

Portfolio Optimization under Partial Information with Expert Opinions

Rüdiger Frey * Abdelali Gabih † Ralf Wunderlich ‡

This paper investigates optimal portfolio strategies in a market with partial information on the drift. The drift is modelled as a function of a continuous-time Markov chain with finitely many states which is not directly observable. Information on the drift is obtained from the observation of stock prices. Moreover, expert opinions in the form of signals at random discrete time points are included in the analysis. We derive the filtering equation for the return process and incorporate the filter into the state variables of the optimization problem. This problem is studied with dynamic programming methods. In particular, we propose a policy improvement method to obtain computable approximations of the optimal strategy. Numerical results are presented at the end.

Keywords: Portfolio optimization, hidden Markov model, dynamic programming

1. Introduction

It is well-known that the drift of asset prices has a crucial impact on the optimal trading strategy in dynamic portfolio optimization problems. At the same time this parameter is notoriously difficult to estimate from historical asset price data: first, drifts tend to fluctuate randomly over time; second, even if drifts were constant, a long time series is needed to estimate this parameter with a reasonable degree of precision as drift effects are usually dominated by volatility. For these reasons practitioners rely mostly on external sources of information such as news, company reports or ratings and on their own intuitive views when determining an estimate for the future growth rate of an asset; these outside sources of information are labelled expert opinions in this paper. The popular Black-Litterman model (see Black and Litterman (1992), Schöttle et al. (2010)) where subjective views are used to update equilibrium-implied returns in a Bayesian way is a typical example for the use of expert opinions in *static* (one-period) models. However, to the best of our knowledge expected utility maximization in *dynamic* portfolio optimization models with expert opinions has so far not been studied.

In the present paper we set out to do exactly that. We consider a hidden Markov model (HMM) where asset prices follow a diffusion process whose drift is driven by an unobservable finite-state Markov chain Y . Information on the hidden chain

*Department of Mathematics, University of Leipzig, D-04009 Leipzig, Germany,
e-mail: frey@math.uni-leipzig.de

†Laboratoire de Technologie de l'Information et Modélisation (TIM), ENSA-Université de Marrakech, Boulevard Abdelkrim El Khattabi BP 575, Marrakech, Morocco, e-mail: gabih@ensa.ac.ma

‡Mathematics Group, Zwickau University of Applied Sciences, PSF 201037,
D-08056 Zwickau, Germany, e-mail: ralf.wunderlich@fh-zwickau.de

2 *R. Frey, A. Gabih and R. Wunderlich*

is of mixed type. First, investors observe stock prices. Moreover, and this is the novel feature of this paper, expert opinions are included in the analysis as a second source of information. Mathematically, expert opinions are represented by a marked point process with jump-size distribution depending on the current state of Y . Standard filtering results for HMMs and Bayesian updating are used to derive a finite-dimensional filter for the state of the hidden Markov chain. This allows us to reduce the portfolio optimization problem to a problem under complete information where the new state variables are the filter distribution and the wealth of the investor. In this model the market is incomplete, as the investor filtration is partly generated by the non-tradable marked point process that models the expert opinions. This makes the application of duality methods and of the martingale approach to portfolio optimization relatively involved. Hence we resort to dynamic programming and work with the associated Hamilton-Jacobi-Bellman (HJB) equation instead. We consider the case of logarithmic and power utilities. In the latter case the HJB equation can be simplified by a change of measure and we end up with a quasi-linear integro-differential equation. Finally we propose a policy improvement method to obtain an approximation of the optimal strategy.

Portfolio optimization under partial information on the drift has been studied extensively over the last years. There are two popular model classes for the drift, linear Gaussian dynamics and HMMs. For Gaussian dynamics explicit solutions for the problem of optimizing the expected utility of terminal wealth are provided for example in Lakner (1998), Brendle (2006), Danilova et al. (2010), where the last paper focuses on additional insider information. Utility maximization for a HMM model is investigated for example in Rieder and Bäuerle (2005), Sass and Haussmann (2004), Sass and Wunderlich (2010) and Gabih et al. (2009). These approaches are generalized in Björk et al. (2010). In the present paper we follow Rieder and Bäuerle (2005) for the setup of the HJB equation in a model with an unobservable drift modelled by a finite-state Markov chain. Moreover, we were inspired by the change of measure technique used among others by Nagai and Runggaldier (2008) and Davis and Lleo (2010).

The paper is organized as follows. In Section 2 we introduce our model of the financial market and formulate the portfolio optimization problem. Section 3 is devoted to the filtering problem for the unobservable drift. Section 4 treats the optimization problem for the special case of logarithmic utility, Section 5 is devoted to the case of power utility. In Section 6 we discuss approximation methods for the optimal strategy; numerical results are presented in Section 7.

2. Financial Market Model

Fix some date $T > 0$ representing the investment horizon. We work on a filtered probability space $(\Omega, \mathcal{G}, \mathbb{G}, P)$, with filtration $\mathbb{G} = (\mathcal{G}_t)_{t \in [0, T]}$ satisfying the usual conditions. All processes are assumed to be \mathbb{G} -adapted. For a generic \mathbb{G} -adapted process H we denote by \mathbb{G}^H the filtration generated by H .

Price dynamics. We consider a market model for one risk-free bond with price $S_t^0 = 1$ and n risky securities with prices $S_t = (S_t^1, \dots, S_t^n)^\top$ given by

$$dS_t^i = S_t^i \left(\mu^i(Y_t) dt + \sum_{j=1}^n \sigma^{ij} dW_t^j \right), \quad S_0^i = s^i, \quad i = 1, \dots, n. \quad (2.1)$$

Here $\mu = \mu(Y_t) \in \mathbb{R}^n$ denotes the mean stock return or drift which is driven by some factor process Y described below. The volatility $\sigma = (\sigma^{ij})_{1 \leq i, j \leq n}$ is assumed to be a constant invertible matrix and $W_t = (W_t^1, \dots, W_t^n)$ is an n -dimensional \mathbb{G} -adapted Brownian motion. The invertibility of σ always can be ensured by a suitable parametrization if the covariance matrix $\sigma \sigma^\top$ is positive definite. The factor process Y is a finite-state Markov chain independent of the Brownian motion W with state space $\{e_1, \dots, e_d\}$ where e_i is the i th unit vector in \mathbb{R}^d . The generator matrix is denoted by Q and the initial distribution by $\pi = (\pi^1, \dots, \pi^d)^\top$. The states of the factor process Y are mapped onto the states μ_1, \dots, μ_d of the drift by the function $\mu(Y_t) = MY_t$, where $M_{lk} = \mu^l(e_k)$, $1 \leq l \leq n$, $1 \leq k \leq d$.

Define the return process R associated with the price process S by $dR_t^i = dS_t^i/S_t^i$, $i = 1, \dots, n$. Note that R satisfies

$$dR_t = \mu(Y_t) dt + \sigma dW_t,$$

so that the quadratic variation of R^i and of $\log S^i$ coincide, $[R^i]_t = [\log S^i]_t$. Since moreover $R_t^i = \log(S_t^i) + \frac{1}{2}[R^i]_t$ we have the equality $\mathbb{G}^R = \mathbb{G}^{\log S} = \mathbb{G}^S$. This is useful, since it allows us to work with R instead of S in the filtering part.

Investor Information. We assume that the investor does not observe the factor process Y directly; he does however know the model parameters, in particular the initial distribution π , the generator matrix Q and the functions $\mu^i(\cdot)$. Moreover, he has noisy observations of the hidden process Y at his disposal. More precisely we assume that the investor observes the return process R and that he receives at discrete points in time T_n noisy signals about the current state of Y . These signals are to be interpreted as expert opinions; a number of examples is given below.

We model expert opinions by a marked point process $I = (T_n, Z_n)$, so that at T_n the investor observes the realisation of a random variable Z_n whose distribution depends on the current state Y_{T_n} of the factor process. The T_n are modelled as jump times of a standard Poisson process with intensity λ , independent of Y , so that the timing of the information arrival does not carry any useful information. The signal Z_n takes values in some measurable space \mathcal{Z} with reference measure dz . Examples are a discrete space with the counting measure or $\mathcal{Z} = \mathbb{R}^N$ with the Lebesgue measure. We assume that the Z_n are conditionally independent given \mathcal{F}_T^Y and that the distribution of Z_n is absolutely continuous w.r.t. the reference measure dz with density $f(Y_{T_n}, z)$. We identify the marked point process $I = (T_n, Z_n)$ with the associated counting measure denoted by $I(dt \times dz)$. Note that the \mathbb{G} -compensator of I is $\lambda dt f(Y_t, z) dz$.

Summarizing, the information available to the investor is given by the *investor filtration* \mathbb{F} with

$$\mathcal{F}_t = \mathcal{G}_t^R \vee \mathcal{G}_t^I, \quad 0 \leq t \leq T. \quad (2.2)$$

Next we give some simple examples for the random variables Z_n that are inspired by the Black-Litterman approach, see for example Schöttle et al. (2010).

Example 2.1. In the Black-Litterman framework one distinguishes so-called absolute and relative views of an investor. An *absolute view* is a prediction on the return of a single asset; it might take the form “asset i has a return of 5%”. Moreover, the investor might specify the confidence in his views. This can be modelled by taking

$$Z_n = (Z_n^{(1)}, Z_n^{(2)}) \text{ with } Z_n^{(1)} = \mu_i(Y_{T_n}) + Z_n^{(2)}\varepsilon_n.$$

Here (ε_n, Z_n) is an iid sequence that is independent of Y ; $\varepsilon_n \sim \mathcal{N}(0, 1)$; $Z_n^{(2)}$ follows a given distribution on $(0, \infty)$ and $Z_n^{(1)}$ and $Z_n^{(2)}$ are independent. A realization $Z_n = (0.05, 0.02)$ means that the investor forecasts a growth rate of 5% and that he believes that the standard deviation of the prediction error of his forecast is 2% (which would correspond to a high level of confidence); a realization $Z_n = (0.05, 0.05)$ on the other hand corresponds to an investor who believes that the standard deviation of his prediction error is 5% reflecting a low level of confidence. A high (low) level of confidence implies that the current view of the investor has a strong (weak) impact on his filter estimate for the drift as can be seen formally from the Bayesian updating formula (3.2) below. The special case where the investors’ confidence does not vary is included by setting $Z_n^{(2)} = \sigma_\varepsilon$ for some constant $\sigma_\varepsilon > 0$.

A *relative view* might take the form “on average asset i outperforms asset j by 2%”. This can be modelled by taking

$$Z_n^{(1)} = \mu_i(Y_{T_n}) - \mu_j(Y_{T_n}) + Z_n^{(2)}\varepsilon_n,$$

where, as before, $Z_n^{(2)}$ is used to model the investors’ confidence.

Finally we remark that in the Black-Litterman approach the measure P should be seen as the subjective probability measure of the investor. We do not claim that his views are in fact correct or that his predictions are actually unbiased. Rather we only assume that the investor *believes* that his predictions are unbiased.

Portfolios and optimization problem. We describe the self-financing trading strategy of an investor by the initial capital $x_0 > 0$ and the n -dimensional \mathbb{F} -adapted process h where h_t^i , $i = 1, \dots, n$, is the proportion of wealth invested in stock i at time t . It is well-known that in this setting the wealth process $X^{(h)}$ has the dynamics

$$\frac{dX_t^{(h)}}{X_t^{(h)}} = \sum_{i=0}^n h_t^i \frac{dS_t^i}{S_t^i} = h_t^\top \mu(Y_t) dt + h_t^\top \sigma dW_t, \quad X_0^{(h)} = x_0. \quad (2.3)$$

Since $\mu(Y_t)$ is bounded and σ is constant, equation (2.3) is well defined if $\int_0^T \|h_s\|^2 ds < \infty$. It will be useful to impose the stronger requirement

$$E\left(\exp\left\{\int_0^T \|h_s\|^2 ds\right\}\right) < \infty. \quad (2.4)$$

A trading strategy satisfying this condition is called admissible; the class of *admissible trading strategies* will be denoted by \mathcal{H} .

We assume that the investor wants to maximize the expected utility of terminal wealth for logarithmic utility $U(x) = \log(x)$ and power utility $U(x) = \frac{1}{\theta}x^\theta$, $\theta < 1$, $\theta \neq 0$. The optimization problem thus reads as

$$\max\{E(U(X_T^{(h)})) : h \in \mathcal{H}\}. \quad (2.5)$$

This is a maximization problem under partial information since we have required that the strategy h is adapted to the investor filtration \mathbb{F} . Note that for $x_0 > 0$ the solution of the SDE (2.3) is strictly positive. This guarantees that $X_T^{(h)}$ is in the domain of logarithmic and power utility.

3. Partial information and filtering

In this section we explain how the control problem (2.5) can be reduced to a control problem with complete information via filtering arguments. We use the following notation: for a generic process H we denote by $\widehat{H}_t = E(H|\mathcal{F}_t)$ its optional projection on the filtration \mathbb{F} , and the filter for the Markov chain Y_t is denoted by $p_t = (p_t^1, \dots, p_t^d)$ with $p_t^k = P(Y_t = e_k|\mathcal{F}_t)$, $k = 1, \dots, d$. Note that for a process of the form $H_t = h(Y_t)$ the optional projection is given by $\widehat{h}(Y_t) = \sum_{k=1}^d h(e_k)p_t^k$. In particular, the projection of the drift is given by

$$\widehat{\mu}(Y_t) = \sum_{k=1}^d \mu_k p_t^k = Mp_t.$$

The following two processes will drive the dynamics of p_t . First, let

$$\widetilde{W}_t := \sigma^{-1}\left(R_t - \int_0^t Mp_s ds\right).$$

By standard results from filtering theory \widetilde{W} is an \mathbb{F} -Brownian motion (the so-called innovations process). Second, define the predictable random measure

$$\nu(dt \times dz) = \lambda dt \sum_{k=1}^d p_{t-}^k f(e_k, z) dz.$$

By standard results on point processes ν is the \mathbb{F} -compensator of I , see for instance Bremaud (1981). The compensated random measure will be denoted by $\gamma(dt \times dz) := I(dt \times dz) - \nu(dt \times dz)$.

6 *R. Frey, A. Gabih and R. Wunderlich*

Filtering. Next we use filtering results in order to derive a stochastic differential equation (SDE) for the filter p_t . We start with the situation where the investor can only observe the return process R . In that case we are in the classical situation of a hidden Markov model and we can use the standard filter for that case (see for example Wonham (1965), Elliott et al. (1994), Liptser and Shiryaev (2001)): It is well-known that p_t solves the SDE system

$$dp_t^k = \sum_{j=1}^d Q^{jk} p_t^j dt + p_t^k \left(\sigma^{-1} (\mu_k - \sum_{j=1}^d p_t^j \mu_j) \right)^\top d\widetilde{W}_t, \quad p_0^k = \pi^k. \quad (3.1)$$

In the presence of the additional information contained in I we have to add an correction term to the above equation. Assume that at time T_n the investor observes Z_n . Denote by $p_{T_n-} = (p_{T_n-}^1, \dots, p_{T_n-}^d)$ the a-priori probabilities before the arrival of this new observation. Then the a-posteriori probabilities $p_{T_n}^k$, $k = 1, \dots, d$, are determined from Bayes formula as follows:

$$p_{T_n}^k = \frac{p_{T_n-}^k f(e_k, Z_n)}{\bar{f}(p_{T_n-}, Z_n)} \quad \text{with} \quad \bar{f}(p, z) := \sum_{j=1}^d p^j f(e_j, z). \quad (3.2)$$

The increment $\Delta p_{T_n}^k := p_{T_n}^k - p_{T_n-}^k$ is thus given by

$$\Delta p_{T_n}^k = p_{T_n-}^k \left(\frac{f(e_k, Z_n)}{\bar{f}(p_{T_n-}, Z_n)} - 1 \right) = p_{T_n-}^k - \int_{\mathcal{Z}} \left(\frac{f(e_k, z)}{\bar{f}(p_{T_n-}, z)} - 1 \right) I(\{T_n\} \times dz). \quad (3.3)$$

By combining (3.1) and (3.3) we arrive at the following result.

Proposition 3.1. *The filter p solves the following d -dimensional SDE system*

$$dp_t^k = \sum_{j=1}^d Q^{jk} p_t^j dt + p_t^k \left(\sigma^{-1} (\mu_k - \sum_{j=1}^d p_t^j \mu_j) \right)^\top d\widetilde{W}_t + p_{t-}^k \int_{\mathcal{Z}} \left(\frac{f(e_k, z)}{\bar{f}(p_{t-}, z)} - 1 \right) \gamma(dt \times dz), \quad (3.4)$$

with initial condition $p_0^k = \pi^k$. and compensated jump measure $\gamma = I - \nu$

Proof. In view of (3.1) and (3.3), all that remains to show is the relation

$$\int_{[0, t] \times \mathcal{Z}} \left(\frac{f(e_k, z)}{\bar{f}(p_{s-}, z)} - 1 \right) \nu(ds \times dz) = 0 \quad \text{for all } t; \quad (3.5)$$

this allows us to replace the integral with respect to I in (3.3) by an integral with respect to the compensated jump measure γ . Using the representation $\nu(ds \times dz) =$

$\lambda ds \bar{f}(p_{s-}, z) dz$ and $\bar{f}(p_{s-}, z) = \sum_{j=1}^d p_{s-}^j f(e_j, z)$ we get

$$\begin{aligned} \int_{[0,t] \times \mathcal{Z}} \left(\frac{f(e_k, z)}{\bar{f}(p_{s-}, z)} - 1 \right) \nu(ds \times dz) &= \int_{[0,t] \times \mathcal{Z}} (f(e_k, z) - \bar{f}(p_{s-}, z)) \lambda ds dz \\ &= \int_{[0,t]} \left(\int_{\mathcal{Z}} f(e_k, z) dz - \sum_{j=1}^d p_{s-}^j \int_{\mathcal{Z}} f(e_j, z) dz \right) \lambda ds \\ &= \int_{[0,t]} \left(1 - \sum_{j=1}^d p_{s-}^j \cdot 1 \right) \lambda ds = 0, \end{aligned}$$

where we used the normalization properties $\int_{\mathcal{Z}} f(e_k, z) dz = 1$ and $\sum_{j=1}^d p_{s-}^j = 1$. \square

It is well-known (see for example Lakner (1998), Sass and Haussmann (2004)) that the \mathbb{F} -semimartingale decomposition of $X^{(h)}$ is given by

$$\frac{dX_t^{(h)}}{X_t^{(h)}} = h_t^\top M p_t dt + h_t^\top \sigma d\widetilde{W}_t. \quad (3.6)$$

Now note that for given $h \in \mathbb{R}^n$ the $(d+1)$ -dimensional process $(X^{(h)}, p)$ is an \mathbb{F} -Markov process as is immediate from the dynamics in (3.4) and (3.6). Hence the optimization problem (2.5) can be considered as a control problem under complete information with the $(d+1)$ -dimensional state variable process $(X^{(h)}, p)$. In Section 5 we will study this problem using dynamic programming techniques.

4. Logarithmic utility

In the case of logarithmic utility the optimization problem can be solved directly.

Lemma 4.1. *Suppose that $U(x) = \log x$, then the optimal strategy for problem (2.5) equals $h_t^* = (\sigma \sigma^\top)^{-1} \widehat{\mu}(Y_s)$.*

Proof. From (3.6) it follows that

$$U(X_T^{(h)}) = \log X_T^{(h)} = \log x_0 + \int_0^T \left(h_s^\top \widehat{\mu}(Y_s) - \frac{1}{2} \|\sigma^\top h_s\|^2 \right) ds + \int_0^T h_s^\top \sigma d\widetilde{W}_s.$$

For $h \in \mathcal{H}$ we have $E\left(\int_0^T h_s^\top \sigma d\widetilde{W}_s\right) = 0$ and hence

$$E[U(X_T^{(h)})] = \log x_0 + E\left(\int_0^T \left(h_s^\top \widehat{\mu}(Y_s) - \frac{1}{2} \|\sigma^\top h_s\|^2 \right) ds\right). \quad (4.1)$$

Since $\widehat{\mu}(Y_t) = M p_t$ is bounded, the strategy h^* given in the lemma is bounded; in particular $h^* \in \mathcal{H}$. Moreover, for all $s \in [0, T]$ the quantity h_s^* maximizes the integrand in (4.1), which implies that h^* is the maximizer of $E(\log(X_T^{(h)}))$. \square

Remark 4.1. If the factor process Y_t and hence the drift $\mu(Y_t)$ is observable then the optimal strategy is well-known - at time t one has to invest the fractions $(\sigma\sigma^\top)^{-1}\mu(Y_t)$ of wealth in the risky stocks. So for logarithmic utility the so-called *certainty equivalence principle* holds, i.e. the optimal strategy under partial information is obtained by replacing the unknown drift $\mu(Y_t)$ by the filter estimate $\widehat{\mu}(Y_t)$ in the formula for the optimal strategy under full information.

5. Dynamic programming equation for the case of power utility

A simplified optimization problem. In Section 3 we have shown that using filtering theory the original problem (2.5) can be transformed into a problem under complete information where the state variables are $X^{(h)}$ and p . In principle this problem could be attacked with dynamic programming methods. However, it turns out that the resulting HJB equation is fully nonlinear. Following Nagai and Runggaldier (2008) we simplify the problem by a change of measure. This will lead to a new problem where the set of state variables is reduced to p and where the HJB equation takes on a simpler quasi-linear form.

First we compute the utility of terminal wealth $U(X_T^{(h)}) = \frac{1}{\theta}(X_T^{(h)})^\theta$. From (3.6) it follows that

$$\frac{1}{\theta}(X_T^{(h)})^\theta = \frac{x_0^\theta}{\theta} \exp \left\{ \theta \int_0^T \left(h_s^\top \widehat{\mu}(Y_s) - \frac{1}{2} \|\sigma^\top h_s\|^2 \right) ds + \theta \int_0^T h_s^\top \sigma d\widetilde{W}_s \right\}. \quad (5.1)$$

Define now the random variable $Z_T^{(h)} = \exp \left\{ \int_0^T \theta h_s^\top \sigma d\widetilde{W}_s - \frac{1}{2} \int_0^T \|\theta \sigma^\top h_s\|^2 ds \right\}$ and the function

$$b(p, h) = b(p, h; \theta) = -\theta \left(h^\top M p - \frac{1-\theta}{2} \|\sigma^\top h\|^2 \right). \quad (5.2)$$

Recall that $\widehat{\mu}(Y_s) = \sum_{k=1}^d \mu_k p_s^k = M p_s$. Hence (5.1) can be written in the form

$$\frac{1}{\theta}(X_T^{(h)})^\theta = \frac{x_0^\theta}{\theta} Z_T^{(h)} \exp \left\{ \int_0^T -b(p_s, h_s; \theta) ds \right\}. \quad (5.3)$$

Since σ is deterministic, the Novikov condition together with (2.4) implies that $E(Z_T^{(h)}) = 1$. Hence we can define an equivalent measure $P^{(h)}$ on \mathcal{F}_T by $dP^{(h)}/dP = Z_T^{(h)}$, and Girsanov's theorem guarantees that $B_t := \widetilde{W}_t - \theta \int_0^t \sigma^\top h_s ds$ is a standard $(P^{(h)}, \mathbb{F})$ -Brownian motion. We set

$$a_k(p) = p^k \sigma^{-1} (\mu_k - \sum_{k=1}^d \mu_k p^k) = p^k \sigma^{-1} M (e_k - p) \in \mathbb{R}^n, \quad (5.4)$$

and let as before $\bar{f}(p, z) = \sum_{k=1}^d p^k f(e_k, z)$. Hence we have the following dynamics for the filter under $P^{(h)}$

$$dp_t^k = \left((Q^\top p_t)_k + \theta a_k^\top(p_t) \sigma^\top h_t \right) dt + a_k^\top(p_t) dB_t + p_t^k \int_{\mathcal{Z}} \left(\frac{f(x_k, z)}{f(p_{t-}, z)} - 1 \right) \gamma(dt \times dz). \quad (5.5)$$

In view of these transformations, for $0 < \theta < 1$ the optimization problem (2.5) is equivalent to the new optimization problem

$$\max \left\{ E \left(\exp \left\{ \int_0^T -b(p_s^{(h)}, h_s; \theta) ds \right\} \right) : h \in \mathcal{H} \right\} \quad (5.6)$$

where for $h \in \mathcal{H}$ the process $p^{(h)}$ has the dynamics (5.5) with initial condition $p_0^{(h)} = \pi$.

For $\theta < 0$ on the other hand (2.5) is equivalent to minimizing the expectation in (5.6). In the sequel we will concentrate on the case $0 < \theta < 1$; the necessary changes for $\theta < 0$ will be indicated where appropriate. Moreover, θ will be largely removed from the notation. The reward function for this control problem equals

$$v(t, p, h) = E_{t,p} \left(\exp \left\{ \int_t^T -b(p_s^{(h)}, h_s) ds \right\} \right) \quad \text{for } h \in \mathcal{H},$$

and the value function is given by $V(t, p) = \sup \{v(t, p, h) : h \in \mathcal{H}\}$. Note that $v(T, p, h) = V(T, p) = 1$.

The HJB equation. As a first step in the derivation of the HJB equation we compute the generator of the process $p_t^{(h)}$ with dynamics (5.5) for a constant strategy $h_t \equiv h \in \mathbb{R}^n$. In that case $p^{(h)}$ is obviously Markovian and a standard application of the Ito-formula shows that the generator \mathcal{L}^h operates on $g \in \mathcal{C}^2(\mathcal{S})$ as follows

$$\begin{aligned} \mathcal{L}^h g(p) &= \frac{1}{2} \sum_{i,j=1}^d a_i^\top(p) a_j(p) g_{p^i p^j} + \sum_{i=1}^d \left\{ \sum_{j=1}^d Q^{ji} p^j + \theta a_i^\top(p) \sigma^\top h \right\} g_{p^i} \\ &\quad + \lambda \int_{\mathcal{Z}} \{g(p + \Delta(p, z)) - g(p)\} \bar{f}(p, z) dz. \end{aligned} \quad (5.7)$$

Here $\mathcal{S} = \{p \in [0, 1]^d : \sum_{i=1}^d p^i = 1\}$ denotes the unit simplex in \mathbb{R}^d and $\Delta(p, z)$ is defined via

$$\Delta^k(p, z) = p^k \left(\frac{f(x_k, z)}{\bar{f}(p, z)} - 1 \right), \quad k = 1, \dots, d. \quad (5.8)$$

Next we turn to a heuristic derivation of the HJB equation for the optimization problem (5.6). Consider an arbitrary admissible strategy h and time points $t \leq u \leq T$. We obtain, by conditioning on \mathcal{F}_u

$$V(t, p) \geq v(t, p, h) = E_{t,p} \left(\exp \left\{ - \int_t^u b(p_s^{(h)}, h_s) ds \right\} v(u, p_u^{(h)}, h) \right). \quad (5.9)$$

Consider now a sequence of strategies h^n on the time interval $[u, T]$ such that $v(u, p_u^{(h^n)}, h^n)$ converges monotonically to $V(u, p_u^{(h)})$. Then we get by passing to the limit in (5.9) that

$$V(t, p) \geq E_{t,p} \left(\exp \left\{ - \int_t^u b(p_s^{(h)}, h_s) ds \right\} V(u, p_u^{(h)}) \right). \quad (5.10)$$

Define $\beta_t^h = \exp \left\{ - \int_0^t b(p_s^{(h)}, h_s) ds \right\}$. It is immediate from (5.10) that $\beta_t^h V(t, p_t^{(h)})$ is a supermartingale. Suppose now that h^* is an optimal strategy. In that case we have equality in (5.9) and $\beta_t^{h^*} V(t, p_t^{(h^*)})$ is a martingale. If we moreover assume that $V(t, \cdot) \in \text{dom } \mathcal{L}^h$ for all $t \in [0, T)$, we get from the Dynkin formula the following HJB equation

$$V_t(t, p) + \sup_{h \in \mathbb{R}^n} \left\{ \mathcal{L}^h V(t, p) - b(p, h; \theta) V(t, p) \right\} = 0, \quad (t, p) \in [0, T) \times \mathcal{S}, \quad (5.11)$$

with terminal condition $V(T, p) = 1$. In case that $\theta < 0$ the equation is similar, but the sup is replaced by an inf. Plugging in \mathcal{L}^h as given in (5.7) and $b(p, h)$ as given in (5.2) into (5.11) the HJB equation can be written more explicitly as

$$\left. \begin{aligned} 0 = & V_t(t, p) + \frac{1}{2} \sum_{k,l=1}^d a_k^\top(p_t) a_l(p_t) V_{p^k p^l}(t, p) + \sum_{k=1}^d \left\{ \sum_{l=1}^d Q^{lk} p^l \right\} V_{p^k}(t, p) \\ & + \sup_h \left\{ \sum_{k=1}^d a_k^\top(p_t) \sigma^\top \theta h V_{p^k}(t, p) + \theta V(t, p) \left(h^\top M p - \frac{1}{2} \|\sigma^\top h\|^2 (1 - \theta) \right) \right\} \\ & + \lambda \int_{\mathcal{Z}} \{ V(t, p + \Delta(p, z)) - V(t, p) \} \bar{f}(p, z) dz. \end{aligned} \right\} \quad (5.12)$$

Now the second line of (5.12) is quadratic in h so that the optimum is attained at the solution h^* of the following linear equation (the first-order condition)

$$\sigma \sum_{k=1}^d a_k(p) V_{p^k}(t, p) + V(t, p) \left(M p - \sigma \sigma^\top h (1 - \theta) \right) = 0.$$

Since σ is an invertible matrix, h^* is given by

$$h^* = h^*(t, p) = \frac{1}{(1 - \theta)} (\sigma \sigma^\top)^{-1} \left\{ M p + \frac{1}{V(t, p)} \sigma \sum_{k=1}^d a_k(p) V_{p^k}(t, p) \right\}. \quad (5.13)$$

If we plug this form of h^* back in the HJB equation (5.12) we obtain a *quasi-linear* integro-differential equation since the resulting equation is linear in the second derivatives $V_{p^k p^l}$ and in the integral part.

Remark 5.1 (On the formal justification of the Bellman equation). So far, our derivation of the HJB equation and of the candidate optimal strategy h^* in (5.13) is purely heuristic. A possible approach to give a precise mathematical meaning to the equation is to use *verification arguments*: if we can find a classical solution $V \in \mathcal{C}^{1,2}([0, T) \times \mathcal{S})$ of equation (5.12) with bounded derivatives, the function h^* introduced in (5.13) is bounded. Hence there exists a solution p^* of the SDE (5.5) with $h_t = h^*(t, p_t^*)$. Then a standard verification result such as Theorem 3.1 of Fleming and Soner (2006) immediately gives that V is the value function of the control problem (5.6) and that $h_t^* := h^*(t, p_t^*)$ is the optimal strategy.

However, the existence of a classical solution of equation (5.12) is an open issue. The main problem is the fact that one cannot guarantee that the equation is *uniformly elliptic*. To see this note that the coefficient matrix of the second derivatives in (5.12) is given by $C = A^\top A$ where the matrix A is given by

$$A = A(p) = (a_1(p), \dots, a_d(p)) \in \mathbb{R}^{n \times d}.$$

By definition equation (5.12) is uniformly elliptic if we can find some $c > 0$ such that for all $\xi \in \mathbb{R}^d$ we have $\xi^\top C \xi \geq c \|\xi\|^2$; in particular the matrix C needs to be strictly positive definite. This is possible only if there are no non-trivial solutions of the linear equation $Ax = 0$ so that we need to have the inequality $n \geq d$ (at least as many assets as states of the Markov chain Y). Such an assumption is hard to justify economically; imposing it nonetheless out of mathematical necessity would severely limit the applicability of our approach.

At present we therefore study an alternative route to justifying equations (5.12) and (5.13), see Frey et al. (2011). First, using results from Pham (1998), it is possible to show that V is a *viscosity solution* of (5.12). Moreover, we are currently working on a homogenization argument that will show that (5.13) can be used to compute an approximately optimal strategy. For this we add a term $\varepsilon d\tilde{B}_t$, with $\varepsilon > 0$ and \tilde{B} a d -dimensional Brownian motion independent of B , to the dynamics of the state equation (5.5). The HJB equation associated with these modified dynamics has an additional term $\varepsilon \sum_{k=1}^d V_{p^k p^k}$ and is therefore uniformly elliptic. Hence the results of Davis and Lleo (2010) apply directly to the modified equation, yielding the existence of a classical solution V^ε . Moreover, the optimal strategy $h^{\varepsilon,*}$ of the modified problem is given by (5.13) with V^ε instead of V . Clearly, one expects that for ε sufficiently small $h^{\varepsilon,*}$ is approximately optimal in the original problem. We are currently working on a formal proof of this statement.

Remark 5.2. Inspection of equation (5.13) shows that in the case of power utility the candidate optimal strategy consists of two parts. The first part

$$h_t^{(0)} = \frac{1}{1-\theta} (\sigma \sigma^\top)^{-1} M p_t \quad (5.14)$$

is the so-called *myopic strategy*; it is obtained by replacing the unknown drift $\mu(Y_t)$ with the filter estimate $M p_t$ in the classical formula for the optimal strategy under full information. Moreover, there is a correction term

$$\frac{1}{(1-\theta)V(t,p)} (\sigma^\top)^{-1} \sum_{k=1}^d a_k(p) V_{p^k}(t,p)$$

which is called *drift risk* in Rieder and Bäuerle (2005). In particular, in the case of power utility the certainty equivalence principle does not hold.

6. Policy Improvement

Solving the Bellman equation (5.12) numerically with finite difference methods is feasible only if d (the number of states of the hidden Markov chain Y) is small. In

12 *R. Frey, A. Gabih and R. Wunderlich*

this section we therefore propose a policy improvement procedure that permits to find an approximation of the optimal strategy which is computable using Monte-Carlo methods. The starting point is the myopic strategy $h_t^{(0)}$ from (5.14) with corresponding reward function

$$v^{(0)}(t, p) := v(t, p, h^{(0)}) = E_{t,p} \left(\exp \left\{ - \int_t^T b(p_s^{(h^{(0)})}, h_s^{(0)}) ds \right\} \right). \quad (6.1)$$

Assume that $v^{(0)}$ is a $\mathcal{C}^{1,2}$ function. Motivated by the derivation of the HJB equation in the previous section we compute a new strategy $h^{(1)}$ by maximizing (for $0 < \theta < 1$) resp. minimizing (for $\theta < 0$) the drift of $\beta_t^h v^{(0)}(t, p_t^{(h)})$. This leads to the following optimization problem

$$\max \left\{ \text{sign}(\theta) \left(\mathcal{L}^h v^{(0)}(t, p) - b(p, h) v^{(0)}(t, p) \right) : h \in \mathbb{R}^n \right\}. \quad (6.2)$$

By analogous arguments as in the derivation of the candidate optimal strategy h^* in the previous section we obtain

$$h^{(1)}(t, p) = h^{(0)}(t, p) + \frac{1}{(1 - \theta) v^{(0)}(t, p)} (\sigma^\top)^{-1} \sum_{k=1}^d a_k(p) v_{p^k}^{(0)}(t, p). \quad (6.3)$$

Note that $h^{(1)}$ has a similar structural form as the optimal strategy h^* , but with $v^{(0)}$ instead of the value function V in the correction term for the drift risk. The reward function corresponding to the strategy $h^{(1)}$ is given by $v^{(1)}(t, p) := v(t, p, h^{(1)})$. The next lemma shows that $h^{(1)}$ is in fact an improvement of $h^{(0)}$.

Lemma 6.1. *Suppose that the reward functions $v^{(0)}$ and $v^{(1)}$ are $\mathcal{C}^{1,2}$ functions. Then it holds that $v^{(1)}(t, p) \geq v^{(0)}(t, p)$, $(t, p) \in [0, T] \times \mathcal{S}$.*

Proof. We concentrate on the case $\theta > 0$. Here we have the following inequalities

$$0 = v_t^{(0)} + \left(\mathcal{L}^{h^{(0)}} v^{(0)} - b(p, h^{(0)}) v^{(0)} \right) \leq v_t^{(0)} + \left(\mathcal{L}^{h^{(1)}} v^{(0)} - b(p, h^{(1)}) v^{(0)} \right)$$

where the first equality follows from the fact that $\beta_t^{h^{(0)}} v^{(0)}(t, p_t^{(h^{(0)})})$ is a martingale while the inequality follows from the fact, that $h^{(1)}$ is the maximum of (6.2). This implies that the process $\beta_t^{h^{(1)}} v^{(0)}(t, p_t^{(h^{(1)})})$ has a non-negative drift, consequently it is a submartingale. Since $v^{(0)}(T, \cdot) = v^{(1)}(T, \cdot) \equiv 1$ we obtain

$$\beta_t^{h^{(1)}} v^{(0)}(t, p) \leq E_{t,p} \left(\beta_T^{h^{(1)}} v^{(0)}(T, p_T^{h^{(1)}}) \right) = E_{t,p} \left(\beta_T^{h^{(1)}} v^{(1)}(T, p_T^{h^{(1)}}) \right) = \beta_t^{h^{(1)}} v^{(1)}(t, p),$$

where the last equality follows since $\beta_t^{h^{(1)}} v^{(1)}(t, p_t^{h^{(1)}})$ is a martingale. Division by $\beta_t^{h^{(1)}}$ yields the assertion. \square

In order to compute $h^{(1)}$ one needs to compute the reward function $v^{(0)}$ and its partial derivatives. Since $h^{(0)}$ is known this can be done with standard Monte-Carlo methods.

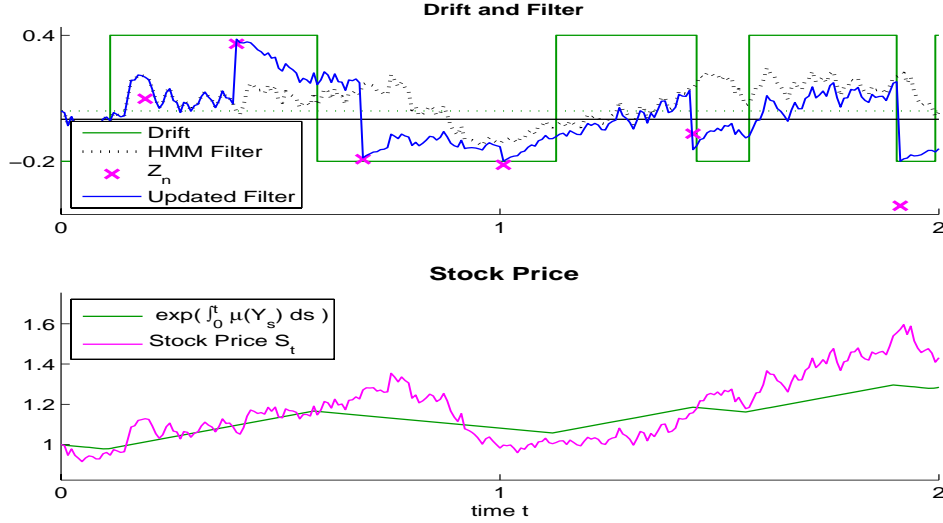


Fig. 1. Top: drift $\mu(Y_t)$, classical HMM filter, predictions Z_n and updated filter $\widehat{\mu}(Y_t)$
 Bottom: stock price S_t and drift component $\exp(\int_0^t \mu(Y_s) ds)$

Remark 6.1. The assumption that $v^{(0)}$ and $v^{(1)}$ are $\mathcal{C}^{1,2}$ functions can be verified for the modified version of the state equation introduced in Remark 5.1, since the Kolmogorov backward equation associated with the modified version of the state process is strictly elliptical. The argument outlined at the end of Remark 5.1 shows that for ε sufficiently small $h^{(\varepsilon,1)}$ (the policy improvement in the modified model) is also a “good” strategy for the original problem.

7. Numerical example

In this section we illustrate the findings of the previous sections. We consider a market with $n = 1$ stock with volatility $\sigma = 0.2$. The drift process $\mu(Y_t)$ is modelled by a Markov chain with $d = 2$ states, $\mu_1 = 0.4$ and $\mu_2 = -0.2$. The transition intensities of Y are set as $\lambda_{12} = 1.5$ and $\lambda_{21} = 1$, yielding the stationary distribution $(0.4, 0.6)$ and the ergodic mean $\lim_{t \rightarrow \infty} E[\mu(Y_t)] = 0.06$. The trading horizon is $T = 2$ years, time discretization works with $M = 250$ time steps.

We consider absolute views for the drift as described in Example 2.1. At the jump times T_n of a Poisson process with intensity λ the views Z_n are generated according to $Z_n = \mu(Y_{T_n}) + \sigma_\varepsilon \varepsilon_n$ where the constant $\sigma_\varepsilon > 0$ describes the confidence of the prediction.

We have simulated a path of the drift process $\mu(Y_t)$, predictions Z_n with confidence parameter $\sigma_\varepsilon = 0.1$ arriving with intensity $\lambda = 2$ and stock returns using $\Delta S_t/S_t = \mu(Y_t)\Delta t + \sigma\Delta W_t$. The upper panel of Figure 1 shows the non-observable




Fig. 2. Initial capital of the non-informed investor which is required to obtain the same maximized expected log-utility at time T as the informed investor who started at time 0 with unit wealth.

drift path $\mu(Y_t)$. The quantities which are observable to the investor are divided into two parts. First, the path of the stock price S_t shown in the lower panel together with the (non-observable) path of $\exp(\int_0^t \mu(Y_s) ds)$ representing the drift component. Second, the predictions Z_n constituting the extra information, which are represented by crosses in the upper panel.

Filter. Given the observations of the stock prices respective returns and the marked point process (T_n, Z_n) the investor can compute the filter p_t by integrating the filter equation (3.4). As initial distribution π we use the stationary distribution of the Markov chain. The resulting filter for the drift $\widehat{\mu}(Y_t) = \sum_{k=1}^d \mu_k p_t^k$ is drawn in the upper panel of Figure 1. For the sake of comparison we also show the classical HMM filter which is based on the observed stock returns only, but not on the extra information.

It can be seen that for the chosen parameters the observed drift predictions Z_n arriving at the information dates T_n are very informative for the investor since for $t = T_n$ the filter is quite close to the actual value of the drift. That means, for $t = T_n$ and for the chosen distribution of Z_n , the investor has nearly full information on the drift. Between the information dates the filter is pushed back towards the ergodic mean.

For the computation of the filter p_t we apply the procedure described in Section 3. We integrate filter equation (3.4) between two jumps, i.e. in the interval (T_{n-1}, T_n) , using a Euler scheme with time step size $\Delta t = T/M$ and add at T_n the correction term resulting from Bayesian updating. In order to reduce time-discretization errors in the integration of the filter equation between the jumps we work instead of the nonlinear Wonham filter equation (3.1) for the normalized filter p_t with a linear filter equation for the unnormalized filter and apply robust filter techniques, see Sass and Wunderlich (2010).

Information gain for log-utility. In order to quantify the value of the additional information from the expert opinions we compare two utility maximizing investors.

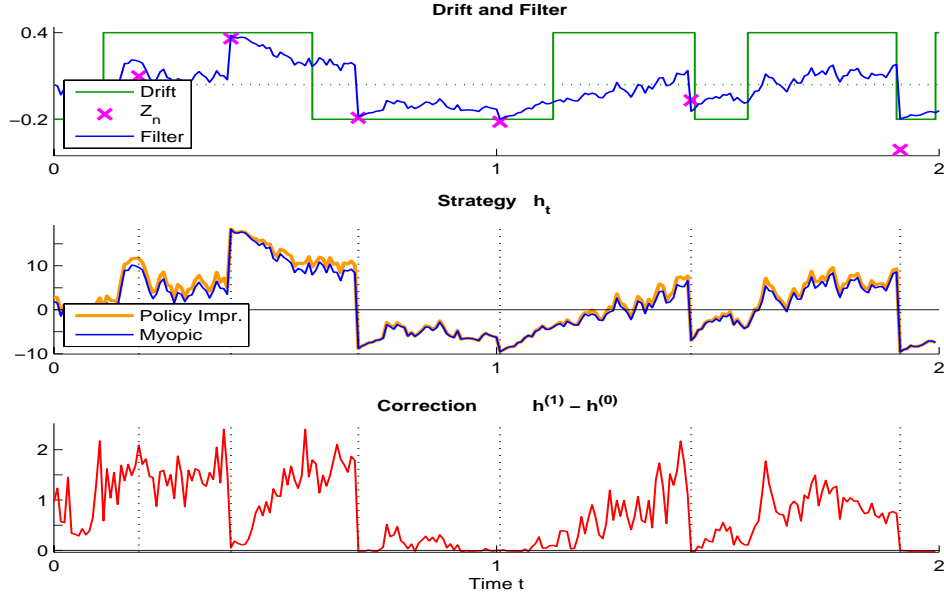


Fig. 3. Top: drift $\mu(Y_t)$, filter $\widehat{\mu(Y_t)}$ and predictions Z_n , Center: myopic strategy $h^{(0)}$ and first approximation $h^{(1)}$, Bottom: correction term $h^{(1)} - h^{(0)}$

First, the “non-informed” investor can only observe stock returns. Second, the “informed” investor additionally has access to expert opinions. Now we consider the initial capital which the non-informed investor needs to obtain the same maximized expected utility at time T as the informed investor who started at time 0 with unit wealth. The difference between this capital and one can be interpreted as information gain for the informed investor. This comparison is restricted to logarithmic utility. Here the optimal strategy h^* given in Lemma 4.1 coincides with the strategy $h^{(0)}$ and expected utility $E(U(X_T^{(h^*)}))$ can be computed easily via Monte-Carlo simulation using representation (4.1).

Denote by X_T^I and X_T^N the optimal terminal wealth of the informed and of the non-informed investor starting both with an initial wealth of one. Then the additional initial capital required by the non-informed investor is given by $x_0^N = \exp\{E(\log(X_T^I)) - E(\log(X_T^N))\}$. Figure 2 shows this initial capital as a function of the intensity λ for different values of the confidence parameter $\sigma_\varepsilon = 0.1, 0.3, 0.5$ representing high, moderate and low confidence. As expected the information gain increases with λ and decreases with σ_ε . The Figure also shows the result for the fully informed investor which is obtained for $\lambda \rightarrow \infty$.

Optimal strategy for power utility: Next, we consider an investor who maximizes expected power utility $U(x) = x^\theta/\theta$ of terminal wealth with $\theta = 2/3$. For computing an approximation of the optimal strategy h^* we apply the policy im-

provement described in the preceding section. We start with the myopic strategy $h^{(0)}$ given in (5.14) which can be computed directly from the filter for the drift $\widehat{\mu}(Y_t) = \bar{\mu}(p_t)$. For the computation of the first improvement $h^{(1)}$ given in (6.3) we need the reward function $v^{(0)}$ and its partial derivatives. Here we use representation (6.1) of $v^{(0)}$ as conditional expectation. Assume that at time t^* based on observations of the stock prices S_t and the marked point process I in $[0, t^*]$ we have obtained the filter $p_{t^*} = p^*$. Then a Monte-Carlo approximation of $v^{(0)}(t^*, p^*)$ is computed from $N = 1000$ paths of $p_s^{h^{(0)}}$ on $[t^*, T]$ which are solutions of the filter equation (5.5) starting at time t^* with p^* . The partial derivatives of $v^{(0)}(t, p)$ with respect to p^k , $k = 1, \dots, d$ are approximated using central differences.

Figure 3 shows in the second panel the myopic strategy $h^{(0)}$ and the first approximation $h^{(1)}$ of the policy improvement. For the chosen parameters the myopic strategy $h^{(0)}$ is very close to the first approximation $h^{(1)}$ of the policy improvement. The lower panel shows the correction term $h^{(1)} - h^{(0)}$ of the policy improvement. It turns out, that this difference vanishes for $t = T_n$ (marked by dotted vertical lines) where additional information allows for quite accurate estimates for the drift which are close to the actual values. So the investor has nearly full information on the drift, and both, the optimal strategy under incomplete information and the myopic strategy are close to the optimal strategy under full information. Moreover, the correction term tends to zero for $t \rightarrow T$ as one would expect.

References

- Björk, T., Davis, M.H.A. and Landén, C. (2010): Optimal investment with partial information. *Mathematical Methods of Operations Research* 71, 371-399.
- Black, F., Litterman, R. (1992): Global portfolio optimization. *Financial Analysts Journal* 48(5), 28-43.
- Bremaud, P. (1981) : *Point Processes and Queues: Martingale Dynamics*, Springer, New York.
- Brendle, S. (2006): Portfolio selection under incomplete information. *Stochastic Processes and Their Applications* 116, 701-723.
- Danilova, A., Monoyios, M. and Ng, A. (2010): Optimal investment with inside information and parameter uncertainty. *Mathematics and Financial Economics*, 3, 13-38
- Davis, M. and Lleo, S. (2010). Risk-sensitive asset management in a jump-diffusion factor model. Working Paper, arXiv:1001.1379v2.
- Elliott, R.J., Aggoun, L. and Moore, J.B. (1994): *Hidden Markov Models*, Springer, New York.
- Fleming, W. H. and Soner, H. M. (2006): *Controlled Markov Processes and Viscosity Solutions*, 2nd edn, Springer, New York.
- Frey, R., Gabih, A. and Wunderlich, R. (2011): Portfolio Optimization under Partial Information with Expert Opinions: Mathematical Aspects. In preparation.
- Gabih, A., Sass, J. and Wunderlich, R. (2009): Utility maximization under bounded expected loss. *Stochastic Models*, 3(25), 2009, 375-407.
- Lakner, P. (1998): Optimal trading strategy for an investor: the case of partial information. *Stochastic Processes and their Applications* 76, 77-97.
- Liptser, R.S. and Shiryaev A.N. (2001): *Statistics of Random Processes: General theory*, 2nd edn, Springer, New York.

- Nagai, H. and Peng, S. (2002): Risk-sensitive dynamic portfolio optimization with partial information on infinite time horizon. *The Annals of Applied Probability* 12, 173-195.
- Nagai, H. and Runggaldier, W.J. (2008): PDE approach to utility maximization for market models with hidden Markov factors. In: *Seminar on Stochastic Analysis, Random Fields and Applications V* (R.C.Dalang, M.Dozzi, F.Russo, eds.). Progress in Probability, Vol.59, Birkhuser Verlag, 493-506.
- Pham, H. (1998): Optimal stopping of controlled jump diffusion processes: a viscosity solution approach. *Journal of Mathematical Systems, Estimation and Control* 8, 1-27.
- Rieder, U. and Bäuerle, N. (2005): Portfolio optimization with unobservable Markov-modulated drift process. *Journal of Applied Probability* 43, 362-378.
- Sass, J. and 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.
- Sass, J. and Wunderlich, R. (2010): Optimal portfolio policies under bounded expected loss and partial information. *Mathematical Methods of Operations Research* 72, 25-61.
- Schöttle, K., Werner, R. and Zagst, R. (2010): Comparison and robustification of Bayes and Black-Litterman models. *Mathematical Methods of Operations Research*, 71, 453-475.
- Wonham, W.M. (1965): Some applications of stochastic differential equations to optimal non-linear filtering, *SIAM J.Control Opt.*, 347-369.