

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. Decision Support

Accepted Manuscript

Dynamic portfolio optimization with transaction costs and state-dependent drift

Jan Palczewski, Rolf Poulsen, Klaus Reiner Schenk-Hoppé, Huamao Wang

PII: S0377-2217(14)01058-3
DOI: [10.1016/j.ejor.2014.12.040](https://doi.org/10.1016/j.ejor.2014.12.040)
Reference: EOR 12699



To appear in: *European Journal of Operational Research*

Received date: 25 April 2014
Revised date: 18 December 2014
Accepted date: 21 December 2014

Please cite this article as: Jan Palczewski, Rolf Poulsen, Klaus Reiner Schenk-Hoppé, Huamao Wang, Dynamic portfolio optimization with transaction costs and state-dependent drift, *European Journal of Operational Research* (2014), doi: [10.1016/j.ejor.2014.12.040](https://doi.org/10.1016/j.ejor.2014.12.040)

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

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

Jan Palczewski

School of Mathematics, University of Leeds, UK.

Rolf Poulsen

Department of Mathematical Sciences, University of Copenhagen, Denmark.

Klaus Reiner Schenk-Hoppé

Leeds University Business School and School of Mathematics, University of Leeds, UK.

Department of Finance, NHH–Norwegian School of Economics, Norway.

Huamao Wang*

School of Mathematics, Statistics and Actuarial Science, University of Kent, UK.

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 nu-

[☆]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.

Email addresses: J.Palczewski@leeds.ac.uk (Jan Palczewski), Rolf@math.ku.dk (Rolf Poulsen), K.R.Schenk-Hoppe@leeds.ac.uk (Klaus Reiner Schenk-Hoppé), H.Wang@kent.ac.uk (Huamao Wang)

merical 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.

Keywords: Dynamic programming, numerical methods, state-dependent drift, transaction costs, Markov Chain approximation

JEL: C61, C63, G11

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 et al. 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 nu-

merically 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 et al. \(1993\)](#) 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 et al. 2013](#)) or to use a genetic programming algorithm to derive approximations of trading strategies ([Lensberg and Schenk-Hoppé 2013](#)). These algorithms work well for short time-horizons, typically less than one year, and when the number of periods is small. By proposing a method that works for 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 trans-

action 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% in one year and by 80% in ten years compared to a naive investor. Learning also reduces the loss in utility caused by transaction costs by 50% 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

$$dS(t) = \mu(t, S(t))S(t)dt + \sigma S(t)dW(t), \quad S(0) = S_0 \quad (1)$$

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

$$dS(t) = mS(t)dt + \sigma S(t)dW(t). \quad (2)$$

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

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

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:

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

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 et al. \(2010, p. 55\)](#). 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. Davis et al. 1993) to describe an investor's trading strategy with two non-decreasing right-continuous processes $L(t)$ and $M(t)$ representing, respectively, the cumulative number of shares bought and sold over $[0, t]$. The dynamics of portfolio positions $(x(t), y(t))$, where $x(t)$ is the value of bonds held and $y(t)$ is the number of shares, is

$$\begin{aligned} dx(t) &= rx(t)dt - (1 + \lambda)S(t)dL(t) + (1 - \lambda)S(t)dM(t), \\ dy(t) &= dL(t) - dM(t). \end{aligned}$$

Given an initial position (x_0, y_0) , the investor maximizes the expected utility of the terminal wealth by following a trading strategy $(L(t), M(t))$:

$$\max_{(L, M)} \mathbb{E}\{U(x(T) + y(T)S(T))\}.$$

We impose two standard assumptions: there are no liquidation costs of the portfolio at the terminal time T and the investor has a utility function with a constant absolute risk aversion (CARA) coefficient α :

$$U(w) = -\exp(-\alpha w). \tag{5}$$

In the cases of an informed investor or a naive investor, this utility maximization problem is classical. For learning investors one can show

that it is optimal to estimate the true drift using (3) and to solve the optimization problem under the stock price dynamics given by (1) with $\mu(t, S(t)) = \mu^L(t, S(t))$.¹ Biased investors' optimization problem mimics behavioral decision making.

Stochastic differential equations with drift of the form (3) or (4) do not satisfy the standard conditions for existence and uniqueness of solution. We therefore provide a result that establishes existence of a unique solution.

Lemma 2.2. *Assume that $\mu : [0, T] \times (0, \infty) \rightarrow \mathbb{R}$ is a continuous function that satisfies a logarithmic growth condition*

$$|\mu(u, S)| \leq M(1 + |\log(S)|), \quad S > 0, \quad u \in [0, T],$$

and a logarithmic Lipschitz condition

$$|\mu(u, S_1) - \mu(u, S_2)| \leq M|\log(S_1) - \log(S_2)|, \quad S_1, S_2 > 0, \quad u \in [0, T]$$

for some constant $M > 0$. Then there is a unique strong solution to the stochastic differential equation (1) for every initial condition $S > 0$.

Proof. Øksendal (2003, Theorem 5.2.1) implies that under the assumptions of the lemma there is a unique strong solution to the stochastic differential equation

$$dZ(u) = \left(\mu(u, e^{Z(u)}) - \frac{\sigma^2}{2} \right) du + \sigma dW(u), \quad Z(t) = 0. \quad (6)$$

¹The justification is based on the separation principle (Fleming and Rishel 1975, Theorem 11.2) and a Kalman-Bucy filter (Øksendal 2003, Chapter 6). The original optimization problem is equivalent to the one with the drift replaced by its filtering estimate (3).

By Itô's formula the process $S(u) = S(t)e^{Z(u)-Z(t)}$, $u \geq t$, satisfies (1), i.e., it is a strong solution to this equation. To prove uniqueness, assume that there is another strong solution to (1), denoted by $\bar{S}(u)$, $u \geq t$, with $\bar{S}(t) = S(t)$ and $\bar{S}(u) \neq S(u)$ for $u > t$. Define $\bar{Z}(u) = \log(\bar{S}(u)/\bar{S}(t))$. Again, by Itô's formula $\bar{Z}(u)$ satisfies (6) and is different from $Z(u)$. This contradicts the uniqueness of the solution to (6). \square

Let us verify that the drifts of the forms (3) and (4) satisfy assumptions of the above lemma. We have

$$|\mu^L(u, S_1) - \mu^L(u, S_2)| \leq \frac{\gamma_0}{\sigma^2} |\log(S_1) - \log(S_2)|,$$

and

$$\begin{aligned} |\mu^L(u, S)| &\leq \sup_{t \geq 0} |\mu^L(t, S)| \\ &\leq \sup_{t \geq 0} \left\{ \frac{\gamma_0 \sigma^2}{\sigma^2 + \gamma_0 t} \left(\frac{\mu_0}{\gamma_0} + \frac{t}{2} + \frac{1}{\sigma^2} |\log(S_0)| \right) \right\} + \sup_{t \geq 0} \left\{ \frac{\gamma_0 \sigma^2}{\sigma^2 + \gamma_0 t} \frac{1}{\sigma^2} \log(S) \right\} \\ &\leq \frac{\gamma_0}{\sigma^2} |\log(S_0)| + \frac{\sigma^2}{2} + \mu_0 + \frac{\gamma_0}{\sigma^2} |\log(S)|. \end{aligned}$$

For a biased investor, we obtain

$$|\mu^a(u, S)| \leq \mu_0 + |a| \frac{\pi}{2},$$

and

$$\begin{aligned} |\mu^a(u, S_1) - \mu^a(u, S_2)| &\leq |a| \sup_{x \in (-\infty, \infty)} |\arctan'(x)| |\log(S_1) - \log(S_2)| \\ &\leq |a| |\log(S_1) - \log(S_2)|. \end{aligned}$$

Denote by $V(t, s, x, y)$ the value function corresponding to the utility optimization problem. This is the highest expected utility achievable by an

investor whose portfolio at time t consisting of x value of bond and y shares of the stock priced at $S(t) = s$:

$$V(t, s, x, y) = \sup_{(L(u), M(u))_{u \geq t}} \mathbb{E}\{U(x(T) + y(T)S(T)) \mid (S(t), x(t), y(t)) = (s, x, y)\}.$$

In the simplest case when the drift function $\mu(t, s) \equiv \bar{\mu}$ (a constant), the value function is characterized as a unique viscosity solution of an HJB equation (Davis et al. 1993):²

$$\max \left\{ V_t + rxV_x + \bar{\mu}sV_s + \frac{\sigma^2}{2}s^2V_{ss}; \right. \\ \left. V_y - (1 + \lambda)sV_x; \quad -V_y + (1 - \lambda)sV_x \right\} = 0 \quad (7)$$

with the terminal condition $V(T, s, x, y) = U(x + ys)$ (subscripts in (7) denote partial derivatives). Solving this equation is usually carried out using numerical approximation. For general drift functions, an HJB representation is not known. We therefore take a different route to study optimal investment when the drift function is time- and state-dependent. In this paper, approximations are designed for the stochastic control problem itself.

3. Numerical Approach

We apply Bellman's dynamic programming principle to solve the control problem with state-dependent drift. The stock price model is discretized in time and space, and the programming works recursively backwards in time.

Let time be discretized in steps of length Δt with $\Delta t = T/N$ where N is the number of time steps. At each time-point the investor has to choose

²This result requires a restriction of the set of available trading strategies $(L(t), M(t))$: the liquidation value at any time must be greater than or equal to a fixed constant.

whether to trade and, if yes, how many units of stock to trade. The bond holdings are determined by the self-financing condition. The expected utility derived from each possible trading choice is determined by the value function. To select the trading decision that maximizes expected utility, the investor solves the maximization problem:

$$\begin{aligned}
 V(t, s, x, y) = \max \bigg\{ & \underbrace{\mathbb{E} (V(t + \Delta t, S(t + \Delta t), e^{r\Delta t}x, y) | S(t) = s)}_{\text{benefit from not trading, } \Delta y = 0}, \\
 & \max_{\Delta y > 0} \underbrace{\mathbb{E} (V(t + \Delta t, S(t + \Delta t), e^{r\Delta t}(x - \Delta y \times s(1 + \lambda)), y + \Delta y) | S(t) = s)}_{\text{benefit from buying } \Delta y > 0 \text{ shares}}, \\
 & \max_{\Delta y > 0} \underbrace{\mathbb{E} (V(t + \Delta t, S(t + \Delta t), e^{r\Delta t}(x + \Delta y \times s(1 - \lambda)), y - \Delta y) | S(t) = s)}_{\text{benefit from selling } \Delta y > 0 \text{ shares}} \bigg\} \quad (8)
 \end{aligned}$$

where the maximization is over the type of trade and the corresponding volume to be traded.

One might conjecture that the spatial discretization of the stock price process is complicated when its drift is state-dependent. However, one can use a standard binomial tree approximation of [Cox et al. \(1979\)](#) and define adjusted probabilities for the up- and down-movement of the discretized stock price. This Markov Chain approximation is provided in, e.g., [Kushner and Dupuis \(1992\)](#) and [Zakamouline \(2005\)](#). The benefit of this representation is that the stock-price model retains the property of being a recombining tree. Specifically, we use the following binomial model. Define the coefficients $u = 1/d = e^{\sigma\sqrt{\Delta t}}$, and set the process

$$S(t + \Delta t) = \begin{cases} uS(t) & \text{with probability } p(t, S(t)) = [e^{\mu(t, S(t))\Delta t} - d]/[u - d], \\ dS(t) & \text{with probability } 1 - p(t, S(t)). \end{cases}$$

A natural discretization of the state space of money and stock holdings

is given by the set $M_x \times M_y$ with $M_x = \{x_j : x_j = \underline{x} + j\delta x \leq \bar{x}, j \in \mathbb{N}\}$ and $M_y = \{y_k : y_k = \underline{y} + k\delta y \leq \bar{y}, k \in \mathbb{N}\}$ with given \underline{x} , \bar{x} , \underline{y} , and \bar{y} , where δx (resp. δy) is the grid spacing in the dimension of money (resp. stock holdings).

A direct algorithm for determining the value function and the optimal trading strategy proceeds as follows.

Define the value function at the terminal time as the realized utility. Set $V(T, s, x_j, y_k) := U(x_j + y_k s)$ for all values s of the discretized stock prices in period T and all portfolio holdings $(x_j, y_k) \in M_x \times M_y$.

For $t = T - \Delta t, \dots, 0$

For all values of the discretized stock price $s = S(t)$ at time t

For all values $(x_j, y_k) \in M_x \times M_y$

Given the functions $V(t + \Delta t, \dots)$, find the highest value in (8) obtained over all values Δy such that $y_k + \Delta y \in M_y$ ³. Denote the maximum value $V(t, s, x_j, y_k)$.

End For

End For

End For

The computational complexity of the direct method is of the order $O(N^2 \times M_x \times M_y \times M_y)$ or $O(N^5)$.⁴ The factor N^2 arises because the algorithm loops through all points on the stock price lattice, the factor $M_x \times M_y$ is due to the loop through the grid of portfolio holdings, and the final factor M_y

³ $V(t + \Delta t, \dots)$ is approximated via a linear interpolation because $\exp(r\Delta t)[x_j \mp \Delta y s(1 \pm \lambda)]$ is typically not an element of M_x .

⁴We let M_x and M_y linearly depend on time steps N to ensure that the grid sizes δx and δy approach 0 when Δt is close to 0 with increasing N , see, e.g. [Zakamouline \(2005\)](#).

comes from the Δy -search. This is slow; doubling the number of steps in all dimensions increases computation time by a factor of 32.

The range of $M_x \times M_y$ is usually large in order to include optimal solutions for all possible states $(t, S(t))$ on the lattice. The above direct numerical method uses a standard equidistant grid and searches for optimal solutions in the whole set.

As a benchmark, suppose the direct algorithm is implemented in a high-level language such as *Matlab* on a typical laptop computer. Pricing an option on a binomial lattice with $T = 1$ year and time steps of 1 day takes 5 - 10 milliseconds, while on the lattice doing optimization over a grid of $M_x \times M_y = 100 \times 100$, i.e. $O(M_x \times M_y^2) = O(10^6)$, takes about 2 hours. This is not a computationally feasible approach since reasonable outputs require high-resolution grids and thousands of simulations of a random drift.

Five measures are employed to reduce running time of simulations:

Reducing dimension. When measuring utility by the negative exponential function (5), the value function V can be written in the form

$$V(t, s, x, y) = H(t, s, y) \exp(-\alpha x \exp(r(T - t))), \quad (9)$$

where $H(t, s, y)$ is defined by $H(t, s, y) := V(t, s, 0, y)$, see, e.g., [Davis et al. \(1993\)](#) or [Monoyios \(2004\)](#). This representation allows reducing the dimension of the optimization problem by one. However, this measure carries a potential cost. Suppose an investor's money and stock holdings are large (in absolute terms) but offsetting in terms of value. Then the exponent of the exponential utility function implied by $H(t, s, y)$ will include the product of a very large stock holding and a large stock price. This can cause nu-

merical over- or underflow errors in the computer program, which are dealt with by our following function $H(t, s, y)$ scale, along with local search and non-equidistant discretization that speeds up the program.

Scaling the function $H(t, s, y)$. To mitigate over- and underflow issues, the value function $H(t, s, y)$ is scaled:

$$G(t, s, y) := V(t, s, -ys, y) = H(t, s, y) \exp(\alpha y s \exp(r(T - t))).$$

Function G satisfies a discrete time dynamic programming equation similar to (8) with the terminal condition $G(T, s, y) = -1$.

Local Δy -search. The solution to $H(t, s, y)$ is known to have a particular structure. The space of stock holdings is split into three regions: buy, no-trade and sell. If the stock position is either in the buy or sell region, a trade is initiated that leads to a stock position on the closest boundary of the no-trade region. If the stock position is inside or on the boundary of the no-trade region, the investor does not change his stock position.

In the case of a constant drift (Monoyios 2004, p. 902) the upper boundary (above which one sells) and the lower boundary (below which one buys) of the no-trade region at a given time t can be both defined by market values of stock positions. It is therefore sufficient to determine the optimal trade in all time- t nodes with a node (t, S) to find the two boundaries.

With a state-dependent drift, this observation no longer holds true: If the forecast of the drift is revised depending on the current stock price, then the no-trade region will depend on this information. One therefore has to determine a no-trade region in each node (t, S) . This is computationally costly. A numerically efficient approach, which we implement, is to deter-

mine the boundaries of the no-trade region through searching over a local range of y . One can also use a binary search algorithm as in [Zakamouline \(2005\)](#) to improve computational speed further. The local range denoted by $[\varphi_b(t, S), \varphi_s(t, S)]$ is determined by an appropriate extension of the boundaries at the successive nodes⁵.

Non-equidistant y -discretization. The structure of optimal trading strategies suggests that it is not efficient to have an equidistant discretization of the y -space. The set of discretization points should be denser close to the boundaries of the no-trade region. We therefore use a symmetric, non-equidistant discretization.

The set is centered at Merton's closed-form solution for the case of a constant drift and no transaction costs, which is denoted by $\varphi_M(t, S)$. The value of drift μ is given by investors (possibly an actual value or an estimate). The non-equidistant grid has larger step-sizes away from the center $\varphi_M(t, S)$. For a given (t, S) -node and the local range $[\varphi_b(t, S), \varphi_s(t, S)]$, we first define the radius

$$\Phi(t, S) := \max \{ \varphi_M(t, S) - \varphi_b(t, S), \varphi_s(t, S) - \varphi_M(t, S) \}, \quad (10)$$

⁵Denoting the buy (resp. sell) boundary by y_b (resp. y_s) we identify the endpoints by

$$\varphi_b(t, S) = \min \{ y_b(t + \Delta t, dS(t)), y_b(t + \Delta t, uS(t)) \} - C_1,$$

$$\varphi_s(t, S) = \max \{ y_s(t + \Delta t, dS(t)), y_s(t + \Delta t, uS(t)) \} + C_2,$$

where C_1 and C_2 are two positive constants selected to ensure the local range is large enough. We check whether the no-trade boundaries obtained numerically hit the endpoints of the local range. If they do, larger values of C_1 and C_2 are chosen and the corresponding computation is repeated.

where

$$\varphi_M(t, S) = \frac{\mu - r}{e^{r(T-t)} \alpha \sigma^2 S}$$

is the Merton solution. Then we define the set of discretization points as:

$$y(t, S, k) = \varphi_M(t, S) + \frac{\Phi(t, S)}{\left(\frac{M_y}{2}\right)^\varpi} \left(k - \frac{M_y}{2}\right) \left|k - \frac{M_y}{2}\right|^{\varpi-1}. \quad (11)$$

The coefficient $\varpi > 1$ controls the level of dispersion.⁶ Numerical experiments (see Wang (2010, Sect. 3.6.4) for details) show that an appropriate choice of the coefficient ϖ is 1.6.

Low-level language. Implementation in a low-level language, e.g., C++, reduces computation time by a factor of approximately 10.

Numerical illustration. We use the following values of parameters as a base case for our numerical results: the actual drift drawn from the normal distribution with mean $\mu_0 = 0.15$ and variance $\gamma_0 = 0.04$, volatility $\sigma = 0.25$, proportional transaction cost rate $\lambda = 0.01$, initial stock price $S_0 = 15$, risk aversion $\alpha = 0.1$, interest rate $r = 0.03$, time-horizon $T = 1$ year, and discretization parameters $\Delta t = 0.01$, $M_y = 3,500$ and $\varpi = 1.6$.

Figure 1 demonstrates the joint effect of transaction costs and state-dependent drift. It shows one realization of the optimal trading strategy over a 40-year time-horizon. The effect is substantial as evidenced by the high variability of the boundaries of the no-trade region. The volatility of

⁶If $\varpi = 1$, the grid degenerates to the equidistant discretization. But when ϖ is large, the points are too concentrated around the center.

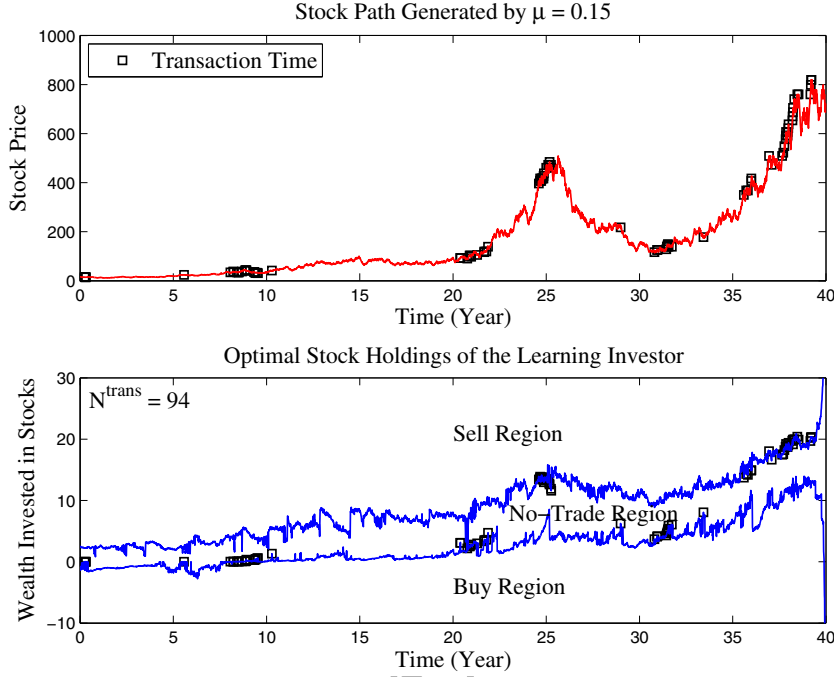


Figure 1: Dynamics of the no-trade region with state-dependent drift $\mu^L(t, S(t))$ within $T = 40$ years horizon. The squares indicate transaction times. N^{trans} is the total number of transactions.

these boundaries reflects changes in the learning investor's estimate of the drift. For instance, both boundaries move downwards around year 25-30 in response to a pronounced fall in the stock price. They move upwards again around year 32 when the stock price recovers. With a known, constant drift, these boundaries (when measured in terms of the amount of wealth invested in stocks) are hyperbola-like curves that are independent of the stock price.

Comparison with Monoyios (2004)'s results. Verification of our method is carried out by comparing numerical results with those reported in Monoyios (2004). The comparison is for the simple case of a known, constant drift

which is considered in the latter paper. Table 1 reports the two boundaries of the no-trade region at the initial time for different transaction costs. We calculate results with our method under both equidistant and non-equidistant discretization. In all three scenarios and for different transaction costs the calculated boundaries coincide up to 3-4 significant digits.

The non-equidistant discretization requires fewer points on the y -grid than the equidistant discretization, which substantially shortens the runtime of the program. Our approach works efficiently because we take state-dependent non-equidistant discretization on a small local range of y -values. In fact, the discretization equation (11) produces a great number of dense points with the precision up to 0.0001 around the area centered at the Merton solution where the no-trade region is most probably located. The discretization points are gradually becoming sparser towards the two end-points of the local range of y -values.⁷ As a result, it suffices to set $M_y = 3,500$ to achieve results similar to those obtained by the standard equidistant discretization that requires 0.27–2.38 million grid points, depending on the full range of y -values, see the last row in Table 1.

We also compare the performance of non-equidistant and equidistant discretizations in the case of the state-dependent drift $\mu^L(t, S(t))$. Figure 2 shows that the most stable results are obtained under the non-equidistant discretization. The precision of the approximation increases gradually as the number of time steps increases. Equidistant discretizations exhibit a more volatile behavior.

⁷See Wang (2010, Figures 3.5 and 3.6) for an example of the frequency histogram and the diagram of varying precision of y -values.

	$\lambda = 0.005$	$\lambda = 0.01$	$\lambda = 0.02$	$\lambda = 0.03$
Monoyios	[0.3866, 0.5780]	[0.3499, 0.6197]	[0.2702, 0.7196]	[0.1813, 0.8243]
Equidistant	[0.3870, 0.5772]	[0.3510, 0.6193]	[0.2708, 0.7137]	[0.1851, 0.8161]
Non-equidistant	[0.3864, 0.5763]	[0.3527, 0.6209]	[0.2720, 0.7177]	[0.1826, 0.8113]
Range of y	[-10.748, 16.213]	[-24.622, 30.443]	[-56.800, 64.677]	[-109.43, 128.71]

Table 1: Boundaries of no-trade region at $t = 0$. The first row is taken from Monoyios (2004, Table 1) using a binomial lattice: $r = 0.1$, $\Delta t = 0.02$, $\mu = 0.15$, $\sigma = 0.25$, $S_0 = 15$, $\alpha = 0.1$, $T = 1$ year. The second row uses the equidistant discretization with $\Delta y = 0.0001$, while the third row uses the non-equidistant discretization (11) with $M_y = 3,500$ and $\varpi = 1.6$. The last row presents the ranges of y grid determined by equations (A.2) and (A.5) in Monoyios (2004).

We finally consider the relationship between computation time and numerical accuracy. Figure 3 shows the log-log scale plot⁸ for the absolute error $|V_i - \hat{V}|$ of the value function $V(t = 0, s = 15, x = 0, y = 0)$ and computation time for the non-equidistant discretization with local search in the case of the state-dependent drift $\mu^L(t, S(t))$. Specifically, the quantities V_i are the results using non-equidistant discretization and $N = 20 + i \times 20$, $i = 0, 1, \dots, 9$, in Figure 2. The benchmark $\hat{V} = -0.9018$ is obtained by using non-equidistant discretization and $N = 420$ in Figure 2. We assume \hat{V} is a reliable approximation of the true value of $V(0, 15, 0, 0)$. Then the difference $|V_i - \hat{V}|$ between V_i and the “true” value \hat{V} is the error of numerical algorithm. An increase in N reduces the error at the cost of longer computation time, as shown in Figure 3.

⁸A log-log plot uses logarithmic scales on both the horizontal and vertical axes.

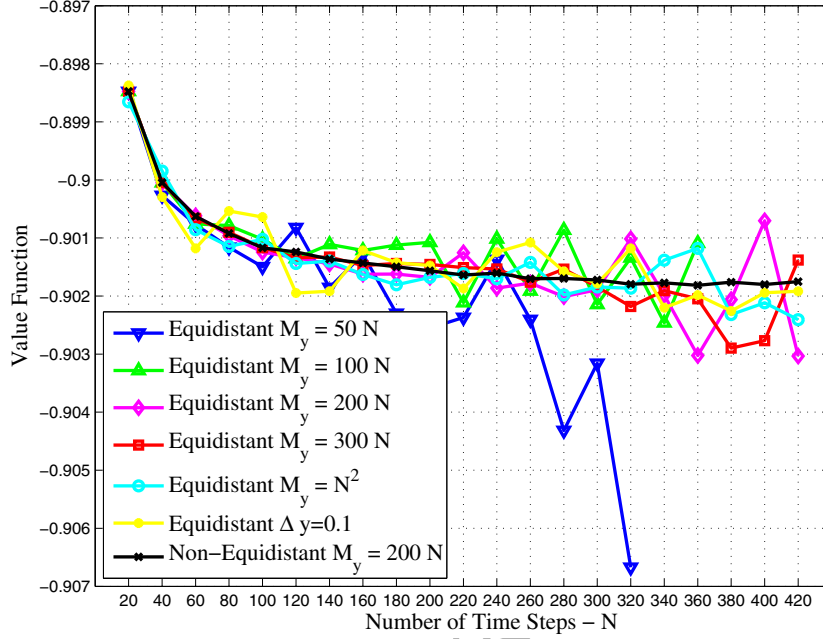


Figure 2: Value functions at initial time versus the number of time steps N with the state-dependent drift $\mu^L(t, S(t))$, where $N = 20 + i \times 20$, $i = 0, 1, \dots, 20$. Other parameters are the same as in the base case.

To estimate the order of time-complexity, we first assume the relationship $|V - \hat{V}| = a \tau^{-b}$, where τ is computation time and we call b the convergence order. We estimate b (and $\log a$) by performing an ordinary least squares regression of $\log(|V - \hat{V}|)$ on $\log(\tau)$.

All observations in Figure 3 are close to a straight line with slope -0.4 (taking logarithms of both variables). This means that to halve the numerical error, computing time is increased by a factor of $2^{(1/0.4)} \approx 5.7$. Note that this is only marginally slower than for standard option pricing calculations in a binomial model (where computing needs to be quadrupled for the error to

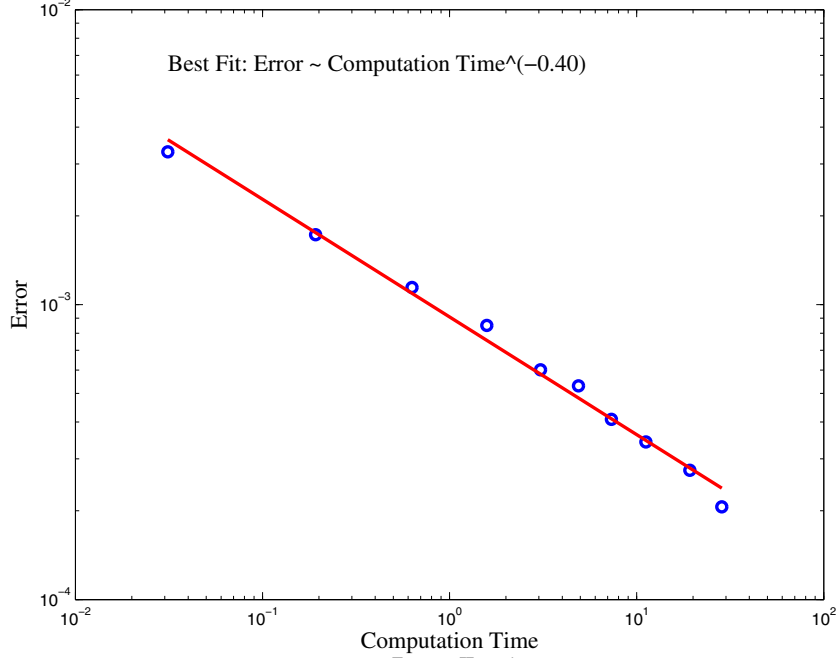


Figure 3: Convergence with non-equidistant discretization and the state-dependent drift $\mu^L(t, S(t))$. The y-axis reports the absolute error $|V_i - \hat{V}|$ of the value function $V(t = 0, s = 15, x = 0, y = 0)$. The benchmark $\hat{V} = -0.9018$ has been obtained by using $N = 420$, and V_i 's are the results with $N = 20 + i \times 20$, $i = 0, 1, \dots, 9$.

be halved) and much faster than for the direct method (8), where a simple halving of all step sizes increases computation time by a factor of 32. In Appendix A, we investigate the convergence order for different random sets of values of parameters⁹. This confirms the robustness of the results reported in Figure 3.

⁹We are grateful to an anonymous reviewer for suggesting this analysis.

4. Results

The numerical solution technique is applied to study the effects of transaction costs and uncertainty over investment time-horizons of up to 10 years. We consider the four types of investors introduced in Definition 2.1.

Our numerical results provide three main insights of practical relevance:

- Not knowing the true stock price dynamics leads to large losses in utility for less risk-averse investors, strongly biased investors, and naive investors (in decreasing order).
- Learning generally reduces the loss in utility caused by uncertainty about the true drift.
- Lower trading volumes due to transaction costs explain about half of the total loss in utility. The other half is caused by transaction cost payments.

When comparing the choices of different investors that are in the same situation or of identical investors who are in different situations, one has to take into account two aspects. First, quantifying an investor's gain or loss should be done using monetary units. This allows expressing differences in utility as the value of contract that, for instance, provides the investor with information about the drift or frees an investor from having to pay transaction costs. These values are defined as the amount of wealth that an investor has to pay (needs to receive) at initial time in order to be indifferent between two situations. Second, naive investors and investors with biases make trading decisions that are not optimal. Such an investor will obtain

a lower average realized utility than expected ex ante. We therefore take realized rather than perceived utility when measuring losses relative to an informed investor.

Section 4.1 considers the value of knowing the realization of the drift and the true stock price dynamics ('value of information') and Section 4.2 analyzes the true (economic) cost of proportional transaction costs.

4.1. Value of information

For each investor type, the average realized utility is given by

$$R(x) := E_\mu \bar{U}_\mu(x) \quad (12)$$

where x is the initial money endowment (the initial share is zero). E_μ denotes expectation with respect to μ which has the distribution $\mathcal{N}(\mu_0, \gamma_0)$. The realized utility \bar{U}_μ is determined by the realized stock price path, the investor's realized trading strategy (L, M) , and the utility function U :

$$\bar{U}_\mu = \mathbb{E}\{U(x(T) + y(T)S(T)) \mid (L, M)\}.$$

The average realized utility cannot be higher than the expected one, i.e.

$$R(x) \leq E_\mu V_\mu(0, S_0, x, 0),$$

where $V_\mu(0, S_0, x, 0)$ is the value of expected utility for a given μ . For naive and biased investors, the inequality will, in general, be strict as these investors make incorrect assumptions about the stock price dynamics. Therefore they overestimate their expected utility. However, an informed investor's average realized utility satisfies

$$R^F(x) = E_\mu V_\mu^F(0, S_0, x, 0),$$

where $V_\mu^F(0, S_0, x, 0)$ is the expected utility which the investor maximizes under knowledge of the value of μ . For a learning investor, who always uses μ_0 as prior for the drift estimate at the initial time, the average realized utility is

$$R^L(x) = V^L(0, S_0, x, 0).$$

The monetary value of being informed rather than having to learn the true drift over time from observations is:

$$IE^L(x) = \sup\{c \geq 0 \mid R^L(x) \leq R^F(x - c)\}. \quad (13)$$

This is the maximum amount a learning investor can pay to obtain the true value of μ without being worse off can be interpreted as an *information equivalent* (IE). If the realization of the randomly drawn drift could be purchased then $IE^L(x)$ were the highest price a learning investor is willing to pay to be certain about the value μ . Since the utility function (5) is CARA, the measure defined in (13) is actually independent of the monetary endowment x .

As the value functions of these two investors satisfy (9), one finds

$$IE^L = \frac{1}{\alpha} \exp(-rT) \log(H^L / E_\mu H_\mu^F),$$

where H is the reduced form value function. An approximation \hat{H}^F of the expected value $E_\mu H_\mu^F$ is calculated as follows:

1. Draw independently M_μ values from the distribution $\mathcal{N}(\mu_0, \gamma_0)$.
2. For each random draw μ_i , calculate the value function $H_{\mu_i}^F$ by solving the portfolio optimization problem (8) with (9).

3. Calculate

$$\hat{H}^F = \frac{1}{M_\mu} \sum_i^{M_\mu} H_{\mu_i}^F.$$

Similar to (13), we can calculate the monetary value of being an informed investor rather than a naive investor or a biased investor. One first needs to solve the optimization problem to determine trading strategies. Using these strategies one can determine realized utility in a Monte Carlo simulation.¹⁰ To obtain the average realized utility one has to repeat this procedure for many draws of μ . In addition, these calculations have to be carried out for different levels of parameters for comparative analysis. The efficient numerical method in Section 3 allows performing these simulations in a matter of hours.

Figure 4 depicts information equivalents for different levels of risk aversion and different investor types. The lowest values are obtained for a learning investor. This confirms that empirical estimation of the drift using the Kalman-Bucy filter (3) is beneficial. The highest values are associated with aggressive trend-followers and contrarian investors while less aggressive ones have information equivalents close to that of the naive investor.

Information equivalents are decreasing in the risk aversion α : more risk-averse investors receive lower benefits from knowing the true drift. For instance, the investors with $\alpha = 0.5$ are only willing to pay from about 17% to

¹⁰We use $M_\mu = 1,001$ during a simulation. Our results show that our sample approximates the normal distribution well (Wang 2010, Sect. 4.8.6). We apply an inverse transformation method with the Beasley-Springer-Moro algorithm (Glasserman 2004, p. 68).

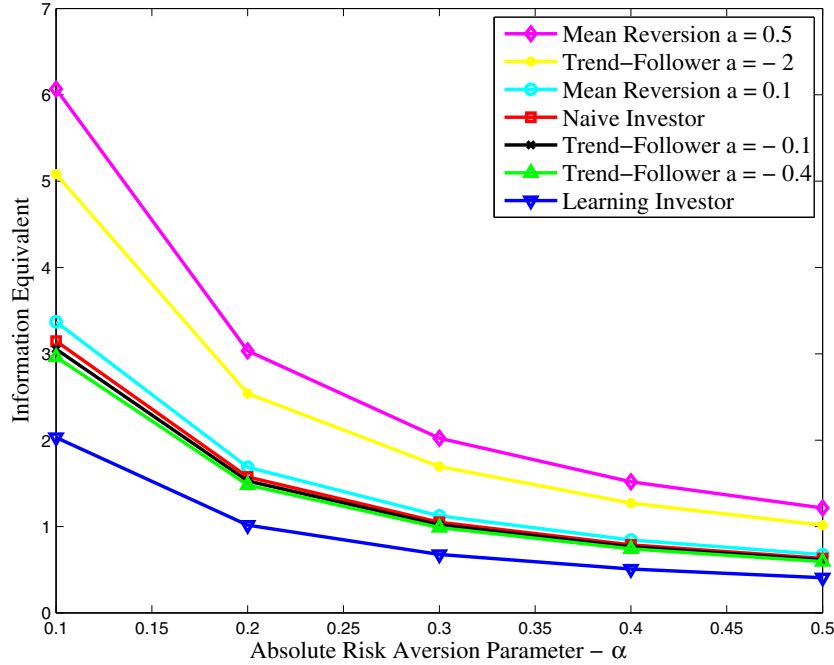


Figure 4: Information equivalents for different levels of risk aversion, $T = 1$ year.

25% as much as the investors with $\alpha = 0.1$ to remove uncertainty about the actual drift. At first sight this might be surprising as higher risk-aversion is generally associated with higher willingness to pay in order to avoid risk. The opposite is true here as higher risk aversion leads to less investment in the stock, see also [Muthuraman and Kumar \(2006\)](#). [Cvitanić et al. \(2006\)](#) also find that the certainty equivalents that they examine achieve the highest values for the lowest risk aversion in different setups.

The sentiment parameter a in (4) has a marked impact on information equivalents. Figure 5(a) shows the information equivalent is a U-shaped function of a varying from -2 (strongly trend-following) to 0.5 (strongly contrarian). The minimum is obtained for $a \approx -0.4$. A mild trend-following

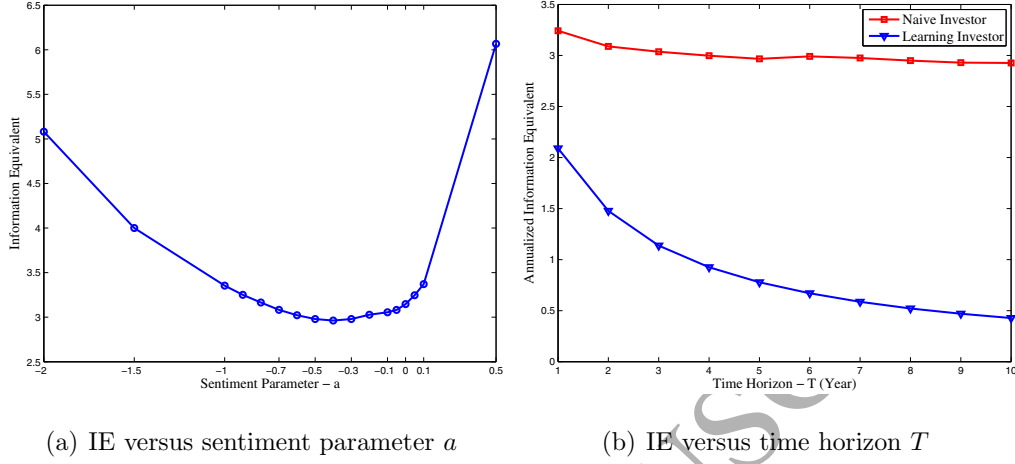


Figure 5: This figure illustrates information equivalents for: (a) biased investors with different values of a (see (4)): naive investor ($a = 0$), trend-follower ($a < 0$) and contrarian investor ($a > 0$); and (b) different investment horizons.

investor therefore mimics the optimal filtering. Hence, the trading strategy of an investor whose estimate of the drift is derived from cautious interpolation of an observed short-term trend, is close to that of a learning investor.

The effect of the investment time-horizon T on the (annualized) information equivalent is studied in Figure 5(b), which is defined as $IE \cdot e^{rT}/T$ since IE is defined at the initial time. First, the naive investor has more to gain from knowing the true drift than the learning investor, and the annualized benefit is fairly constant across different T . In contrast, the decreasing information equivalent of a learning investor reflects the gain in knowledge from filtering which reduces conditional variance when T increases. It also provides a hedge against unfavorable realizations of the drift (Brennan 1998).

The information equivalent is positive even at a 10-year investment hori-

zon. The lesson is that the true drift is difficult to estimate and one cannot eliminate uncertainty about the drift. Thus, learning via filtering has benefits even in the long run. A naive investor with a 1-year (10-year) horizon could reduce the loss by 33% (80%) when adopting a filtering strategy. Previous studies using filtering without transaction costs also find substantial utility gains from 2.93% to 215.73%, see [Cvitanic et al. \(2006\)](#).¹¹

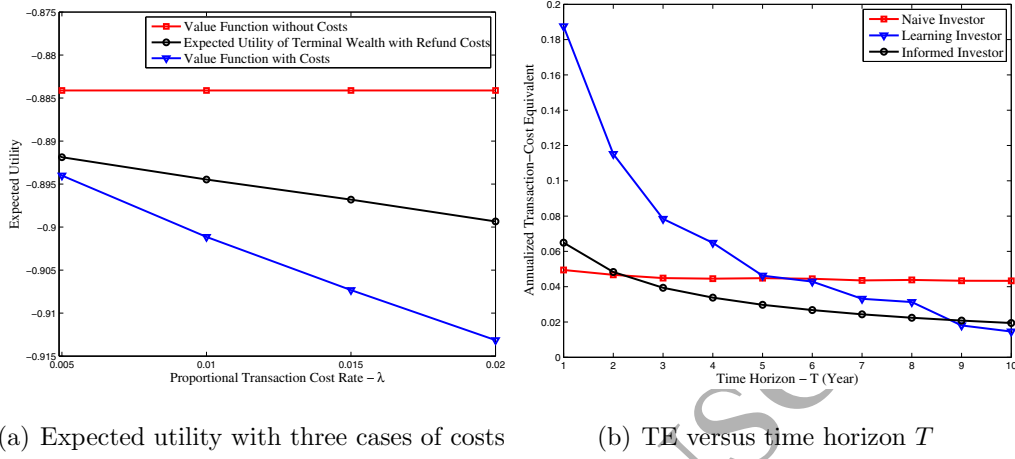
4.2. Transaction costs

Trading strategies are sensitive to transaction costs. Figure 6(a) shows the utility of a learning investor under different scenarios. The top line is the benchmark case of no transaction costs. The bottom line is the utility with transaction costs, which is decreasing as the proportional transaction cost increases. This coincides with previous studies (see, e.g. [Gennotte and Jung 1994](#)). In the range 0.5% to 2% the loss in utility is approximately linear.

This loss is caused by two effects of transaction costs: (a) a direct effect due to the additional expense incurred and (b) an indirect effect due to less trading. We strip out the first one by reimbursing all transaction costs (with interest) at the final period. The investor optimizes his strategy without knowing about this reimbursement. The result is the middle line in Figure 6(a) which is about halfway (except those for the small $\lambda < 0.01$) between the zero-cost and positive-cost without reimbursement case.

The difference between the reimbursement and the zero-cost case is the deadweight loss from the transaction costs. It measures the true economic

¹¹A quantitative comparison between the results of [Cvitanic et al. \(2006\)](#) and ours is inappropriate since the models and values of parameters are substantially different.



(a) Expected utility with three cases of costs

 (b) TE versus time horizon T

Figure 6: This figure depicts (a) maximum expected utility of a learning investor under three situations of transaction costs; and (b) transaction-cost equivalents of naive / learning / informed investors within different investment horizons.

cost of this friction. We find that the total effect of the transaction cost is about twice (except $\lambda < 0.01$) as large as the loss in utility due to less trading resulting from transaction costs. The implications are that freely re-balancing portfolio significantly contributes to expected utilities, and less re-balancing brings about half of the total loss.

To capture the value from investing in a market without transaction costs, we denote the gain to an investor of type \cdot as

$$TE(\lambda) = \sup\{c \geq 0 \mid E_{\mu} V_{\mu,\lambda}(0, S_0, x, 0) \leq E_{\mu} V_{\mu,\lambda=0}(0, S_0, x - c, 0)\}, \quad (14)$$

where $V_{\mu,\lambda}(0, S_0, x, 0)$ is the value of expected utility. In contrast to IE in Section 4.1, we compare here one investor (rather than two) in two situations with or without costs irrespective of his opinion about the drift. The transaction-cost equivalent $TE(\lambda)$ is the maximum price an investor is will-

ing to pay to avoid transaction costs. The CARA utility function (5) implies that the measure is independent of the monetary endowment x . As V satisfies (9), one has

$$\text{TE}(\lambda) = \frac{1}{\alpha} \exp(-rT) \log(E_{\mu} H_{\mu, \lambda} / E_{\mu} H_{\mu, \lambda=0}).$$

We express $\text{TE}(\lambda)$ in a consistent way with μ as one of the subscripts in H without specifying an investor. In fact, only for an informed investor, does the value function depend on $\mu \sim \mathcal{N}(\mu_0, \gamma_0)$. For all other types, one can drop E_{μ} and the subscript μ .

Figure 6(b) shows the welfare effect of transaction costs on three investor types. The annualized transaction-cost equivalents are approximately constant for the naive investor but slowly decreasing for the informed investor and rapidly decreasing for the learning investor. For time-horizons of up to 5 years, the learning investor is the one most strongly affected because the estimate of the drift is inaccurate and can vary drastically in the short run (e.g. [Lundtofte 2008](#)).¹² This increases the learning investor's incentive to trade and leads to higher transaction costs.

At longer time horizons, the naive investor has the most to gain from the absence of transaction costs as the misspecification of the drift leads to excess trading compared to the investors who either know or have learned enough about the actual drift. For a learning investor, trading is slightly contrarian, which leads to the lowest transaction-cost equivalent. For instance, a sudden sharp drop (rise) in the stock price leads to a stock purchase (sale) from the

¹²At short time-horizons, the conditional variance of the filter, which decreases with time, is relatively large compared to those in the long run.

informed investor. A learning investor at the same time lowers (increases) the estimate of the drift and therefore tends to make a smaller trade, incurring lower transaction costs. As a result, the learning investor reduces the loss in utility by about 50% over a 10-year time-horizon compared with the naive investor. The benefit of learning mirrors the substantial utility gains found by Cvitanić et al. (2006) without considering transaction costs.

5. Conclusion

The efficient algorithm introduced in the paper allows us to solve portfolio optimization problems with state-dependent drift and long time-horizons in the presence of proportional transaction costs. We apply the method to explore scenarios in which investors (a) use past stock prices to learn about the true drift, (b) react to stock price movements as trend-followers or contrarians, or (c) are naive and ignore information revealed over time.

The numerical results show that forecasting behavior has a strong impact on trading. We quantify the value of information and the welfare effect of transaction costs. Information is most valuable to the least risk-averse investor, and transaction costs are most detrimental to naive investors. The total loss in utility from transaction costs is generally about twice as large as the direct cost incurred. Learning reduces the utility losses due to the uncertain drift and transaction costs, especially for medium and long horizons.

Appendix A Test of Convergence Order

In Figure 3 with the values of parameters of the base case, we obtained that the estimate of b in the relationship $|V - \hat{V}| = a \tau^{-b}$ is 0.4. This means

Panel A: comparative analysis of the convergence order b					
	$\lambda = 0.0$	$\lambda = 0.005$	$\lambda = 0.01$	$\lambda = 0.015$	$\lambda = 0.02$
Order	0.4571	0.4260	0.4083	0.5244	0.5920
	$r = 0.03$	$r = 0.06$	$r = 0.09$	$r = 0.12$	$r = 0.15$
Order	0.3946	0.4042	0.4604	0.5810	0.6020
	$\sigma = 0.25$	$\sigma = 0.3$	$\sigma = 0.35$	$\sigma = 0.4$	$\sigma = 0.45$
Order	0.4010	0.4158	0.4337	0.4833	0.5477
	$\gamma_0 = 0.04$	$\gamma_0 = 0.09$	$\gamma_0 = 0.16$	$\gamma_0 = 0.25$	$\gamma_0 = 0.36$
Order	0.4034	0.4419	0.4591	0.4622	0.4540
	$\mu_0 = 0.15$	$\mu_0 = 0.2$	$\mu_0 = 0.25$	$\mu_0 = 0.3$	$\mu_0 = 0.35$
Order	0.3993	0.3998	0.4293	0.4564	0.4540
	$\alpha = 0.1$	$\alpha = 0.2$	$\alpha = 0.3$	$\alpha = 0.4$	$\alpha = 0.5$
Order	0.4025	0.4014	0.4012	0.4020	0.4016
Panel B: the convergence order b for random parameter sets					
	Set 1	Set 2	Set 3	Set 4	Set 5
Order	0.4216	0.5352	0.4546	0.4463	0.5708

Table 2: This table displays the convergence order b with non-equidistant discretization and the state-dependent drift $\mu^L(t, S(t))$. Panel A shows the comparative analysis of b to six parameters. Other parameters are the same as the base case in Section 3. Panel B lists the convergence order for five random sets of values of parameters in Table 3, assuming each parameter satisfies a continuous uniform distribution within the region bounded by the 2nd and 6th column of Panel A.

	λ	r	σ	γ_0	μ_0	α
Set 1:	0.0055	0.1390	0.3711	0.2319	0.1662	0.2746
Set 2:	0.0146	0.0720	0.3369	0.3199	0.2893	0.4320
Set 3:	0.00011	0.0109	0.3477	0.0513	0.1846	0.3493
Set 4:	0.0067	0.0517	0.2635	0.1226	0.2221	0.2350
Set 5:	0.0175	0.0392	0.2873	0.1833	0.3351	0.1615

Table 3: The random sets of values of parameters.

that for the state-dependent drift, the computing time of our algorithm is increased by a factor of $2^{(1/0.4)} \approx 5.7$ in order to halve the numerical error. Here we test the convergence order b for different sets of values of parameters.

In general, the test results show that the convergence order b is at least around 0.4. In fact, it is faster for the most values of parameters tested. The speed of convergence is mainly affected by volatile no-trade regions of stock holdings due to our local search along with other improvements. Since we appropriately enlarge the no-trade regions of the successive nodes as the range of local search, we can locate the no-trade region faster if these regions do not change dramatically between adjacent nodes.

In contrast to a global search, here the size of no-trade region does not impede the search speed. Indeed wider no-trade regions usually indicate that they are less volatile which in turn decreases running time of our algorithm. Panel A of Table 2 indeed shows that the convergence order b rises gradually with a large transaction cost rate λ . Similarly, the convergence order b is highest for random Set 2 and Set 5 in Panel B where λ is large. Furthermore, our algorithm also benefits from the no-trade regions with small sizes. For

instance, a large risk-free r (a less attractive stock) or a large volatility σ (a riskier stock) implies substantially narrow and low no-trade region, which reduces computation time. In addition, the order slightly fluctuates around 0.40 to 0.46 for another three parameters since merely varying one of these values does not significantly change the volatility of the no-trade region.

References

- Björk, T., Davis, M., Landén, C., 2010. Optimal investment under partial information. *Mathematical Methods of Operations Research* 71, 371–399.
- Brennan, M. J., 1998. The role of learning in dynamic portfolio allocations. *European Finance Review* 1, 295–306.
- Clelow, L., Hodges, S., 1997. Optimal delta-hedging under transactions costs. *Journal of Economic Dynamics and Control* 21, 1353–1376.
- Cox, J. C., Ross, S. A., Rubinstein, M., 1979. Option pricing: A simplified approach. *Journal of Financial Economics* 7, 229–263.
- Cvitanić, J., Lazrak, A., Martellini, L., Zapatero, F., 2006. Dynamic portfolio choice with parameter uncertainty and the economic value of analysts' recommendations. *Review of Financial Studies* 19, 1113–1156.
- Davis, M. H. A., Panas, V. G., Zariphopoulou, T., 1993. European option pricing with transaction costs. *SIAM Journal on Control and Optimization* 31, 470–493.
- Fleming, W. H., Rishel, R. W., 1975. *Deterministic and Stochastic Optimal Control*. Springer, New York.

- Gennotte, G., Jung, A., 1994. Investment strategies under transaction costs: The finite-horizon case. *Management Science* 40, 385–404.
- Glasserman, P., 2004. *Monte Carlo Methods in Financial Engineering*. Springer, New York.
- Herzog, R., Kunisch, K., Sass, J., 2013. Primal-dual methods for the computation of trading regions under proportional transaction costs. *Mathematical Methods of Operations Research* 77, 101–130.
- Kushner, H., Dupuis, P., 1992. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer, New York.
- Lensberg, T., Schenk-Hoppé, K. R., 2013. Hedging without sweat: a genetic programming approach. *Quantitative Finance Letters* 1, 41–46.
- Luderer, B., Nollau, V., Vethers, K., 2010. *Mathematical Formulas for Economists*, 4th Edition. Springer, New York.
- Lundtofte, F., 2008. Expected life-time utility and hedging demands in a partially observable economy. *European Economic Review* 52, 1072–1096.
- Monoyios, M., 2004. Option pricing with transaction costs using a Markov chain approximation. *Journal of Economic Dynamics and Control* 28, 889–913.
- Muthuraman, K., Kumar, S., 2006. Multidimensional portfolio optimization with proportional transaction costs. *Mathematical Finance* 16, 301–335.
- Øksendal, B., 2003. *Stochastic Differential Equations: An Introduction with Applications*, 6th Edition. Springer, New York.

Shefrin, H., 2008. A Behavioral Approach to Asset Pricing Theory, 2nd Edition. Academic Press, San Diego.

Wang, H., 2010. Optimal Portfolio Choice under Partial Information and Transaction Costs. PhD thesis, University of Leeds, <http://ssrn.com/abstract=1664427>.

Zakamouline, V., 2005. A unified approach to portfolio optimization with linear transaction costs. Mathematical Methods of Operations Research 62, 319–343.