European Journal of Operational Research 243 (2015) 921–931

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Decision Support

Dynamic portfolio optimization with transaction costs and state-dependent drift ☆, ☆☆

Jan Palczewski [a], Rolf Poulsen [b], Klaus Reiner Schenk-Hoppé [c,d], Huamao Wang [e,*]

[a] School of Mathematics, University of Leeds, UK
[b] Department of Mathematical Sciences, University of Copenhagen, Denmark
[c] Economics, School of Social Sciences, University of Manchester, UK
[d] Department of Finance, NHH–Norwegian School of Economics, Norway
[e] School of Mathematics, Statistics and Actuarial Science, University of Kent, UK

ARTICLE INFO

Article history:
Received 25 April 2014
Accepted 21 December 2014
Available online 31 December 2014

Keywords:
Dynamic programming
Numerical methods
State-dependent drift
Transaction costs
Markov Chain approximation

ABSTRACT

The problem of dynamic portfolio choice with transaction costs is often addressed by constructing a Markov Chain approximation of the continuous time price processes. Using this approximation, we present an efficient numerical method to determine optimal portfolio strategies under time- and state-dependent drift and proportional transaction costs. This scenario arises when investors have behavioral biases or the actual drift is unknown and needs to be estimated. Our numerical method solves dynamic optimal portfolio problems with an exponential utility function for time-horizons of up to 40 years. It is applied to measure the value of information and the loss from transaction costs using the indifference principle.

1. Introduction

Numerical methods for dynamic portfolio optimization under proportional transaction costs typically assume that the drift of the risky asset is constant. However, a state-dependent drift enters the optimization problem in many scenarios. For instance, if the drift is unobservable, it can be estimated with the Kalman–Bucy filter. This leads to an optimization problem where the drift depends on the currently observed stock price (e.g. Björk, Davis, & Landén, 2010). The drift is also state-dependent when contrarian investors optimize portfolios under the assumption that prices are mean-reverting; for instance when an investor is a victim of the Gambler's fallacy, see, e.g., Shefrin (2008). Similarly, investors who aim at following market trends will include a state-dependent drift in optimization.

In these cases an investor's optimal trading strategy strongly depends on the forecasting function used to predict asset prices. This poses a numerically demanding problem. Our paper proposes an efficient numerical method to solve finite-horizon portfolio optimization problems with transaction costs and state-dependent drift. The method has time-complexity of $O(N^{2.5})$, where N is the number of time steps in the discrete approximation of the investment interval. In contrast, a discrete-time dynamic programming algorithm (see (8) in Section 3) that directly solves the problem has time-complexity $O(N^5)$. Our method allows us, for instance, to study 40-year investment horizons with time steps of 4-day length on a basic laptop computer.

There are several numerical methods for solving the optimization problem with a constant drift under transaction costs. Davis, Panas, and Zariphopoulou (1993) proposes a backward recursive method which has seen a number of improvements in the past 20 years. For instance, Monoyios (2004) provides an approximation to the optimal decision in the final period which allows searching over a smaller range of stock holdings. Zakamouline (2005) proposes bounds on stock holdings. Another method is to solve the Hamilton–Jacobi–Bellman (HJB) equations of optimization problems by finite differences (e.g. Herzog, Kunisch, & Sass, 2013) or to use a genetic programming algorithm to derive approximations of trading strategies (Lensberg & Schenk-Hoppé, 2013). These algorithms work well for short time-horizons, typically less than 1 year, and when the number of periods is small. By proposing a method that works for

☆ The authors are grateful to the two anonymous reviewers and the editor, Professor Immanuel Bomze, for their helpful comments and advice.

☆☆ The paper benefitted from the first and third authors' visits to the Hausdorff Research Institute for Mathematics at the University of Bonn in the framework of the Trimester Program *Stochastic Dynamics in Economics and Finance*. The second author gratefully acknowledges support from the Danish Strategic Research Council under contract number 10-092299, the research center 'HIPERFIT'. The fourth author acknowledges support from Kent Faculty of Sciences Research Fund.

* Corresponding author. Tel.: +44 1227 823593; fax: +44 1227 827932.
E-mail addresses: J.Palczewski@leeds.ac.uk (J. Palczewski), Rolf@math.ku.dk (R. Poulsen), Klaus.Schenk-Hoppe@manchester.ac.uk (K. R. Schenk-Hoppé), H.Wang@kent.ac.uk (H. Wang).

http://dx.doi.org/10.1016/j.ejor.2014.12.040
0377-2217/

non-constant drift and long time-horizons, our paper fills this gap in the literature.

The main challenges arising from a state-dependent drift are that the search for the optimal strategy has to be carried out for all nodes of a binomial tree, and that the state-dependent strategy results in a larger range of stock holdings. This increases the likelihood of over- and underflow arising for the exponential utility function as pointed out by Clewlow and Hodges (1997). For a constant drift, in contrast, the optimal strategy is independent of stock prices at time t. One only needs to search for the optimal strategy at a node at time t, see Monoyios (2004, p. 902).

To overcome the challenges, we develop a fast numerical method to approximate the optimal solution well. The approach combines four aspects: (a) reducing dimensionality, (b) scaling the objective function, (c) carrying out local searches for optimal trading decisions, and (d) non-equidistant discretization of the state space.

We apply the numerical method to a study of the true costs of market frictions using the indifference principle. The analysis reaps the full benefit of the approach because measuring these costs requires taking averages over many realizations of the drift. For each realization, one has to calculate trading strategies and carry out Monte Carlo simulations. In general, a state-dependent drift is observed to make the strategy more variable than a constant drift. This, in turn, entails more aggressive trading.

The indifference principle yields the following results.

First, the value of information is measured by comparing realized utilities of different types of investors. We find that information is most valuable to the least risk-averse investor. It also turns out that cautious trend-followers do almost as well as investors who estimate the drift from observations.

Second, the utility loss due to transaction costs is measured as the maximum amount of money an investor is willing to pay up front to avoid transaction costs. The loss is observed to be about twice as large as the direct expense incurred. Transaction costs are most detrimental to naive investors (who do not revise their initial estimates of the drift) when investing over a medium or long time horizon. It implies that in the long run naive investors are the most active traders and usually hold wrong beliefs. At short time-horizons, transaction costs strongly affect the learning investor as his estimate of the drift varies drastically in the short run.

Third, we examine the impact of the investment time horizon. The main finding is that, although uncertainty about the true drift cannot be removed completely, learning about the drift reduces the loss in utility due to the uncertain drift by 33 percent in 1 year and by 80 percent in 10 years compared to a naive investor. Learning also reduces the loss in utility caused by transaction costs by 50 percent over a 10-year time-horizon.

Section 2 presents the model. The numerical method is explained in Section 3 and applied in Section 4 to quantify the economic costs under various assumptions on the state-dependent drift. Section 5 concludes.

2. Model

We consider an investor who maximizes utility from wealth by trading in a risk-free bond with a constant interest rate r, and a risky stock. The randomness of the stock price is modelled on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ which supports a one-dimensional Brownian motion $(W(t))$ and an independent random variable m whose role will be explained later. The investor assumes that the dynamics of the stock price $S(t)$ is given by

$$dS(t) = \mu\big(t, S(t)\big)S(t)dt + \sigma S(t)dW(t), \qquad S(0) = S_0 \tag{1}$$

with a constant volatility $\sigma > 0$. The function $\mu(t, S)$ is a time- and state-dependent drift of the stock price.

We consider a situation in which the true dynamics of the stock price is unknown: The actual drift is a random variable m which is determined at the initial time and fixed over the horizon (recall that it is independent of the Brownian motion $(W(t))$). Hence the true price dynamics is

$$dS(t) = mS(t)dt + \sigma S(t)dW(t). \tag{2}$$

The drift m is not observed by investors with an exception of an informed investor (a benchmark) who additionally knows the drift m. If the structure of the price dynamics is known, one can use observed stock prices to estimate m. Assume throughout the paper that m is normally distributed with mean μ_0 and variance $\gamma_0 > 0$:

$$m \sim \mathcal{N}(\mu_0, \gamma_0).$$

Then the Kalman–Bucy filter gives that the best estimate of m given an observation of the stock price trajectory up to time t is

$$\mu^{\mathrm{L}}(t, S(t)) = \frac{\gamma_0\sigma^2}{\sigma^2 + \gamma_0 t}\left(\frac{\mu_0}{\gamma_0} + \frac{t}{2} + \frac{1}{\sigma^2}\log(S(t)/S_0)\right). \tag{3}$$

This estimate takes the form $\mu(t, S(t))$, and hence entails a dynamics as defined in (1).

Investors who are not aware of the characteristics of the random variable m and/or the dynamics (2) make suboptimal decisions. We consider two types of such investors. The first one is a naive investor who assumes that the dynamics is given by (2) with $m = \mu_0$, i.e., $\mu(t, S(t)) = \mu_0$ in (1). The second type of investor suffers from a behavioral bias and estimates the drift as:

$$\mu^{a}(t, S(t)) = \mu_0 + a \arctan\Big((\mu_0 - \sigma^2/2)\,t - \log(S(t)/S_0)\Big). \tag{4}$$

The second item of (4) characterizes the investor's adjustment to his initial estimate μ_0. The arctangent function is a symmetric about the origin and increasing function taking values within $(-\pi/2, \pi/2)$ on the domain $(-\infty, +\infty)$, see, e.g. Luderer, Nollau, and Vetters (2010, p. 55). The adjustment vanishes when the logarithmic return $R(t) := \log(S(t)/S_0)$ equals $(\mu_0 - \sigma^2/2)t$ which was the expected value $\mathbb{E}[R(t)]$ if the drift of the stock price was a known constant μ_0. In this case, it is known that, see, e.g. Øksendal (2003, p. 64)

$$R(t) := \log(S(t)/S_0) = (\mu_0 - \sigma^2/2)\,t + \sigma\, W(t).$$

We refer to the parameter 'a' as the investor's *sentiment*. It measures the investor's confidence in his initial estimate μ_0. If the parameter a is positive then the investor believes that the price will revert to the predicted mean: A higher than predicted return is forecast to lead to a drift smaller than μ_0. The investor's decision is contrarian. It can be interpreted as the result of overconfidence about the ability to predict the stock price dynamics. If the parameter a is negative, the investor will revise the initial estimate upwards if the returns are higher than predicted (resp. downwards if returns are lower than μ_0). The investor is a trend follower; he places more trust in the market's view about stock price dynamics than in his own view.

Definition 2.1. **Informed** investors observe the realization of the random drift m at the initial time.
Learning investors use (3) to estimate the realization of the random drift m.
Naive investors assume that the drift is constant $m = \mu_0$.
Biased investors use (4) as their estimate of the drift.

Trading in the stock incurs proportional transaction costs with the rate $\lambda \in [0, 1)$. Purchasing y shares costs $y(1 + \lambda)S(t)$ at time t while selling y shares brings in $y(1 - \lambda)S(t)$. It is customary (e.g.