Journal of Economic Dynamics & Control 34 (2010) 207–230

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

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

Numerical solution of the Hamilton–Jacobi–Bellman formulation for continuous time mean variance asset allocation

J. Wang, P.A. Forsyth *

David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, ON, Canada N2L 3G1

ARTICLE INFO

Article history:
Received 20 December 2008
Accepted 31 August 2009
Available online 15 September 2009

JEL classification:
C63
G11

Keywords:
Optimal control
Mean variance tradeoff
HJB equation
Viscosity solution

ABSTRACT

We solve the optimal asset allocation problem using a mean variance approach. The original mean variance optimization problem can be embedded into a class of auxiliary stochastic linear-quadratic (LQ) problems using the method in Zhou and Li (2000) and Li and Ng (2000). We use a finite difference method with fully implicit timestepping to solve the resulting nonlinear Hamilton–Jacobi–Bellman (HJB) PDE, and present the solutions in terms of an efficient frontier and an optimal asset allocation strategy. The numerical scheme satisfies sufficient conditions to ensure convergence to the viscosity solution of the HJB PDE. We handle various constraints on the optimal policy. Numerical tests indicate that realistic constraints can have a dramatic effect on the optimal policy compared to the unconstrained solution.

1. Introduction

Continuous time mean variance asset allocation has received considerable attention over the years (Zhou and Li, 2000; Li and Ng, 2000; Nguyen and Portrai, 2002; Leippold et al., 2004; Bielecki et al., 2005). Financial applications include hedging futures (Duffie and Richardson, 1991), insurance (Chiu and Li, 2006; Wang et al., 2007), pension asset allocation (Gerrard et al., 2004; Hojgaard and Vigna, 2007) and optimal execution of trades (Lorenz and Almgren, 2007). In its simplest formulation, an investor can choose to invest in a risk-free bond or a risky asset, and can dynamically alter the proportion of wealth invested in each asset, in order to achieve a mean variance efficient result.

The continuous time mean variance problem does not lend itself easily to a dynamic programming formulation. There have been two main approaches to this problem. The original mean variance optimal control problem can be embedded into a class of auxiliary stochastic linear-quadratic (LQ) problems, which can then be solved in terms of dynamic programming (Zhou and Li, 2000; Li and Ng, 2000). Alternatively, Martingale techniques can be used (Bielecki et al., 2005). In the case of the LQ method, previous papers use analytic techniques to solve the nonlinear Hamilton–Jacobi–Bellman (HJB) PDE for special cases. In order to obtain analytic solutions, the authors typically make assumptions which allow for the possibility of unbounded borrowing and infinite negative wealth (bankruptcy). However, some analytic solutions have been developed for handling specific constraints: no stock shorting (Li et al., 2002) (but shorting the bond is still allowed) and the no bankruptcy case (Bielecki et al., 2005) (but again allowing for shorting the bond).

A popular approach for optimal stochastic control problems in finance is to use utility functions (Cairns et al., 2006; Damgaard, 2006; Muthuraman, 2007; Chellathurai and Draviam, 2007; Munk, 2000). This approach usually results in

* Corresponding author.
E-mail addresses: j27wang@uwaterloo.ca (J. Wang), paforsyt@uwaterloo.ca, paforsyt@cs.uwaterloo.ca (P.A. Forsyth).

0165-1889/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.jedc.2009.09.002

financial models which also generate HJB PDEs. Although investment policy based on mean variance optimization has its critics, an advantage of this approach compared to power-law or exponential utility maximization is that the results can be easily interpreted in terms of an efficient frontier.

The objective of this paper is to develop a numerical method for solving the continuous time mean variance optimal asset allocation problem. We will use a fully numerical scheme based on solving the HJB equation resulting from the LQ formulation. Our scheme can easily handle any type of constraint (e.g. nonnegative wealth, no shorting of stocks, margin requirements).

Although the methods developed in this paper can be applied to any asset allocation problem, such as those discussed in Nguyen and Portrai (2002), Bielecki et al. (2005), Chiu and Li (2006), and Wang et al. (2007), we will focus on asset allocation problems which are relevant to defined contribution pension plans as discussed in Hojgaard and Vigna (2007) and Cairns et al. (2006). In Hojgaard and Vigna (2007), the objective is to determine the mean variance efficient strategy in terms of final wealth. In Cairns et al. (2006), the pension plan model includes a stochastic salary component. In Cairns et al. (2006), the problem was formulated in terms of maximizing the utility of the wealth-to-income ratio. Here, we consider the same model, but solve for the optimal continuous time mean variance efficient frontier. Note that by setting the contribution rate to zero, the pension plan problem reduces to the classical continuous time (multi-period) portfolio selection problem (Zhou and Li, 2000; Li and Ng, 2000; Li et al., 2002; Bielecki et al., 2005; Li and Zhou, 2006).

The main results of this paper are:

- Based on the methods in Forsyth and Labahn (2008) and Wang and Forsyth (2008), we develop a fully implicit method for solving the nonlinear HJB PDE, which arises in the LQ formulation of the mean variance problem. Under the assumption that the HJB equation satisfies a strong comparison property, our methods are guaranteed to converge to the viscosity solution of the HJB equation. In addition, the policy iteration scheme used to solve the nonlinear algebraic equations at each timestep is globally convergent. Note that an explicit method would have timestep restrictions due to stability considerations. In the case of an unbounded control, the maximum stable timestep would be difficult to estimate. This problem does not arise if an implicit method is used.
- By solving the HJB PDE and a related linear PDE, we develop a numerical method for constructing the mean variance efficient frontier (in continuous time). Any type of constraint can be applied to the investment policy.
- We pay particular attention to handling various constraints on the optimal policy. In particular, in order to compare the numerical solution with the known analytic solution in special cases, it is necessary to allow for negative wealth and unbounded controls. This requires careful attention to the grid construction and form of the control as the mesh and timesteps shrink to zero.
- From a practical point of view, we observe that the addition of realistic constraints can completely alter some of the properties of the mean variance solution compared to the unconstrained control case (Li and Zhou, 2006).

We should point out here that the optimal mean variance strategy in Zhou and Li (2000), Li and Ng (2000), Nguyen and Portrai (2002), Leippold et al. (2004), Bielecki et al. (2005), Chiu and Li (2006), Wang et al. (2007), and Hojgaard and Vigna (2007) is the *pre-commitment* strategy, i.e. once the initial strategy has been determined (as a function of the state variables) at the initial time, the investor commits to this strategy, even if the mean variance policy, computed at a later time would differ from the pre-commitment strategy. This contrasts with the *time-consistent* policy whereby the investor optimizes the mean variance tradeoff at each instant in time, assuming optimal mean variance strategies at each later instant. The subtle distinction between these two approaches is discussed in Basak and Chabakauri (2007). We note that the efficient frontier for the pre-commitment strategy must always lie above the efficient frontier for the time-consistent strategy. However, there is some economic controversy about the meaning of these two approaches. We will focus on the pre-commitment policy in this paper, leaving the time-consistent problem for future work.

2. Mean variance efficient wealth case

We first consider the problem of determining the mean variance efficient strategy in terms of the investor's final wealth. We will refer to this problem in the following as the *wealth* case. This will allow an explanation of the basic approach for construction of the efficient frontier, without undue algebraic complication. For certain special cases, there are also some analytic solutions available (Hojgaard and Vigna, 2007) for this problem. This will enable us to compare with the numerical solution.

Suppose there are two assets in the market: one is risk free (e.g. a government bond) and the other is risky (e.g. a stock index). The risky asset S follows the stochastic process,

$$dS = (r + \xi_1\sigma_1)S\,dt + \sigma_1 S\,dZ_1, \tag{2.1}$$

where dZ_1 is the increment of a Wiener process, σ_1 is volatility, r is the interest rate, ξ_1 is the market price of risk (or Sharpe ratio) and the stock drift rate can then be defined as $\mu_S = r + \xi_1\sigma_1$. Suppose that the plan member continuously pays into the pension plan at a constant contribution rate π in the unit time. Let $W(t)$ denote the wealth accumulated in the pension plan at time t, let p denote the proportion of this wealth invested in the risky asset S, and let $(1-p)$ denote the fraction of