sum_{k=1}^n \tilde{Z}_k = Z_{nh}$ representing different observation scenarios. The same plot contains also the approximating strategy $\underline{\alpha}_n$ using the same $s = Z_{nh}$. One can easily see from (3.3) that for $s = Z_{nh} = 0$ the approximating strategy converges to $\frac{\mu_0 - r}{\sigma^2}$ for $h \rightarrow 0$. Moreover, for fixed $s = Z_{nh} \neq 0$, $\hat{\mu}_{nh}$ converges to $\pm\infty$ for $h \rightarrow 0$ and the sign does depend on the value of s . Hence, since short-selling is excluded, the corresponding strategies will converge to either 0 or 1 for $h \rightarrow 0$, if $s \neq 0$. Note, that if $s = Z_{nh}$ was not fixed we would always find that $\hat{\mu}_{nh} \rightarrow 0$ as $h \rightarrow 0$. The main purpose of this figure, however, is to show that the two strategies are almost identical even for larger values of h . Hence, in practice, the

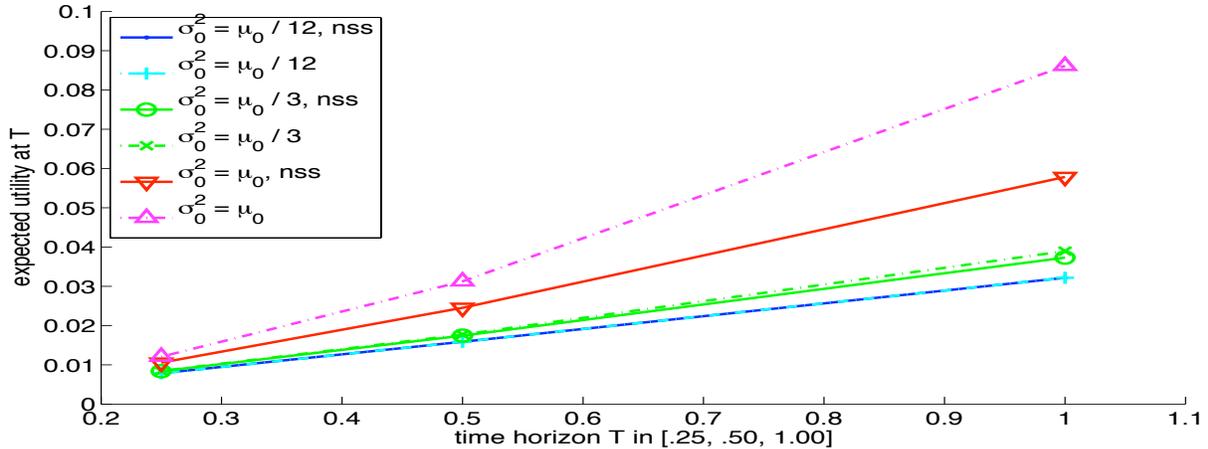


FIGURE 2. Expected utilities $J(1)$, $J^{nss}(1)$ of the continuous investor for $T \in \{\frac{1}{4}, \frac{1}{2}, 1\}$, $\mu_0 = \log(1.05)$ and $\sigma_0^2 \in \{\frac{1}{12}\mu_0, \frac{1}{3}\mu_0, \mu_0\}$, represented by the colours blue, green and red.

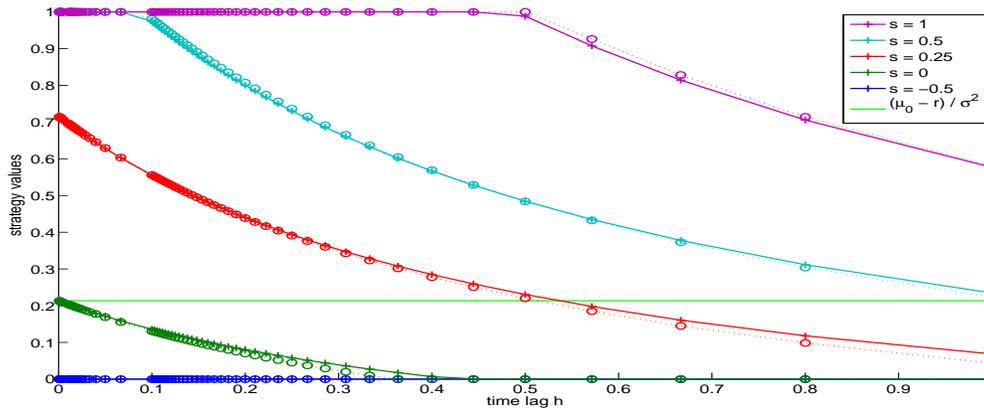


FIGURE 3. Comparison of the h -investor's optimal strategy $\alpha_n^*(s)$ and the approximating strategy α_n for $n = 10$, for different values of the sufficient statistic $s = \sum_{k=1}^n \tilde{Z}_k = Z_{nh}$ and different time lags h . For every colour-coded combination the straight line gives the value of $\alpha_n^*(s)$, while the dotted curve represents α_n .

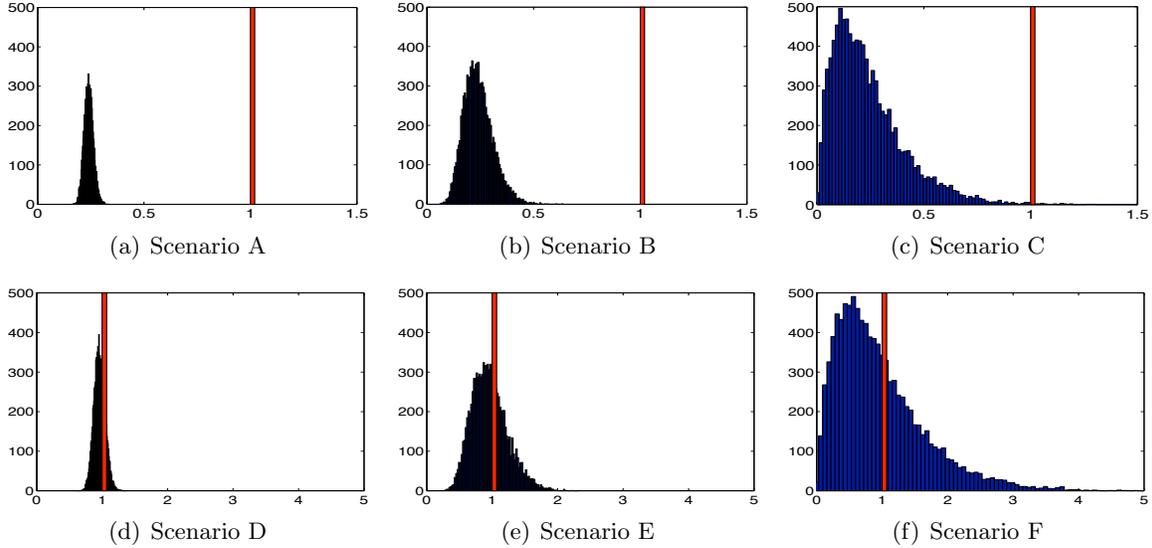
approximating strategy can be very useful since it is available in closed form and is reasonably close to the optimal strategy.

7.2. Unknown Volatility. In contrast to the previous setting, we now assume that the drift is a known constant and the volatility is the unobservable random variable. We set $\mu = \log(1.05)$ and $\Sigma^{-1} \sim \mathcal{W}_1(\nu, \Sigma_0^{-1})$. We choose six different parameter combinations for ν, Σ_0^{-1} as specified in Table 2. In scenario A, B and C we have $\mathbb{E}[\Sigma^{-1}] = \nu \Sigma_0^{-1} = 12.5$, while in D, E and F $\mathbb{E}[\Sigma^{-1}] = 50$.

Figure 4 illustrates the distribution of $\Sigma^{-1}(\mu - r)$ by showing corresponding histograms for all 6 scenarios. If $\Sigma^{-1}(\mu - r)$ is used as a trading strategy within the first 3 scenarios, the no short-selling restriction is satisfied with a very high probability. In the scenarios D, E, F, however, being constrained to no short-selling or not makes a notable difference. The red vertical line in Figure 4 indicates the upper bound 1 on $\Sigma^{-1}(\mu - r)$ which is in place as soon as short-selling constraints are imposed. For D, E, F we observe a significant proportion of the probability mass outside the $[0, 1]$ interval.

TABLE 2. Market scenarios for the prior distribution of Σ^{-1} .

scenario	ν	Σ_0^{-1}
A	250	1/20
B	25	1/2
C	4	25/8
D	250	1/5
E	25	2
F	4	25/2

FIGURE 4. Histogram obtained by simulating $\Sigma^{-1}(\mu - r)$ for all 6 scenarios.

Again, we are interested in the discretization gap. We can easily compute J^{vol} , the expected utility of the continuous investor without short-selling constraints, since this is available in closed form, see (6.2).

The optimal solution to the h -investor's problem with unknown volatility is difficult to calculate, but we can use some approximations:

When trading continuously, the volatility can be determined with any desired precision over an arbitrary small compact interval, thus the continuous investor will invest optimally when holding Merton's ratio $\Sigma^{-1}(\mu - r)$ in the stock. Applying this strategy in discrete time is denoted by $\alpha^{vol,known,N}$, where $N = T/h$. Restricting it to no short-selling leads to a strategy denoted by $\underline{\alpha}^{vol,known,N}$. These strategies are obviously not \mathcal{F}^R -adapted and therefore not admissible for the h -investor, but they are useful for our comparisons.

We also consider the following strategies which are admissible for the h -investor since the unknown Σ^{-1} is replaced by the estimator defined in (6.6):

$$\alpha_n^{vol,est,N}(s) := \widehat{\Sigma}^{-1,h}(s, \lfloor \frac{n}{2} \rfloor)(\mu - r),$$

$$\underline{\alpha}_n^{vol,est,N}(s) := \max\{0, \min\{1, \widehat{\Sigma}^{-1,h}(s, \lfloor \frac{n}{2} \rfloor)(\mu - r)\}\},$$

where $n = 0, 1, \dots, N - 1$ and the latter obeys short-selling restrictions. Here s denotes again the value of the sufficient statistic. By Monte Carlo simulation (2 Million replications) we have found the results given in Table 3.

Clearly and as expected, taking into account additional information gives an advantage to the investor. Additional information is particularly useful if the estimator is not very precise yet due to a high variance of the prior.

TABLE 3. Expected terminal utility for strategies with either estimated or known parameter Σ^{-1} for different N and different scenarios.

α	$N = 4$		$N = 200$		$N = 2000$		continuous
	$J_{0,\alpha}^{h,vol}(1, (0, 0))$	$J_{0,\alpha}^{h,vol}(1, (0, 0))$	$J_{0,\alpha}^{h,vol}(1, (0, 0))$	$J_{0,\alpha}^{h,vol}(1, (0, 0))$	$J_{0,\alpha}^{h,vol}(1, (0, 0))$	$J_{0,\alpha}^{h,vol}(1, (0, 0))$	$J^{vol}(1)$
	$\underline{\alpha}^{vol,est,4}$	$\underline{\alpha}^{vol,known,4}$	$\underline{\alpha}^{vol,est,200}$	$\underline{\alpha}^{vol,known,200}$	$\underline{\alpha}^{vol,est,2000}$	$\underline{\alpha}^{vol,known,2000}$	π
A	0.03181	0.03183	0.03182	0.03184	0.03185	0.03186	0.03187
B	0.03164	0.03182	0.03177	0.03184	0.03185	0.03186	0.03187
C	0.02984	0.03183	0.03165	0.03183	0.03182	0.03184	0.03187
Standard deviations of $J_{0,\alpha}^{h,vol}(1, (0, 0)) \in [0.000048, 0.000065]$ for A, B, C.							-
D	0.03868	0.03874	0.03868	0.03873	0.03874	0.03876	0.03880
E	0.03799	0.03853	0.03835	0.03853	0.03852	0.03856	0.03880
F	0.03144	0.03763	0.03723	0.03759	0.03757	0.03761	0.03880
Standard deviations of $J_{0,\alpha}^{h,vol}(1, (0, 0)) \in [0.000076, 0.000127]$ for D, E, F.							-

TABLE 4. Discretization gap $J^{vol}(1) - J_{0,\alpha}^{h,vol,known,2000}(1, (0, 0))$, difference between continuous and discrete trading for continuous observations $J_{0,\alpha}^{h,vol,known,2000}(1, (0, 0)) - J_{0,\alpha}^{h,vol,known,4}(1, (0, 0))$, and difference between continuous and discrete observations for discrete trading $J_{0,\alpha}^{h,vol,known,4}(1, (0, 0)) - J_{0,\alpha}^{h,vol,est,4}(1, (0, 0))$.

scen.	discretization gap	cont. - disc. trading with cont. observations	cont. - disc. observations with disc. trading
A	0.00001	0.00003	0.00002
B	0.00001	0.00004	0.00018
C	0.00003	0.00001	0.00199
D	0.00004	0.00002	0.00006
E	0.00024	0.00003	0.00054
F	0.00119	0.00098	0.00619

Table 4 shows the discretization gap for unknown volatility. Since $J^{vol,nss}(1)$ is not available in closed form, we approximate it by $J_{0,\alpha}^{h,vol,known,2000}(1, (0, 0))$ which can be justified by Theorem 6.3. We find that the discretization gap is small in scenarios A-C. Moreover, it is still within the size of the standard deviation of our estimator. This result is not too surprising, given that the no short-selling restriction in these scenarios is satisfied with a very high probability anyway and hence the difference between $J^{vol}(1)$ and $J^{vol,nss}(s)$ is expected to be small.

The situation is very different in scenarios D-F. Even if we account for the standard deviation of the estimator, the discretization gap is clearly there and becomes more obvious when we move from D to F . In these scenarios the no short-selling restriction is usually not automatically satisfied and therefore these results seem natural.

Next we compare the difference between continuous and discrete trading assuming that we have continuous observations. We compare $J_{0,\alpha}^{h,vol,known,2000}(1, (0, 0))$ as an approximation for continuous trading with $J_{0,\alpha}^{h,vol,known,4}(1, (0, 0))$ which represents discrete trading. We find that the difference is rather small.

The story is very different if we account for the different information of the discrete and the continuous investor. We compare $J_{0,\alpha}^{h,vol,known,4}(1, (0, 0))$, i.e., the expected utility of the h -investor who was told the volatility, with $J_{0,\alpha}^{h,vol,est,4}(1, (0, 0))$ which is the expected utility she gains from using the 2-step-volatility estimator defined in (6.6). The difference in expected utility

is severe. Given that $N = 4$, the estimator for the volatility is only based on two observations and therefore not very precise yet. Therefore the loss in expected utility is large.

Generally, we find that if the volatility is unknown the information loss is the main problem for the h -investor.

8. CONCLUSION

We have derived optimal trading strategies for an investor who can only trade at discrete time points and does not observe all market parameters. We have provided a detailed comparison between the h -investor and the continuous investor regarding their optimal strategies and their expected utilities.

For unknown drift, we have proved that the estimators of the drift of the continuous and the h -investor coincide at the discrete trading times. The expected utility of the h -investor, however, does not converge to the expected utility of the continuous investor for $h \rightarrow 0$. This is due to the fact that the h -investor's optimal strategy excludes short-selling whereas for the continuous investor this is not the case. The arising discretization gap is just the difference between the expected utility of the continuous investors with and without short-selling constraints.

For unknown volatility, already the estimators between the two types of investors differ. The continuous investor can perfectly estimate the unknown volatility from a very small time frame of available stock prices and effectively works under full information. The h -investor, however, loses relevant information by observing the asset price at discrete time and therefore has a disadvantage. We still observe a very similar limiting behaviour for $h \rightarrow 0$. Again the expected utility of the h -investor only converges to the expected utility of the continuous investor satisfying short-selling constraints, i.e., again the discretization gap occurs.

For both situations (unknown drift, unknown volatility, respectively) we have provided simple strategies for the h -investor which are almost optimal, simple to implement and ensure that the corresponding expected utility converges to the optimal expected utility.

Of course, it is also possible to combine the situation of unknown drift and unknown volatility, but we leave this for future research.

Our results under incomplete information are in contrast to those under full information, see Rogers (2001), where the discretization gap does not occur. Hence, we have found a natural example in which a discretization of a continuous problem does not converge to the continuous solution when the size of the time step approaches zero.

PROOFS

Proof of Theorem 5.1. The second equality holds by definition, therefore we only need to prove the first equality. The expected log-utility of the continuous investor can be written as:

$$J_{\pi}^{nss}(x) = \mathbb{E}_x \log X_T^{\pi} = \log(x) + \mathbb{E}_x \left[\int_0^T (r + \pi_s^{\top}(\mu - r\mathbf{1}) - \frac{1}{2}\pi_s^{\top}\Sigma\pi_s) ds \right]$$

and the expected log-utility of the h -investor is

$$J_{0,\underline{\alpha}}^h(x, 0) = \mathbb{E}_x \log X_N^{\underline{\alpha}} = \log(x) + rT + \sum_{n=1}^N \mathbb{E}_x \log (1 + \underline{\alpha}_{n-1}^{\top} R_n).$$

So it is obviously sufficient to show that

$$\lim_{h \rightarrow 0} \sum_{n=1}^N \mathbb{E}_x \log (1 + \underline{\alpha}_{n-1}^{\top} R_n) = \mathbb{E}_x \left[\int_0^T (\pi_s^{\top}(\mu - r\mathbf{1}) - \frac{1}{2}\pi_s^{\top}\Sigma\pi_s) ds \right].$$

We first prove:

$$\mathbb{E} \left[\log(1 + \underline{\alpha}_{n-1}^{\top} R_n) \right] = \mathbb{E} \left[\pi_{(n-1)h}^{\top}(\mu - r\mathbf{1})h - \frac{1}{2}\pi_{(n-1)h}^{\top}\Sigma\pi_{(n-1)h}h \right] + o(h).$$

For a fixed time point $t = (n-1)h$ write $\alpha = \underline{\alpha}_{n-1} = \underline{\pi}_t$. Using the tower property of the conditional expectation we treat the \mathcal{F}_{n-1}^R -measurable $\underline{\alpha}_{n-1}$ as given and write \mathbb{E}_t for the conditional expectation. Consider the function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ defined by:

$$f(x_1, \dots, x_d) = \log \left(1 + \sum_{i=1}^d \alpha_i (e^{x_i} - 1) \right).$$

If we set $\tilde{Y}_n^i := (\mu_i - r)h - \frac{1}{2}\|\sigma^i\|^2 h + \sqrt{h}\sigma^i \varepsilon_n$ then

$$\mathbb{E}_t \left[\log(1 + \alpha^\top R_n) \right] = \mathbb{E}_t f(\tilde{Y}_n^1, \dots, \tilde{Y}_n^d).$$

Next use a Taylor series expansion of f around $\underline{0} = (0, \dots, 0)$. We use the notation $D_i f$ for the partial derivative w.r.t x_i and $\theta \in (0, 1)$:

$$f(x) = f(\underline{0}) + \sum_{i=1}^d D_i f(\underline{0})x_i + \frac{1}{2} \sum_{i,j=1}^d D_i D_j f(\underline{0})x_i x_j + \frac{1}{6} \sum_{i,j,k=1}^d D_i D_j D_k f(\theta x)x_i x_j x_k.$$

Computing the derivatives we obtain: $D_i f(\underline{0}) = \alpha_i$, $D_i D_j f(\underline{0}) = -\alpha_i \alpha_j$ for $i \neq j$ and $(D_i)^2 f(\underline{0}) = \alpha_i(1 - \alpha_i)$. Further note that $|D_i D_j D_k f(\theta \tilde{Y}_n)| \leq 1$. Thus, our statement follows, since

$$\begin{aligned} \mathbb{E}_t \left[\log(1 + \alpha^\top R_n) \right] &= \sum_{i=1}^d \alpha_i \mathbb{E}_t \tilde{Y}_n^i - \frac{1}{2} \sum_{i,j=1}^d \alpha_i \alpha_j \mathbb{E}_t \tilde{Y}_n^i \tilde{Y}_n^j + \frac{1}{2} \sum_{i=1}^d \alpha_i \mathbb{E}_t (\tilde{Y}_n^i)^2 + o(h) \\ &= \alpha^\top (\mu - r\mathbf{1})h - \frac{1}{2} \alpha^\top \Sigma \alpha h + o(h). \end{aligned}$$

We obtain with dominated convergence, the fact that the constants in front of $o(h)$ can be made independent of n (due to stationarity) and the \mathbb{P} -a.s. continuity of $t \mapsto \underline{\pi}_t$ that

$$\begin{aligned} &\lim_{h \rightarrow 0} \sum_{n=1}^{T/h} \mathbb{E}_x \log(1 + \underline{\alpha}_{n-1}^\top R_n) = \\ &= \lim_{h \rightarrow 0} \left(\sum_{n=1}^{T/h} \mathbb{E}_x \left[\underline{\pi}_{(n-1)h}^\top (\mu - r\mathbf{1})h - \frac{1}{2} \underline{\pi}_{(n-1)h}^\top \Sigma \underline{\pi}_{(n-1)h} h \right] + \sum_{n=1}^{T/h} o(h) \right) \\ &= \mathbb{E}_x \left[\lim_{h \rightarrow 0} \sum_{n=1}^{T/h} \underline{\pi}_{(n-1)h}^\top (\mu - r\mathbf{1})h - \frac{1}{2} \underline{\pi}_{(n-1)h}^\top \Sigma \underline{\pi}_{(n-1)h} h \right] + \lim_{h \rightarrow 0} o(1) \\ &= \mathbb{E}_x \left[\int_0^T \underline{\pi}_s^\top (\mu - r\mathbf{1}) - \frac{1}{2} \underline{\pi}_s^\top \Sigma \underline{\pi}_s, ds \right] \end{aligned}$$

which implies the result. \square

Proof of Theorem 5.3. Using the same Taylor expansion as in the previous proof we can write the optimization problem in (4.4) as

$$\begin{aligned} &\int \int \log(1 + a^\top \hat{z}) Q^{\hat{Z}}(d\hat{z}|m) \mu(dm|s, n) = \\ &= \int \int \left(\sum_{i=1}^d a_i (\hat{z}^i - \frac{1}{2}\|\sigma^i\|^2 h - rh) - \frac{1}{2} \sum_{i,j=1}^d a_i a_j \hat{z}^i \hat{z}^j + \frac{1}{2} \sum_{i=1}^d a_i \|\sigma^i\|^2 h \right) Q^{\hat{Z}}(d\hat{z}|m) \mu(dm|s, n) + o(h) \\ &= \left(a^\top (\hat{\mu}_t - r\mathbf{1}) - \frac{1}{2} a^\top \Sigma a \right) h + o(h). \end{aligned}$$

If we define $f(a) := a^\top (\hat{\mu}_t - r\mathbf{1}) - \frac{1}{2} a^\top \Sigma a$ then

$$f_h(a) = \frac{1}{h} \int \int \log(1 + a^\top \hat{z}) Q^{\hat{Z}}(d\hat{z}|m) \mu(dm|s, n)$$

has the property that $f_h \rightarrow f$ for $h \rightarrow 0$ and the constant in front of $o(h)$ can be chosen independent of a . Since D is compact, $a \mapsto f_h(a)$ is continuous and strictly concave, i.e., the maximum point of f_h is unique, we obtain our statement with Theorem A.1.5 in Bäuerle & Rieder (2011). \square

Proof of Theorem 6.2. First observe that the strong law of large numbers implies

$$\lim_{h \rightarrow 0} \hat{\Sigma}^h \left(\sum_{k=1}^{\lfloor \frac{t}{2h} \rfloor} \tilde{Y}_k \tilde{Y}_k^\top, \lfloor \frac{t}{2h} \rfloor \right) = \lim_{h \rightarrow 0} \frac{1}{\nu - d - 1 + \lfloor \frac{t}{2h} \rfloor} \left(\Sigma_0 + \frac{1}{2h} \sum_{k=1}^{\lfloor \frac{t}{2h} \rfloor} \tilde{Y}_k \tilde{Y}_k^\top \right) = \Sigma, \mathbb{P}\text{-a.s.}$$

Since the limit is independent of t , the convergence is also uniform on compact intervals of the form $[\epsilon, T]$ for $\epsilon > 0$. Hence, when we denote the objective function of (6.3) by $f(a) := a^\top(\mu - r\mathbf{1}) - \frac{1}{2}a^\top \Sigma a$ and by

$$f^h(a) := a^\top(\mu - r\mathbf{1}) - \frac{1}{2}a^\top \hat{\Sigma}^h \left(\sum_{k=1}^{\lfloor \frac{t}{2h} \rfloor} \tilde{Y}_k \tilde{Y}_k^\top, \lfloor \frac{t}{2h} \rfloor \right) a$$

the objective function of (6.7), we obviously have $f^h \rightarrow f$ for $h \rightarrow 0$. Since D is compact and $a \mapsto f^h(a)$ is continuous and strictly concave, we obtain the statement with Theorem A.1.5 in Bäuerle & Rieder (2011). \square

Proof of Theorem 6.3. As in the proof of Theorem 5.1 we have

$$\mathbb{E}_x \log(1 + \underline{\alpha}_{n-1}^{h, \text{vol}^\top} R_n) = \mathbb{E}_x \left[\pi_{(n-1)h}^{h, \text{vol}^\top} (\mu - r\mathbf{1})h - \frac{1}{2} \pi_{(n-1)h}^{h, \text{vol}^\top} \Sigma \pi_{(n-1)h}^{h, \text{vol}} h \right] + o(h).$$

Since the convergence of the portfolio strategy is uniform on intervals $[\epsilon, T]$ the statement finally follows as in Theorem 5.1. \square

Proof of Corollary 6.4. With Theorem 6.3 and

$$J_{0, \underline{\alpha}^{h, \text{vol}}}^{h, \text{vol}}(x) \leq J_0^{h, \text{vol}}(x) \leq J^{\text{vol}, \text{nss}}(x)$$

we obtain the result. \square

REFERENCES

- Bain, A. & Crisan, D. (2009). *Fundamentals of stochastic filtering*, vol. 60 of *Stochastic Modelling and Applied Probability*. Springer, New York.
- Bäuerle, N. & Rieder, U. (2007). Portfolio optimization with jumps and unobservable intensity process. *Mathematical Finance* **17**, 205–224.
- Bäuerle, N. & Rieder, U. (2011). *Markov Decision Processes with applications to finance*. Springer.
- Björk, T., Davis, M. H. A. & Landén, C. (2010). Optimal investment under partial information. *Mathematical Methods of Operations Research* **71**, 371–399.
- Brendle, S. (2006). Portfolio selection under incomplete information. *Stochastic Processes and their Applications* **116**, 701–723.
- Brendle, S. (2008). On a problem of optimal stochastic control with incomplete information. *Applied Mathematics & Optimization* **58**, 257–274.
- Corsi, M., Pham, H. & Runggaldier, W. J. (2008). Numerical approximation by quantization of control problems in finance under partial observation. In *Mathematical modelling and numerical methods in finance*, 325–360, North Holland.
- Cvitanović, J. & Karatzas, I. (1992). Convex duality in constrained portfolio optimization. *Annals of Applied Probability* **2**, 767–818.
- DeGroot, M. H. (1970). *Optimal Statistical Decisions*. McGraw-Hill Book Company.
- Duffie, D. & Protter, P. (1992). From discrete- to continuous-time finance: Weak convergence of the financial gain process. *Mathematical Finance* **2**, 1–15.

- Frey, R. & Runggaldier, W. J. (1999). Risk-minimizing hedging strategies under restricted information: The case of stochastic volatility models observable only at discrete random times. *Mathematical Methods of Operations Research* **50**, 339–350.
- Fristedt, B., Jain, N. & Krylov, N. (2007). *Filtering and prediction: A primer*, vol. 38 of *Student Mathematical Library*. American Mathematical Society, Providence, RI.
- Gandy, A. & Veraart, L. A. M. (2012+). The effect of estimation in high-dimensional portfolios. *Forthcoming in Mathematical Finance*.
- He, H. (1991). Optimal consumption-portfolio policies: A convergence from discrete to continuous-time models. *Journal of Economic Theory* **55**, 340–363.
- Honda, T. (2003). Optimal portfolio choice for unobservable and regime-switching mean returns. *J. Econ. Dyn. Contr.* **28**, 45–78.
- James, M. R., Krishnamurthy, V. & Gland, F. L. (1996). Time discretization of continuous-time filters and smoothers for HMM parameter estimation. *IEEE Transactions on Information Theory* **42**, 593–605.
- Karatzas, I. & Zhao, X. (2001). Bayesian adaptive portfolio optimization. In *Option Pricing, Interest Rates and Risk Management* (eds. E. Jouini, J. Cvitanic & M. Musiela), 1st ed., chap. 17, 632–669, *Handbooks in Mathematical Finance*, Cambridge University Press.
- Lakner, P. (1995). Utility maximization with partial information. *Stochastic Processes and their Applications* **56**, 247–273.
- Lakner, P. (1998). Optimal trading strategies for an investor: The case of partial information. *Stochastic Processes and their Applications* **76**, 77–97.
- Pham, H. & Quenez, M. C. (2001). Optimal portfolio in partially observed stochastic volatility models. *Annals of Applied Probability* **11**, 210–238.
- Pham, H., Runggaldier, W. J. & Sellami, A. (2005). Approximation by quantization of the filter process and applications to optimal stopping problems under partial observation. *Monte Carlo Methods and Applications* **11**, 57–81.
- Prigent, J.-L. (2003). *Weak convergence of financial markets*. Springer-Verlag, Berlin.
- Rieder, U. & Bäuerle, N. (2005). Portfolio optimization with unobservable Markov-modulated drift process. *Journal of Applied Probability* **42**, 362–378.
- Rishel, R. (1999). Optimal portfolio management with partial observations and power utility. In *Stochastic analysis, control, optimization and application*, 605–619, Birkhäuser, Boston.
- Rogers, L. C. G. (2001). The relaxed investor and parameter uncertainty. *Finance and Stochastics* **5**, 131–154.
- Sass, J. & Haussmann, U. G. (2004). Optimizing the terminal wealth under partial information: The drift process as a continuous time Markov chain. *Finance and Stochastics* **8**, 553–577.
- Taksar, M. & Zeng, X. (2007). Optimal terminal wealth under partial information: Both the drift and the volatility driven by a discrete-time Markov chain. *SIAM Journal on Control and Optimization* **46**, 1461–1482.
- von Rosen, D. (1988). Moments for the inverted Wishart distribution. *Scandinavian Journal of Statistics* **15**, 97–109.
- Zohar, G. (2001). A generalized Cameron-Martin formula with applications to partially observed dynamic portfolio optimization. *Mathematical Finance* **11**, 475–494.