Optimal Investment under Partial Information

Tomas Björk [*]

Department of Finance

Stockholm School of Economics

Box 6501, SE-113 83 Stockholm

SWEDEN

e-mail: tomas.bjork@hhs.se

Mark H.A. Davis

Department of Mathematics

Imperial College

London, SW7 2AZ

England

e-mail: mark.davis @imperial.ac.uk

Camilla Landén,

Department of Mathematics

Royal Insitute of Technology

SE-100 44 Stockholm, SWEDEN

camilla@math.kth.se

April 4, 2008 (First version August 2006)

Abstract

We consider the problem of maximizing terminal utility in a model where asset prices are driven by Wiener processes, but where the various rates of returns are allowed to be arbitrary semimartingales. The only information available to the investor is the one generated by the asset prices and, in particular, the return processes cannot be observed directly. This leads to an optimal control problem under partial information and for the cases of power, log, and exponential utility we manage to provide a surprisingly explicit representation of the optimal terminal wealth as well as of the optimal portfolio strategy. This is done without any assumptions about the dynamical structure of the return processes. We also show how various explicit results in the existing literature are derived as special cases of the general theory.

1 Introduction

We consider a financial market consisting of the risk free bank account and n risky assets without dividends. The risky assets are driven by a multi

1

dimensional Wiener process which is adapted to some filtration $\mathbf{F}$. This filtration is "big" in the sense that it properly includes the filtration generated by the Wiener process. The return processes of the risky assets are allowed to be general $\mathbf{F}$-adapted semimartingales, and in particular we make no Markovian assumptions. The information available to the investor is the filtration $\mathbf{F}^S$ generated by the asset price processes only, so the investor can in general not observe the return processes directly. The problem to be solved is that of maximizing expected utility of terminal wealth over the class of portfolio strategies adapted to the observable information $\mathbf{F}^S$. This leads to a optimal control problem under partial information.

There exists a considerable literature on investment problems of this kind, and the standard approach is more or less as follows.

- Assume that the return processes have a "hidden Markov model" structure.

- Project the asset price dynamics on the observable filtration, thereby obtaining a completely observable model.

- Write down the filtering equations for the return processes and adjoin the filter estimate processes as extended state variables.

- Apply standard dynamic programming techniques to the reformulated completely observable problem, and solve the associated Bellman equation to obtain the optimal control. Alternatively, use martingale techniques.

In this literature the two typical assumptions have been to model the returns as either linear diffusions, leading to the Kalman filter, or as functions of a finite state Markov chain, leading to the Wonham filter. The two basic references are Lakner (1995, 1998). In Lakner (1995), the author studies the problem in general terms and derives, using martingale methods, the structure of the optimal investment and consumption strategies. Explicit results are then obtained for log and power utility in a model where the rates of returns are constant random variables. In Lakner (1998) the same methodology is applied to the case when the return process is a linear diffusion.

The linear diffusion model is studied further, and from a slightly different perspective, in Brendle (2004, 2006), where explicit results for the value of information are derived for power and exponential utility. The effects of learning on the composition of the optimal portfolio is studied in Brennan (1998) and Xia (2001), and Brennan and Xia (2001) discusses (apparent)

asset price anomalies in the context of partially observed models. In Gennotte (1986) the linear diffusion case is studied within a dynamic equilibrium framework and in the recent paper Cvitanic et al. (2006), the authors use a model with constant random returns to analyze the value of professional analyst's recommendations.

Using dynamic programming arguments, the case of an underlying Markov chain is studied in Bäuerle and Rieder (2005, 2004), where also the optimal investment in the partially observable model is compared to the one in the case of a fully observable model.

The Markov chain model is also studied, using Martingale methods and Malliavin calculus, in Haussmann and Sass (2003, 2004b). In Haussmann and Sass (2004a) this analysis is extended to stochastic volatility, and in Sass (2007) convex constraints are added to the model. In Nagai and Runggaldier (2006) the Markov chain model is studied using dynamic programming methods and a new stochastic representation formula is provided.

Two interesting alternative modeling approaches are presented in the recent papers Bäuerle and Rieder (2007) and Callegaro et al. (2006) where the asset prices are driven by jump processes instead of the usual Wiener process.

The structure of the present paper is as follows.

In Section 2 we present our model, which is basically a non-Markovian asset price model with invertible volatility matrix, where the asset prices are driven by a multi dimensional Wiener process. The return processes are, however, not necessarily adapted to the filtration generated by the asset prices. Instead they are allowed to be general semimartingales adapted to a larger filtration, thus leading to a model with partial information. Section 3 is devoted to a fairly detailed study of the optimal investment problem in the special (and rather well known) case of complete observations, using martingale techniques. We have two reasons for including this well studied problem in our paper. Firstly, we need the results for the later parts of the paper. Secondly, in our non-Markovian setting we obtain new and surprisingly explicit results for the optimal wealth and investment processes in the cases of log, power, and exponential utility. In Section 4 we turn to the partially observable case. The main (and standard) idea is then to project the asset price dynamics onto the observable filtration using results from non-linear filtering. We thus reduce the partially observable problem to a completely observable (non-Markovian) problem, and to solve this we only have to copy the results from the previous section. In Section 5 we study the special case when the return processes are generated by a "hidden Markov process". By adjoining the filter equation for the conditional density of the hidden Markov process as a new state variable we can compute

3

the optimal investment strategy explicitly up to the solution of a PDE with infinite dimensional state space. For the cases when the return processes are driven by a finite state Markov chain or a linear SDE, we recover most of the known results from the literature, the exceptions being Bäuerle and Rieder (2007) and Callegaro et al. (2006).

The contributions of the present paper are as follows.

- We add to the literature on the completely observable case by deriving explicit expressions of the optimal wealth and investment processes in a non-Markovian setting.

- For the general partially observed non-Markovian case, our results are considerably more explicit than those obtained in Lakner (1995).

- By the non-Markovian approach we manage to provide a unified treatment of a large class of partially observed investment problems.

- In particular, we obtain most previously known results for models, where the returns are driven by a linear diffusion of by a finite state Markov chain, as special cases of the general theory.

- On the didactic side, we feel that one of the main contributions of the paper is that it shows how simple the problem is, given the proper perspective.

2   Setup

We consider a financial market living on a stochastic basis $(\Omega, \mathcal{F}, \mathbf{F}, P)$, where the filtration $\mathbf{F} = \{\mathcal{F}_t\}_{0 \leq t \leq T}$ satisfies the usual conditions, and where P is the objective probability measure. The basis carries an n-dimensional P-Wiener process W, and the filtration generated by the W process is denoted by $\mathbf{F}^W$.

The financial market under consideration consists of n non-dividend paying risky assets with price processes $S^1, \ldots, S^n$, and a bank account with price process B. The formal assumptions concerning the price dynamics are as follows.

Assumption 2.1 .

1. *The risky asset prices have P-dynamics given by*

$$dS_t^i = \alpha_t^i S_t^i dt + S_t^i \sigma_t^i dW_t, \quad i = 1, \ldots, n. \tag{2.1}$$

   Here $\alpha^1, \ldots, \alpha^n$ are assumed to be $\mathbf{F}$-adapted scalar processes, and $\sigma^i, \ldots, \sigma^n$ are $\mathbf{F}^W$-adapted d-dimensional row vector processes.

2. *The short rate r_t is assumed to be a bounded $\mathbf{F}^S$-adapted process; then the bank account has dynamics given by*

$$dS_t^0 = r(t)S_t^0 dt.$$

For $a < b$ we will denote $D_{a,b} = \exp\left(-\int_a^b r(t)dt\right)$.

We note that, by the quadratic variation properties of W, the assumption that $\sigma^i, \ldots, \sigma^n$ are $\mathbf{F}^W$-adapted is essentially without loss of generality. Defining the stock vector process by[1] $S = [S^1, \ldots, S^n]'$, the return vector process by $\alpha = [\alpha^1, \ldots, \alpha^n]'$ and the volatility matrix by $\sigma = [\sigma^1, \ldots, \sigma^n]'$, we can write the asset price dynamics as

$$dS_t = D(S_t)\alpha_t dt + D(S_t)\sigma_t dW_t,$$

where $D(S)$ denotes the diagonal matrix with $S^1, \ldots, S^n$ on the diagonal.

We will need an important assumption concerning the volatility matrix.

Assumption 2.2 *We assume that with probability one the volatility matrix σ_t is non-singular for all t.*

The interpretation of the model above is that we do not have access to the full information contained in the filtration $\mathbf{F}$, but that we are only allowed to observe the evolution of the asset price vector process S. A special case of the model above would for example be that α is of the form $\alpha(t) = \alpha(t, Y_t)$ where Y is a "hidden Markov process" which cannot be observed directly. It is, however, important for the rest of the paper that we do **not** make any such assumption of a Markov structure in our model. As regards the short rate, it may seem artificial to assume that r_t is adapted to $\mathbf{F}^S$, and indeed this would be artificial if we were to interpret all the components of S^i as equity prices. Our assumption is however natural if some components are interest-rate-related securities such as zero-coupon bonds. Our model thus covers mixed investment strategies in equity and fixed-income assets.

To state the formal problem to be studied, we define the **observable filtration G** by $\mathbf{G} = \mathbf{F}^S$, and consider a fixed utility function U, satisfying the usual Inada conditions. The problem to be solved is that of maximizing expected utility over the class of observable portfolio strategies. Denoting the initial wealth by x, and portfolio value process by X, we thus want to maximize.

$$E^P\left[U(X(T))\right],$$

---

[1]Prime denotes transpose.

over the class of self financing **G**-adapted portfolios with the initial condition $X(0) = x$. We are thus facing a stochastic control problem with partial information.

Our strategy for attacking this problem is by first solving the corresponding (and much simpler) problem with complete information. By using results from non-linear filtering theory, we will then show that our partially observed problem can be reduced to a related problem with complete information and we are done.

3   The completely observable case

For the rest of this section we assume that we are in the completely observable case, i.e. that $\mathbf{F} = \mathbf{F}^W = \mathbf{F}^S$. Given our assumption about the invertibility of the volatility matrix σ this implies that the model is complete, so we can separate the determination of the optimal wealth profile from the determination of the optimal strategy. More precisely, we can proceed along the following scheme.

- Find the optimal wealth profile at time T by solving the static problem

$$\max_{X \in \mathcal{F}_T} \quad E^P[U(X)] \tag{3.1}$$

subject to the budget constraint

$$E^Q[D_{0,T}X] = x, \tag{3.2}$$

where x is the initial wealth, and Q is the unique (because of the assumed completeness) martingale measure.

- Given the optimal wealth profile X_T^*, we can (in principle) compute the corresponding generating portfolio using martingale methods.

As is well known, the static problem above can easily be solved using Lagrange relaxation. Since we need the formulas, we briefly recall the basic technique. We start by rewriting the budget constraint (3.2) as

$$E^P[D_{0,T}L_T X] = x,$$

where L is the likelihood process between P and Q, i.e.,

$$L_t = \left.\frac{dQ}{dP}\right|_{\mathcal{F}_t}.$$

We now relax the budget constraint to obtain the Lagrangian

$$\mathcal{L} = E^P[U(X)] - \lambda\left(E^P[D_{0,T}L_T X] - x\right),$$

so

$$\mathcal{L} = \int_{\Omega} \left\{ U(X(\omega)) - \lambda \left[D_{0,T}(\omega) L_T(\omega) X(\omega) - x \right] \right\} dP(\omega).$$

This is a separable problem, so we can maximize for each ω. The optimality condition is

$$U'(X) = \lambda D_{0,T} L_T$$

so the optimal wealth profile is given by

$$X^* = I \left(\lambda D_{0,T} L_T \right), \qquad (3.3)$$

where $I = (U')^{-1}$. The Lagrange multiplier is as usual determined by the budget constraint (3.2).

We do in fact have an explicit expression for the Radon-Nikodym derivative L_T above. From the price dynamics (2.1) and the Girsanov Theorem it is easily seen that the L dynamics are given by

$$\begin{aligned}
dL_t &= L_t \left\{ \sigma_t^{-1} \left(\mathbf{r}_t - \alpha_t \right) \right\}' dW_t, \\
L_0 &= 1,
\end{aligned}$$

where $\mathbf{r}_t$ denotes the n-column vector with r_t in each component. We thus have the explicit formula

$$L_t = \exp \left\{ \int_0^t \left\{ \sigma_s^{-1} \left(\mathbf{r}_s - \alpha_s \right) \right\}' dW_s - \frac{1}{2} \int_0^t \| \sigma_s^{-1} \left(\mathbf{r}_s - \alpha_s \right) \|^2 ds \right\}. \qquad (3.4)$$

We will treat three special cases in detail: power utility, log utility and exponential utility.

3.1 Power utility

The most interesting case is that of power utility. In this case the utility function is of the form

$$U(x) = \frac{x^\gamma}{\gamma},$$

for some non-zero $\gamma < 1$. We have

$$I(y) = y^{-\frac{1}{1-\gamma}}.$$

3.1.1 The optimal terminal wealth profile

From (3.3), and the expression for I above we obtain the optimal wealth profile as

$$X_T^* = \left(\lambda D_{0,T} L_T \right)^{\frac{-1}{1-\gamma}}. \qquad (3.5)$$

The budget constraint (3.2) becomes

$$\lambda^{\frac{-1}{1-\gamma}} E^P \left[D_{0,T}^{-\beta} L_T^{-\beta} \right] = x, \tag{3.6}$$

with β defined by

$$\beta = \frac{\gamma}{1 - \gamma}. \tag{3.7}$$

Solving for $\lambda^{-\frac{1}{1-\gamma}}$ in (3.6) and inserting this into (3.5) gives us the optimal wealth profile as

$$X_T^* = \frac{x}{H_0} D_{0,T}^{\frac{-1}{1-\gamma}} L_T^{\frac{-1}{1-\gamma}},$$

where

$$H_0 = E^P \left[D_{0,T}^{-\beta} L_T^{-\beta} \right].$$

The optimal expected utility $V_0 = E^P \left[U \left(X_T^* \right) \right]$ can easily be computed as

$$V_0 = \frac{1}{\gamma} E^P \left[(X_T^*)^\gamma \right] = H_0^{1-\gamma} \frac{x^\gamma}{\gamma}. \tag{3.8}$$

We will now study H_0 in some detail. From (3.4) we obtain

$$L_T^{-\beta} = \exp \left\{ \int_0^T \beta \left\{ \sigma_t^{-1} \left(\alpha_t - \mathbf{r}_t \right) \right\}' dW_t + \frac{1}{2} \int_0^T \beta \| \sigma_t^{-1} \left(\alpha_t - \mathbf{r}_t \right) \|^2 dt \right\}.$$

This expression looks almost like a Radon Nikodym derivative, and this observation leads us to define the P-martingale L^0 by

$$L_t^0 = \exp \left\{ \int_0^t \beta \left\{ \sigma_s^{-1} \left(\alpha_s - \mathbf{r}_s \right) \right\}' dW_s - \frac{1}{2} \int_0^t \beta^2 \| \sigma_s^{-1} \left(\alpha_s - \mathbf{r}_s \right) \|^2 ds \right\}$$

i.e. with dynamics

$$dL_t^0 = L_t^0 \beta \left\{ \sigma_t^{-1} \left(\alpha_t - \mathbf{r}_t \right) \right\}' dW_t.$$

We can thus write

$$L_T^{-\beta} = L_T^0 \exp \left\{ \frac{1}{2} \int_0^T \frac{\beta}{1 - \gamma} \| \sigma_t^{-1} \left(\alpha_t - \mathbf{r}_t \right) \|^2 dt \right\}, \tag{3.9}$$

to obtain

$$H_0 = E^0 \left[\exp \left\{ \beta \int_0^T \left(r_t + \frac{1}{2(1 - \gamma)} \| \sigma_t^{-1} \left(\alpha_t - \mathbf{r}_t \right) \|^2 \right) dt \right\} \right]. \tag{3.10}$$

where we integrate over the measure Q^0 defined through the likelihood process L^0.

For easy reference we collect the definitions of Q and Q^0.

Definition 3.1

- *The risk neutral martingale measure Q is defined by*

$$\frac{dQ}{dP} = L_t, \quad on \ \mathcal{F}_t, \tag{3.11}$$

with L given by

$$dL_t = L_t \left\{ \sigma_t^{-1} \left(\mathbf{r}_t - \alpha_t \right) \right\}' dW_t. \tag{3.12}$$

- *The measure Q^0 is defined by*

$$\frac{dQ^0}{dP} = L_t^0, \quad on \ \mathcal{F}_t,$$

with L^0 given by

$$dL_t^0 = L_t^0 \beta \left\{ \sigma_t^{-1} \left(\alpha_t - \mathbf{r}_t \right) \right\}' dW_t.$$

with

$$\beta = \frac{\gamma}{1-\gamma}.$$

We also collect our results so far.

Proposition 3.1 *With definitions as above, the following hold.*

- *The optimal terminal wealth is given by*

$$X_T^* = \frac{x}{H_0} D_{0,T}^{\frac{-1}{1-\gamma}} L_T^{\frac{-1}{1-\gamma}}, \tag{3.13}$$

where H_0 is defined by (3.10) above.

- *The optimal utility V_0 is given by*

$$V_0 = H_0^{1-\gamma} \frac{x^\gamma}{\gamma}.$$

3.1.2 The optimal wealth process

We have already computed the optimal terminal wealth profile X_T^* above. and we can in fact also derive a surprisingly explicit formula for the entire optimal wealth process X^*.

Proposition 3.2 *The optimal wealth process X^* is given by*

$$X_t^* = \frac{H_t}{H_0} (D_{0,t} L_t)^{-\frac{1}{1-\gamma}} x, \tag{3.14}$$

where

$$H_t = E^0 \left[\exp \left\{ \beta \int_t^T \left(r_s + \frac{1}{2(1-\gamma)} \| \sigma_s^{-1} (\alpha_s - \mathbf{r}_s) \|^2 \right) ds \right\} \middle| \mathcal{F}_t \right]. \tag{3.15}$$

Proof. From general theory we know that the wealth process (normalized with the bank account) of any self financing portfolio will be a Q martingale, so we have

$$D_{0,t}X_t^* = E^Q\left[D_{0,T}X_T^*\middle|\mathcal{F}_t\right]. \tag{3.16}$$

Using the expression (3.13) for X_T^* and the abstract Bayes' formula, we obtain

$$
\begin{aligned}
X_t^* &= \frac{x}{H_0}E^Q\left[D_{t,T}D_{0,T}^{\frac{-1}{1-\gamma}}L_T^{\frac{-1}{1-\gamma}}\middle|\mathcal{F}_t\right] \\
&= \frac{x}{H_0}D_{0,t}^{\frac{-1}{1-\gamma}}E^Q\left[D_{t,T}^{-\beta}L_T^{\frac{-1}{1-\gamma}}\middle|\mathcal{F}_t\right] \\
&= \frac{x}{H_0}D_{0,t}^{\frac{-1}{1-\gamma}}\frac{E^P\left[D_{t,T}^{-\beta}L_T^{-\beta}\middle|\mathcal{F}_t\right]}{L_t}.
\end{aligned} \tag{3.17}
$$

Using (3.9) we have

$$
\begin{aligned}
E^P\left[D_{t,T}^{-\beta}L_T^{-\beta}\middle|\mathcal{F}_t\right] &= L_t^0\frac{E^P\left[D_{t,T}^{-\beta}L_T^0\exp\left\{\frac{1}{2}\int_0^T\frac{\beta}{1-\gamma}\|\sigma_t^{-1}(\alpha_t-\mathbf{r}_t)\|^2dt\right\}\middle|\mathcal{F}_t\right]}{L_t^0} \\
&= L_t^0 E^0\left[D_{t,T}^{-\beta}\exp\left\{\frac{1}{2}\int_0^T\frac{\beta}{1-\gamma}\|\sigma_t^{-1}(\alpha_t-\mathbf{r}_t)\|^2dt\right\}\middle|\mathcal{F}_t\right] \\
&= L_t^0\exp\left\{\frac{1}{2}\int_0^t\frac{\beta}{1-\gamma}\|\sigma_s^{-1}(\alpha_s-\mathbf{r}_s)\|^2ds\right\}H_t \\
&= L_t^{-\beta}H_t,
\end{aligned} \tag{3.18}
$$

where H_t is defined by (3.15). The expression (3.14) for the optimal wealth X_t^* now follows from (3.17) and (3.18). ∎

3.1.3 The optimal portfolio

We can also derive a reasonably explicit formula for the optimal portfolio. We need some notation so let us denote by μ_H and σ_H the relative drift and diffusion parts of the process H defined in (3.15), i.e.

$$dH_t = H_t\mu_H(t)dt + H_t\sigma_H(t)dW_t. \tag{3.19}$$

Furthermore, for any self financing portfolio we denote by $u_t = (u_t^1, \ldots u_t^n)$ the vector process of portfolio weights on the risky assets. This of course implies that the weight, u^0, on the bank account is given by $u_t^0 = 1 - u_t\mathbf{1}$, where $\mathbf{1}$ denotes the n-column vector with a unit in each position.

Proposition 3.3 *The optimal portfolio weight vector process u^* is given by*

$$u_t^* = \frac{1}{1-\gamma} (\alpha_t - \mathbf{r}_t)' \left(\sigma_t \sigma_t'\right)^{-1} + \sigma_H(t)\sigma_t^{-1} \qquad (3.20)$$

Proof. From standard portfolio theory it follows that the wealth process X of any self financing portfolio has the dynamics

$$dX_t = X_t u_t \alpha_t dt + X_t(1 - u_t\mathbf{1})r_t dt + X_t u_t \sigma_t dW_t, \qquad (3.21)$$

implying that the discounted wealth process $Z_t = D_{0,t} X_t$ has dynamics

$$dZ_t = Z_t u_t \left(\alpha_t - \mathbf{r}_t\right) dt + Z_t u_t \sigma_t dW_t, \qquad (3.22)$$

the point being that the portfolio u can be determined from the diffusion part of the Z-dynamics. From (3.14) we see that the optimal discounted wealth process has the form

$$Z_t^* = A H_t D_{0,t}^\beta L_t^c,$$

where $A = x H_0^{-1}$ and $c = -(1-\gamma)^{-1}$. Using (3.12) and (3.19) we easily obtain

$$dZ_t^* = Z_t^* (\ldots) dt + Z_t^* \left\{ c \left\{\sigma_t^{-1} (\mathbf{r}_t - \alpha_t)\right\}' + \sigma_H(t) \right\} dW_t,$$

where we do not care about the exact form of the dt term. Post-multiplying the diffusion part by the term $\sigma_t^{-1}\sigma_t$ and comparing with (3.22) shows that

$$u_t^* = \left\{ \frac{1}{1-\gamma} \{\sigma_t^{-1}(\alpha_t - \mathbf{r})\}' + \sigma_H(t) \right\} \sigma_t^{-1},$$

which is equivalent to (3.20). ∎

In (3.20) we recognize the first term as the solution to the classical (completely observable) Merton problem. The second term represents the "hedging demand for parameter risk".

3.2 Log utility

In this case the utility function is given by

$$U(x) = \ln(x),$$

which implies that

$$I(y) = \frac{1}{y}.$$

From the point of view of local risk aversion, log utility is the limiting case of power utility when the risk aversion parameter γ tends to zero. We would thus intuitively conjecture that the solution to the log utility problem is obtained from the power utility case by setting γ to zero, and in fact this turns out to be correct. We have the following result, and since the calculations in this case are very simple we omit the proof.

Proposition 3.4 *For the log utility case, the following hold.*

- *The optimal portfolio process X^* is given by*

$$X_t^* = (D_{0,t} L_t)^{-1} x,$$

  where, as before, the likelihood process L is given by (3.12).

- *The optimal portfolio weight vector process u^* is given by*

$$u_t^* = (\alpha_t - \mathbf{r_t})' \left(\sigma_t \sigma_t' \right)^{-1}$$

In this case $H_t \equiv 1$, so that $\sigma_H = 0$ and there is no hedging demand for parameter risk in the optimal portfolio. This is intuitively expected from the interpretation of log utility as myopic.

In particular we see that results from the power case trivialize in the log case, in the sense that $L^0 \equiv 1$, $Q^0 = P$, and $H \equiv 1$.

3.3   Exponential utility

In this case we have
$$U(x) = -\frac{1}{\gamma} e^{-\gamma x},$$

and
$$I(y) = -\frac{1}{\gamma} \ln(y).$$

3.3.1   The optimal wealth process

From (3.3) the optimal terminal wealth profile is given by

$$X_T^* = I(\lambda D_{0,T} L_T) = -\frac{1}{\gamma} \ln \lambda - \frac{1}{\gamma} \ln(D_{0,T} L_T),$$

and the Lagrange multiplier is easily determined by the budget constraint

$$E^Q \left[D_{0,T} X_T^* \right] = x,$$

giving the following expression for the optimal terminal wealth:

$$X_T^* = \frac{x + \frac{1}{\gamma}J_0}{B_{0,T}} - \frac{1}{\gamma}\ln(D_{0,T}L_T).$$

Here

$$J_0 = E^Q[D_{0,T}\ln(D_{0,T}L_T)]$$

and $B_{t,T}$ is the zero-coupon bond value

$$B_{t,T} = E^Q\left[D_{t,T}|\mathcal{F}_t\right].$$

Proposition 3.5 *For $t \in [0, T]$ define*

$$J_t = E^Q[D_{t,T}\ln(D_{t,T}L_T)|\mathcal{F}_t]. \tag{3.23}$$

Then the optimal terminal wealth process is given by

$$X_t^* = \left(\frac{x + \frac{1}{\gamma}J_0}{B_{0,T}} - \frac{1}{\gamma}\ln D_{0,t}\right)B_{t,T} - \frac{1}{\gamma}J_t. \tag{3.24}$$

Proof. As usual,

$$
\begin{aligned}
X_t^* &= E^Q\left[D_{t,T}X_T^*|\mathcal{F}_t\right] \\
&= \frac{x + \frac{1}{\gamma}J_0}{B_{0,T}}B_{t,T} - \frac{1}{\gamma}E^Q\left[D_{t,T}\ln(D_{0,T}L_T)|\mathcal{F}_t\right].
\end{aligned}
$$

Writing $\ln(D_{0,T}L_T) = \ln(D_{0,t}) + \ln(D_{t,T}L_T)$, we see that

$$E^Q\left[D_{t,T}\ln(D_{0,T}L_T)|\mathcal{F}_t\right] = B_{t,T}\ln(D_{0,t}) + J_t.$$

The result follows. ∎

3.3.2   The optimal portfolio

As in the power utility case, we will identify the optimal portfolio from the (discounted) wealth dynamics (3.21)-(3.22). Indeed, the discounted portfolio value $Z_t = D_{0,t}X_t^*$ is a Q-martingale satisfying

$$dZ_t = Z_t u_t \sigma_t dW_t^Q,$$

where u is the optimal portfolio trading strategy. Define

$$A(t) = \left(\frac{x + \frac{1}{\gamma}J_0}{B_{0,T}} - \frac{1}{\gamma}\ln D_{0,t}\right)D_{0,t}. \tag{3.25}$$

13

Then from (3.24) we have

$$Z_t = A(t)B_{t,T} - \frac{1}{\gamma}D_{0,t}J_t,$$

and we note that both $A(t)$ and $D_{0,t}$ are bounded variation processes. Let $\sigma_B(t)$ and $\sigma_J(t)$ be, respectively, the zero-coupon bond volatility and the volatility of the J_t process, i.e. we have the semimartingale decompositions

$$\begin{aligned} dB_{t,T} &= B_{t,T}\sigma_B(t)dW_t^Q + (\cdots)dt, \\ dJ_t &= J_t\sigma_J(t)dW_t^Q + (\cdots)dt. \end{aligned}$$

Then

$$dZ_t = \left(A(t)B_{t,T}\sigma_B(t) - \frac{1}{\gamma}D_{0,t}J_t\sigma_J(t)\right)dW_t^Q.$$

We thus obtain the following result.

Proposition 3.6 *The optimal portfolio strategy in the exponential utility case is*

$$u(t) = \left(\frac{A(t)B_{t,T}\sigma_B(t) - \frac{1}{\gamma}D_{0,t}J_t\sigma_J(t)}{A(t)B_{t,T} - \frac{1}{\gamma}D_{0,t}J_t}\right)\sigma^{-1}(t). \qquad (3.26)$$

3.3.3   Exponential utility with constant interest rate

The above expressions for X_t^* and u_t are complicated and not easy to interpret. They simplify considerably, however, in the case of constant interest rate $D_{t,T} = B_{t,T} = e^{-r(T-t)}$. First, note from (3.4) that

$$\ln(L_t) = \int_0^t \varphi_s' dW_s - \frac{1}{2}\int_0^t \|\varphi_s\|^2 ds, \qquad (3.27)$$

where

$$\varphi_t = \sigma_t^{-1}\{\mathbf{r} - \alpha_t\}. \qquad (3.28)$$

Furthermore, the Girsanov Theorem tells us that we can write

$$dW_t = \varphi_t dt + dW_t^Q,$$

where W^Q is Q-Wiener. Hence

$$\ln(L_t) = \int_0^t \varphi_s' dW_s^Q + \frac{1}{2}\int_0^t \|\varphi_s\|^2 ds \qquad (3.29)$$

and

$$\begin{aligned} E^Q\left[\ln(L_T)|\,\mathcal{F}_t\right] &= \ln(L_t) + E^Q\left[\int_t^T \varphi_s' dW_s^Q + \frac{1}{2}\int_t^T \|\varphi_s\|^2 ds\,\Big|\,\mathcal{F}_t\right] \\ &= \ln(L_t) + H_t, \end{aligned}$$

14

where

$$H_t = \frac{1}{2} E^Q \left[\int_t^T \|\varphi_s\|^2 ds \,\Big|\, \mathcal{F}_t \right]. \tag{3.30}$$

Proposition 3.7 *With exponential utility and constant interest rate* r,
(i) The optimal wealth process is given by

$$X_t^* = e^{rt} x + e^{-r(T-t)} \frac{1}{\gamma} \{H_0 - H_t - \ln(L_t)\}, \tag{3.31}$$

(ii) The optimal portfolio investment strategy is

$$u_t^* = e^{-r(T-t)} \frac{1}{\gamma X_t} \left(\{\sigma_t^{-1} [\alpha_t - \mathbf{r}]\}' - \sigma_H(t) \right) \sigma_t^{-1}, \tag{3.32}$$

where σ_H *is obtained from the* H *dynamics as*

$$dH_t = \mu_H(t)dt + \sigma_H(t)dW_t. \tag{3.33}$$

Proof. From the definition (3.23), the process J_t is given in the case of constant interest rate by

$$
\begin{aligned}
J_t &= E^Q \left[e^{-r(T-t)} \{-r(T-t) + \ln(L_T/L_t) + \ln L_t\} \,\Big|\, \mathcal{F}_t \right] \\
&= \{-r(T-t) + \ln L_t + H_t\} e^{-r(T-t)}.
\end{aligned}
$$

In view of (3.29) and (3.33) we have

$$dJ_t = e^{-r(T-t)} \{\varphi_t' + \sigma_H(t)\} dW_t^Q + (\cdots)dt,$$

so that

$$\sigma_J(t) = \frac{1}{J_t} \{\varphi_t' + \sigma_H(t)\} e^{-r(T-t)}.$$

In this case $\sigma_B = 0$, so from (3.26) we obtain

$$u^*(t) = \frac{-\frac{1}{\gamma} D_{0,t} J_t \sigma_J(t)}{D_{0,t} X_t^*} \sigma_t^{-1} = -\frac{e^{-r(T-t)} \{\varphi_t' + \sigma_H(t)\}}{\gamma X_t^*} \sigma_t^{-1}.$$

4   The partially observable case

We now go back to the original partially observable model and recall that the stock price dynamics are given by

$$dS_t^i = \alpha_t^i S_t^i dt + S_t^i \sigma_t^i dW_t, \quad i = 1, \ldots, n.$$

where $\alpha^1, \ldots, \alpha^n$ are assumed to be $\mathbf{F}$-adapted, whereas $\sigma^1, \ldots, \sigma^n$ are assumed to be $\mathbf{F}^S$-adapted. We again stress that there is no assumption of a

Markovian structure. The interesting case is of course when the observable filtration $\mathbf{F}^S$ is strictly included in the "big" filtration $\mathbf{F}$. As before we write the S dynamics on vector form as

$$dS_t = D(S_t)\alpha_t dt + D(S_t)\sigma_t dW_t, \tag{4.1}$$

and we recall that σ is assumed to be invertible. Our problem is to maximize

$$E^P\left[U(X_T)\right]$$

over the class of $\mathbf{F}^S$ adapted self financing portfolios, subject to the initial wealth condition $X_0 = x$.

4.1   Projecting onto the observable filtration

The idea is to reduce the partially observable problem above to an equivalent problem with complete observations. To this end we define the process Z by

$$dZ_t = \sigma_t^{-1} D(S_t)^{-1} dS_t, \tag{4.2}$$

i.e.

$$dZ_t = \sigma_t^{-1}\alpha_t dt + dW_t.$$

Now we define, for any $\mathbf{F}$-adapted process Y, the filter estimate process $\hat{Y}$ as the optional projection of Y onto the $\mathbf{F}^S$ filtration, i.e.

$$\hat{Y}_t = E^P\left[Y_t \mid \mathcal{F}_t^S\right].$$

We go on to define the **innovations process** $\bar{W}$ by

$$d\bar{W}_t = dZ_t - \left(\widehat{\sigma_t^{-1}\alpha_t}\right)dt,$$

which, by the observability assumption on σ, can be written as

$$d\bar{W}_t = dZ_t - \sigma_t^{-1}\widehat{\alpha}_t dt, \tag{4.3}$$

From non linear filtering theory (see e.g. Liptser and Shiryayev (2004)) we recall the following result.

Lemma 4.1 *The innovations process $\bar{W}$ is a standard $\mathbf{F}^S$- Wiener process.*

We now write (4.3) as

$$dZ_t = \sigma_t^{-1}\widehat{\alpha}_t dt + d\bar{W}_t, \tag{4.4}$$

and note that this is the semimartingale representation of Z w.r.t the filtration $\mathbf{F}^S$. Replacing the dZ term in (4.2) by the expression given in (4.4), and rearranging terms, gives us

$$dS_t = D(S_t)\widehat{\alpha}_t dt + D(S_t)\sigma_t d\bar{W}_t. \tag{4.5}$$

This equation represents the dynamics of the S process w.r.t. to its internal filtration $\mathbf{F}^S$. Note that the S occurring in (4.5) is exactly the same process (omega by omega) as the one occurring in (4.1). The only difference is that in (4.1) we have the semimartingale representation of S with respect to the filtration $\mathbf{F}$, whereas in (4.5) we have the $\mathbf{F}^S$ semimartingale representation.

4.2 Solving the optimal control problem

Going back to our partially observed optimal control problem, we wanted to maximize

$$E^P[U(X_T)]$$

over the class of self financing $\mathbf{F}^S$ adapted portfolios, given the initial condition $X_0 = x$. The problematic feature was that the S dynamics were given by

$$dS_t = D(S_t)\alpha_t dt + D(S_t)\sigma_t dW_t,$$

where α was not observable.

However, we have just derived an alternative expression for the S dynamics, namely

$$dS_t = D(S_t)\widehat{\alpha}_t dt + D(S_t)\sigma_t d\bar{W}_t.$$

The point of this is that, since $\widehat{\alpha}$ is by definition adapted to $\mathbf{F}^S$, and $\bar{W}$ is $\mathbf{F}^S$-Wiener, we now have a *completely observable* investment problem in a *complete* market. Since this problem is exactly of the form treated in Section 3, this means that we only need to copy and paste from Section 3 in order to obtain the solution to the partially observable problem. The only difference will be that whenever we have an expression involving α in the results from Section 3, we have to replace α by $\widehat{\alpha}$. We then need some new notation in order to see the difference between the formulas in the completely observable and the partially observable cases.

Remark 4.1 *Concerning the notation below. We reserve $\widehat{\cdot}$ exclusively to denote filter estimates, whereas $\bar{\cdot}$ us used to indicate objects in the partially observed model. For example: $\bar{H}_t$ denotes the process which in the partially observed model has the role that H had in the fully observable model. In contrast, the process $\widehat{H}_t$ would denote the filter estimate of H.*

We start by defining the appropriate martingale measure on $\mathbf{F}^S$.

Definition 4.1 *The $\mathbf{F}^S$ martingale measure $\bar{Q}$ is defined by*

$$\frac{d\bar{Q}}{dP} = \bar{L}_t, \quad on \ \mathcal{F}_t^S,$$

with $\bar{L}$ dynamics given by

$$d\bar{L}_t = \bar{L}_t \left\{ \sigma_t^{-1} \left(\mathbf{r}_t - \widehat{\alpha}_t \right) \right\}' d\bar{W}_t.$$

4.2.1 Power utility

Going to power utility we need to define the $\mathbf{F}^S$ analogues of the measure Q^0 and the process H.

Definition 4.2

- *The measure $\bar{Q}^0$ is defined by*

$$\frac{d\bar{Q}^0}{dP} = \bar{L}_t^0, \quad on \ \mathcal{F}_t^S, \tag{4.6}$$

  with $\bar{L}^0$ given by

$$d\bar{L}_t^0 = \bar{L}_t^0 \beta \left\{ \sigma_t^{-1} \left(\mathbf{r}_t - \widehat{\alpha}_t \right) \right\} d\bar{W}_t. \tag{4.7}$$

- *The process $\bar{H}$ is defined by*

$$\bar{H}_t = E^0 \left[\exp \left\{ \beta \int_t^T \left(r_s + \frac{1}{2(1-\gamma)} \| \sigma_s^{-1} \left(\widehat{\alpha}_s - \mathbf{r}_s \right) \|^2 \right) ds \right\} \Bigg| \mathcal{F}_t \right]. \tag{4.8}$$

We now have the following results. They all follow directly from the corresponding results for the completely observable case.

Proposition 4.1 (Power utility)
With notation as above, the following hold.

- *The optimal wealth process $\bar{X}^*$ is given by*

$$\bar{X}_t^* = \frac{\bar{H}_t}{\bar{H}_0} (D_{0,t} \bar{L}_t)^{\frac{-1}{1-\gamma}} x,$$

  where $\bar{H}$ is given above and the expectation is taken under $\bar{Q}^0$.

- *The optimal portfolio weight vector $\bar{u}^*$ is given by*

$$\bar{u}_t^* = \frac{1}{1-\gamma} \left(\widehat{\alpha}_t - \mathbf{r}_t \right)' \left(\sigma_t \sigma_t' \right)^{-1} + \sigma_{\bar{H}}(t) \sigma_t^{-1}$$

  where $\sigma_{\bar{H}}$ is the volatility term of $\bar{H}$, i.e. $\bar{H}$ has dynamics of the form

$$d\bar{H}_t = \bar{H}_t \mu_{\bar{H}}(t) dt + \bar{H}_t \sigma_{\bar{H}}(t) d\bar{W}_t.$$

4.2.2 Log utility

For log utility we immediately have the following result.

Proposition 4.2 (Log utility) *For the log utility case, the following hold.*

- *The optimal portfolio process $\bar{X}^*$ is given by*

$$\bar{X}_t^* = (D_{0,t}\bar{L}_t)^{-1}x,$$

  where the likelihood process $\bar{L}$ is given above.

- *The optimal portfolio weight vector process $\bar{u}^*$ is given by*

$$\bar{u}_t^* = (\widehat{\alpha}_t - \mathbf{r}_t)' \left(\sigma_t\sigma_t'\right)^{-1}$$

4.2.3 Exponential utility

We readily have the following result.

Proposition 4.3 (Exponential utility)

- *The optimal wealth process is given by*

$$\bar{X}_t^* = \left(\frac{x + \frac{1}{\gamma}\bar{J}_0}{B_{0,T}} - \frac{1}{\gamma}\ln D_{0,t}\right)B_{t,T} - \frac{1}{\gamma}\bar{J}_t.$$

  where

$$\bar{J}_t = E^{\bar{Q}}[D_{t,T}\ln(D_{t,T}\bar{L}_T)|\mathcal{F}_t], \quad t \in [0,T].$$

- *The optimal portfolio, in terms of the optimal weights on the risky assets, is given by*

$$\bar{u}^*(t) = \left(\frac{\bar{A}(t)B_{t,T}\sigma_B(t) - \frac{1}{\gamma}D_{0,t}\bar{J}_t\sigma_{\bar{J}}(t)}{\bar{A}(t)B_{t,T} - \frac{1}{\gamma}D_{0,t}\bar{J}_t}\right)\sigma^{-1}(t).$$

  where $\sigma_{\bar{J}}$ is obtained from the $\bar{J}$ dynamics as

$$d\bar{J}_t = (\cdots)dt + \sigma_{\bar{J}}(t)\bar{J}_td\bar{W}_t,$$

  and $\bar{A}(t)$ is given by (3.25) with J_0 replaced by $\bar{J}_0$.

5 The Markovian case

In order to obtain more explicit results, and to connect to the earlier literature we now make the assumption that we have a Markovian system. We start with the general setting and then go on to the concrete cases of power, log, and exponential utility. For these special cases, we show that the optimal strategy can be computed explicitly, up to the solution of a *linear* PDE with infinite dimensional state space, in contrast to the nonlinear infinite-dimensional Hamilton-Jacobi-Bellman equation required for general stochastic control problems—see Bensoussan (2004).

5.1 Generalities and the DMZ equation

The model is specified as follows.

Assumption 5.1 *We assume that the asset price dynamics are of the form*

$$dS_t^i = \alpha_t^i(Y_t)S_t^i dt + S_t^i \sigma_t^i dW_t, \quad i = 1, \ldots, n.$$

Here $\alpha_t^i(y)$ is assumed to be a deterministic function of t and y, whereas σ_t^i is a deterministic function of t. The process Y is assumed to be a time homogeneous Markov process, independent of W, living on the state space $\mathcal{Y}$ and having generator $\mathcal{A}$. We assume that the short rate r_t is deterministic.

The two most typical special cases of the setup above is that either Y is a (possibly multi dimensional) diffusion or that Y is a Markov chain living on a finite state space. The independence assumption between S and Y is not really needed, but leads to simpler calculations. The assumed time invariance for Y is only for notational simplicity and can easily be relaxed. We could also allow α^i and σ^i to be adapted to the filtration $\mathbf{F}^S$. It might seem natural to suppose that the short rate is a function $r(t) = r(t, Y_t)$ of the factor process, but this does not fit into the nonlinear filtering framework since we would then have a noise-free observation $r(t, Y_t) = \frac{d}{dt}B(t)$ of Y_t. On the other hand, if we continue to assume that r_t is a general $\mathbf{F}^S$-adapted process then we will not obtain the 'Markovian' results below.

In this setting we may apply standard non-linear filtering theory (see Liptser and Shiryayev (2004)) and to do so we need a regularity assumption.

Assumption 5.2 *We assume that the conditional P-distribution of Y_t given $\mathcal{F}_t^S$ admits a density $p_t(y)$ on $\mathcal{Y}$, relative to some dominating measure $m(dy)$.*

If Y is a diffusion on R^n, the measure $m(dy)$ will be n-dimensional Lebesgue measure, and in the case of Y being a finite state Markov chain,

$m(dy)$ will be the counting measure on the (finite) set Y. Given the assumption above we can thus write conditional expectations as integrals w.r.t. $p_t(y)$. More precisely, for any function $f : \mathcal{Y} \to R$ we have

$$E^P \left[f(Y_t) | \mathcal{F}_t^S \right] = \int_{\mathcal{Y}} f(y) p_t(y) m(dy).$$

In the language of non linear filtering we thus have the *signal process Y*, and the *observation process Z*, where Z is given by

$$dZ_t = \sigma_t^{-1} \alpha_t(Y_t) dt + dW_t.$$

We can now write down the Kushner-Stratonovich (KS) equation for the conditional density p_t in our model.

Theorem 5.1 (The KS equation) *With assumptions as above the dynamics of the conditional density are given by*

$$dp_t(y) = \mathcal{A}^* p_t(y) dt + p_t(y) \left\{ \sigma_t^{-1} \left[\alpha_t(y) - \hat{\alpha}_t(p_t) \right] \right\}' d\bar{W}_t. \qquad (5.1)$$

Here, $\mathcal{A}^*$ is the adjoint to $\mathcal{A}$, $\hat{\alpha}$ is given by

$$\hat{\alpha}_t(p_t) = \int_Y \alpha_t(y) p_t(y) m(dy), \qquad (5.2)$$

and

$$d\bar{W}_t = dZ_t - \sigma_t^{-1} \hat{\alpha}_t(p_t) dt.$$

We note that
$$\hat{\alpha}_t(p_t) = \widehat{\alpha_t(Y_t)}$$

where $\hat{\alpha}$ in the left hand side denotes a deterministic function, whereas the $\hat{}$ sign in the right hand side denotes a filter estimate. More precisely, if we denote the convex space of densities on Y by $\mathcal{H}$, then $\hat{\alpha}$ is a deterministic mapping $\hat{\alpha} : \mathcal{H} \times R_+ \to R$, where

$$(p, t) \longmapsto \int_Y \alpha_t(y) p(y) m(dy).$$

We can now view the KS equation as a single infinite dimensional SDE, writing it as

$$dp_t = \mu_p(t, p_t) dt + \sigma_p(t, p_t) d\bar{W}_t,$$

with μ_p and σ_p defined by the KS equation above, so the conditional density process p will be Markovian. The main point of this is that, in a Markovian setting, a conditional expectation like the ones defining $\bar{H}$ in the power and exponential cases above, will be a deterministic function of the state variable

p, and it will also satisfy a PDE (the Kolmogorov backward equation) on an inherently infinite dimensional state space. We will thus be able to provide an explicit solution to the optimal investment problem, up to the solution of a PDE. We note that instead of using the standard conditional density process $p_t(y)$ and the KS equation we could instead use the un-normalized density $q_t(y)$ and the Zakai equation. See Section 5.5 below for details.

5.2 Power utility

For the power case we of course still have Proposition 4.1, but in the present Markovian setting we can obtain more explicit formulas for the processes $\bar{H}$ and u^*. We start by noticing that the measure $\bar{Q}^0$ in (4.6)-(4.7) has likelihood dynamics given by

$$d\bar{L}_t^0 = \bar{L}_t^0 \beta \left\{\sigma_t^{-1}\left[\mathbf{r}_t - \widehat{\alpha}_t(p_t)\right]\right\}' d\bar{W}_t. \tag{5.3}$$

where $\hat{\alpha}_t(p)$ is given by (5.2). The $\bar{Q}^0$ dynamics of the conditional density process are thus, by Girsanov, given by

$$dp_t = \mu_p^0(t, p_t)dt + \sigma_p(t, p_t)d\bar{W}_t^0, \tag{5.4}$$

where

$$\mu_p^0(t, p) = \mu_p(t, p) + \beta\sigma_p(t, p)\sigma_t^{-1}\left[\mathbf{r}_t - \widehat{\alpha}_t(p)\right] \tag{5.5}$$

Furthermore, the process $\bar{H}$ defined by (4.8) will now have the more specific form

$$\bar{H}_t = E^0\left[\exp\left(\int_t^T \left\{r_s + \frac{\beta}{2(1-\gamma)}\|\sigma_s^{-1}\left(\hat{\alpha}_s(p_s) - \mathbf{r}_s\right)\|^2\right\} ds\right)\bigg|\mathcal{F}_t^S\right],$$

and the Markovian structure of the p process will allow us to write the $\bar{H}$ process as

$$\bar{H}_t = \bar{H}(t, p_t),$$

where $\bar{H}$ on the right hand side denotes a deterministic function of the variables t and p, defined by

$$\bar{H}(t, p) = E_{t,p}^0\left[\exp\left(\int_t^T \left\{r_s + \frac{\beta}{2(1-\gamma)}\|\sigma_s^{-1}\left(\hat{\alpha}_s(p_s) - \mathbf{r}_s\right)\|^2\right\} ds\right)\right]. \tag{5.6}$$

Thus the function $\bar{H}$ will solve the following Kolmogorov backward equation (see Da Prato and Zabzcyk (1996)), where $\mathbf{Tr}$ denotes the trace.

$$\frac{\partial\bar{H}}{\partial t} + \frac{\partial\bar{H}}{\partial p}\mu_p^0 + \frac{1}{2}\mathbf{Tr}\left[\sigma_p'\frac{\partial^2\bar{H}}{\partial p^2}\sigma_p^2\right] + \left(r + \frac{\beta}{2(1-\gamma)}\|\sigma^{-1}\left(\hat{\alpha} - \mathbf{r}\right)\|^2\right)\bar{H} = 0 \tag{5.7}$$

$$\bar{H}(T, p) = 1 \tag{5.8}$$

22

Note that this is a PDE in the infinite dimensional state variable p, so the partial derivatives terms w.r.t. p above are Frechet derivatives.

In the Markovian setting, the optimal portfolio weight vector process $\bar{u}^*$ will have the form of a feedback control, i.e. it will be of the form

$$\bar{u}_t^* = \bar{u}^*(t, p_t),$$

where $\bar{u}*$ in the right hand side denotes a deterministic function defined by

$$\bar{u}^*(t, p) = \frac{1}{1 - \gamma} \left(\widehat{\alpha}_t(p) - \mathbf{r}_t \right)' \left(\sigma_t \sigma_t' \right)^{-1} + \frac{1}{\bar{H}(t, p)} \frac{\partial \bar{H}}{\partial p}(t, p) \sigma_p(t, p) \sigma_t^{-1}.$$

This follows directly from the fact that $\sigma_{\bar{H}}$ in Proposition 4.1 can, in the Markovian case, be computed explicitly by using the Itô formula.

5.3 Log utility

The log utility case is trivial. The optimal strategy given in Proposition 4.2 can, in the Markovian framework, be written in feedback form as

$$\bar{u}^*(t, p) = \left(\widehat{\alpha}_t(p) - \mathbf{r}_t \right)' \left(\sigma_t \sigma_t' \right)^{-1}.$$

5.4 Exponential utility

Since in this section $r(t)$ is deterministic, we can use the simpler expressions of Section 3.3.3. The process $\bar{H}_t$, i.e. H_t defined by (3.30) with α replaced by $\widehat{\alpha}$, can be written as

$$\bar{H}_t = \bar{H}(t, p_t),$$

where the function $\bar{H}(t, p)$ is defined as

$$\bar{H}(t, p) = \frac{1}{2} E_{t,p}^Q \left[\int_t^T \| \sigma_s^{-1} (\widehat{\alpha}_s(p_s) - \mathbf{r}_s) \|^2 ds \right],$$

and where $\bar{H}(t, p)$ will satisfy the Kolmogorov backward equation

$$\frac{\partial \bar{H}}{\partial t} + \frac{\partial \bar{H}}{\partial p} \mu_p^Q + \frac{1}{2} \mathbf{Tr} \left[\sigma_p' \frac{\partial^2 \bar{H}}{\partial p^2} \sigma_p^2 \right] + \frac{1}{2} \| \sigma^{-1} (\widehat{\alpha} - \mathbf{r}) \|^2 = 0 \qquad (5.9)$$

$$H(T, p) = 0. \qquad (5.10)$$

Here μ_p^Q is defined by

$$\mu_p^Q(t, p) = \mu_p(t, p) + \sigma_p(t, p) \sigma_t^{-1} \left[\mathbf{r}_t - \widehat{\alpha}_t(p) \right]$$

The optimal portfolio, in terms of the optimal weights on the risky assets, is given in feedback form as

$$\bar{u}^*(t, p, \bar{x}) = e^{-r(T-t)} \frac{1}{\gamma \bar{x}} \left(\left\{ \sigma_t^{-1} \left[\widehat{\alpha}(t, p) - \mathbf{r}_t \right] \right\}' - \frac{\partial \bar{H}}{\partial p}(t, p) \sigma_p(t, p) \right) \sigma_t^{-1}.$$

$$(5.11)$$

5.5 The Zakai equation

An alternative to using the KS equation in (5.1) above, and the related PDEs (5.7)-(5.8) and (5.9)-(5.10) is to use the Zakai un-normalized conditional density process $q_t(y)$. This density satisfies the Zakai equation

$$dq_t(y) = \mathcal{A}^* q_t(y) dt + q_t(y) \sigma_t^{-1} \alpha_t(y) dZ_t. \qquad (5.12)$$

and the advantage of using the Zakai equation is that it is much simpler than the DMZ equation. It is driven directly by the observation process Z, and the drift and diffusion terms are linear in q. The relation between q and p is given by

$$p_t(y) = \frac{q_t(y)}{\int_{\mathcal{Y}} q_t(u) du},$$

and the results on Sections 5.2, 5.3, and 5.4 can easily be transferred to the q formalism.

5.6 Finite dimensional filters

As we have seen above, for the general hidden Markov model the optimal investment strategy u^* is a deterministic function $u^*(t, p, \bar{x})$ of running time t, the conditional density p_t, and the wealth $\bar{X}_t$. Furthermore, we can compute the optimal investment strategy u^* explicitly up to the solution of a PDE (the Kolmogorov backward equation) on an infinite dimensional state space. For a general model, we then have two closely related computational problems.

- The DMZ filter equations (5.1) describes an infinite dimensional SDE, driven by the innovations process. In a concrete application, the filter could thus never be implemented exactly, so one would have to construct an approximate finite dimensional filter.

- As a consequence of the infinite dimensionality of the filter, the Kolmogorov equations above are generically PDEs on an infinite dimensional state space and thus very hard to solve.

In order to simplify the situation it is thus natural to study models where the state space is of finite dimension. This occurs if and only if the DMZ equation evolves on a finite dimensional sub manifold of the inherently infinite dimensional convex space of probability densities, in other words if and only if the associated filtering problem has a finite dimensional filter. It is furthermore well known from filtering theory that the existence of a finite dimensional filter is a very rare phenomenon, related to the finite dimensionality of the Lie algebra generated by the drift and diffusion operators of

the Zakai equation. The two main cases where the filter is finite dimensional are the following:

- The case when Y is a finite state Markov chain, leading to the Wonham filter.

- The case when Y is the solution of a linear SDE, leading to the Kalman filter.

These cases are (apart from Bäuerle and Rieder (2007) and Callegaro et al. (2006)) precisely the cases studied previously in the literature. The linear diffusion case is studied in Brendle (2004, 2006), Brennan (1998), Brennan and Xia (2001), Cvitanic et al. (2006), Gennotte (1986), Lakner (1995), Lakner (1998), and Xia (2001), whereas the Markov chain model is treated in Bäuerle and Rieder (2005, 2004), Haussmann and Sass (2003, 2004a,b), Nagai and Runggaldier (2006), and Sass (2007).

References

Bäuerle, N., Rieder, U., 2004. Portfolio optimization with Markov-modulated stock prices and interest rates. IEEE Trans. Automat. Control 49 (3), 442–447.

Bäuerle, N., Rieder, U., 2005. Portfolio optimization with unobservable markov-modulated drift process. Journal of Applied Probability 42 (2), 362–278.

Bäuerle, N., Rieder, U., 2007. Portfolio optimization with jumps and unobservable intensity. Mathematical Finance 17 (2), 205–224.

Bensoussan, A., 2004. Stochastic Control of Partially-Observable Systems, 2nd Edition. Cambridge University Press.

Brendle, S., 2004. Portfolio selection under partial observation and constant absolute risk aversion, working paper, Princeton University.

Brendle, S., 2006. Portfolio selection under incomplete information. Stochastic processes and their Applications 116 (5), 701–723.

Brennan, M., 1998. The role of learning in dynamic portfolio decisions. European Finance Review , 295–306.

Brennan, M., Xia, Y., 2001. Assessing asset pricing anomalies. Review of Financial Studies 14 (4), 905–942.

Callegaro, G., Di Masi, G., Runggaldier, W., 2006. Portfolio optimization in discontinuous markets under incomplete ionformation. Asia-Pacific Financial Markets 13, 373–394.

Cvitanic, J., Lazrak, A., Martinelli, L., Zapatero, F., 2006. Dynamic portfolio choice with parameter uncertainty and the economic value of analysts' recommendations. Review of Financial Studies 19, 1113–1156.

Da Prato, G., Zabzcyk, J., 1996. Ergodicity for Infinite Dimensional Systems. Cambridge University Press.

Gennotte, G., 1986. Optimal portfolio choice under incomplete information. Journal of Finance 41, 733749.

Haussmann, U. G., Sass, J., 2003. Optimal portfolio choice for unobservable and regime-switching mean returns. Journal of Economic Dynamics and Control 28, 45–78.

Haussmann, U. G., Sass, J., 2004a. Optimal terminal wealth under partial information for HMM stock returns. In: Mathematics of Finance (Contemp. Math. 351). AMS.

Haussmann, U. G., Sass, J., 2004b. Optimizing the terminal wealth under partial information: The drift process as a continuous time markov chain. Finance and Stochastics 8, 553–577.

Lakner, P., 1995. Utility maximization with partial information. Stochastic Processes and their Applications 56, 247–249.

Lakner, P., 1998. Optimal trading strategy for an investor: the case of partial information. Stochastic Processes and their Applications 76, 77–97.

Liptser, R., Shiryayev, A., 2004. Statistics of Random Processes, 2nd Edition. Vol. I. Springer Verlag.

Nagai, H., Runggaldier, W., 2006. PDE approach to utility maximization for market models with hidden markov factors. In: 5th Seminar on Stochastic Analysis, Random Fields and Applications. Birkhauser Verlag.

Sass, J., 2007. Utility maximization with convex costraints and partial information. Acta Appl Math 97, 221–238.

Xia, Y., 2001. Learning about predictability: the effects of parameter uncertainty on dynamic asset allocation. Journal of Finance 56, 205–246.