European Journal of Operational Research 234 (2014) 561–570

Contents lists available at SciVerse ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Investment strategies and compensation of a mean–variance optimizing fund manager

Georgios Aivaliotis, Jan Palczewski *

School of Mathematics, University of Leeds, Leeds LS2 9JT, UK

ARTICLE INFO

Article history:
Available online 3 May 2013

Keywords:
Mean–variance
Continuous-time stochastic control
Viscosity solutions
Investment strategy
Managerial compensation

ABSTRACT

This paper introduces a general continuous-time mathematical framework for solution of dynamic mean–variance control problems. We obtain theoretical results for two classes of functionals: the first one depends on the whole trajectory of the controlled process and the second one is based on its terminal-time value. These results enable the development of numerical methods for mean–variance problems for a pre-determined risk-aversion coefficient. We apply them to study optimal trading strategies pursued by fund managers in response to various types of compensation schemes. In particular, we examine the effects of continuous monitoring and scheme's symmetry on trading behavior and fund performance.

1. Introduction

Markowitz's seminal paper [29] introduced the mean–variance criterion into portfolio optimization. Single-period problems, which are mathematically tractable, have enjoyed popularity both in the academia to model investor preferences and behavior (see, e.g., Epstein [16], Ormiston and Schlee [30], Tobin [35]) and among practitioners (see, e.g., Bodie et al. [12], Litterman [28]). An extension of this theory to continuous-time models proved to be difficult due to fundamental problems introduced by the variance term. A natural approach to continuous-time optimization is to use dynamic programming, which relies on markovianity of functionals. The variance is, however, not markovian. There are three main alternatives. The first involves the study of risk-sensitive functionals (see, e.g., Bielecki et al. [10]), whose second order Taylor expansion has the form of a mean–variance functional

$$E(H) - \frac{\gamma}{2} Var(H), \tag{1}$$

where H is a random outcome of the investment and γ is the risk-aversion coefficient. The second alternative to dynamic programming hinges on the use of martingale methods (see, e.g., Bielecki et al. [9]). Although these methods can be used to obtain closed-form solutions for a class of mean–variance problems, they turned out to be unsuitable as a basis for efficient numerical algorithms for general mean–variance problems.

A substantial progress in the theory for mean–variance functionals was due to a third approach, closely related to the one we employ in this paper. This approach, introduced by Li and Ng [26] in a discrete-time setting, embeds the mean–variance problem into a class of auxiliary stochastic control problems that can be solved by dynamic programming methods (see also Leippold et al. [25]). An extension of this method to a continuous-time framework is presented in Zhou and Li [38], and further employed by Fu et al. [18] and Lim [27]. These papers put several constraints on the optimization problem in order to obtain auxiliary control problems in a linear-quadratic form. In particular, the random variable H in the mean–variance functional (1) is assumed to be a linear function of the portfolio wealth process. Wang and Forsyth [36] design numerical schemes for auxiliary linear-quadratic problems formulated in [38] and construct an efficient frontier.

In this paper we present a mathematical framework for the solution of general mean–variance stochastic control problems in continuous time. This framework extends the continuous-time theory of Zhou and Li [38] in two aspects. First, we allow the random variable H in the mean–variance functional (1) to be specified either as a continuous function of the portfolio wealth at a terminal time (in general, the value of the controlled process) or as an integral of a continuous function of the portfolio wealth (in general, the value of the controlled process) over time. A particular case when H depends linearly on the portfolio wealth at terminal time is covered theoretically in Zhou and Li [38] and numerically in Wang and Forsyth [36]. Second, we relax assumptions on the dynamics of the controlled process to cover non-homogeneous degenerate diffusions with Lipschitz coefficients of linear growth.

To the best of our knowledge, mean–variance optimization problems based on the integral of a function of the value of the controlled process over time have not been widely studied. A closely related paper by Aivaliotis and Veretennikov [4] provides theoretical approximation results via regularization; their solution

* Corresponding author. Tel.: +44 1133435180.
E-mail addresses: G.Aivaliotis@leeds.ac.uk (G. Aivaliotis), J.Palczewski@leeds.ac.uk (J. Palczewski).

0377-2217/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ejor.2013.04.038

leads to randomized strategies. In our paper, the optimization problem is solved directly using the theory of viscosity solutions to Hamilton–Jacobi–Bellman equations (see Fleming and Soner [17], Pham [33]). In particular, our results enable computation of (non-randomized) optimal strategies in a feedback form. The justification of their optimality – the verification theorem – requires very restrictive assumptions (for the latest results see Gozzi et al. [21]) that our control problem does not satisfy. We, therefore, resort to numerically testing the optimality of strategies extracted from numerical solutions of the HJB equations.

Our theoretical results are used to develop numerical algorithms to maximize functional (1) for a given (pre-determined) risk-aversion coefficient γ. Wang and Forsyth [36] solve auxiliary markovian optimization problems which are parametrised neither by risk-aversion nor by the expectation of terminal value. Once the optimal strategy is known, they can compute the risk-aversion, the expectation and the variance. This proves to be sufficient if one aims at graphing an efficient frontier. Our approach is different as we endeavor to find an optimal strategy for a pre-determined risk-aversion coefficient. We reformulate the mean–variance problem as a superposition of a static and a dynamic optimization problem, which is equivalent to solving a set of parametrized HJB equations and maximizing the resulting value functions over a compact interval valued parameter. We demonstrate that, for practical applications, our approach leads to an efficient numerical algorithm.

Recently, Basak and Chabakauri [7,8] proposed another view on mean–variance optimization. They introduced a notion of optimality in an intra-personal game theoretic sense. This has the advantage of turning the optimization problem markovian. It should, however, be noticed that strategies optimal in a game-theoretic sense might not be optimal in a classical sense; and vice versa.

Theoretical results of this paper are applied to a study of a delegated portfolio management problem. It is a common practice in the asset management industry to use mean–variance preferences for choosing portfolios (see Bodie et al. [12] and Littermann [28]). We assume that fund managers apply the same type of preferences to their compensation and tend to follow trading strategies that maximize their satisfaction from compensation. The mathematical framework introduced in this paper allows us to study trading strategies pursued by fund managers in response to various types of compensation (incentive) schemes. We also analyze implications of complex schemes on distributional properties of the fund's wealth process. We consider symmetric (e.g., co-ownership) and asymmetric (with a hurdle rate provision) schemes based on the terminal wealth and on the continuously monitored wealth.

Incentives have been proven to be a significant factor influencing the behavior and performance of fund managers. Agarwal et al. [2] examine, in an empirical study, the influence of incentives and managerial discretion on the performance of hedge funds. They find that managers with performance-related incentives – the inclusion of hurdle rate provisions, or co-ownership – are associated with a better performance. We study numerically the implications of the above incentives on trading decisions of fund managers. Managers with symmetric (co-ownership) compensation schemes show a superior performance over those remunerated by schemes with hurdle-rate provisions: the resulting Sharpe ratio of the terminal wealth is higher.

Incentives also influence the riskiness of trading strategies pursued by fund managers. Elton et al. [15] find that managers with asymmetric incentive contracts tend to follow riskier strategies than those with symmetric compensation schemes.[1] In particular, they observe that asymmetric schemes encourage large variations in the riskiness of portfolios over time: a poor performance at any time triggers a sharp increase in the risk taking. Our numerical results show that such behavior is optimal for a fund manager with mean–variance preferences.

Our numerical study contributes also to the discussion about the frequency of portfolio monitoring (see, e.g., Agarwal et al. [2] and Goetzmann et al. [20]). We analyze trading strategies and portfolio performance when the manager's compensation is based on her performance sampled continuously over the whole investment period. We observe a fall in Sharpe ratios for symmetric and asymmetric schemes. This agrees with the empirical findings of Agarwal et al. [2]. A continuous examination of the fund's wealth diminishes managerial discretion, which, according to [2], impacts on the fund performance. One would, however, expect that the closer scrutiny offered by such compensation schemes lowers the riskiness of investment decisions. We demonstrate that the opposite is true: the variance of excess returns increases.

A classical but more challenging problem is the design of compensation schemes that align preferences of an investor and a fund manager. Existing literature offers results in the case of preferences represented by utility functions (see, e.g., Carpenter [13] and Ou-Yang [31]). Mean–variance optimality criterion has only been used in a static (single-period) framework [6,14].

The outline of the remaining part of the paper is as follows. Section 2 introduces a general mathematical framework for the solution of mean–variance stochastic control problems and prepares the ground for design of efficient numerical schemes. The problem of managerial compensation in a continuous-time market model alongside with various types of compensation schemes and discussion of numerical methods used for computation of optimal investment strategies is presented in Section 3. Analysis of the trading strategies is performed in Section 4. Section 5 concludes. In the Electronic Supplement, Section A introduces numerical schemes and verifies their convergence, Section B collects proofs.

2. Theoretical framework and results

In this section we present a general framework for the solution of mean–variance dynamic optimization problems. Section 2.1 studies functionals depending on the value of controlled process at the terminal time. In Section 2.2, these ideas are extended to functionals based on the whole trajectory of the controlled process. Our exposition is geared towards numerical computations, necessary for practical applications.

The state is described by a d-dimensional non-homogeneous stochastic differential equation (SDE) driven by a d_1-dimensional Wiener process (W_t)

$$dX_t = b(\alpha_t, t, X_t)dt + \sigma(\alpha_t, t, X_t)dW_t, \quad X_{t_0} = x, \tag{2}$$

where $b : A \times [0,\infty) \times \mathbb{R}^d \to \mathbb{R}^d$ and $\sigma : A \times [0,\infty) \times \mathbb{R}^d \to \mathbb{R}^{d\times d_1}$. The process $(\alpha_t, t_0 \leqslant t \leqslant T)$ is from the class $\mathcal{A}$ of all progressively measurable processes (with respect to the filtration generated by the Wiener process (W_t)) with values in a compact set $A \subset \mathbb{R}^\ell$. Its role is to control the dynamics of the diffusion (X_t). Eq. (2) has a pathwise unique (weak) solution if the following conditions are satisfied (see, e.g., Gikhman and Skorokhod [19, Section Remark V.8.1] or Fleming and Soner [17, Appendix D]):[2]

[1] The fixed fee in the paper by Elton et al. [15] can be represented in our framework as a symmetric compensation contract.

[2] These conditions are superficial for the uniqueness of solutions. We will, however, need them later in the study of the value function of the mean–variance optimization problem.